Evolution, 2023, **77(12)**, 2699–2702
https://doi.org/10.1093/evolut/qpad167
Advance access publication 13 October 2023
Book Review

OXFORD

Does science evolve?

Review of: *How Knowledge Grows: The Evolution of Scientific Development*, by Chris Haufe, 2022. The MIT Press, Cambridge MA. 352 pp. ISBN: 9780262544450

Cultural evolution is big business these days. A diverse range of phenomena, from the ontogeny of human psychology (Heyes, 2018), to archaeological assemblages (O'Brien & Lyman, 2002), to differentiation in kinship terminology (Keen, 2022) are analyzed using cultural evolutionary models. But, as many cultural evolutionists are quick to point out, any change over time in a cultural system *can* be modeled using an evolutionary framework. Whether or not it is worth doing so, however, depends on what we learn about the system by modeling it this way.

In "How Knowledge Grows," Chris Haufe applies an evolutionary framework to the development of scientific knowledge. And he is particularly alive to the above issue. Analogies between science and evolution have been drawn by some of the biggest names in the philosophy of science, but, claims Haufe, what we learn about science from drawing these analogies has never been clear.

The book aims to remedy this situation. Haufe sets himself a bold task: develop a Darwinian model of the production of knowledge in scientific communities that explains how that knowledge is generated. He draws on current research from a diverse range of fields—including evolutionary biology, cultural evolutionary theory, the history of science, and the philosophy of biology—to make a case that the evolutionary analogy is worth another look.

Haufe's book is well-informed, enjoyable to read, and well-argued. It provides a much-needed update to an important line of thought in 20th century philosophy of science, and will be useful for researchers working on a diverse range of issues: from traditional epistemological questions such as demarcating science from nonscience; to the application of cultural evolutionary theory in the philosophy of science in practice; to the history of science. However, I do have some reservations regarding Haufe's project. In what follows, I provide a brief overview of the text and then outline my concerns.

The book begins by making a general case for why a rethink of the evolutionary analogy is needed. Haufe characterizes the challenge facing proponents of the analogy via Peter Godfrey-Smith: "Though the analogy between science and Darwinian evolution is something that people keep coming back to, the analogy has not yielded a lot of new insights so far" (Godfrey-Smith, 2003: 166; quoted in Haufe, 2022: 11). In the aptly named "Pummelling of the Predecessors" (Haufe, 2022: 13 & §1.2), previous work by Thomas Kuhn, Karl Popper, Donald Campbell, Steven Toulmin, and David Hull is subject to this criticism. According to Haufe, while these thinkers note similarities between developments in the history of science and concepts from Darwinian evolution, they do not attempt to use the analogy to solve or even re-frame philosophical questions.

However, the Predecessors are not to be faulted for this omission; rather, they were simply victims of their time. There have been significant advances in the philosophy of biology, in the sociology of science and in evolutionary biology that mean the evolutionary analogy can meet Godfrey-Smith's challenge and "yield some insights." And, as it turns out, Godfrey-Smith is key to meeting his own challenge. Haufe's framework and the structure of his argument leans heavily on *Darwinian Populations and Natural Selection* (Godfrey-Smith, 2009); and in particular the idea that the extent to which evolution by natural selection can influence a population is a function of the population's degrees of heredity, intrinsicality, and continuity. On Godfrey-Smith's formulation *heredity* is the extent to which offspring resemble parents due to the causal role of parents; *intrinsicality* is the extent to which differences in reproductive output in a population depend on intrinsic features of the members of the population; and *continuity* is the extent to which the magnitude of trait differences tracks the magnitude of their fitness differences (see Haufe, 2022: 38). The importance of Godfrey-Smith's framework is that it generalizes beyond biology to cultural practices. Haufe's claim is that if you find "populations" of scientific practices that score highly on these metrics, and if the cultural selection at play is appropriately epistemic, then you have found instances where knowledge "grows."

The book is split into two parts. The first part (chapters 2–6) outlines the general framework, and the second part (chapters 7–9) offers an extended case study illustrating the framework. Chapters 2–4 examine aspects of the evolutionary process, framed through Godfrey-Smith's three features of paradigmatically Darwinian populations, and applies them to scientific practice. Chapter 2 looks at heredity, and how experimental and methodological techniques are reliably copied from one generation of scientists to the next. Chapter 3 examines fitness, and how intrinsic features of certain practices contribute to their prevalence within a population. Chapter 4 is concerned with the phenomenon of multiple discoveries and the analogy with convergent evolution, and how this produces continuity.

Received August 11, 2023; revisions received September 8, 2023; accepted September 18, 2023
Associate Editor: James Mallet; Handling Editor: Jason Wolf

In these chapters, Haufe is motivating the claim that increasingly specialized scientific problem-solving is analogous to increasing vertical degrees of adaptive phylogenetic refinement. But the propensity of natural selection to influence populations also produces horizontal branching and divergence. In chapter 5, we see how the evolutionary analogy—and in particular, speciation—can account for the expanding breadth of scientific knowledge. Chapter 6 takes a different line again. Here Haufe is concerned with legitimizing Kuhn's notion of scientific revolutions via analogy with mass extinction events.

The second part of the book offers an extended case study designed to illustrate aspects of the framework developed in the first part of the book. Haufe looks in detail at Stephen Jay Gould's attempts to establish evolutionary paleontology as a new sub-discipline in the evolutionary sciences. As such, this example mainly bolsters the work of chapter 5. And while the case study is both illuminating and entertaining, by my lights there is something of a missed opportunity here. My reservations regarding the book mainly fall on claims developed in the first three chapters, so I was disappointed the case study did not deal more directly with this part of Haufe's framework.

OK—onto those reservations shortly. First, however, a quick caveat is in order. The scope of Haufe's book is incredibly broad: he synthesizes historical, philosophical, sociological, biological, and anthropological research in order to motivate his case. There are always "reservations" to be had with any work of philosophy, let alone one attempting to do substantive work on a philosophical problem via the synthesis of such a diverse range of disciplines. And, in general, I agree with Haufe's approach: understanding the development of scientific knowledge will require synthesizing the various disciplines he appeals to. However, I am concerned about some aspects of this synthesis. So my reservations are pitched at this level: I agree with the general project; I disagree with how Haufe has gone about it.

Here's the tricky thing about applying evolutionary thinking to cultural change: *there are so many ways to make the analogy*. The question is not "can we map evolutionary concepts onto processes of cultural change?," but "which mapping should we use, and why?." Haufe has chosen a particular way of formulating the analogy, but there are other ways. And my worry is that there is another formulation that does a better job of explaining the *production* of scientific knowledge.

The first thing to note about Haufe's approach is that it is a "meme's eye-view" (Shennan, 2011) model of cultural evolution, where selection occurs in the cultural domain, and the units of selection and bearers of fitness are cultural attributes themselves. To illustrate, consider the process of making a tool. If, for whatever reason, individuals in a population tend to favor one way of producing a tool over others, then that technique, and tool, will proliferate. The meme's eye-view approach contrasts with agent-based models, where selection is understood as good old-fashioned natural selection, and the units of selection and bearers of fitness are culturally equipped human agents. For instance, if making a tool in a particular way increases my chance of survival and causes me to have more children, and via social learning those children inherit my toolmaking technique, then that technique and tool will proliferate.[1]

Two important questions face anyone employing a meme's eye-view model: (1) What is the cultural unit that undergoes selection? (2) Why do the selection pressures involved explain cultural change over time? For instance, consider the toolmaking case. What is the cultural unit? It could be the tool itself, or it could be the set of techniques used to make the tool. Both practices and their outputs can be understood as cultural units. Moreover, while in many cases practices and their outputs may co-vary, there is no necessary mapping between change in one and change in the other. We can make similar tools using very different techniques, and we can use similar techniques to make very different tools.[2] This means we cannot always use practices as a proxy for outputs. The practice/output distinction also changes how we understand the relationship between selection and cultural change over time. For instance, if we are interested in change over time in practices, then we might look at how sociological pressures or natural selection favor certain techniques. On the other hand, we might be interested in outputs. Here we could follow the same procedure just outlined, but treat changes in practice as a proxy for changes in output. However, as we have seen, this is risky. Alternatively, we could take outputs as our cultural unit and treat practices as the selection process driving change over time. This may not be particularly palatable when it comes to toolmaking, as we might think this practice is more analogous to developmental processes than evolutionary processes (though I do think there are ways of making the latter analogy work). But, as we shall see, I think it is a very plausible framing for Haufe's project.

So the answers to (1) and (2) matter. However, I found Haufe a bit unclear here. What is the cultural unit under investigation? On one hand, as the title of the book indicates, he is interested in the growth of scientific knowledge. On the other hand, Haufe adopts a "practice-oriented" (2022, 38) approach, where "practices" include "… such apparently varied notions as experimental techniques, concepts, metaphysical commitments about the fundamental constituents of reality, canons of inference, substantive scientific theories, and research problems…" (2022, 39). So, again, we have two cultural units at play: practices (of the scientific variety) and the outputs they produce (scientific knowledge). And again, it is not clear that we can use practices as a proxy for outputs. At least in principle (and, I suspect, in many actual cases) we could have increases in the depth and breadth of scientific practices without any associated increase in scientific knowledge.

What about the kind of cultural selection in operation? Here it is clear that Haufe is interested in cultural—and more specifically, *sociological*—selection operating on practices. Moreover, he thinks that practices can operate as a proxy for knowledge, given that the right kind of sociological selection pressures—namely, epistemic sociological selection pressures—are at play (2022, §1.3.3). However, remember that we also need to specify why the cultural selection process involved explains cultural change over time. It seems clear to me that the kind of sociological pressures Haufe outlines (2022, §3.3) would indeed select for successful scientific practices; that is, practices that produce knowledge. But this explains the *production* of knowledge only at a superficial level. The thought is that sociological pressures—such

[1] A further view counts the cultural descendants of culturally equipped agents; there are, as I say, many ways to make the analogy.

[2] Stout (2011) provides a nice example of how an overly strong focus on tool morphology, rather than toolmaking practices, might have led us to underestimate rates of cultural change in the Lower Paleolithic.

as disciplinary gatekeeping against novel problem-solving techniques and problems that lack "significance"—would select for certain practices *because* they produce knowledge. However, if we want to explain *how* that knowledge is produced, we would have to look at the way scientific practices themselves select accurate hypotheses over inaccurate hypotheses. The idea is that we begin with a population of hypotheses which vary in accuracy. These are then exposed to scientific practice. Given a set of experimental techniques, concepts, theories, inferential methods, research problems etc., certain hypotheses will gain traction and others will fall by the wayside. Knowledge is generated to the extent that this process is epistemic; that is, has a bias towards retaining more accurate hypotheses over less accurate hypotheses. On this formulation, the cultural selection forces that explain the production of scientific knowledge operate in practices. The challenge would be to show how "traits" of truth or accuracy in populations of hypotheses are conserved over time, and how this produces vertical and horizontal increases in knowledge. Sociological pressures that promote successful practices over less-successful practices would be only a secondary concern. I think Haufe is right to point out that a commitment to the view that science is a cultural practice does not entail a commitment to the view that we cannot rank the epistemic merits of different cultural practices (2022: 91). However, I think vindicating this thought will require focusing on scientific practice as a source of selection, rather than as a unit of selection. The difference between Haufe's formulation and the one I have just sketched may be analogous to well-trodden disputes in evolutionary biology over whether selection acts on genes or organisms.[3] It would be interesting to explore these differences in further detail.

Another, more philosophical, worry lurks here. Those of a more pragmatic bent might think that solving scientific problems does not require much in the way of knowledge. Theories that are strictly speaking false can nonetheless be very useful. We still use Newtonian mechanics in many contexts despite being replaced by relativity, because the former gets the job done and the latter is unwieldy. Hence we might deny that selection for successful scientific practice equates to selection for scientific knowledge. Haufe's own characterization of the relationship between practices and knowledge—that the former is "implicated in" the latter—is unlikely to allay concerns of this kind. Some of this might be resolved if Haufe offered a definition of knowledge. Of course, you can't do everything. What counts as knowledge is a doozy of a philosophical debate: despite a history stretching back 2 millennia in the Western tradition there's no sign of resolution on the horizon, so you cannot fault Haufe for avoiding the issue. Nonetheless, some indication of the basic concept he has in mind would be helpful for clarifying the relationship between scientific practice and scientific knowledge. For instance, the know-how vs know-that distinction may be relevant here. If the concept of knowledge Haufe has in mind is more at the practical, know-how end of the spectrum then perhaps these concerns would be resolved.

One of the great strengths of Haufe's framework is that it explains both epistemic success *and* failure. The level of vertical refinement and horizontal breadth in knowledge produced by any given population will depend on how well that population meets the criteria of heredity, intrinsicality, and continuity, and on the extent to which the selection pressure acting on that population is epistemic. Haufe's model predicts that scientific communities corresponding to these conditions will make epistemic progress, but it also predicts that communities failing to meet these conditions will fail to make progress. This is important. A danger to any project like this is that you produce a model that is too inflexible with regard to failure, which subsequently makes it hard to see how science can do anything other than make progress. But science is messy. Sometimes it moves forward and sometimes it gets stuck. Haufe's model captures this aspect of the social reality of science nicely.

However, it also means the model must be rejected or accepted according to these predictions. If we find cases where a scientific community appears to score highly on heredity, intrinsicality, and continuity, and to be subject to epistemic selection pressure, yet is nonetheless failing to make progress, then Haufe's model is in trouble. It is likewise in trouble if a community is making progress in the absence of these conditions. So, how do the model's predictions fare? Well, abstract models like this are always going to find counterexamples in the messy real-world, so I don't think there is much point to rattling off a list. However, one does wonder how representative the nice illustrative cases that Haufe selects—such as the advances in chemistry produced in Justus Liebig's lab in the 1830's—are of science as a whole. Do these cases generalize to a more robust pattern? It would be interesting to see how Haufe's model holds up to testing over a large, randomized sample-set. This would be a great sociology of science project.

Haufe's book raises some important implications for the future of scientific knowledge, and for philosophical attempts to come to grips with it. One in particular struck me. If Haufe is right, and if science continues its current broadly successful trajectory, then the breadth and depth of scientific knowledge will continue to increase and diversify. As sub-disciplines and sub-sub-disciplines become ever more specialized, we might start to wonder whether the project of producing a general overarching synthesis of scientific knowledge is even possible. Wilfrid Sellars once characterized philosophy as the attempt to "… understand how things in the broadest possible sense of the term hang together in the broadest possible sense of the term" (Sellars, 1991: 35). If we project Haufe's model two or three centuries into the future, would achieving this goal look even remotely feasible? Of course, the worry here isn't unique to Haufe's project—rather it is a general concern about increasing specialization in the sciences. But his model does bring the issue into sharp focus. Will scientific knowledge "grow" to the point where putting it "all back together" is an impossible task? If so, what does this mean for attempts to produce a general understanding of reality?

A more practical implication concerns the kinds of conditions Haufe's model predicts would create the most productive research ecosystems. Consider, for instance, the issue of isolating an evolving population of scientific practices from the detrimental impact of too much variation. Just as biological evolution requires clusters within a population to become reproductively isolated from one another, so the emergence of sub-disciplines within the sciences requires a limit on, to use Kuhn's memorable phrase, the "range of possible partners for fruitful intercourse" (Kuhn et al., 2002: 8; quoted in Haufe, 2022: 82–83). Here it is important that scientific communities operate as "epistemic niche constructors" (Haufe, 2022: 72), shielding populations from the introduction of too much novelty and ensuring that practices adapt to epistemic

[3] Thanks to an editor at *Evolution* for pointing this out.

rather than non-epistemic pressures. Haufe's example of this process comes from the funding guidelines of the National Institutes of Health, which provide highly specific problems to be addressed and explicitly warns applicants against "… seeming too innovative" (Haufe, 2022: 89). And while this example is well-taken, it is just one case. And in the current tech-driven climate it is very common to see funding bodies encouraging scientific innovation and novelty. For instance, the European Union's "Horizon Europe" programme and associated European Innovation Council seeks to promote "ground-breaking, high-gain/high-risk research" and "frontier research and breakthrough scientific ideas" (Directorate-General for Research & Innovation, 2021). On Haufe's framework, should we worry that programs such as these are obstructing the horizontal expansion of knowledge by introducing too much variation into certain populations of scientific practices? If he is right, do we need to rethink the incentives these kinds of funding programs prioritize?

I want to finish on a positive note. To reiterate, in my view Haufe's synthesis of evolutionary biology, cultural evolutionary theory, philosophy of biology, and the history of science is the right approach to explaining the development of scientific knowledge. And it's the right way to meet Godfrey-Smith's challenge of showing how the evolutionary analogy can be philosophically insightful. Indeed, I think Haufe's program is stronger than this, insofar as it meets a stronger challenge, which I'll call *Sterelny's challenge*. In his review of Daniel Dennett's "Darwin's Dangerous Idea," Kim Sterelny (1999) takes issue with meme theory for the following reason. The big question facing biology in the 19th century was: "where did all this (apparent) design come from if there was no designer?" The answer to that question was: "evolution by natural selection." No such explanatory challenge faces the development of scientific knowledge (or any other cultural practice). The answer to the question: "where did all this (actual) design come from?" is simply: "Us." Sterelny puts it as follows: "In human culture, the paradigm examples of organized complexity are the results of conscious design; of sighted rather than blind watchmakers. So what needs explaining?" (1999: 259). Here we have a very important *dis*analogy between biological evolution and cultural change, and one that should be addressed by proponents of the analogy.

With Haufe's framework in hand, we can respond to Sterelny's challenge. What needs answering in the case of science is as follows: how can a socially embedded cultural practice carried out by biased and fallible agents achieve the kind of incredible advances in knowledge that science clearly has? (see Haufe, 2022: 2). In other words, we can be skeptical whether human designers are, *on their own*, capable of producing the kind of design we see in science. Haufe's response is that, at their best, scientific communities maintain a set of population-level dynamics that ensure knowledge increases over time despite the shortcomings of individual practitioners. In other words, the whole (scientific knowledge) is greater than the sum of its' parts (individual scientists). This aspect of Haufe's work is importantly similar to recent work in biological anthropology on the evolution of cumulative technological culture. Here the question is: "how do humans maintain a level of technology that outstrips the abilities of any one individual?" And the answer is: "culture" (Boyd, 2017; Boyd & Richerson, 1988; Boyd et al., 2013; Richerson & Boyd, 2008). It would be interesting to explore crossovers between Haufe's project and this body of work.

Ross Pain

Department of Philosophy, University of Bristol, Bristol, United Kingdom

Corresponding author: Department of Philosophy, Cotham House, University of Bristol, Bristol BS6 6JL, United Kingdom.

Email: ross.pain@bristol.ac.uk

Funding

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 101018523).

Conflict of interest: The author declares no conflict of interest.

References

Boyd, R. (2017). *A different kind of animal: How culture transformed our species*. Princeton University Press.

Boyd, R., & Richerson, P. J. (1988). *Culture and the Evolutionary Process*. University of Chicago Press.

Boyd, R., Richerson, P. J., & Henrich, J. (2013). The cultural evolution of technology: Facts and theories. In *Cultural evolution*. The MIT Press. https://doi.org/10.7551/mitpress/9780262019750.003.0007

Directorate-General for Research and Innovation (European Commission). (2021). *Horizon Europe, pillar I – Excellent science: Driving scientific excellence and supporting the EU's position as a world leader in science*. Publications Office of the European Union. https://doi.org/https://data.europa.eu/doi/10.2777/456952

Godfrey-Smith, P. (2003). *Theory and reality: An introduction to the philosophy of science*. University of Chicago Press.

Godfrey-Smith, P. (2009). *Darwinian populations and natural selection*. Oxford University Press.

Haufe, C. (2022). *How knowledge grows: The evolutionary development of scientific practice*. MIT Press.

Heyes, C. (2018). *Cognitive gadgets: The cultural evolution of thinking*. Harvard University Press.

Keen, I. (2022). The evolution of Australian Kin Terminologies: Models, conditions, and consequences. *Current Anthropology*, *63*(1), 31–67. https://doi.org/10.1086/719365

Kuhn, T. S., Conant, J., & Haugeland, J. (Eds.). (2002). *The road since structure: Philosophical essays, 1970–1993, with an autobiographical interview*. University of Chicago Press. https://press.uchicago.edu/ucp/books/book/chicago/R/bo3639271.html

O'Brien, M. J., & Lyman, R. L. (2002). Evolutionary archeology: Current status and future prospects. *Evolutionary Anthropology: Issues, News, and Reviews*, *11*(1), 26–36. https://doi.org/10.1002/evan.10007.

Richerson, P. J., & Boyd, R. (2008). *Not by genes alone: How culture transformed human evolution*. University of Chicago Press.

Sellars, W. (1991). Philosophy and the scientific image of man. In *Science, perception and reality* (pp. 7–40). Ridgeview Publishing Company.

Shennan, S. (2011). Descent with modification and the archaeological record. *Philosophical Transactions of the Royal Society of London, Series B: Biological Sciences*, *366*(1567), 1070–1079. https://doi.org/10.1098/rstb.2010.0380

Sterelny, K. (1999). Dawkins' Bulldog. *Philosophy and Phenomenological Research*, *59*(1), 255–262. https://doi.org/10.2307/2653476

Stout, D. (2011). Stone toolmaking and the evolution of human culture and cognition. *Philosophical Transactions of the Royal Society of London, Series B: Biological Sciences*, *366*(1567), 1050–1059. https://doi.org/10.1098/rstb.2010.0369