

5 Branching and the Growth of Scientific Knowledge

Our primary task in this book has been to investigate the extent to which the evolutionary process might serve as an apt framework for modeling the historical development of scientific knowledge. The larger purpose of this investigation is twofold. First, there is the question of whether and to what degree the evolutionary process of directional refinement is adequate as a description of the causal process that generates scientific knowledge in certain well-described cases. Second, the ability of the evolutionary process to successfully capture the dynamics of scientific knowledge is to serve as a philosophical explanation for why scientific knowledge often exhibits a pattern of development that distinguishes itself from other forms of human activity in a way that is taken to be exemplary of the phenomenon of epistemic progress.

At the core of our efforts lies the concept of a Darwinian population. The degree to which a group satisfies the criteria for being a Darwinian population is to serve as a measure for how effective natural selection is (or could be) at modifying it; or, to put it in slightly different terms, it is a measure of natural selection's particular contribution to the overall modification of a group. Throughout the last few chapters we have used these criteria to make various claims regarding what it would look like for scientific practice to approximate the epistemic equivalent of a paradigmatically Darwinian population—what I have called an idealized epistemic community. We've also looked at aspects of scientific practices in some communities that often seem to closely instantiate the relevant properties of a paradigmatically Darwinian population. Chapter 2 showed that there are several good reasons for thinking that scientific practice often appears to satisfy to a high degree the *heredity* criterion. Chapter 3 likewise argued that there are good

reasons to think that scientific practice often appears to satisfy the *intrinsicity* condition. Finally, chapter 4 argued that a population undergoing optimization by directional natural selection will display a characteristic sequence of modifications of related sizes, and that the historical development of scientific practice often appears to fit that characteristic sequence very well indeed.

This last step allowed us to connect the evolutionary model's potential success as a causal explanation for the historical development of scientific knowledge on the one hand with the model's promise as a normative explanation for the special epistemic status of some instances of scientific inquiry on the other. We take science's ability to solve problems as partly constitutive of what makes science epistemically special. And its ability to solve problems in the way that it does is explained by the fact that it possesses versions of the same properties that cause biological populations to solve the problems posed by their environments. It thus turns out that part of our conception of what makes science epistemically special is an inevitable consequence of the causal process responsible for the historical development of scientific knowledge.

We are now faced with a new explanatory task which, like the one previously discharged, contains a descriptive dimension, a normative dimension, and a dimension that seems to be at once descriptive and normative. In general, the success of the evolutionary model as a normative account of scientific knowledge depends on whether we can derive science's epistemically special features as inevitable consequences of the fact that scientific practice, because of its approximations to a paradigmatically Darwinian population, is often highly susceptible to modification through natural selection. We have done this for science's ability to solve problems. Now we must do the same thing for the perpetually increasing breadth of scientific knowledge. That is, we need to show (1) that the modification of scientific practice through natural selection in idealized scientific communities inevitably results in the broadening of scientific knowledge; (2) that the broadening of scientific knowledge, like problem solving, has a direct biological analogue; and (3) that this biological analogue is also the direct and inevitable result of natural selection's modification of paradigmatically Darwinian biological populations. For simplicity's sake, I use the term "growth" throughout this chapter to refer specifically to the increasing breadth of scientific knowledge.

5.1 The Epistemic Significance of Growth

Some historians of science believe that Isaac Newton probably possessed all of the scientific knowledge that existed during his lifetime (for example, Westfall 1980). As copious as Newton's mind surely was, not even he could achieve mastery over more than a fraction of the body of knowledge to which science can now lay claim—and not just because he is dead (although that is a big part of the problem). Since Newton's time, scientific knowledge has expanded to such an extent that almost all practitioners are well versed only in the literature that pertains directly to the highly esoteric area of research in which they've chosen to work. Although practitioners probably have a general knowledge of all the branches of a field that they need to study before getting to their esoteric corner of research, their knowledge will tend to be fairly basic. They will usually not know much about the frontiers of current research in areas other than their own.

The unwieldy expansion of scientific knowledge since the late seventeenth century is perhaps just one particularly strong illustration of the fundamental fact that knowledge grows over time. But why is the fact that knowledge grows over time epistemologically significant? That is, why should we believe that the mere phenomenon of epistemic growth might be able to tell us something about the nature of knowledge? I think we can articulate at least four independent reasons for its significance, each of which suggests a separate lesson for the study of knowledge itself.

The first reason is simply that any historical trend exhibited by knowledge is *prima facie* epistemologically significant. The formulation, development, and refinement of a theory involves responding to the factors that function as constraints on a theory in the relevant domain. This is as true of philosophical theories as it is of scientific ones. Philosophical theories of knowledge have, in the past century, taken as their primary set of constraints some intuitions about descriptions of cases in which certain allegedly necessary conditions on knowledge are satisfied (e.g., Gettier cases) (Gutting 2009). Now, regardless of whether or not that particular species of constraint is effective, the intention behind its use is to sort potentially viable philosophical theories of knowledge from the weak, the nonfunctional, and the generally unpromising ones. The viable theories obey the constraints that are recognized by a research community at a given time. Some forms of progress in epistemology (as in other areas of theorizing)

involve refining our theories in response to such constraints, uncovering new constraints, and discovering that some constraints had been adopted erroneously.¹

The epistemological significance of knowledge's historical trends emerges through the ability of these trends to function as constraints on a theory of knowledge. A historical trend offers us a picture of how knowledge behaves (1) on large time scales and (2) at the community level, both of which reflect in their own way the close association between knowledge and some form of *stability* (Kitcher 2015, 477). Trends that hold across relatively large chunks of history reflect the insensitivity of certain aspects of knowledge to perturbations of any kind, because these histories are assumed to include a panoply of different sociocultural, intellectual, and environmental conditions. Because variation in these conditions evidently does not disrupt the relevant epistemic phenomena, we impute to such phenomena an ability to persist despite radical changes of context. We are thus given some sense of what knowledge behaves like regardless of the conditions that prevail at any given time. Whether or not these trends resonate with our intuitions about knowledge, they seem as good a candidate as anything for serving as constraints that a viable theory of knowledge would have to obey. To echo Kuhn (yet again), "how could history of science fail to be a source of phenomena to which theories about knowledge may be legitimately asked to apply?" (Kuhn 1962b, 9).

The second reason for the epistemological significance of the growth of knowledge over time is that growth is often treated as an index of success. Success, in turn, typically demands a causal explanation; we are reluctant to accept accounts of success that are grounded in brute luck. The success of scientific knowledge, manifested in its spectacular growth since the seventeenth century, is no accident; the growth of knowledge is epistemologically significant in that it reflects the ability of knowledge to cause successful outcomes.

To make the point sound somewhat less trivial, consider a counterfactual history in which human knowledge had a tendency to remain more or less static, or perhaps even to shrink. Would the conceptual connection between knowledge and success still hold under these conditions? It is not clear to me that it would. If, for example, knowledge remained static for a relatively large chunk of history regardless of intense investment in rational inquiry, it would then appear to be the case that knowledge is not particularly notable

for its contribution to the success of our intellectual endeavors. Thus, despite the intuitive appeal of the idea that knowledge facilitates successful outcomes, this counterfactual history would offer *prima facie* evidence to the contrary. From this vantage point, the historical fact that knowledge does (thankfully) grow over time is deeply significant. It provides strong justification for the use of the causal connection between knowledge and success as a constraint on philosophical theories of knowledge.

Taking the basic association between growth and success as our starting point, we can distinguish a couple of different forms of success, only for one of which is growth a symptom. Over the last few chapters we have been looking at how ideas might be directionally refined and perfected in a manner consistent with varying degrees of influence of natural selection. The relevant sense of *success* in that context was one where researchers were able to develop better answers—and better questions—by building on the contributions made by previous researchers to the understanding of a given phenomenon. But success in this sense is perfectly consistent with zero growth. Indeed, under perfectly normal circumstances our total corpus of knowledge might actually contract during—and possibly as a consequence of—the process of refinement. For example, in chapter 3 we discussed how there can often be selective pressure to reduce the number of different kinds of phenomena one is studying because too much variation can disrupt the optimization process. Thus, it is conceivable—even expected—that our knowledge might sometimes cease to grow or even contract as we attempt to deepen our mastery over a particular phenomenon, because we might have either (1) temporarily suspended the investigation of new phenomena (stasis) or (2) abandoned some of our working posits regarding certain phenomena in lieu of a less general but more precise solution to a particular problem (contraction) (Kuhn 1962b, 170). Yet there is a clear and distinct meaning to the notion that we have achieved success in this instance; we had a good solution to a problem, and we've improved on that solution.

By contrast, the sense of *success* for which growth is a symptom needn't indicate any sort of refinement. Rather, success of this kind involves the broadening of our understanding through the investigation of hitherto unexplored or unknown phenomena. In such cases, our success resides not in the development of a completely satisfying account of a new phenomenon but in *the development of a new line of inquiry*. The mere existence of this new line of inquiry provisions future generations of researchers with

an opportunity for the kind of refinement-oriented success described in the preceding paragraph. Research communities who achieve success of this kind will often have only a vague impression (if any) of how to approach the object of inquiry. But their initially tenuous grip can (and in many instances does) strengthen over time, fueled by the directional pressures detailed in chapter 4.

As promised, that chapter attempted to draw a connection between researchers' tightening grip over a phenomenon on the one hand and an instantly recognizable sense in which science makes progress on the other. This sense is one of two distinct but overlapping varieties of scientific progress—namely, the movement toward an optimum in practice space that we earlier labeled “vertical progress.” If we think of progress in this sense as a kind of movement from one location in practice space to another location, the relevant notion of *success* is that of occupation of such a location. In general, success is something that happens at a specific time, whereas progress is what happens over time. The diachronic trend denoted by “progress” is composed of a series of synchronic events, each denoted by the term “success.”

This relation between the diachronic phenomenon of progress and the synchronic phenomenon of success generalizes to the present chapter because the growth-centered species of success is precisely the reason that growth and progress are treated as conceptually linked in some contexts. If the growth-centered species of success is the addition of a new line of inquiry, there is a related growth-centered species of *progress*, constituted by a positive trend in the number of phenomena of which we have at least a dim understanding. This form of progress—which I earlier called “horizontal”—is associated with the persistent tendency of rational inquiry to develop new branches of knowledge, to add “twigs” to those branches, and to add yet further, nubby outgrowths to those “twigs,” invoking the growth of the proverbial tree of knowledge over time. Horizontal progress—the overall tendency to increase the number of things we understand or investigate—is the third reason for the epistemological significance of the fact that knowledge grows over time.

Armed with this new species of progress, we can articulate a fourth and final way in which the growth of knowledge bears on a philosophical theory of knowledge itself. We observed above that the growth of scientific knowledge in some cases appears on the whole to be insensitive to contingencies of time and place, for which we awarded epistemic growth the highly desirable property of *stability*. Of course, some forms of stability are

more desirable than others; not all forms are epistemologically significant. For example, some of the temporal variation over which the growth of knowledge remains stable involves radical changes in which styles of dress are popular. That's not very interesting or surprising from an epistemological perspective, because our (warranted) default presupposition is that styles of dress are not all that relevant to knowledge.

But there is another sort of change across which the growth of knowledge is historically stable, one with which any philosophical theory of knowledge must reckon—namely, the content of our scientific theories. In Ian Hacking's words, "we know more about polyhedra or atomic weights than we once did, even if future times plunge us into quite new, expanded reconceptualisations of those domains" (1979, 384). Over the past two centuries, we have refined an early insight due to Dalton regarding what elements are and how they combine, but we have also since that time greatly expanded the number of phenomena in this domain that fall within our grasp. This latter trend holds despite radical changes in our views about the nature of atoms and of subatomic particles (e.g., from classical to quantum mechanical). Our ability to increase the number of things we know (or are on the cusp of knowing) has not been impeded by even as drastic a shift as that involved in the transition from discrete to continuous conceptions of matter. Indeed, that ability has almost certainly been enhanced.

What this observation suggests is that one form of progress—horizontal progress—is a higher-order phenomenon that is robust across changes in its constituents. Once we've added a line of inquiry, we have made horizontal progress, even if we eventually abandon everything we initially believed about the phenomena associated with that line of inquiry. Being concerned strictly with the cardinality of the set of phenomena under investigation, horizontal progress generally does not require that our views about its constituents remain fixed. Compare this content independence of horizontal progress with the content dependence of vertical progress. Because the core of vertical progress is the refinement of practice, there must be some causal continuity between earlier and later successes; later successes build on earlier successes. By contrast, progress in the horizontal dimension does not (or, at least, need not) occur by building on previous results. Lines of inquiry can be added quite independently of one another.

To summarize our discussion so far, the historical growth of knowledge is epistemologically significant because (1) historical trends in knowledge

function as constraints on a theory of knowledge; (2) growth is causally associated with success, and the success of science is a large part of what makes it appear epistemically special; (3) the accumulation of instances of the kind of success that constitutes growth is associated with a form of progress; and (4) this form of progress is not directly dependent on the content of our practices.

The preceding considerations establish the growth of knowledge as a datum that any credible theory of scientific knowledge needs to imply. Because of its manifest epistemological significance, epistemic growth must function as a constraint on what kinds of philosophical theories of knowledge promise to deliver the most complete understanding of epistemic phenomena. Theories of knowledge that imply epistemic growth as an inevitable consequence of the presence of knowledge are, therefore, better positioned than others to illuminate the nature of knowledge itself.

Our analysis of epistemic growth suggests that we think of it in terms of the closely allied (and conveniently epistemic) notion of *success*, and that the growth-oriented notion of success be understood in terms of the addition of a refineable line of inquiry. This, in turn, led to a conception of *horizontal* progress, which centered around the accumulation of such lines over time (as opposed to vertical progress, which is consistent with the nonproliferation of lines of inquiry). This analysis has raised the bar for successfully linking the growth of scientific knowledge to some biological phenomenon, because it requires that we be more precise about what we mean by “growth” in biological contexts. In particular, the relevant form of biological growth should bear a close resemblance to the sense of success associated with the addition of a refineable line of inquiry.

In section 5.2 I show that the conception of growth operative in evolutionary biology closely mirrors the notion of epistemic growth as analyzed above. This will establish the *prima facie* case for thinking of growth-oriented epistemic success as an evolutionary phenomenon. I then move on in section 5.3 to consider the evolutionary processes by which lineages undergo lineage addition—what I will henceforth refer to as *branching*. It will turn out that only certain kinds of branching events follow inevitably from the process of natural selection. Since part of what makes idealized scientific communities exemplary from an epistemic point of view is the way in which their distinctive epistemic properties follow from the process of

natural selection, we need to be able to model at least some kinds of scientific growth as branching events of these kinds. Following our discussion of evolutionary branching in general, in section 5.4 I examine the arguments for thinking that scientific practice does, at least sometimes, branch in this epistemically desirable manner. In part II of the book, I use an extended case study to illustrate the nature of epistemic branching.

5.2 Evolutionary Success

To establish the claim that the growth-oriented conception of scientific success is a kind of evolutionary success, let us begin by examining the concept of *evolutionary success* as it is (and has historically been) used in paleontology.² Warning: it is a bit of a mess.

Two dominant themes emerge from a survey of discussions of evolutionary success in the paleontological literature, both of which can be discerned in the following three excerpts picked at random from the writings of the century's most prominent paleontologists:

If sheer weight of numbers is the deciding factor, this is the most spectacularly successful of all the phyla, for it includes the insects. (Simpson 1949, 28)

We know, for example, that some biologic groups have been much more successful than others, as measured by taxonomic or ecologic diversity or by temporal persistence. (Raup et al. 1973, 526)

This radiation spawned highly successful lineages of woody plants from the standpoint of their [species] longevity, structural diversity, and species numbers. (Donoghue 2005, 85)

One theme is that paleontologists are well aware of the fact that there is no consensus statement on what the term "success" means in evolutionary contexts; most uses of "success" are accompanied by some kind of clarificatory hedge or other. The other theme that emerges is that, despite the known absence of a precise definition, paleontologists seem to more or less agree on what counts as evolutionary success; one rarely encounters disputes over whether the term "success" fits a particular case, or whether certain conceptions of success are incorrect from an evolutionary standpoint. Thus, "success" and its more frequent cognate, "successful," are used fairly liberally, but always with a statement describing the conception of success being invoked.

Paleontologist Pete Ward has recently tried to capture the variegated applications of the evolutionary concept of success in a list composed of eight distinct notions, each of which corresponds to what has been meant by “evolutionary success” at different places in the paleontological literature:

1. *Individual longevity*: how long an individual organism lives
2. *Species longevity*: for a higher-level taxon, the average duration of time between when a species first appears in the fossil record and when it disappears
3. *Species fecundity*: the number of daughter species produced by a given species
4. *Individual abundance*: the number of organisms that are members of a given lineage
5. *Percentage of the planet's biomass*: the proportion of the organic material comprising a lineage, relative to the total amount of organic material on Earth
6. *Species that co-opt other species for their betterment*: species that either manipulate or exterminate other species, resulting in an increase in individual abundance
7. *Wide geographic range*: a lineage for which the members that are farthest apart from one another are far apart in an absolute sense
8. *Surviving mass extinctions*: lineages that were not driven to extinction during at least one of the major extinction events in the history of life

There are a few observations we can make about the list. First is that it doesn't capture every sense of “successful” employed in the passages quoted above. Simpson's “sheer weight of numbers” corresponds to *individual abundance*. Donoghue's “species number” and “longevity,” as well as Raup's et al. “diversity or . . . persistence,” are all captured by *species fecundity* and *species longevity*, respectively. However, Donoghue's “structural diversity”—that is, the number of different morphotypes in a lineage—cannot be accommodated directly by any conception on Ward's list.

The second thing to notice is that all conceptions (save one)³ are defined in terms of some form of *relative growth*: relative growth in (1) years (*longevity*); (2) diversity (*lineages* or *structures*); (3) members (*abundance* and *biomass*);

(4) area (*geographic range*); or (5) habitats (*geographic range*). Underlying these distinct conceptions is a notion of evolutionary success that seems fundamentally to be about the tendency of a clade to exhibit relative growth along a dimension that is evolutionary significant on long timescales.⁴

Third, each subcategory of success describes a property associated with the *stability of a clade across environmental change*. Longevity is a measure of stability across time, which is itself a proxy for degree of environmental change; the longer the time, the more environmental change there is likely to have been. Abundance is associated with stability across change in a different way, in that it reduces the overall probability of extinction through either sheer weight of numbers⁵ or increasing the amount of variation available for responding to environmental change.⁶ Diversity and geographic range are both associated with stability across different environments and hence across environmental change. The centrality of stability across environmental change also allows us to capture Ward's inclusion of "surviving mass extinctions" on the list of properties that are closely allied with evolutionary success.

Putting these second and third observations together, it is probable that the reason why these forms of relative growth are evolutionarily significant on large time scales is precisely because they are associated with stability across environmental change. The more growth a clade undergoes in one of these categories, the more likely the clade is to persist across changes in the conditions to which its members are adapted. A stability-centered generalized conception of evolutionary success would then follow: *evolutionary success* involves a clade's relative growth along any dimension associated with stability across environmental change.

We can further improve on this conception by reflecting on a fourth trend—namely, the hierarchical relationship between the different kinds of evolutionary success identified thus far. It will be noticed that each kind of evolutionary success fits into a nested hierarchy of sets of environmental challenges. At the lowest level sits the *abundances*, which correlate to the narrow set of environmental challenges faced by a local population and which, to an extent, define it as a population (see chapter 2). Above that sits *species diversity* and *longevity*, which group populations into distinct lineages within a clade, each of which is associated with a distinct set of environmental challenges that distinguish its evolutionary fate from that of its

sister lineages. Moving up a level we find *geographic range*, which includes not only the environmental challenges that a clade's lineages actually encounter, but also those that they could or did encounter—that is, each of the environmental factors that either could potentially affect the selective advantageousness of any trait in the clade or did affect advantageousness at some point as the clade's geographic range expanded. At the very top of the hierarchy rests *surviving a mass extinction*, which includes every possible set of environmental challenges for many effectively unrelated clades.

Each level in the hierarchy corresponds to a qualitative distinction in the degree of difference in the nature of the environmental challenges faced by denizens of a given rank. The problems faced by individual members of a local population will be maximally similar to one another; this is why competition is evolutionarily significant primarily at the population level. Increasing abundance shows that the population is getting better and better at solving the problems that make a difference to survival and reproduction. Next are the sets of problems faced by sister species, which are expected to be different enough to have kept them reproductively distinct, but similar enough to have made branching possible. Following that, the breadth of a geographic range and the time taken to expand it are both metrics for the degree of difference between the problems encountered by members of a clade, with wide ranges indicating potentially radical differences in the kinds of problems affecting members. Finally, those lineages that survive mass extinctions have been exposed to the most radical differences in environmental challenges of which we are currently aware.

Viewed in this way, each kind of success involves the expansion either in degree or in kind of a group's capacity for solving evolutionary problems. Increasing abundance reflects growth in the ability to solve the problems of local adaptation. Increasing diversity reflects growth in a clade's capacity to solve problems that are distinct enough to promote or maintain divergence. Increasing geographic range reflects growth in a clade's capacity to solve qualitatively different kinds of selective problems. And surviving a mass extinction reflects a clade's capacity to rapidly adapt to entirely distinct selective regimes. Incorporating this fourth trend into our generalized conception of evolutionary success suggests the following picture: at any given level, *evolutionary success is distinguished by relative growth in problem-solving capacity*.

5.3 Assessing the Analogy between Evolutionary Success and Epistemic Success

The formulation of evolutionary success as growth in problem-solving capacity highlights the close connection between it and the growth-oriented sense of epistemic success described in section 5.1 and makes plausible the idea that growth-oriented epistemic success is a type of evolutionary success. Fundamentally, both conceptions of success are centered on the idea that adding members to a group increases the group's capacity to solve problems. In the case of scientific practice, the "group members" are lines of inquiry (or a superset of lines of inquiry; see later in this section). In the case of biological groups, the members are individual organisms, lineages, or groups of lineages.

I think there is cause to take this analogy very seriously. First, the depiction of knowledge as a tree—a collection of "groups subordinate to groups"—dates back to ancient Greece and continues steadily to the present day. These depictions are indistinguishable in form from representations of the tree of life, a tradition that grew out of the effort to classify organisms in a nested hierarchy. In both domains, "growth" has consistently involved the addition of "branches." Second, both biological and epistemic branches are *perpetuated by descent*. In chapter 2 we saw a variety of mechanisms that can work effectively to reproduce scientific practice with a high degree of fidelity. As the emphasis on historical relationships became increasingly important in both domains—histories of science and natural histories—the branching scheme took on a new significance, in that it was intended to represent lines of descent. In the same way that the tree of life grows not by merely adding new lineages but by new lineages emerging from other lineages, so too does the growth of knowledge routinely proceed by the emergence of new lines of inquiry from existing lines. Relatedly, in both cases resemblances between branches are most commonly explained in terms of degree of genetic relatedness. As with species, two branches of knowledge most closely resemble each other whenever they are immediate descendants of a common ancestor.

Third, recall that the growth of knowledge was associated with a recognizable conception of progress in the following way: The addition of a refinable line of scientific inquiry qualifies as an epistemic *success* in that it constitutes a broadening of our understanding through new (or new kinds

of) investigations; the accumulation of such instances of success over time constitutes (horizontal) progress—that is, a positive *trend* in the number of (kinds of) phenomena of which we have at least some minimal understanding or on which we now have some degree of traction. How, if at all, does the accumulation of biological lineages correspond to progress in this sense?

In chapter 1 I used some of Darwin's ideas on evolutionary progress to formulate a basic picture of what progress looks like when the historical development of a lineage is governed by natural selection. In practical terms, Darwin imagined a population whose members, despite continuing to improve their ability to profit from the resources for which they were already fairly well suited, had failed to make use of a particular endemic resource—say, a source of potential nutrition—due to the absence of variation capable of exploiting it. If a suitable variant were to arise, the population could expand its resource base, allowing it to grow through its ability to feed more members.

We can, alternatively, envision the expansion of a population's geographic range through similar means: a variant arises capable of exploiting an available resource whose range is broader than the current resource profile for that population. This variant is thus able to venture beyond the population's current geographic range, where it will encounter yet further novel sources of sustenance that it might someday chance to exploit as well. In this way, the original population expands both in abundance and in geographic range. According to Darwin's picture, the descendants of newly emergent variants will—after a sustained period of pursuing hitherto unexploited opportunities—have diverged morphologically from their parent population to such an extent that they eventually comprise a distinct lineage whose evolutionary fate is no longer directly tied to its parent population. He labeled this effect “the Principle of Divergence” and took it to illustrate how natural selection, along with explaining adaptation, is also able to account for the fact that there are so many different kinds of organisms, and that those kinds form “groups subordinate to groups” (Darwin 1859, chap. 4).

Driven by this process of “divergence,” the accumulation of biological lineages amounts to a positive trend in the number of environmental problems that a clade is involved in solving. The opportunity to diverge begins with the appearance of a variant that partially solves a problem that has never been solved—namely, how to exploit an available resource. Once the exploitation begins, given enough pressure or time, that variant's descendants will have evolved into their own stable, self-perpetuating lineage.

Through the emergence of a stable lineage in this way, the larger clade of which that new lineage is now part has thus increased the number of problems that it is involved in solving.

At the core of Darwin's image of progress is the phenomenon of *increasing specialization*. Vertical progress is defined by incremental increases in the degree of specialization along a trajectory established by previous generations. Horizontal progress, in contrast, consists of the founding of a new trajectory; where there had previously been no opportunity for specialization, one now exists. Both forms of progress are characterized by improvement in the population's ability to exploit potential resources.

By now it should be clear that the general conception of evolutionary success that was articulated above bears the same relationship to horizontal progress in biological contexts that growth-oriented epistemic success bears to horizontal progress in epistemic contexts. In the same way that scientific communities often appear to progress horizontally by increasing the number of problems that practitioners are involved in solving through the *addition of lines of inquiry*, biological populations make horizontal progress through the addition of new lineages, which increases the number of problems that clades are involved in solving. Fundamentally, the growth of a clade is driven by the emergence of a new opportunity to specialize. And so it is with the growth of knowledge. I believe that there is much to recommend the image of a scientific problem as a resource that scientific communities use to sustain themselves. Problems are the fuel that drives inquiry.

5.4 Evolutionary Branching

This last dimension of the analogy points to the next step in my argument for the idea that part of what makes scientific knowledge epistemically distinctive rests on the fact that the growth of scientific knowledge is driven by natural selection. Up to now, we've been discussing the notions of *growth*, *lineage addition*, and *problem* in a fairly general way, in order to make it easy to highlight the connections among them in both biological and scientific contexts. We now need to introduce a single refinement, common to each notion, that will clarify both their evolutionary and epistemic importance.

Let's begin with the more basic notion of lineage addition. By the addition of a lineage, I mean a *speciation* or *branching event*—that is, the emergence

of a distinct, autonomous lineage whose evolutionary fate is not necessarily tied to that of any other lineage.⁷ We're not really sure how often this occurs in nature, but it is almost certain that any measurement would underestimate its frequency because most newly emergent autonomous lineages are likely to be of very short duration.

Intuitively, the addition of an ephemeral lineage does not correspond to a kind of growth. This intuition is confirmed by the fact that species longevity is taken to be one of the components of evolutionary success, a notion that we found to be grounded in the concept of growth. We would not count a clade as evolutionarily successful in the case where, although it gave rise to a great many species, each of them quickly went extinct. Rather, in order for the addition of a lineage to count as an evolutionary success—and thus, to contribute to a clade's growth—it needs to last for an appreciable duration on a macroevolutionary timescale.

There is an obvious parallel here with the growth of knowledge. Just as the addition of an ephemeral biological lineage fails to qualify as evolutionary progress, we would not count the addition of a line of inquiry as epistemically progressive unless it made a lasting contribution to human knowledge. When a line of inquiry is judged by a community of practitioners to make a lasting contribution to knowledge, we say that it is *significant* (Haufe ms-b):

Def: a line of inquiry is *significant* just in case it is judged by a community of practitioners to make a lasting contribution to human knowledge.

Now, if horizontal epistemic progress is fundamentally rooted in the growth of knowledge, and if the growth of knowledge is grounded in the addition of a significant line of inquiry, then the accumulation of significant lines of inquiry must be at the heart of horizontal epistemic progress. In other words, one of science's epistemically important attributes is its propensity to generate significant lines of inquiry. Thus, not only is *significance* essential to our understanding of how the growth of knowledge contributes to the epistemic power of science, but the very notion of significance is grounded in the very same property that connects growth to evolutionary success—namely, that in both cases, the addition of a new line constitutes growth only if its effects persist for an appreciable duration.

This is an important result. Yet the centrality of natural selection to our philosophical explanation for the epistemic power of science requires us

to go even further. Because (1) if the perpetual emergence of significant lines of inquiry is one of science's most epistemically important attributes, and (2) if the epistemic importance of an attribute is explained in terms of natural selection's causal responsibility for it, then (3) it had better be the case that a line of inquiry's probability of significance—that is, of its making a lasting contribution to scientific knowledge—increases directly with natural selection's influence on the modification of practice.

I want to begin the argument for natural selection's direct effect on the significance of a line of inquiry by looking at analogous phenomena in a biological context, where we already have a relatively firm grasp of the evolutionary dynamics of branching. The essence of the evolutionary branching process is *genetic divergence*. In order for a lineage to undergo an evolutionary branching event—that is, in order for it to add a lineage—the genome of some subgroup within that lineage needs to become different enough to reliably prevent successful breeding with members of the larger parent group of which it is an offshoot.

We can decompose natural selection's role in the branching process into two independent components. First, natural selection might play an important role in branching events by causing certain variants within a population to diverge for ecological reasons. This is Darwin's "Principle of Divergence" discussed in section 5.3, which is known in contemporary evolutionary theory as *ecological speciation*. Second, natural selection might play an important role in branching events when gene flow between two subgroups has been obstructed. I will discuss each of these components in turn.

5.4.1 The Continuum of Ecological Disruption

We can begin to understand natural selection's role in branching events by revisiting the phenomenon of ecological speciation. At the core of the ecological speciation process is, as Darwin clearly articulated, a source of "ecological disruption" (Bolnick and Fitzpatrick 2007). Imagine a population whose sole source of nutrition is grapes. Through directional selection, the ability of members to procure and metabolize Biscuits steadily improves. Meanwhile, a local resource—toast—widely available and rich in potential nutrition, goes unused by the population due to their inability to metabolize it. Suppose that at some point, variants emerge within the population and are able to metabolize toast. Suddenly, there is a decrease in the selective pressure on these variants to improve their ability to procure

and metabolize grapes.⁸ After all, whatever they lack with respect to the population's average when it comes to grape-eating can be made up with a few easily obtained toast.

At this point, the population begins to undergo what is called *disruptive* or *divergent* selection—two distinct regimes of directional selection, each operating on a different subgroup within the population. While grape-selection continues to directionally improve typical members of the original population, the newly emergent toast-eaters have essentially opted out of grape-selection and are being pushed further and further down the path of toast-eating optimization.

Several factors will affect the chances that an ecological disruption like this will lead to the emergence of an evolutionarily autonomous lineage. The dominant factor is the toast-eating optimum's distance in morphospace. As each generation of toast-eaters continues to improve its toast-eating abilities, the toast-eating subgroup continues to diverge morphologically from its founding fathers and, a fortiori, the grape-eaters. This divergence ceases once the toast-eating optimum has been reached. If the optimum is relatively near in morphospace, toast-eaters will most likely have undergone relatively little modification. And being now only slightly different from grape-eaters, the toast-eating genome will most likely not have diverged enough to persist as a self-contained autonomous lineage. But if (1) the toast-eating optimum is quite far away, and if (2) selection for improved toast-eating is allowed to continue unabated, the chances of permanent divergence begin to look more promising. In short, the further the optimum, the more opportunity selection has to contribute to divergent modification.

Independently of the morphological optimum, another factor affecting the probability of divergence is the expected utility of toast-eating. This is a function of the reproductive benefit of (say) eating one piece of toast combined with the availability of toast. For the first toast-prone variants, divergence will not be possible unless it is at least as beneficial to eat toast as it is to eat grapes. If toast-eating causes a net loss in reproductive success, toast-eating will gradually disappear from the population. But if toast-prone variants are not reproductively penalized for toast-eating, they may exit the grape-eating selection regime and transition to eating toast without negatively impacting their reproductive success.

Since expected utility is a product, we can decompose it into two equally important selective components. One thing that will affect the expected

utility is the magnitude of the benefit of eating a piece of toast (not to be underestimated, in my experience). If eating a piece of toast causes an increase in reproductive success, toast-eating variants will become more widespread, other things being equal. But since they are *ex hypothesi* not competing against grape-eaters, the prevalence of grape-eaters remains unaffected. The toast-eating subgroup is simply allowed to grow due to its ability to exploit an unused resource. *Ceteris paribus*, the greater the benefit per piece of toast, the faster the toast-eating population grows. The more it grows, the more effective selection will be at optimizing toast-eating morphology, which (as has been mentioned) raises the probability of divergence.

But benefit magnitude can only perform this service if the benefits can be realized. Thus, the other component involves the availability of toast more generally—a proxy for how likely it is that a toast-eater will be able to procure a piece of toast (or several). If there are few toast-eaters and many toast, odds are it's toast for breakfast. If there are many toast-eaters fighting over a single piece of toast, the probability of procuring a piece of toast is very low. Thus, so is the expected utility of toast-eating. And so the probability of divergence varies directly with the probability of toast procurement.

These two components—benefit magnitude and availability—modulate the probability of divergence through their independent effects on how effectively selection can modify toast-eaters. Even if the net gain of eating a piece of toast is relatively low, toast-eating will still be a rewarding strategy if there is ample supply. Equally, if the net gain per piece of toast is relatively high but toast is in short supply, toast-eating will be similarly rewarded. In either case, selection is given the opportunity to optimize toast-eaters on account of the fact that the expected utility of eating toast is at least that of eating grapes.

5.4.2 Gene Flow

The foregoing factors are sufficient to cause a full-scale branching event in sexually reproducing populations under random mating, but only if the value of each factor is on the high end. Otherwise, if grape-eaters and toast-eaters continue to mate despite their divergent lifestyles, the population of toast-eating variants will generally be perpetually pulled back from the toast-eating trajectory and toward the population mean for the *ex hypothesi* larger parent population of grape-eaters.

There are a few ways of preventing the swamping of toast-eaters through random mating. Most models of ecological divergence get around this obstacle by positing a nonrandom component to mating, typically in the form of mate preferences (Bolnick and Fitzpatrick 2007, 469). A preferable, more conservative approach⁹ would be to introduce nonrandom mating through spacial considerations, along the following lines: toast-eaters like hanging out where there are toast to eat. If toast territory and grape territory share less than 50% of their respective spaces, toast-eaters are more likely to encounter (and mate with) other toast-eaters. Either way, the toast-eating subgroup is able to preserve the progress it has made toward the toast-eating optimum, because its mean phenotype can evolve unencumbered by the effects of the orthogonal optimizing forces governing grape-eaters (Fry 2003).

The simplest and most intuitive model of ecological speciation is, in fact, the process described in section 3.1. Here, assortative mating preferences evolve in response to selection against mating with individuals who might drag one's offspring toward intermediate phenotypic values and thus reduce the fitness of those offspring. It pays toast-eaters to mate with other toast-eaters because that is more likely to advance their offspring toward the toast-eating optimum than mating with grape-eaters. Likewise, it pays grape-eaters to mate with other grape-eaters. In general, "individuals who mate randomly risk producing intermediate offspring with lower fitness, indirectly favoring individuals that mate with their own ecotype" (Bolnick and Fitzpatrick 2007, 470). Those with sufficiently strong mate preferences will be favored by selection, because their offspring will tend to have higher fitness.

Regardless of which model one favors, each of them attempts to devise some way of *interrupting gene flow between subgroups*. It is generally recognized that the cessation of gene flow is a critical element in the development of a variant subgroup into a full-blown new species. Canonical models of speciation are "allopatric," meaning that genetic divergence occurs because subgroups are confined to different areas and mating between subgroups is impossible. Because ecological models of speciation aim at describing conditions under which speciation occurs even when members of divergent subgroups could still mate with each other in principle, such models generally look for ways in which mating between subgroups is permitted but discouraged.

The reason why cessation of gene flow is thought to be so important for speciation has to do with the homogenizing effects of random mating in

large populations. This principle is reflected in our discussion above as well as in section 3.1. In such populations, effectively random mating ensures that members' genomes tend to stay relatively similar to one another. Whatever minor genetic deviations arise in a generation will typically be lost due to chance, and so the next generation will more or less resemble its parents.

Sometimes, though, a subgroup becomes geographically isolated from its parent group, cutting off access to the parent's gene pool and the homogenizing effects of breeding within it. Alone in the wilderness, the isolated subgroup's genome begins to diverge from its parent. In such circumstances, there are two distinct gene pools, and the random morphological effects of sampling error play out differently in each pool. In addition, the isolated subgroup—the so-called founder population—is customarily hypothesized to be relatively small. And in a small population, the effects of random genetic drift are more pronounced; a rare mutation that disappears in a large population is less likely to do so in a small population. Thus, in contrast to the ecological process of speciation with which we began, these canonical allopatric models of speciation do not require natural selection to drive populations apart. The amplifying effects of small population size on genetic drift are sufficient to generate permanent divergence between a subgroup and its parent. After enough time has gone by, the isolated subgroup's genome has—purely through drift—become so different from its parent's that successful interbreeding is a low-probability occurrence. A new species is born.

Of course, natural selection also plays a role in fueling the divergence of genetically isolated subgroups. We saw in our discussion of ecological speciation how assortative mating increases the efficacy of natural selection by restricting gene flow between subgroups. Geographic barriers restrict gene flow to similar effect, thus allowing the two subgroups to pursue their distinct evolutionary optima. Although natural selection is not essential for speciation in allopatric models, it is a standard feature of such models because (1) it is ever-present; and (2) its modifying effects are typically much stronger and much more persistent than those of random drift (see section 3.1 for more details).

To summarize, the contribution that natural selection is able to make to the divergent modification of a new lineage increases with an increase in any of the following: (1) distance of the variant optimum, (2) magnitude of the

reproductive gain provided by the new resource, (3) relative abundance of the new resource, or (4) restriction of gene flow between subgroups.

As we have seen, another major contributor to the probability of permanent divergence is the small size of founder populations. But unlike (1)–(4), small population size does not increase the efficacy of natural selection; quite the reverse. Thus, the standard picture (right or wrong) of speciation is one in which gene flow is curtailed, thus allowing drift to dominate early on. As the isolated subgroup grows, the probability of sampling error dwindles commensurately, allowing ecological factors to properly dominate the modification of the lineage.

5.5 The Emergence of Specializations

I see the relation that *significant lines of inquiry* bear to *specialization* as that of *well-marked varieties* to *species* in the *Origin*. In general, to say that the emergence of a significant line of inquiry is analogous to the process of ecological speciation in biological contexts is to say that the kinds of Darwinian, selection-promoting factors that generate new species also work to generate specializations of lasting significance to scientific knowledge. What I would like to do in the remainder of this chapter is to provide a schematic description of a selection-driven branching process in science, a process in which the propensity to generate new specializations is modulated by the degree to which the development of scientific practice is dominated by selection—that is, the degree to which scientific practice approximates a paradigmatically Darwinian population.

The aim of this section is to illustrate a phenomenon of central epistemic importance, which often appears to arise in the historical development of scientific knowledge—namely, the inevitable emergence of *significant* lines of inquiry as a by-product of the influence of selection on the modification of scientific practice. The emphasis on “significant” is essential here. I do not maintain that every line of inquiry’s existence can be attributed to the influence of natural selection in scientific practice. Scientific knowledge does not necessarily benefit from the proliferation of lines of inquiry; it is not a weeping willow, which thrives on the unchecked multiplication of branches per se. Its likeness is rather that of a carefully curated bonsai tree. It must be constantly pruned, staked, and encouraged in very specific ways so that it might have the best chance of satisfying our purposes.

Nor I am suggesting that significant new lines of inquiry can only arise in science through the influence of the sorts of factors that facilitate the efficacy of natural selection. Just as no one working today would argue that new biological species can only arise through ecological speciation, we should embrace a plurality of circumstances out of which significant lines of inquiry have emerged. For example, I acknowledge the presence of serendipity in science, although I think its centrality to discovery has been vastly overplayed. More to the point, I fail to see how anything of epistemic importance might accrue to science by virtue of the fact that it is susceptible to the chance occurrence of significant discoveries. If these discoveries are indeed as chancy as they are often portrayed, why is modern scientific inquiry so disproportionately blessed with them? Why, for example, was the much longer empirical tradition of alchemical investigation not similarly favored by chance?

Ultimately, we must come to grips with and account for the fact that the growth of modern scientific knowledge through the proliferation of significant lines of inquiry is, for some reason, an inevitable consequence of the nature of modern scientific inquiry. This fact, along with the problem-solving capacity of modern science, are its most distinctive epistemic characteristics. The reason for this inevitable proliferation, I argue, is that scientific practice is—at least at times—a nearly paradigmatic Darwinian population and is thus highly susceptible to the influence of natural selection. And one of natural selection's inevitable by-products is the emergence of independent, refinable lineages. We call them “species” in biology. In science, they are called “specializations,” or “disciplines.”

I am going to argue that a line of inquiry's probability of significance—that is, of its making a lasting contribution to scientific knowledge—increases directly with natural selection's influence on the modification of scientific practice. The case for this claim develops along the following lines: first, I argue for the definition of “significance” stated in section 5.4—in particular, for the part of the definition that makes practitioners' esteem essential to a line of inquiry's significance. I then argue that the sorts of factors that often increase the probability of practitioners' esteem are closely analogous to those that increase the probability of selectively driven (ecological) speciation. What this will show is that the same properties that affect natural selection's influence over the emergence and endurance of new species also affect natural selection's influence over the emergence of significant

lines of inquiry. Were we to show this, we will have established that the probability of a significant line of inquiry emerging is directly tied to how susceptible scientific practice is to the influence of natural selection—that is, how closely it approximates a paradigmatically Darwinian population. This supports our contention that natural selection’s governance of the historical development of scientific practice is what accounts for its most epistemically significant attributes.

5.5.1 Achieving Significance in Science

We’ve already offered some discussion of the topic of *significance* in chapter 3. There we made two points that are relevant to the current section. First, we observed that a scientific community sometimes functions as a constructed niche, where the addition of new practices—be they new problems or new problem-solving approaches—is carefully controlled by the preferences of the scientific community. This suggests that, at a very basic level, a line of inquiry is significant when the community decides that it is significant. Indeed, “what better criterion than the decision of the scientific group could there be?” (Kuhn 1962b, 170). Because of the many possible directions in which inquiry could potentially proceed, the decision is not taken lightly. Where a community decides that a line of inquiry is worth pursuing, it has invested in the belief that its pursuit will result in a lasting contribution to scientific knowledge.

We also saw in chapter 3 how science funding agencies attempt to gauge the potential significance of proposed research by examining the strength, proximity, and number of connections it possesses with ongoing scientific research. Although this gauge does not necessarily tell us what significance is *per se*, it does have two important consequences. First, it provides a constraint on a viable theory of significance—namely, that any philosophical theory of significance should imply that significance is something that can be reliably measured by examining a line of inquiry’s connections with ongoing research. Ideally, the reliable detection of significance through an examination of connections with ongoing research should be an obvious, natural consequence of a philosophical account of significance in science.

Second—and related to this constraint—it fits precisely with the Kuhnian notion that significance is an internal matter. Although a community employs some objective criteria when it comes to gauging significance,

those criteria are themselves ultimately rooted in the community's prior decisions about which lines of inquiry count as significant—that is, those lines that the community had previously decided would make a lasting contribution to scientific knowledge. This is entirely appropriate from a Kuhnian perspective. Because if (1) the community is better positioned than any other party to judge current significance, and if (2) the relation to current significance is the most reliable guide to future significance, then (3) the community is best positioned to judge future significance. But why, from the community's perspective, should a line of inquiry's probability of making a lasting contribution increase with its relatedness to lines of inquiry that have made, or are poised to make, lasting contributions already?

I think there are a variety of ways in which one could account for the perceived relationship between current and future significance. In keeping with earlier discussions in the book, I will try to explicate this phenomenon within Kuhn's general framework. One of the most powerful consequences that follows from the adoption of a paradigm is that practitioners are furnished with a guide for discerning which kinds of problems are important, which kinds of phenomena are important to understand, which facts will be "particularly revealing of the nature of things" (Kuhn 1962b, 25). In many cases, the paradigm will immediately implicate certain quantities directly, as Darwinian theory did with respect to the mechanisms of heredity, or as Newtonian theory did with the gravitational constant. In both cases, adopting the paradigm consisted partly of accepting that these are questions whose answers are critical for advancing our understanding of nature.

But it is part of the character of scientific inquiry that, as time goes on, new questions emerge out of the investigation of those whose significance was directly highlighted by the adoption of the paradigm. In the case of Darwinian theory, this includes the many questions now pursued by molecular biology—questions that, while directly relevant to the issue of heredity, could not possibly have been conceived of in the mid-nineteenth century. Here, practitioners converged on the significance of these questions because of their direct connection to questions whose significance had been previously implicated. Whether the question's significance is partly constitutive of a paradigm's adoption, or whether it is instead a downstream consequence of a paradigm's adoption, the fundamental

factor influencing perceptions of significance is how closely the focal question can be linked to issues that are significant according to the paradigm.

What this reasoning suggests is that to be significant *just is* to be significant according to the paradigm. Significance is, in Carnap's terms, an "internal question"; it makes no sense for members of a scientific community to ask whether the governing paradigm has correctly identified some phenomenon as significant (Carnap 1950). If what allows a group of practitioners to constitute a scientific community is the acceptance of a shared paradigm, and if accepting a paradigm implies accepting certain questions as significant, then one cannot both be a community member and challenge the significance of a paradigm-validated question. In other words, community members must, virtually by definition, accept the significance of paradigm-validated questions.

This, then, would explain both why there can be no higher standard than the assent of the relevant community and why communities judge potential for significance on the basis of relation to past significance. If being significant just means being paradigm-validated, and if paradigm-validated significance only becomes paradigm-validated through the assent of the community, then a question can only achieve significance if the community is convinced that it is significant according to the paradigm. To say that there is no higher standard than the assent of the relevant community is just another way of saying that there is no higher standard than the paradigm. To the extent that we conceive of practitioners as epistemic niche constructors, things could hardly be otherwise.

There is no prescribed route by which practitioners must become convinced of a problem's significance. Its significance may be immediately obvious to all, given its centrality to or ubiquity within paradigm-driven research. Or practitioners may gradually, independently, and quite on their own come to see the relevance of the problem to their ongoing research, whose significance (of course) can be derived from its relation to paradigm-validated questions. Or, as I will explore at length in part II, practitioners may come to be convinced of a problem's significance through explicit argument aimed at establishing that its significance follows by implication from the paradigm. In the final analysis, all routes seem to result in the same basic effect: the community accepts the problem's significance, and it thereby becomes a question worthy of pursuit.

5.5.2 How Significant Problems Give Rise to New Specializations

Once a community of practitioners recognizes the question as a potential object of legitimate scientific inquiry—that is, as a paradigm-validated question—research on that question becomes part of the broader communal effort to understand nature. At this stage, the community is ready to add a *line of inquiry*. But it is not clear yet whether an entirely new *specialization* will emerge. Whether the new object of inquiry will give rise to a new specialization depends on the same sorts of factors that determined whether a source of potential nutrition would ultimately give rise to a new species.

The disciplinary branching process described below can be understood as a special case of Pagel and Mace's (2004) speciation-inspired model of the production of human cultural diversity. In that model, branching is driven by two primary influences:

One is the drive to secede from larger groups whenever possible, the better to control some defensible resource; this is what gives rise to the geographical patterns of diversity. The second set of forces is social and behavioural. They maintain cooperation within groups and create cultural identity and coherence, causing barriers to gene flow and meaning that vertical cultural transmission dominates. (Pagel and Mace 2004, 278)

In science, the “defensible resource” is a certain supply of scientific problems. The “social and behavioural [forces that] create cultural identity and coherence” are the replication and isolation mechanisms that allow scientific communities to persist over time. In the remainder of this chapter, I show how these two forces conspire to produce a version of the Pagel-Mace cultural branching process in science.

Whose job is it to solve the new and important problem that stands before the community? We can think about this question in the context of the grape- and toast-eaters we met earlier in the chapter. Whose job is it to eat the toast? Presumably everyone would *like* to be able to eat the toast. But in that population, only certain variants are equipped with the capacity to metabolize toast, and so the toast bonanza is left for them alone.

Toast vs. grape is an extreme case; the canonical Darwinian process is not one in which qualitatively new abilities arise in a single generation. But it is not an impossibility, and it illustrates an important evolutionary lesson—namely, branching events are most likely to occur when there are no intermediate phenotypes (Bolnick and Fitzpatrick 2007, 469). When

toast-eating ability is something that one either has or doesn't have, it becomes possible for two selection regimes to emerge, one pushing some variants toward a toast-eating optimum, the other pushing some variants toward a grape-eating optimum. In contrast, if each member of the population can eat both toast and grapes, there will be just one selective process (at least, with respect to toast- and grape-eating)—the selective process that drives the population to the optimum toast-grape balance (Bürger 2005).

The same lesson applies in the context of a scientific community's new research problem. Even if the new problem is important enough to warrant a new line of inquiry, it might not be distinct enough to constitute a new kind of research problem; it may simply be a hitherto unencountered version of the kind of problem practitioners are adept at solving. Indeed, enormous amounts of intellectual energy are often spent trying to transform queer new problems into the kind of problem for which we have well-worn methods of solving. If the new problem is something that could be solved with the community's existing set of practices, or with a variation on existing practices, then the emergence of a new specialization is improbable. When an existing problem-solving method (or some variation thereon) can be applied with success to the new problem, no practitioner can claim privileged access to the methods required to solve the problem. Anyone who can add even numbers together can just as easily add odd numbers together.

We can put the point in evolutionary terms: If the methods of solving the new problem are close enough in practice space to reach via a sufficiently plastic application of existing methods, there won't be the sort of disruptive pressure on the community that is required to split it into ecologically distinct subgroups. But if qualitatively distinct problem-solving methods are required, an autonomous new specialization may be able to emerge. For, at that point, each practitioner will have to choose to acquire one set of methods or another.

Now, because disruptive pressure emerges as a result of the unreachability of certain points in practice space, and because problems are themselves an element of practice, it ought to be the case that sufficiently distinct research problems are capable of driving the emergence of new specializations all by themselves. That is, even if a new problem is solvable with existing methods, it still might be capable of fueling the emergence of a new specialization. This phenomenon is easiest to illustrate in the context of very general problem-solving methods like mathematics and statistics.

Two statisticians who receive exactly the same training may nevertheless find themselves in two different specializations; one might be, say, a data scientist, while the other is a biostatistician. Despite drawing on the same problem-solving toolkit, the problems involved are themselves distinct enough to be unrelated. Progress in data science could easily have no bearing on progress in biostatistics, because there might be very little overlap in their respective domains. This would be a version of the model of ecological divergence discussed above that uses spatial considerations to interrupt gene flow: a population can branch into two distinct species, so long as there are two subgroups whose territorial overlap is sufficiently small.

Mulkay (1975) observed that the introduction of new technologies can have a similar disruptive effect. Given an instrument whose competent operation is sufficiently recondite, it can sometimes happen that entire specializations grow around the use of that technology per se. This is becoming increasingly common, as the variety of tools necessary for answering certain kinds of questions becomes increasingly disparate. Under these circumstances, it will often be impracticable to have a single practitioner develop competence in the use of a certain technology in addition to mastering the subject matter to which that technology is to be applied.

In sum, as with biological populations, the ability of natural selection to drive the emergence of a new autonomous scientific specialization is directly related to how specialized one needs to be in order to extract a benefit from the hitherto unexplored domain. Whether it is a new problem, a new technology, or a new theoretical framework, the less common ground there is between practitioners who engage in one practice and practitioners who engage in the other practice, the more probable it is that the single community of practitioners will fissure into two specializations.

In section 5.4.1 we saw how the magnitude of the reproductive gain provided by toast would have to exceed that provided by grapes in order to lure toast-prone variants away from grape-eating. When the benefit of toast is less than that of grapes, toast-eating is selected against; if it is equal to grapes, toast-prone variants will be indifferent between toast and grapes, and the probability of the lineage splitting is minimal. As the benefit of eating toast increases, natural selection will increase its bias in favor of toast-eating, and the prospects for permanent ecological disruption begin to rise.

What should we expect in science in light of this dynamic? Let us begin with the observation that, *ceteris paribus*, the probability of ecological disruption increases as the expected utility of improving one's grape-eating ability declines. As grape-eating optimization becomes less beneficial, the relative benefit of eating toast will climb, causing increased selection for toast-eating. This suggests the following principle for the emergence of new scientific specializations: *the probability that a new specialization will emerge varies inversely with the marginal utility of progress in an established specialization.*

Earlier efforts to analyze the growth of scientific knowledge describe the emergence of new specializations in precisely these terms. Mulkay's (1975) account, for example, begins with the following observation:

The growth of a new area starts with the perception, by scientists already at work in one or more existing areas, of unsolved problems, unexpected observations or unusual technical developments, the pursuit of which lies outside their current field. Thus the exploration of a new area is often set in motion of a process of scientific migration. . . . Scientific migration is not a random process, for the scientists moving into a new field tend to come from other areas with specifiable characteristics. In particular, they come from research areas which have experienced a pronounced decline in the significance of current results; from areas where there are few or no avenues of research easily available; from areas whose members have special competence in or knowledge of techniques which appear to have wider application; and from areas which have been disrupted, often by events originating outside the research community, and whose members have consequently no firm commitment to an established field. (Mulkay 1975, 5)

A similar dynamic is depicted in the graph produced by Holton (1962) (figure 5.1). Holton analyzed the emergence of new specializations in terms of the change in the "number of undiscovered interesting ideas" as a function of time. Although this is not, strictly speaking, the same thing as "marginal utility of progress," both notions attempt to capture the same idea: the probability of making major progress in a given specialization declines with time. It will be noted that, in his representation, branching events occur right at the point where the slope of the "parent" curve starts to level off—that is, at the point where the probability of making major progress begins to decline. As we will see in part II's case study, the branching point for paleobiology as a new specialization occurs at the precise moment when the magnitude of contributions to evolutionary theory starts to plateau.

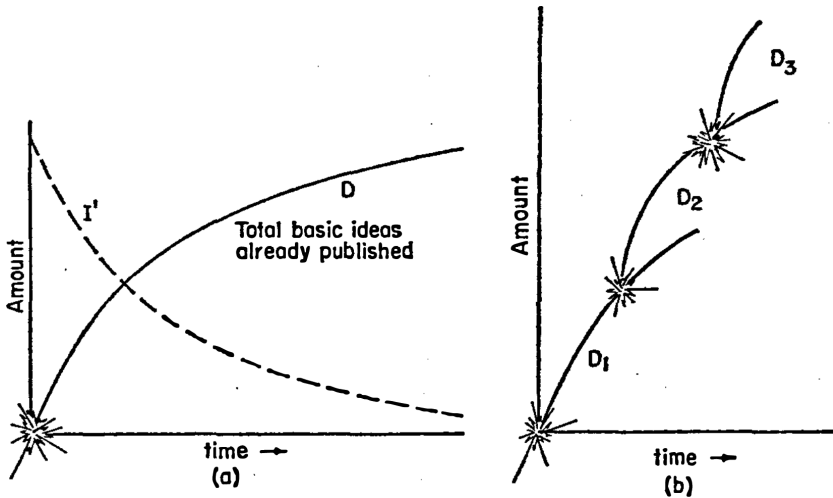


Figure 5.1

The escalation of discovery lines. From Holton (1962).

Both Holton and Mulkey discuss the emergence of specializations as a branching process, one that explains how scientific knowledge grows over time. Mulkey, for example, concludes his account by stating, “This account of the processes of scientific development emphasises the way in which science grows through the *branching of new lines of research*” (Mulkey 1975: 5–7; emphasis added). Holton (who literally uses a drawing of a large, leafy tree to illustrate the process) concludes that “the growth of scientific research proceeds by the *escalation of knowledge*—or perhaps rather new areas of ignorance” (Holton 1962, 126). For Holton, “escalation” occurs through “the applicability of the early techniques, and the originating of new questions now suggesting themselves in neighboring parts of the same fields, [which] provoked a rapid branching into several new directions” (118).

The key point for our argument lies in the fact that, according to these models (and soon to be corroborated by part II), new scientific specializations often appear to arise at the exact moment at which natural selection becomes most effective—that is, when the marginal utility of progress in an established discipline goes into decline. Once the magnitude of the benefit of making headway in a given specialization starts to level off, practitioners who “have special competence in or knowledge of techniques which

appear to have wider application” begin to look for opportunities to opt out of the race to squeeze the last few morsels of progress out of a declining specialization. These opportunities can be found most easily in research areas that offer benefits comparable to those available in their current specialization but have yet to see an influx of capable researchers.

This process, in turn, raises the related issue of whether the hitherto unexploited source of problems is sufficiently abundant to support a group of practitioners large enough to propagate itself. On the face of it, this problem is just that of toast-abundance, instantiated in a scientific context: it doesn’t matter how important the new type of problem is if there are not enough tokens to allow the community to grow. But scientific problems can behave somewhat differently in the sense that working on them will often generate more problems, such that the growth of knowledge can often really properly be characterized as a steady decrease in the ratio of solved problems to unsolved problems. This is what Holton alludes to when he describes the growth of scientific research “by the *escalation* of . . . new areas of ignorance.”

The upshot of this important difference in how the supply of problems behaves is that problem supply will rarely affect natural selection’s power to modify scientific practice to the same extent that nutrient or mate supply will affect selection’s power to modify biological populations. Even if opportunities in another domain are initially limited, their pursuit can often result in an increase in the supply of problems, thus laying the foundation for the migration of a far greater number of researchers, in Kuhn’s words, “away from competing modes of scientific activity.” Thus, at least with respect to resource abundance, the influence of natural selection can be counted on to be less constrained in scientific contexts than in biological ones.

The scientific community’s ability to restrict the inflow of new practices was discussed extensively in chapter 3, as a way of arguing that new practices often become widespread in science chiefly by virtue of the goodness of fit between those practices on the one hand and the epistemic challenges that scientific communities set for themselves on the other. Thus, much of what I will have to say about the flow of variation per se has already been argued. Here I will simply highlight a few points about the isolation of communities specifically as it pertains to the emergence of new specializations.

As Kuhn made clear at various points, the isolation of a scientific community was the most important factor for promoting its ability to specialize. By definition, specialization requires limiting the scope of inquiry and

the tools employed therein. The normative constraints that paradigm adoption places on acceptable problems and methods achieve this isolation. But Kuhn's view that isolation was caused primarily by lexical divergence was overly narrow (see chapter 3 for a critique). If Kuhn's view were correct, there should be far fewer specializations than there in fact are, since it is false that the only way to effectively isolate practitioners is through communication breakdowns. But this is an entirely separate matter from that of whether specialization accounts for the directional refinement of scientific practice, which it arguably does.

Because, *pace* Kuhn, scientific communities have so many mechanisms for isolating themselves from the influx of migratory practices, they can often be rather ideal systems for generating new specializations. Unlike biological populations, which are usually at the mercy of insurmountable physical obstacles to prevent gene flow, scientific communities can police the flow of practices by arbitrary fiat. This gives each community virtually unlimited power to refine methods, problems, and the like, solely in response to the epistemic challenges they choose to confront. The freedom to pursue their own epistemic ends in relative isolation accelerates the divergence of subgroups of practitioners from one another, eventually resulting in the creation of independent, autonomous research communities. I should emphasize that Kuhn largely embraced this image of the branching process, the one difference being that he seemed to think isolation was more difficult than it probably is.

The goals of this section have been (1) to present an interpretation of what it would look like for natural selection to drive the emergence of new scientific specializations and (2) to argue that selection's actual influence on the emergence of new specializations is probably significant in some cases. The argument for (2) is based on two premises, the first of which was our interpretation of adaptive branching in scientific contexts. This interpretation consists of drawing a direct connection between the variety of conditions that independently contribute to selection's influence over the branching process in biological populations on the one hand and the factors thought to be of the most general causal relevance to the emergence of new scientific specializations on the other. The second premise is simply that this connection is no accident. In other words, the reason why factors that often cause the emergence of new scientific specializations are directly analogous to

factors that amplify natural selection's role in the emergence of new species is because the emergence of new scientific specializations is often driven by natural selection. Chapter 4 emphasized how the process of refinement of scientific practices is analogous in character to the process of refinement of structures by natural selection. We now see that effect of selection on the growth of scientific knowledge is equally important. When selection's role in scientific inquiry is less pronounced—that is, when scientific practice is somewhat less than a paradigmatically Darwinian population—the practitioners' practices will tend not to diverge from one another with the pace, magnitude, and shape that they characteristically exhibit.

5.6 Conclusion

Under the overwhelming influence of natural selection, the emergence of new scientific specializations is as inevitable as the emergence of new species. Now, if the influence of natural selection on the development of practice makes new specializations inevitable, and if scientific practice becomes more susceptible to selection's influence as its approximation of a paradigmatically Darwinian population increases, then the fact that scientific practice is a nearly paradigmatic Darwinian population would be a very good explanation for its unrelenting propensity to generate new specializations.

We are then forced to take even more seriously the philosophical claim that part of what gives modern scientific inquiry its epistemic power is the fact that its development over time seems often to be dominated by the influence of natural selection. As with the directional refinement of practices, the growth of scientific knowledge is a datum with which any philosophical account of scientific knowledge must reckon. If, as I have argued, the perpetual growth of scientific knowledge consists of the perpetual emergence of new scientific specializations, then the growth of scientific knowledge is inevitable if scientific practice resembles a paradigmatically Darwinian population. There is no available alternative theory that makes science's two most significant epistemic features an inevitable consequence of the process of scientific inquiry under a few specifiable conditions.

And yet there is a further story to tell. So far, our discussion has remained very close to two traditions of thought. At its most general level, our goal has been to shed philosophical light on the historical pattern of development in some instances of scientific practice, and we have relied largely

on Thomas Kuhn's descriptions of the various elements of that historical pattern. There have been a few instances in which it seemed necessary to amend Kuhn's views to a degree, but none of those emendations called into question a broadly Kuhnian perspective on what the historical pattern of scientific development actually looks like in certain contexts. Parallel to this discussion have, of course, been our repeated appeals to the tradition of thought known as Darwinian evolutionary theory, whose wide exploration was made possible by our attempt to follow the implications of thinking about scientific practice as a Darwinian population.

These two traditions have, up to now, lined up very nicely indeed—much more tidily than I would have ever dreamed when I first began thinking about these issues. But there are major features of each tradition that we have thus far neglected. It will no doubt have occurred to the reader that, even though Kuhn's model is specifically intended to describe the process of revolutionary scientific change, I have made virtually no mention of any such change, nor have I attempted to find an evolutionary analogue to scientific revolutions.

Does my silence on this matter stem from skepticism about the existence of Kuhnian revolutions? No. I am mostly convinced that the kind of revolutionary change described by Kuhn—incommensurability and all—actually exists. Is it because scientific revolutions lack epistemic significance? Again, no. They are as much a part of the historical pattern of scientific development as either directional refinement or branching. More importantly, major revolutions are viewed by practitioners as epistemically unique, in that there are no other occasions in which so much scientific progress is made in so little time.

However, in contrast to the other epistemically significant phenomena we've analyzed, major scientific revolutions do not follow inevitably from the extent to which scientific practice is a paradigmatically Darwinian population. That is, a population of practices that is maximally susceptible to modification by natural selection does not thereby increase its chances of undergoing a major scientific revolution. All that being a paradigmatically Darwinian population guarantees is that specializations will emerge and that they will be directionally refined through a process Darwin labeled *adaptation* and Kuhn labeled *normal science*. And although Kuhn clearly thought that normal science was a necessary condition for a scientific revolution (more on that in chapter 6), major revolutions require more than

normal science. And that extra stuff does not follow as a consequence of paradigmatic Darwinian populationhood.

Which brings us to the feature of evolutionary theory that we've thus far neglected: extinction—mass extinction, in particular. Some of the most significant developments in empirical and theoretical work on evolutionary change since the publication of the *Origin* have been the very recent growth in our understanding of the nature of mass extinction and its fundamental role in shaping evolutionary history. But the subject has yet to come up in this book because, like scientific revolutions, mass extinctions are not an inevitable consequence of the Darwinian evolutionary process. A paradigmatically Darwinian population is no more or less susceptible to a mass extinction event than a paradigmatically un-Darwinian population is.

In chapter 6 I argue that these two phenomena—major scientific revolutions and mass extinctions—are two names for the same thing: a major scientific revolution is a mass extinction of scientific practices. Hence, their simultaneous neglect has not been a coincidence. We will then have to confront the epistemological question of whether major scientific revolutions can be part of what accounts for the distinctive power of modern science, given that they do not follow from the fact that scientific practice is a Darwinian population.

This is a section of [doi:10.7551/mitpress/14461.001.0001](https://doi.org/10.7551/mitpress/14461.001.0001)

How Knowledge Grows

The Evolutionary Development of Scientific Practice

By: Chris Haufe

Citation:

How Knowledge Grows: The Evolutionary Development of Scientific Practice

By: Chris Haufe

DOI: [10.7551/mitpress/14461.001.0001](https://doi.org/10.7551/mitpress/14461.001.0001)

ISBN (electronic): 9780262371599

Publisher: The MIT Press

Published: 2022

The open access edition of this book was made possible by generous funding and support from MIT Press Direct to Open



The MIT Press

© 2022 Massachusetts Institute of Technology

This work is subject to a Creative Commons CC-BY-NC-ND license. Subject to such license, all rights are reserved.



The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Stone by Westchester Publishing Services, Danbury, CT.

Library of Congress Cataloging-in-Publication Data

Names: Haufe, Chris, author.

Title: How knowledge grows : the evolutionary development of scientific practice / Chris Haufe.

Description: Cambridge, Massachusetts : The MIT Press, [2022] | Includes bibliographical references and index.

Identifiers: LCCN 2021057609 | ISBN 9780262544450 (paperback)

Subjects: LCSH: Science—Philosophy. | Science—Methodology. | Knowledge, Theory of.

Classification: LCC Q175 .H35198 2022 | DDC 501—dc23/eng20220405

LC record available at <https://lcn.loc.gov/2021057609>