

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>

Unprecedented Worlds

A large pterosaur¹ sporting a teardrop fedora impatiently waves a notepad and pen clutched in the claws at the end of their enormous wings.

“Prof. Ichthy . . .” they begin, speaking at a rapid pace, “You’ve said that knowing about the past is important because it can help us think about the future. Do you really think that? Can you back that up? And can you say something concrete? Can you go on record for *Eons’* readers, explaining how all this great progress about knowing how things were might tell us about how things will go? Would you be willing to take any bets on the future? Will you, Professor, put your money where your mouth is?”

“The thing is . . .” the pterosaur’s speaking pace slows, and she sheepishly shrugs, “I find my readers are worried about the future; they worry about it a whole lot. And I don’t just mean small-time, personal worries for their fortunes and families and such, but big-time worries about where we’ll be in ten years, or fifteen years, or a hundred years. People feel like the social world, the political world, the technological world, and even the natural world are all spinning out of control.”

The pterosaur puts pen to notepad expectantly as they finish their question: “In light of all of this anxiety, is there any advice you can give for facing that future? What has all your learning about the past given you in terms of lessons for thinking about the future?”

The professor grimaces, takes a pause.

“Well, you’re calling me out a bit here. We academics are often encouraged to frame what we’re doing in as earth-shattering a way as possible. And this can sometimes lead to us overdrawing

1 A member of the *Dearc* genus, to be taxonomically exact if metaphysically dubious.

the bow on topics that exceed our expertise. I'm afraid I am no expert on the future. Advising your readers in specific, concrete ways is somewhat above my pay grade. But I'll try nonetheless to corral some thoughts. I think I have three things to say."

Alert readers will have noticed that my definitions of "past world" and "denizen" are only restricted to the past by conceptual stipulation. While past worlds are, by definition, subsets of the actual past, nothing is stopping us from considering denizens occupying subsets of the current world or denizens occupying subsets of the future. This is a feature, not a bug, as similar epistemic challenges arise for "lost" denizens in the present and in the future as for those in the past. For instance, a major challenge in understanding denizens of the contemporary deep sea comes from that environment's oddities—a lack of sunlight, extremely high pressures—which lead many of the properties of those denizens to be "lost" relative to the environments we're used to or can easily examine. When biologists are trying to understand denizens of the deep, such as the colossal squid, they adopt strategies that rely on the examination of proxies, such as smaller and less pressure-adapted squid (Levy and Currie 2015). I think it is reasonable to consider such practices as examples of grafting toward understanding a denizen of the "lost world" of today's deep oceans.

What about the future? What can we say about future worlds, particularly *unprecedented* worlds, in light of my discussion of past lost worlds?

Concerns about rapid, global, previously unseen changes occurring now or in the nearish future, and these escalating, have become so prevalent that it is difficult to talk about them without relying on a dystopian stock of near-clichés involving tipping points and catastrophic collapse. The increasing pace of technological and social change, coupled with our increasing impact on the environment around us, makes it increasingly possible, even plausible, that we are calling forth an unprecedented future. That is, we can perhaps expect a future where previous certainties can no longer be taken for granted. A future that obeys different rules from the present. Further, even putting aside an anthropogenic perspective, if something like chapter 8's arguments hold water, then there's no particular reason to think the often-shifting, peculiar past shouldn't continue to surprise us into the future.

Let's define an *unprecedented future*, then, symmetrically with lost worlds.

As a past world is composed of past denizens that do not exist in the present, an unprecedented future is composed of future denizens that do not exist in the present.²

This conceptual symmetry suggests an epistemic symmetry between lost and unprecedented worlds. Worlds are lost when their denizens are—that is, when some denizen, under some characterization, existed then but does not exist now. So, lost and unprecedented worlds only differ in their temporal orientation. This suggests that, perhaps, the strategies historical scientists adopt in trying to uncover lost worlds might apply to unprecedented futures. In this chapter, I'll suggest that this epistemic symmetry only partially holds.

Metaphysically speaking, you might be committed to a determined, fixed future, and thus think that—in principle at least—sufficient knowledge from the present could track the future; you might be committed to an open, probabilistic future, and thus think that—again in principle—there is simply no way of knowing some future outcomes. In keeping with the methodology of the last chapter, I don't think the analysis of scientific strategies I've provided can ground commitments either way on this score. Further, I'm not convinced that these metaphysical views generate such epistemic conclusions, at least at the scales that concern us here. I'll briefly return to metaphysical asymmetries at the chapter's close, but, for my purposes, I'm primarily interested in analysis not only at an abstract epistemological level, but in terms of the strategies, methods, and practices available to us for uncovering the future. As such, the partial asymmetry I'll try to establish concerns epistemic practice, not metaphysical questions about determinism or otherwise.

I'll start by discussing various ways in which the knowledge we've generated about the past might be helpful for thinking about the future. Don't expect any big surprises, although my emphasis on the modal nature of historical scientific knowledge perhaps lets me push the boat out for the importance of past knowledge a little further than others have. I'll then ask whether my answers concerning knowledge of loss carry over to unprecedented futures. Do the epistemic strategies that enable historical scientists to uncover lost worlds apply to unprecedented worlds as well? I'll provide an example of how careful characterization can open research into improbable but catastrophic

2 As with lost worlds, unprecedented futures are profligate, due to a lack of restriction on what kinds of entities, processes, and dynamics can be "denizens." I don't see this as a problem. As with the past, we'll be interested in *significantly* unprecedented futures, where a lack of denizens makes epistemic trouble for our knowledge of the future.

global risks in ways analogous to the approaches of historical science. So, I'll suggest that very similar strategies apply to unprecedented worlds as to lost worlds: both are amenable to strategic perspectivalism.

However, I'll then turn to an epistemic asymmetry between investigations of lost and unprecedented worlds. I've argued that the success of the science of lost worlds doesn't just turn on clever characterization, but on *re-contextualization*, that is, the results of investigations being situated within an actual history as revealed by, for instance, the fossil record. Further, I've argued that traces play a central role in the productively iterative investigations characteristic of the science of lost worlds. It is unclear whether there are "records" of the future. Without such "records," the power of iteration and recontextualization is *prima facie* undermined. So—perhaps surprisingly—at least under some circumstances and about some questions, we should be more pessimistic about our knowledge of the unprecedented future than the lost past.

I'm on record as not thinking there's anything that in principle distinguishes historical from non-historical research. The epistemic differences within those categories are greater than the differences between them: there's nothing epistemically distinctive about historical research *per se* (Currie 2019a). My claim about our knowledge of unprecedented futures requires a slight revision to this way of thinking: I think access to records does mean that a scientist of the past differs from a scientist of the future in a particular way. Specifically, the existence of records allows scientists of the past to map out the past's surprising peculiarities and to explore the possibilities closely aligned with those. These are not available, or at least not to the same extent, for the future, and—unless other approaches and epistemic tricks are discovered—something like traces are required for rich knowledge of unprecedented futures. As we'll see, I consider this conclusion revisable, based as it is on our current methods and approaches to uncovering the future.

Before diving in, it is worth noting a rather obvious further disanalogy between our knowledge of the past and our knowledge of the future. We can change the future, but we cannot change the past (except in rather meta-physically inconsequential ways: Currie and Swaim 2021). I'll incorporate this further disanalogy into the discussion as we go.

1. Past-Derived Knowledge and the Future

“The first, and I think quite obvious thing to point out is that the past can be a guide to the future. And not just because the past provides us relevant data or analogues, but because the past gives us knowledge of what is possible.”

This chapter primarily concerns whether the knowledge-generating strategies applicable to lost worlds can be applied to unprecedented worlds, so this is not the place to develop a particularly long-winded account of how the past might epistemically inform us about the future, but it is nonetheless important to sketch these ways. I’ll draw on a geological example to provide some illustrative, slightly caricatured examples.

In 1982, Christopher Scotese (at the time, a graduate student at the University of Chicago) speculated about the future formation of a supercontinent he called *Pangea Proxima* (Scotese 2018).³ The name indicates the new supercontinent’s similarity to Pangea, the supercontinent that broke apart at the close of the Jurassic period. This new Pangea would be formed by the closing of the Atlantic and Indian oceans, thus unifying the Old and New Worlds (that is, the Americas along with Europe, Africa, and Asia). This is postulated to occur about 250 million years from now (Romano and Cifelli 2015).

Scotese’s speculation about continental arrangement 250 million years in the future was based on the *supercontinent cycle*, the idea that there is a regular cycle of supercontinental assembly and break-up with something like a 500-million-year turnover (Nance, Murphy, and Santosh 2014; Mitchell et al. 2021). The last few billion years have seen three supercontinents, from Columbia to Rodinia to the most recent, Pangea. Continental movement is largely governed by subduction processes, where one plate slides under another to be consumed by the mantle. When the mechanisms of plate tectonics began is disputed, with estimates ranging from when indisputably subducted rocks appear during the Neoproterozoic just under a billion years ago to the Mesoarchean around 3 billion years ago (Stern 2018; Palin and Santosh 2021). So far as I can tell, earlier dates are mostly favoured (although there remains plenty of uncertainty, Harrison 2024). Prior to subduction taking off, it is thought that there were patterns of continental formation for “supercratons,” but as these plates were segregated, the pattern seen in supercontinental

3 Originally *Pangea Ultima*, but he changed the name later.

assembly wasn't detected. Scotese more or less took current continental positions and extrapolated forward based on the supercontinent cycle.

With our case on the table, we can ask: how might what we know about the past, especially lost worlds, help us understand the future? Items in the following list overlap in various ways, but I think it useful to articulate them separately.

First, we can do so by identifying regularities and establishing their modal profile by testing them against a wide range of conditions. The supercontinent cycle is a long-term pattern, but is nonetheless one that has a history—it began at some point—and emerged due to a range of fragile conditions, such as the formation of a melted mantle and subduction processes kicking off. Our understanding of this regularity and how it might play out in the future depends on our studying long temporal scales. Understanding a process' modal profile by testing it against many conditions enables us to understand how that process might work in the future (for discussion of the philosophical point, see Jeffares 2008; Currie 2018a, chapter 7; Page 2021).

Second, this can be done by tracing local dynamics in a historical individual that may be projected into the future. In addition to general regularities and mechanics, historical individuals sometimes themselves unfold in predictable patterns. Scotese took the current position of the continents as a starting point, say, North America's current position and trajectory, and extrapolated into the future. Insofar as historical individuals follow peculiar trajectories, knowing their past and present patterns and projecting forward is a route into the future.

Third, this can be done by providing partial analogues to particular possible future scenarios or conditions. Scotese's name for the postulated continent hints at the importance of Pangea as a model for imagining the kinds of conditions the formation of *Pangea Proxima* would herald. For instance, more exposed land might increase surface weathering, thus storing more carbon in geological processes, which seems to have occurred for Pangea. As such, Pangea acts as an analogue for the hypothetical future continent. It may be, then, that at least some features of the Jurassic would return if we returned to Jurassic continental conditions (for discussion of the philosophical point, see Wylie 1985; Currie 2018a, chapter 8; Wilson 2023; Watkins 2024c).

Fourth, both the second and third strategies—extrapolating local trends into the future and identifying partial analogues—can aid in discovering the future by suggesting new things to look for, new ways of conceptualizing the entities populating the present, and how they might lead us into the future.

For instance, understanding that the continents shift on tectonic plates underwrites characterizing continents as (if you want) “tectonic individuals.” Another set of examples are studies of mass extinctions—uncovering these in the deep past opens the door to asking whether we are currently undergoing one, for instance.

Fifth, providing information for calibrating measurements and models can help us understand the future. We might be interested in rates of continental dispersal at various scales. For instance, Condie et al. (2015) examine long-term trends across the Earth’s history to argue that rates of supercontinental assembly and dispersal diverge from the velocity of typical plate movement. To do this, they needed to calibrate measurements of subduction and then compare them to large-scale trends. The development and calibration of such measurements require a long-term historical perspective, often involving identifying and developing proxies for the intended measurement (for discussion of the philosophical point, see Bokulich 2020; Bocchi et al. 2022; Wilson and Boudinot 2022; Watkins 2024b).

Sixth, providing a range of conditions relevant to global or local situations can help us understand the future. By comparing various global supercontinental states, the transitions between them, and various rates of dispersal, arrangement, and an understanding of the mechanisms driving them, we build a rich set of resources for imagining future possible Earths and comparing our system to others. Earth is the only planet we know of with active plate tectonics (although it has been postulated for Europa; see Gaidos and Nimmo 2000), and this potentially has major implications for the kinds of peculiar dynamics on our planet—particularly for making sense of the particular stabilities and features seen between the lithosphere, atmosphere, and biosphere that are (so far as we know) unique to our planet. Further, these rich resources can be used to imagine our planet in various states, actual or otherwise.

Seventh, studies of the past can provide a guide and inspiration for imagining possible futures. In chapter 7, I argued that scientific imagination is constrained, enabled, and enacted by the artifacts scientists construct. These artifacts range from particular specimens to data collection, tabulation, and analysis to the models and other devices constructed to represent and examine data, hypotheses, and specimens. There is a wide variety of models and simulations constructed and being constructed to understand the continental future of our planet (e.g., Yoshida 2016; Broussolle 2022), all of which draw on information from past continental positions, as well as tectonic mechanisms informed by historical data. More ephemerally, understanding both the

range of ways the Earth's lithosphere has been, and objective modal knowledge concerning the ways it could have been, underwrites our capacity to imagine how it could be.

Overall, I take this list to demonstrate that past-derived knowledge, both about actual past worlds and possible worlds, has critical and kaleidoscopic upshots for thinking about the future. If you're called upon to defend the epistemic importance of understanding the past for knowing the future, you have bountiful things to say. Some of this knowledge, but not all, concerns significantly lost worlds, so let's turn to ask whether the strategies used to uncover lost worlds might apply to unprecedented worlds.

2. Strategic Perspectivalism and Unprecedented Worlds

"The second thing to say is that many of the ways that we scientists try to find out about the past, the strategies we adopt, if you want, are also applicable to thinking about the future. Many of the ways I've figured out surprising things about Broadfoot and their kin might be turned to thinking about what the future is like."

Throughout this book, I've argued that historical scientists, through devious use of strategic perspectivalism, gain rich knowledge of lost worlds. Can the same strategy be applied to unprecedented worlds? In this section, I'll draw on some examples from existential and catastrophic risk to say:

Yes.

However, in the next section, I'll pivot to a:

Yes, but . . .

Let's begin with some common-sense ideas about *prediction*. A prediction takes a set of information, some concerning initial conditions, and some concerning the laws, regularities, and rules governing how those initial conditions might unfurl. These are brought together to show what state of affairs or outcomes we should expect to follow. In simple contexts, the combination of regularities and initial conditions guarantees the predicted outcome, or near enough: epistemic possibility and actuality converge. In trickier contexts, the relationship between prediction, regularities, and initial conditions becomes

more complex: ranging from being less guaranteed, that is, providing a set of possible outcomes with associated probabilities, to difficult cases involving ribbons of possible outcomes with uncertain relationships, possibilities, and probabilities. In complex cases, even knowing which variables might be relevant becomes extremely hard to determine.

Our capacity to intervene is a double-edged sword in these contexts. Controlling initial conditions and intervening as events unfurl can enable us to follow desired paths to bring about the outcomes we desire—as is paradigmatically the case in experimental contexts. But this relies on us both understanding the effects of our actions and having some handle on what our future actions will be like. Often, the things we care about in the future cannot or should not be controlled like an experiment. And ongoing intervention often requires ongoing coordination, political will, and economic and resource capacity. These can be fragile indeed.⁴

All of this, then, requires deep modal knowledge not only of the relationship between initial conditions and the various regularities governing them, but also of our own agency in those conditions and regularities. While (unless our technological prowess becomes impressive indeed) it is unlikely we will have the ability to make much difference to the tectonic fates of continents, there are plenty of global processes that we have plenty of capacity to interfere with.

One area where consideration of unprecedented worlds is prominent concerns *existential* and *catastrophic risk*. These risks often come in familiar lists, ranging from events we have limited control over, such as asteroid impacts, megavolcanoes, and so on, to events that we can take pretty direct blame for, such as nuclear catastrophes, malicious (or indifferent) artificial superintelligence, runaway synthetic biology, the downstream effects of extreme climate change, and the degradation of the planet's biosphere and atmosphere more generally (Bostrom 2002; Bostrom and Cirkovic 2011; Rhodes 2024). Existential risk has been characterized in a variety of ways (Torres 2023): from relatively straightforward human extinction (e.g., Tonn and Stiefel 2013), to including civilizational collapse (Beard, Rowe, and Fox 2020), to the more esoteric, such as closing off the possibility of our descendants' flourishing in an intergalactic post-human utopia.

⁴ This, I take it, is what lies behind the difference between a “prediction” and a “projection.”

The events and processes associated with existential and catastrophic risk are often considered to be low probability but high impact (e.g., Bussiere and Fratzscher 2008), although these days the “low probability” assumption might be questionable for at least some of the risks. The chances that the effects of climate change are so bad that they fundamentally destroy our species’ capacity to survive on this planet, or that a megavolcanic eruption will occur over the next few centuries, are low. But the occurrence of such events would be so bad that we should nonetheless put significant effort into understanding, avoiding, or mitigating them.

I tend to think “existential risk” is particularly interesting because the risk pertains to the *existence of something*, as opposed to, say, risk as understood in terms of financial costs, opportunity costs, or other measures (Currie, under review). I take a personal existential risk by riding a bike rather than driving a car; languages or even whole cultures often face existential risks; our actions have led many species to face existential risk; our putting various features of the biosphere and atmosphere under existential risk puts our own species at existential risk. Despite its wide applicability, the concept is almost always applied to thinking about risks to the extinction of our own species. This is partly because existential risk is often considered in terms of the risk to our species failing to achieve its potential and the loss of future generations.

There are a fair few avenues for criticizing the perspectives often adopted by existential and catastrophic risk researchers (see, for instance, Cremer and Kemp 2021; Schuster and Woods 2022; Sundaram, Maas, and Beard 2024; Corsico 2024; McLaughlin 2025; Currie, under review). First, there are many outcomes from processes like climate change, the global adoption of large-scale computing, and environmental degradation that fall short of global catastrophe (and certainly of species-level extinction), that nonetheless involve an unacceptable amount of suffering, death, and other significant badness. An overemphasis on species-level extinction, especially when backed up by a simplistic moral calculus, can lead to recommendations ranging from mildly embarrassing to horrific. Second, the vast majority of that much more likely suffering will be the burden of more vulnerable people. To some extent, worrying about existential risk at the level of our species is the domain of the privileged, and that shows. Third, the philosophical underpinnings of existential risk work often draw on fairly fringe ideas from the perspective of academic philosophy, what Cremer and Kemp characterize as “techno-utopian” (Cremer and Kemp 2021). Fourth, the focus on catastrophes as caused by calamitous events has at least two deleterious and, I think, related upshots:

a focus on the long term rather than the medium term, and a tendency to conceive of these impacts as one-off “bolts from the blue.” Let’s focus on the latter point.

Much work on existential risk in the first fifteen years of this century focused on particular, hard-to-predict events and their avoidance. This is most stark in the discussion of the possibility of runaway effects of the development of artificial intelligence. Generally, we might think of these as existential risks caused by a single, precipitating event. One advantage of this perspective is the capacity to carefully uncover various causal features of those events (I make a point like this regarding historical science in Currie 2019e). There is, however, a major disadvantage, as Liu et al. point out: “A high profile ‘one-hit-KO’ existential risk such as global nuclear war or a pandemic may constitute only one avenue towards that ‘adverse outcome,’ and concentrating predominantly upon (ways to intervene in) its origin and direct pathway, risks overshadowing other potential paths or disaster interaction effects that functionally converge towards that same disastrous outcome” (Liu et al. 2018, 7).

An existential calamity might not come about due to a particular powerful trigger but due to a slow decline or a perfect storm (Kuhlemann 2018). To take an infamous palaeontological example, the collision occurring around the K-Pg event may have been a “one-hit” event, which more or less alone triggered the subsequent mass extinction, or it could have acted as the “last straw” to an already weakened biota, or it could have been the impact’s interaction with other events, such as continental break-up and a period of increased volcanism, that together led to the extinction. There are, then, multiple kinds of routes to catastrophic outcomes. These slower, more complex, and less sexy catastrophes raise further epistemic issues (Baum 2020). Thus, in addition to narrow perspectives on particular risks, we need perspectives that can productively think about risks in more general ways that are both empirically and politically informative.

Many existential risks to our own species are unprecedented, at least in terms of scale and target. Some are due to unprecedented features of our own species: the extent to which we (intentionally or otherwise) engineer global systems, global interconnectedness, the speed of information exchange, and so on. So, the targets of these investigations are highly peculiar and, to varying levels of characterization, unprecedented—they concern denizens

not currently present in our world in an epistemically deleterious way.⁵ As such, trying to understand them has at least some analogies to how we might understand lost worlds.⁶ According to strategic perspectivalism, when faced with a challenge from a lost, or putatively lost, denizen, scientists will adopt—or construct—perspectival tools that productively characterize their target. We see something similar, I think, in some studies of existential and catastrophic risk. An example is Shahar Avin and colleagues’ classification of global catastrophic risks.

Similarly to Liu et al., Avin et al. characterize earlier work as “identifying individual GCR [global catastrophic risk] scenarios, and in compiling lists of the scenarios of greatest concern,” but they don’t yet see “a comprehensive, interdisciplinary view of severe global catastrophic risks” (Avin et al. 2018, 20). This is the aim of their classification. Instead of focusing on particular sources of threat—big rocks hitting the earth, rampant artificial intelligence, etc.—they instead draw our attention to the kinds of systems that could break down and how we might fail to stop these.

Avin et al.’s framework has three dimensions.

First, they provide an analysis of *critical systems* that operate within particular safety boundaries. Here, inspiration is drawn from the literature on “planetary boundaries” (Rockstrom et al. 2009; Steffen et al. 2015). The thought behind that literature is that human life depends on various systems remaining sufficiently stable: on a global level, the biosphere and atmosphere are dependent on various processes, such as the oxygen cycle we discussed at the beginning of chapter 5. This perspective takes human life as dependent on various, often-interdependent, fragile regularities. It emphasizes the peculiarity of global systems. Avin et al. extend the notion of safe planetary boundaries to include various intertwined hierarchical levels. These include, for example, ecological levels such as trophic webs relying on primary producers; social levels like food production and dissemination; and individual levels including cardiovascular, immunological, and other systems. All of these processes might be disrupted by potential threats.

5 Lalitha Sundaram and I (2024) have argued that a science of existential and catastrophic risk shouldn’t trade in quantified probabilities and prediction narrowly understood, but instead walk in the ambiguities of possible scenarios, the relationships between these, heuristics guiding actions, and warning signs heralding bad outcomes. Even this less ambitious picture of our knowledge of unprecedented worlds requires modal understanding.

6 In Currie (2019e), I compare the epistemic situation of palaeontology to existential risk science.

Second, Avin et al. point out, to reach global scales, system failures require mechanisms of spread. Some large-scale events, such as significant meteor impacts or megavolcanoes, have fairly immediate spread mechanisms. Others, such as pandemics, might rely on human-built dissemination structures such as global travel or cyberattacks through information networks. Avin et al. identify three sets of such mechanisms: “natural” processes on a global scale, such as air or water-based dispersal, anthropogenetic networks such as communication systems and via various forms of replication, such as biological, digital, or cultural.

Finally, third, Avin et al. include possible failures of prevention and mitigation. The approach here is to identify “potentially fragile areas in the global risk prevention and mitigation system” (Avin et al. 2018, 23). As in their analysis of critical systems, this takes the form of an intertwined hierarchy. Individual failures, such as failures in risk perception or cognitive bias, stand alongside interpersonal failures in, say, trust, reputation, or conflict resolution. On higher levels, institutional failures include inappropriate incentive structures, adaptability, and decision-making processes, while extra-institutional factors might include governance, diversity, and enforcement.

A major benefit of the framework is that it allows some threats to be bundled together: “Scenarios with significantly different primary causes could manifest their GCR [Global Catastrophic Risk] potential through a similar mechanism” (Avin et al. 2018, 24). This both simplifies and helps guide research and resource efforts toward critical systems that might be the target of multiple disruptions. Avin et al. also highlight how such a framing makes space for interdisciplinary collaboration and the identification of relevant expertise for mitigating and tackling risks. Their final application is of particular import to us here: “The system is a tool to highlight highly uncertain or neglected corners of the GCR possibility space, and guide research efforts towards these corners, in the hope of discovering unknown unknowns. The combinatorial nature of the classification systems provides a natural way of progressing from well-known systems and mechanisms to a vast and as-yet largely unexplored space of possible GCRs” (Avin et al. 2018, 25).

The framework captures a complex modal space, intended to convey the various ways in which critical systems might fail, and those failures might spread, and how we might fail to stop them. Some spaces within that framework are familiar and well-studied; others are not. If it works well, the framework would help identify these unknowns and reveal some unknown unknowns, guiding research into those areas.

We can think of Avin et al.'s classification as a perspectival tool: it *characterizes* catastrophic risk scenarios in a way intended to unify various discipline-bound theories and data, such that it helps identify and focus research and policy efforts. In its exploratory guise, we can understand it in terms of the artifactualist account of imagination defended in chapter 7. The framework is a mode, consisting of three complex dimensions, that may be instantiated across various media. I've used a casual linguistic medium here, but Avin et al. also present it using various diagrams and tables, which enable comparisons as well as tracing possible routes from system failure to spread mechanism to mitigation failure, and considering other possible GCR sources. It thus enables exploration of the space of GCRs, scaffolding our otherwise unsystematic or isolated work on particular threats. It is, in other words, a tool for the imagination.

There are further similarities between Avin et al.'s framework and many of the artifacts we've seen in the historical sciences. Mitchell et al. make use of a model of ecological space differentiation, which brings out the complex relationships among population size, niche stability, and other factors to map out possible conditions under which we would expect neutrally structured ecosystems. Both Mitchell et al. and Avin et al.'s possibility spaces should not be understood as targeting mere epistemic possibility. We're not just interested in what will happen, but what might. Further information might shape or complexify the model, but it won't shrink toward the actual. Further, Avin et al. explicitly emphasize the potentially iterative relationship between their framework and the various studies that could be carried out, for instance, that it could "both inform, and be informed by, different 'foresight' tools" (Avin et al. 2018, 25). These "foresight" tools are themselves ways of imagining and monitoring future possibilities, typically involving various expert opinion-gathering and amalgamating activities, as well as various kinds of simulations (Conway 2006; Kemp 2024).

So, in tackling the unprecedented future, we see researchers strategically adopting perspectives that enable fruitful, iterative studies and constrain and enable epistemically relevant imagining.

You might worry this conclusion is unsurprising, as strategic perspectivism is so general a strategy that it is no wonder that it applies across cases. If we define the strategy in highly abstract terms, specifically, as the practice of characterizing a target and constructing a tool to understand it, then this is right. However, I've understood strategic perspectivism more narrowly, construed as a way of uncovering lost worlds. As such, I've emphasized the

role of artificing, grafting, and anchoring, and the iterative studies they generate, toward understanding denizens unfamiliar to our present. It is in this more narrow sense that Avin et al. exemplify many of the features of strategic perspectivalism. By identifying critical systems, they are able to integrate—or graft—current systems with future possible states. Mechanisms of spread can be understood via various artificing strategies, particularly simulations. The framework can work as a structure supporting various iterative studies.

In developing their perspectival tool, then, Avin et al.'s approach makes use of many of the strategies discussed in the first half of the book. Artificing, especially of global-level models, but also of various scenarios used in policy planning, enables us to create partial simulacra of unprecedented worlds, and, as we've seen, extending trends from previous records into the future enables us to graft unprecedented worlds, and various anchors can be found throughout history. In short, faced with potentially unprecedented denizens, we seek characterizations that link them in various ways with what we do know.

Our potential knowledge of unprecedented worlds, then, is great. And I think the analysis I've provided in this book can make some sense of that knowledge. So that's my "Yes." However, it's time for the "But."

3. Records of the Future?

"The third, and I'm afraid, much less happy point, involves thinking about the crucial importance of things like this," Ichthy gestures to the skull in the glass. "I've said a few times that I couldn't have imagined Broadfoot without this specimen, that nothing about today other than this specimen speaks to something like this animal existing. It is fossils, and records built of objects like them, that help us find surprises in the past. So, if the future is supposed to be as surprising as the past, how are we supposed to detect those surprises? That is, what could play the role of a fossil for our investigation of the future? I don't think there is such a thing. And if there isn't," Ichthy sighs, "then I suppose I cannot provide as optimistic a story about our ability to understand the future as we can the past. Unless. Unless there could be something like a record of the future . . .?"

I've argued that many of the tools of strategic perspectivalism are available for scientists of the unprecedented, so many of the lessons of the first half of this book carry over to future-oriented projects. However, I have a hunch—one

I can't quite kick—that there is nonetheless an asymmetry concerning our knowledge of the past and the future. This difference is grounded in the importance of records. There are different ways of establishing the asymmetry, and I'll sketch a few in this section. Even if the negative point doesn't hold water—and for our sake, I hope it doesn't—the discussion nonetheless serves to capture what is so important about traces for understanding lost worlds.

Recall the (hopefully) commonsensical account of prediction I introduced earlier: predictions work by taking some initial conditions and, in combination with regularities governing what unfolds given those conditions, estimating some later set of conditions. Trace-based reasoning, one might think, is simply a temporally-reversed example of this. One takes some current conditions, the existence of a fossil, say, and based on regularities pertaining to the formation of that kind of thing, estimates some set of earlier conditions. A “trace” is some current observation that, in combination with a justified theory capturing the dependency between the observation and its causal or informational descendants, licenses an inference to the past (Currie 2018a, chapter 3). So, a “future trace” would simply be a current observation that, in combination with the relevant theory, licenses an inference into the future. At a certain level of description, that's correct, and it suggests a symmetrical treatment: just as fossils and other traces allow us to build records of the past, we can treat the current state of the world as “traces” or “records” of the future. Let's call that principle *symmetry*: epistemically speaking, knowledge of both the past and the future is gained by making use of contemporary observations, and extrapolating either forward or backward from those observations using regularities about how those phenomena transform over time.

I don't think this symmetry works *at the level of epistemic practice*. First, traces can tell us *what is surprising* about the *actual* past, while extrapolating from present observations can, at best, provide some *possible* surprises in the future. Second, traces enable *recontextualization* with an actual past, while extrapolation from current observations at best enables probabilistic integration of possible futures. Third, traces are perspectival tools that enable iterative, constrained imaginings tied to an actual past; observations of current phenomena don't seem to enable that kind of investigation.

Let's start by considering a few features of trace-based reasoning.

In this book's first chapter, I suggested that an ironic downside of over-emphasizing traces in the philosophy of the historical sciences, that is, focusing on information decay at the expense of loss, is that the value of

traces themselves is not sufficiently captured. Why do I think that? Consider (again!) some of the cases of loss throughout the book: rangeomorph spatial neutrality, morphology, and feeding strategies; the unique Avalonian conditions—bacterial mats, heterogeneous oxygenation—that led to ivesheadiomorph fossils; the sequence of changes that took us from the Avalonian to the more familiar Phanerozoic. These are all circumstances that depart from the usual running of things in our current world. They are surprising in a particular kind of way. If we took the world around us as representative of the past, we wouldn't expect to find those denizens. Plausibly, our knowledge of these surprising denizens is grounded in traces. As Prof. Ichthy emphasized throughout their lecture, were it not for the fossil of Broadfoot, they wouldn't have imagined that creature. Without fossils, I don't think we would have imagined many dinosaurid forms, and certainly wouldn't have thought those imaginings to track an actual past.

So, traces matter because they tell us what is surprising about the past. This is what makes them so crucial for driving the otherwise impoverished imaginations of scientists.

I've further emphasized the importance of *recontextualization*. Strategic perspectivalists often characterize their targets in ways amenable to artificing, anchoring, and various forms of data analysis. This typically involves isolating and abstracting various properties from traces; for instance, Mitchell et al. isolated the spatial and taxa-specific properties of Avalonian rangeomorphs. While this isolation is powerful, it's also necessary to resituate those properties within the past world. Consideration of features of rangeomorph environments, such as resource availability, and further information about their ecological situation, such as their small populations, was required to align the various anchors and models Mitchell et al. used to the actual past world, and thus generate their explanation of ecological neutrality in the Avalonian. So, recontextualization brings the fruits of various more-isolated perspectives into contact with an actual history and provides further opportunities for integration and iteration.

So, traces matter because they provide a link, often multiple links, to the actual past, which enables recontextualization.

Compare now the sources of uncertainty we might have of the past and of the future. In trace-based reasoning, we must be sure that we've characterized the trace correctly (say, that it is in fact a fossil) and that the various theories required to interpret it are well-founded, thus allowing us to infer that the fossil formed at a particular time, is of a particular taxon, and so on. When

these are relatively secure, we can say with relative confidence that *this taxon* existed in *this location* at *this time*—we can treat it as a fixed point.

Switching to the future, from some current observation, and our understanding of how things unfold, we are often able to lay out possible outcomes. Sometimes this is highly constrained: if a trajectory continues, the outcome is certain—a species will go extinct so long as its population continues to decline. Some constraints on the future are fairly powerful: for instance, wooden structures won't last as long as metallic structures. Our knowledge of the future is based on (1) the current states, (2) extrapolating from historical trends, and (3) constraints from the regularities we've identified. In trying to understand future continent formations, we rely on the mechanisms of tectonics we've managed to generate (and these not breaking down), past patterns in tectonics and current continent positions (the initial conditions), and what we're able to glean from various artificing and anchoring procedures.

On the face of it, for prediction, there doesn't seem to be epistemic objects that concretely play the specific roles I've emphasized for traces. That is, how does analysis of the present reveal surprising—i.e., unprecedented—things about the future? And further, how does analysis of the present provide iterative links with an actual future?

Our knowledge of the future is strongest when targets unfurl in regular, patterned, and convergent ways, and when the records we have of the past can be extended into the future. For plate tectonics, we have an enormous amount of historical data, a good understanding of the mechanisms, and so on. This allows us to model the future using “conservative extensions” of existing theoretical and empirical knowledge. But if the world is often peculiar, those regularities will be highly fragile, and what new regularities might arise will be extremely hard to predict. Many of the various scenarios of existential and catastrophic risk are highly peculiar. Food security, for instance, depends on production methods, the economics, legal structures, and material conditions of trade, local and global weather patterns, soil conditions, and so on. All of these are deeply interdependent and in constant flux. Predicting the future of food security to any degree of accuracy requires getting a handle on all of these highly peculiar systems. In contrast, our knowledge of the past—deeply fallible though it might be—sometimes includes records that can ground knowledge of what we are often unable to detect in the future, that is, the surprising and the peculiar. Even in the case of tectonics, we are left less with a concrete hypothesis of what will happen in 250 million years,

and more with a range of options, often highly conditional on what comes before—and that's without worrying about unprecedented surprises.

Further, the artifacts that ground and direct historical science are not exhausted by extrapolating from existing trends, the models of artificing, or the various analogues of anchoring. They also include intimate, productively iterative analysis of specimens and traces. As I said earlier, records are, from one perspective, current objects construed as data relevant to past inferences. However, in chapter 7, I argued that they are also objects that we have constructed to be revelatory of times and possibilities in the past. Mitchell et al.'s exploration of neutral evolution couldn't have occurred without careful data gathering and recording, and then using this to reconstruct a past world grounded in the Mistaken Point fossils. These frozen(ish) time-slices of lost denizens partially ensured the relevance of their exploration of the possible to an actual past. For future records, we might examine contemporary objects and imagine them as revelatory of the future. Some metals degrade more slowly than others, and as such, some of our metallic constructions are likely to survive into the future, and we can imagine how those might constrain and affect the peculiar futures we might move into. But these are typically unsurprising futures. While current objects can be iteratively examined to understand how the future might be, the specimens that typically enliven historical research stand in often obstinate testament to various past events and, as such, strongly link investigation to the actual past.

These arguments rely on an approximate distinction between two kinds of evidence from the present concerning other times. Sometimes, we might take the state of the world as *representative* of other times. This is roughly uniformitarian: take the denizens of today and anticipate them as continuing to exist in future times. Sometimes, we might take particular components of the current world as processually connected with components of the world at other times. This is what we see in trace-based reasoning and grafting more generally. Considering unprecedented and lost worlds, the former approach cannot be taken as a steady guide. The latter is, or so it seems, more available for the past than for the future. If that is right, our knowledge of past worlds is more powerful than our knowledge of unprecedented worlds. The symmetry breaks in that sense. Our imaginations lack the material constraints required to make them relevant to the actual future or possibilities relevant to it.

We further might distinguish between two kinds of *surprise*, implied above. In scientific contexts, we might standardly think of surprise as relative to our knowledge—something is surprising if it bucks our empirical and

theoretical expectations. Triceratops, at least concerning our most well-established knowledge of them, are no longer surprising in this sense. However, there is another sense of surprise—the sense of being *unfamiliar*—in which triceratops are still surprising. This sense of surprise takes the present and compares it to other times, possibilities, or circumstances. Traces enable knowledge of things that are surprising in this latter sense. Without traces, our capacity to identify such future surprises is reduced.⁷

A single trace, then—when background theory is sufficient—nails down a part of a past world and provides a pivot point for developing modal knowledge. A pattern of traces, linked together, provides a trajectory—a story—through the actual past. These patterns provide powerful access to aspects of the past that are no longer represented in the present—that is, to lost worlds. While something like this might sometimes be available for the future, when our theoretical constraints are sufficient to extrapolate to highly constrained paths beyond the present, or when future paths converge in cases of low peculiarity. But these are significantly more tied to our current world and offer significantly fewer opportunities to explore possibility in a way grounded in, relevant to, and guided by future actuality.

So, in some limited ways, the metaphor of “records” might be applied to our knowledge of the future in some circumstances, mostly when we’re extending a set of historical patterns from one location into that location’s future, or when we adopt an extremely abstract perspective. But what concerns us often goes beyond a location or takes the location into futures where the intertwining possibilities stretch epistemic possibility beyond what we can handle; and what concerns us is often the particular, the finer-grained, rather than what can be delivered by abstract modelling of possibility or extrapolated from current patterns.

Depending on how central we take traces to be in the strategy of uncovering lost worlds, this might lead us to conclude that, in fact, strategic perspectivalism is *not* a strategy adoptable for much future-oriented work. This turns on whether we think trace-based reasoning is essential to strategic perspectivalism. I’m not sure what exactly hangs on that discussion in the current context, but I am open to being convinced of this point.

7 There’s likely a lot more to be said about “surprise” as discussed here, especially regarding its phenomenological and aesthetic senses. See Currie (2018b) for epistemic surprise, and French and Murphy (2023) for surprise in science more generally.

One objection to the asymmetry I've identified comes from *epistemic function*. You might argue that while at least sometimes for past knowledge we care about the actual past—we want something like a retrodiction—for the future all we need are projections, that is, what would happen under certain policies or interventions. Because we only want projections concerning the future, the distinction I've drawn is irrelevant. In response, I'm not convinced the difference between a retrodiction or a prediction on the one hand, and projections on the other, is so stark. A projection is a *conditional prediction*, and although this means the prediction coming true depends on that conditional occurring, it isn't obvious to me why this modal character makes such a profound epistemic difference. Even if it does, and that exploring possibilities is all we need for (some) future-oriented science, this doesn't undermine the idea that there is knowledge we can have of the past that we cannot have of the future.

In his discussion of betting on science's progress, Derek Turner distinguishes between *methods neutral* bets and *current methods* bets (Turner 2016). The former bets that some knowledge won't be gained regardless of whatever progress we might make on methods. The latter bets that, given our current technological methodological capacities, some knowledge is unlikely to be gained. This provides a few ways of construing my arguments. On a *current-methods* reading, I'm arguing that, given the current epistemic resources available to us, absent either different ways of uncovering the future or some way of developing records of the future, our past knowledge will outstrip our future knowledge. On a *methods-neutral* reading, I'm further claiming that we *won't* develop further ways of understanding the future or records of the future. I take myself to be making something like the weaker bet, and why is revealing. To see this, let's consider what might underwrite the epistemic asymmetry.

What accounting might we give of the epistemic asymmetry? One might underwrite this asymmetry via a metaphysical claim: if the past is truly closed and the future truly open, then as a matter of fact what has happened is a series of fixed points and processes and what will happen is in some sense indeterminate (Lewis 1979); if, as Carol Cleland has argued (2002), the laws of physics have temporal asymmetries embedded in them, and these laws in some regard govern everything else we might care about, then epistemic consequences might follow. I don't think those sorts of claims are appropriate here. I'm sceptical that we have good grounds to commit to a metaphysically open future, am sceptical that the derivations of physical theory must apply

so uniformly, and (especially) doubt that such temporal asymmetries make a profound difference to the knowledge of historical scientists like palaeontologists and archaeologists.

However, the *peculiarity* of the past (and the future!) is a metaphysical claim I'm happy to entertain. In a peculiar world, whatever regularities we encounter depend on complex sets of generating and maintaining processes. This can ground local asymmetries in various ways. Peculiar patterns could be *path-dependent* in an *accumulative* way. That is, the previous set of states might strongly constrain the direction of future states (this, I think, is captured by Bill Wimsatt's notion of "generative entrenchment" (2007)). In that context, access to previous states and some handle on how these constrain the future could lead to something like a trace of the future. Because a current state is necessary but insufficient for later states, with the right understanding of those constraints, we can study that current state as bearing a mark of the future. However, peculiar patterns can also be wildly disrupted or follow chaotic dynamics. Here, previous states do not constrain later states in any strong way. Under these circumstances, present states are nothing like traces of the future, as there isn't a clear inferential path between that putative trace and the future. As we saw in chapter 1, philosophers sometimes discuss trace-based reasoning in terms of whether various historical processes are information-preserving or destroying. These two kinds of peculiarity provide us with a future-oriented framing of this distinction. If peculiarity is accumulative, then processes carrying signals of the future often preserve information; if peculiarity is disruptive, then those processes will often be information-destroying.

On this kind of framing, whether peculiarity metaphysically grounds local temporal asymmetries turns on a further question concerning the kind of peculiarity we face. There's a lot more to say on this score, which I'll leave for now. For our purposes here, suffice to say, appeal to peculiarity alone won't ground a sufficiently systematic asymmetry.

Instead, we could explain an imperfect and revisable temporal asymmetry by appeal to our epistemic situation. That is, rather than pointing to the structure of the world we study, we point to features of ourselves as knowers. There are weaker and stronger versions of this thought. On weaker versions, we point to the current scientific toolkit we happen to possess. We have powerful theories geared toward understanding the processes by which traces form. But we lack powerful theories, at least often, geared toward understanding how current states might unfurl into future states—at least

surprising future states. On stronger versions, we might say that regardless of the general structure of the world, the knowledge that we generate depends in part on our temporal nature. Regardless of the metaphysics of time's arrow, our *experience* and *agency* operate in one narrow temporal direction. We are thus epistemically bound by our subjective temporal straightjacket.

I'm not sure if ultimately much sense can be made of either of these suggestions, but regardless, it strikes me that both may be overcome in principle. Weaker views require new approaches and techniques to thinking about the future. Stronger views require us to find ways to epistemically transcend the limitations of our temporal agency. This isn't as obscure as it sounds: we overcome our spatial limitations through various detection and visualization techniques that bring things too big or too small to human levels; we do the same with durations that are too long or too short (see Rheinberger's discussions in 1997, 2010, 2023).

Just as I'm not confident in our ability to get a good grip on a peculiar future generally, I also doubt we're in a good position to predict the path of peculiar science.⁸ Thus, I'm left with a *current methods* bet. Crucially, this transforms an apparently pessimistic argument into a challenge.

The pterosaur scribbles on her notepad, disappointment creeping across her face. "So," she says, "the future is just unknowable?"

Prof. Ichthy waves their flippers conciliatorily, "No, no, not at all! It is really hard to know about the future, but I have great faith in our ingenuity, if perhaps not as much in our wisdom. As I said, I'm no expert on the future, so cleverer and more informed heads than my own might be able to provide a more reassuring answer. It's worth noting that some things we can be pretty sure of, but some other things we can't." Ichthy switches from talking directly to the reporter to the audience more generally.

"I'm sometimes asked to make predictions about the future of science. What new discoveries might we make? What questions will remain unanswered? Here, I think we have good reason to be optimistic that we'll discover a lot of things. The history of science, the ingenuity of scientists, and so forth, all provide good reason for optimism. I think this is true even of my subject, the deep past. Although fossils provide only partial information, our ongoing capacity to rejig our perspectives via new sources of evidence, new finds, and new technologies is likely to continue

8 That is, science itself is contingent (Kidd 2011; Rheinberger 2023, chapter 8).




ILLUSTRATION 9.1 Robo-Ichthy

into the future. If someone were to say that we won't know much about the past, or that some particular mystery will never be resolved, they're placing a pretty risky bet." Ichthy nods emphatically, their posture indicating a debate long-ago played out from their perspective.

"So, I think our relationship with the future—and scientific ingenuity—makes it an incredibly risky bet to say that some particular question will go unanswered. However, trying to tell what exactly will be discovered—how future science will play out—is a different question altogether. Another question I'm asked is about how we scientists will do our science in the future. Against the backdrop of technological advancement, what new ways of doing science will emerge? To be a little topical: as technology advances, will scientists like me even be required? We can imagine scientific futures dominated by automation, where fancy machines whir up answers that now rely on scientific thinking, or great machines do our fieldwork rather than our eyes and flippers. I admit to finding such futures deeply unattractive, but I also think it is very difficult to make predictions about them. Partly this is due to our influence over the future—do we want a future with automated science?—but that is not the only source of difficulty. I think we know so much about the past in part

because we have records of it. So, can we develop records of the future? Or can we show that we don't need them?"

Still unsatisfied, the pterosaur responds, "Okay, well, I'm sure *Eons'* readers will want to know, what should we do? If there's such uncertainty, not just in the future generally but in science's future, what policies or paths should we follow?"

Ichthy again grimaces, "We're once again far outside of my wheelhouse. Decision-making under uncertainty, the details of policy, and certainly the ethics and justice of the whole question bring us well outside of what I could say anything sensible about. I don't really see myself as the kind of academic who speaks confidently outside of my expertise. I'm sorry to let you down."

A further silence stretches across the audience. One attendee whispers to another, "It'd be a shame to finish on such a bummer . . ."

