Dappled Science in a Unified World

Michael Strevens

To appear in *Philosophy of Science in Practice: Nancy Cartwright and the Nature of Scientific Reasoning*, edited by Hsiang-Ke Chao and Julian Reiss, Springer.

Abstract

Science as we know it is "dappled". Its picture of the world is a mosaic in which different aspects of the world, different systems, are represented by narrow-scope theories or models that are largely disconnected from one another. The best explanation for this disunity in our representation of the world, Nancy Cartwright has proposed, is a disunity in the world itself: rather than being governed by a small set of strict fundamental laws, events unfold according to a patchwork of principles covering different kinds of systems or segments of reality, each with something less than full omnipotence and with the possibility of anomic indeterminism at the boundaries. This paper attempts to undercut Cartwright's argument for a dappled world by showing that the motley nature of science, both now and even at the completion of empirical inquiry, can equally well be explained by proponents of the "fundamentalist" view that the universe's initial conditions and fundamental physical laws determine everything that ever happens.

1. Fundamentalism and Unity in Science

Everything is made up of a single kind of physical stuff and everything that happens is directed solely by fundamental laws of physics that, depending on the configuration of stuff at one moment, determine its configuration at the next. Thus, all the complexities of life, all human striving and failing, all happenings throughout the universe's spatial and temporal extent, can be predicted from, and understood in terms of, the physical makeup of the things involved as acted upon by the fundamental laws. This is the doctrine of *fundamentalism*.

Fundamentalists have had to loosen or revise the description given in the previous paragraph. "Physical stuff" must be understood broadly enough to include the values of quantum fields at space-time points. The conception of time as a sequence of moments must be made more sophisticated, and the notion of temporally asymmetric dependence—that the state of the universe at earlier times determines its state at later times—has been called into question. Perhaps most important, it has to be allowed that the fundamental laws are stochastic. In that case, they perhaps do not predict and explain everything, but they predict and explain everything predictable and explicable.

Even as these allowances have been made, the empirical evidence for fundamentalism has accumulated swiftly, so that in the middle of the twentieth century logical empiricist philosophers of science were able to proclaim as a central working hypothesis the "unity of science", a doctrine premised on a version of the fundamentalist world picture adjusted to accommodate modern physics as just described.

Fundamentalism proclaims a unified structure for the world, not for science, but it is natural to suppose that it has immediate implications for scientific methodology.

One such implication concerns the structure of completed science: a comprehensive representation of the structure of the world would articulate the fundamental laws of physics, describe at least schematically how complex objects such as viruses, brains, and democracies are built from fundamental physical stuff, and derive from the fundamental laws the characteristic behaviors of these higher-level systems, so predicting and explaining what they do, to the extent that it can be predicted and explained.

The universities at the end of inquiry will contain more than just a department of physics: there will be departments or sub-disciplines for each structure or kind of structure, specializing in the derivation of behavior in each case. Further, the majority of the members of these departments will likely know little or nothing about physics, as they will specialize in deriving behavior from laws of behavior one level down—laws of economics from laws of psychology, say. This is more or less the picture sketched by Oppenheim and Putnam (1958), with each unit of specialists bridging adjacent tiers, or more likely small parts of adjacent tiers, in the great layer cake of being.

Perhaps the cake is not the best metaphor, however, as it has interesting structure only in the vertical dimension, whereas our world is, even at a single level—at the level of organisms, of minds, of cultures, or whatever—extraordinarily diverse. This horizontal structure, too, scientific unifiers wish to bring together into an integrated whole. Fundamentalism does not guarantee horizontal unification in quite the straightforward way that it guarantees vertical unification, but the fundamentalist world picture holds out hope all the same, either of a grand derivation pattern that yields all parts at a given level, or of a causal unification that ties the parts together into a single narrative.

A metaphor for a completed science ought to have structure in every direction, then, but it ought to bind that structure into a totality driven by a single underlying principle. The best choice is perhaps a shopping mall. It lacks something, however, in poetry. A university? Not metaphorical enough. Let me suggest the Hindu and Buddhist conception of a mandala, a complex but unified geometric figure that represents the universe in its totality (figure 1). Fundamentalism's implication for future science—for science at the end of

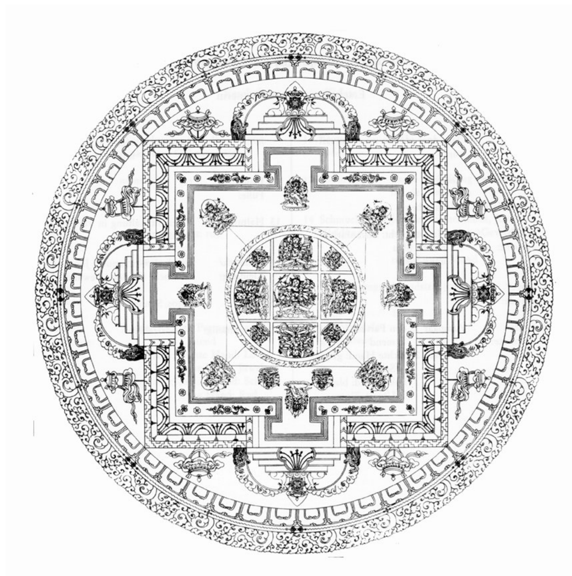

Figure 1: Guhyasamaja mandala (Tibet)

inquiry—is that it has, in its intricacy, the order and integration of a mandala.

The other methodological implication of fundamentalism is for science now: that it should aim at the ideal of the mandala, seeking to amalgamate islands of expertise into continents of systematic theory, and then deriving those theories in steps from the fundamental theory to create the unified whole.

That science should not yet have succeeded in drawing the mandala is hardly surprising. Many departments of the university have, after all, only gotten underway in the last century; it took the department of physics three centuries to discover and then to replace Newtonian theory, and it is still working on some important details.

But scientists should be, at the very least, pursuing integration, expressing dissatisfaction where connections between their own theories and others nearby are missing or incomplete and working to develop them. Derivation from the next lower level should be always on their minds, for its own sake and because of the relationships it reveals with the neighbors.

Unity is as much aspirational as actual. How do those aspirations stand?

2. Disunity and the Dappled World

Real science is not only largely disunified; it is largely content to be disunified. That is what I have learned from the work of Nancy Cartwright, who argues persuasively that, for every E. O. Wilson or Steven Weinberg painting majestic pictures of unification in the popular press and urging their colleagues to put together the pieces, there are tens of thousands of scientists happy to focus on their own limited canvases with no ambition to follow some grand blueprint. To put it more plainly, in most areas of research, most scientists spend their time building models with tightly circumscribed scope—models that only apply, and are only intended to apply, to a small set of behaviors of a small set of systems. Deriving those models from lower-level principles or integrating them into a broader theory is far down the agenda, and perhaps not on the agenda at all. This is the disunity of science.

Cartwright at first developed the picture of science as disunified with reference to physics—not fundamental physics, but the physics of structured systems such as lasers. She later turned to economics, where wide scope and accuracy seem to have great difficulty cohabiting—where, as Cartwright puts it, *useful* models tend to apply only to a narrow range of economic structures (Cartwright 1999, 2003, 2007). I will make the same point with respect to evolutionary theory in section 5.

Not every scientist neglects unity; indeed, a drive to unity is apparently a key motivation in some scientific sub-specialties, such as cognitive neuroscience. The popular press is full of calls to unity. But real science, as it unfolds in a million labs around the world, seems on the whole to be rather indifferent. Real science is disunified.

This disunity, this manifest contentment with science as a mosaic rather than a mandala, issues a challenge to the unity of science and thence to fundamentalism. The challenge to unity: do scientists see something that we philosophers do not see, which has led them to abandon unity as an ideal? Or have they encountered, over and over, difficulties in establishing intertheoretic connections that have led them to abandon unity as a practical goal? The challenge to fundamentalism: is what scientists see, or what they have inadvertently run into, a world structure that is profoundly different from the fundamentalist structure?

To both questions Cartwright answers yes, concluding that the world we live in is not fundamentalist but "dappled": it contains small islands of regular behavior, each orderly in its own way, floating in a sea of complexity. There are no fundamental laws of physics that predict and explain what goes on either at sea or on the vast majority of the islands. What we call the fundamental laws apply only to a few small islands of their own, consisting in the behavior of fundamental particles in simple environments such as oil-drop experiments or particle accelerators and their collision detectors.[1]

The Cartwright argument for a dappled world works by demonstrating disunity and then reversing the logic of the advocates of scientific unity. Because fundamentalism is true, the unifiers argue, we have the opportunity and motivation to pursue a unified science. But scientists apparently lack either one or both of opportunity and motivation, Cartwright argues, and so fundamentalism is false.

The aim of this paper is to provide an explanation for modern science's disunity in fundamentalist terms. Given what we know already about the workings of the world, I will show, we ought to expect that the majority of scientists will find great difficulty in building connections to other theories, and that they will see little motivation to do so. It is not merely hard to find the connections in most cases—it is unrewarding.

1. In earlier work Cartwright seemed content to argue that it is a real possibility that the world is dappled, that the nomological mosaic is a prospect worthy of serious philosophical consideration. Now, however, she is happy to talk about her "belief in the dappled world" (Cartwright 1999, 1).

* * *

I can see five explanations of disunity compatible with the fundamentalist world structure.

The first explanation is simple: scientists are too dull or too blinkered or too professionalized to care about the project of unification. They could and should be drawing connections between models and theories, but they are busy publishing snippets of data. Plausible? There might be something to this acidulous characterization, but there are more interesting accounts of disunity to be found.

According to the second explanation—an optimistic variant on the first?—the work of unification requires only a few inspired thinkers. It is right and proper that 99% of scientific effort is expended on local, disconnected projects; the 1% will weave it all together into the mandala. This is a possibility only if unification is not much work. But few unifiers, I think, believe that. The best fundamentalist explanations of disunity assume quite the opposite.

The three remaining explanations—from complexity, antireductionism, and contingency—will be discussed at greater length, with contingency playing the leading role.

3. The Explanation from Complexity

The explanation from complexity is straightforward in its general form: much unification, though possible in principle as fundamentalism dictates, is rendered impossible in practice by the complexity of the necessary derivations. The mandala can be drawn, but it is too difficult for us to draw.

That is surely the correct story to tell about some pieces of science. Quantum chemists would very much like to develop accurate models, based on fundamental physics, of atomic and molecular structure. What impedes them is the fiendish mathematics involved; nevertheless, they strive to build approximation techniques and faster algorithms to do the job better and better. There is a temporary failure of reduction wherever those models fail to replicate structural detail, but it is clear that scientists have every interest in bridging such gaps; further, their steady progress suggests that there are no in-principle obstacles to success (Hoefer 2008).

Complexity, in the form of sensitive dependence on numerous hard-to-measure boundary conditions, is also surely what accounts for our having no science of the trajectories traced by falling leaves or wafting banknotes—to take a famous example discussed by Neurath and Cartwright. Modeling the leaf is not only laborious but also unremunerative and uninteresting. Why expend the resources? So we don't.

When it really matters, however, scientists will devote a great deal of attention to problems of this form. Accurate weather forecasts are practically important in many ways; meteorologists consequently have the motivation and resources to press on with the reductive strategy despite the complications. They divide the atmosphere into smaller and smaller cells, they include more and more details of the processes occurring in those cells, such as type and variety of cloud cover, in many cases using models of the processes derived from fluid mechanics and other branches of fairly low-level physics, and they are rewarded with forecasts that are more and more accurate.

Push hard enough against the complexity, it seems, and nature will yield. But often there are not the resources to push hard enough; in these cases, you will see a science that leaves phenomena at different levels unconnected—not because there are no connections, but because they are too expensive to reveal. Science is, in these cases, disunified for economic rather than for metaphysical reasons.

A closely related reason for disunity, especially important in the historical sciences, is what I will call epistemic erasure: some of the facts necessary to carry out a unifying derivation have become, over time, buried or eroded, and are now extremely difficult to access.

Complexity of derivation, along with epistemic erasure, explains some of our disunified science's missing links. There is a certain kind of gap that it does not, however, address: an absence of motivation, a lack of interest, in closing the theoretical gap. If it is only complexity that stands in the way of unification, we should find in the sciences a regret wherever unity is too expensive to pursue. In the case of evolutionary theory, I will shortly argue, such regret is conspicuously absent. First, however, let me consider an explanation for disunity that does propose to account for the missing motivation.

4. The Explanation from Antireductionism

On behalf of fundamentalism, Hoefer (2008) appeals to philosophical antireductionism. His argument is brief, but it could be understood as follows.[2] Suppose, as many antireductionists argue, that there is a great deal of explanatory autonomy in the sciences. Among the consequences of autonomy, it is supposed, is that many or all low-level facts are explanatorily irrelevant to facts at a sufficiently high level, even in cases where the high-level facts can be derived from the low-level facts—as exemplified by Philip Kitcher's claim that, for the purposes of understanding independent assortment, "it's irrelevant whether the genes are made of nucleic acid or of Swiss cheese" (Kitcher 1999, 200). If that were the case, scientists would have a principled reason to ignore the connections made available by the fundamentalist world structure. The connections would be there for the taking, but they would not be taken. Disunity would be the methodological rule, despite the world's metaphysical unity.

I myself cannot seriously believe that a derivation of a phenomenon from lower levels, or even from the fundamental level, can be entirely explanatorily

2. I am unsure whether Hoefer's antireductionism denies the possibility of reductive unification in principle, in practice, or merely the explanatory relevance of the enterprise. At one point he writes (p. 317) "I suspect most fundamentalists have no wish to argue that such a reduction is possible, for us at least", suggesting the intermediate view, with the delineation of the mandala perhaps blocked by complexity, but it is the last of the three possibilities that I attribute to him in the main text.

irrelevant to that phenomenon—that there is literally no insight whatsoever to be gained from understanding the lower-level implementation of higher-level patterns of behavior.

There is a weaker sort of irrelevance, however, that certainly does play a part in accounting for apparent disunity in science. Science, like the rest of the university, operates according to a grand division of labor. The biologists do the biology; the physicists do the physics; the economists do the economics—and within biology, the evolutionary biologists, the ecologists, the molecular biologists and so on divide the scientific to-do list into smaller pieces. In the context of this division, what the physicists do is irrelevant to the biological enterprise just because the biologists have agreed to leave it to the physicists.

Insofar as the division is temporary—insofar as there is, notionally at least, a plan to convene an assembly at the end of inquiry to stitch together the patches of knowledge woven by science's departments and sub-departments (with many interim conferences, presumably, to be held before that time)—the atomization of scientific inquiry for these practical reasons is not genuine disunity, not a riposte to the advocates of the unity of science.

It has, nevertheless, considerable power to explain the extent to which the different divisions of science ignore one another, and in particular the extent to which the higher-level sciences ignore questions of lower-level implementation. Why, for example, do evolutionary theorists explaining sex ratios pay little attention to the mechanics by which genetic material is transmitted from one generation to the next, though the relative amounts transmitted by males and females stand to make a significant difference to sex ratios (Kitcher 1999)? Simply because that is someone else's job, someone far better qualified to do the research. The evolutionary and genetic research programs appear to be epistemic islands—they have their own conferences, journals, agendas, explanatory styles—but this segregation of production in no way implies disunity of the product. Indeed, by making possible efficiencies of specialization and scale, it only speeds the arrival of an integrated science of life.[3]

The antireductionist explanation, strong or weak, is well able to explain scientists' lack of concern for, and thus science's (temporary) lack of, vertical integration. What it does not account for is a horizontal gap, an absence not of reduction but of amalgamation.

The explanation from complexity and epistemic erasure, then, accounts for both vertical and horizontal gaps but not for the lack of will to fill them. The explanation from antireductionism explains the apparent neglect of the vertical gaps. What is needed is something that does the same for the horizontal gaps.

5. The Mosaic Aspect of Evolutionary Theory

Evolutionary theory has at its core a small set of tenets: that all life on Earth shares a common ancestry; that the primary agent of adaptive change is natural selection; that adaptations can to a large degree be passed down from parents to offspring. Along with the tenets comes technical machinery of similarly broad scope, most notably the mathematics of population genetics.

At first glance, it might seem that the work of evolutionary biologists is to develop models of specific types of evolution wholly or at least largely derived from an enhanced and sophisticated version of the core tenets, in the same way that the work of physicists is (apparently) to develop models of specific physical processes derived from the fundamental physical laws. The core tenets would in that case unify evolutionary theory both in content and in method.

And yet, though the core merits its name, it is not the grand unifier that you might suppose. To see why, consider the structure of what I will call a *simple selection model*: an evolutionary model describing and explaining the replacement of one variant in a population by another fitter variant, or the

3. This argument is developed at greater length in Strevens (2016).

attainment of an equilibrium between the two variants. A simple selection model derives its logical and mathematical structure from the core. The very idea of explaining a change in the relative frequency of the traits by appealing to fitness differences comes from the core; so too do strategies such as the use of population genetics to track the changes in the frequencies of alleles that cause the appearance of the traits under examination. (For simplicity's sake, I assume genetic determinism.) In its framework, then, the model is a faithful reflection of the core, and all models like it—all simple selection models—are in this respect unified by the core.

Much of the work put into simple selection models is not, however, unified at all. To see why, take a closer look—in the way that Cartwright has taught us—at the empirical content that transforms a mere framework or template into an informative model of a real evolutionary process. In a simple selection model, that content falls into two classes. First, there are facts about the genetic determination of the traits—which alleles, assembled into which genotypes, determine the appearance of what features. Second, there are facts about the differential rates of proliferation (determined by facts about viability and reproduction) of the two traits.

These facts are clearly not derived from evolutionary theory's core. There is no way that the core tenets, or any extension of the tenets, could entail the specifics of genetics and development in some particular organism or the specifics of that organism's survival and reproduction in some particular habitat. The facts constitute, therefore, what you might call *local* as opposed to core content: content that is specific to some particular evolutionary scenario or small group of scenarios.

Whereas core content, being common to a wide range of evolutionary models, is a unifier of evolutionary research, local content is a disunifier, splitting such research both methodologically and content-wise into many small islands, with each islet corresponding to a particular piece of a particular selective history of a particular life form.

To what extent, then, is an evolutionary model's functional content of the unifying sort, and to what extent is it of the disunifying sort? In a simple selection model, the local content alone—the genetics of the relevant traits and their differential contribution to rates of proliferation—entails pretty much everything that the model predicts or explains.[4] It appears, then, that the local content exhausts the model's empirical content; what the core supplies is purely logical and mathematical.[5]

Consider as a concrete example one of the best known models in population genetics, that of the maintenance of the gene *S* for sickled hemoglobin in populations exposed to malaria. Two copies of *S* cause sickle-cell anaemia, a disease that without modern medical care significantly shortens lifespan. Because the alternative to *S*, the normal gene *A*, has no ill effects, you would therefore expect the sickle gene to be expunged from the gene pool by natural selection. But in some populations, such as West Africans, it is quite common (around 20% of adults have one copy of *S*). How can that be?

The answer is that having a single copy of *S* and a single copy of *A*—having the heterozygous genotype *AS*—confers powerful resistance against the worst effects of malaria, especially the strain caused by the *Plasmodium falciparum* parasite. In populations where *P. falciparum* was rife, the heterozygote advantage was sufficiently strong to maintain the sickle gene at a certain equilibrium level in the population.

The location of the equilibrium point—whether what's optimal is to have 5%, 15%, 25%, or some other proportion of hemoglobin genes be of the *S* variety—depends on the details, such as the virulence of the parasite, the degree of resistance conferred by the heterozygous genotype, and the severity

4. More exactly, what is entailed is a high physical probability for these descriptive and explanatory consequences.

5. The core has some empirical content, most obviously the thesis of common ancestry. But in simple selection models, that content motivates a certain method without contributing to the method's products—it motivates a search for models of population change driven by natural selection, but it does not appear in those models as a working part.

of the harm caused by sickle-cell anemia. Given values for these parameters, a model can be constructed to predict and explain a certain rate of prevalence of the sickle gene and so of sickle-cell disease—in both cases, surprisingly high.

The local content of the model comprises the genetics of hemoglobin and the parameters just mentioned, such as the virulence of *P. falciparum*—facts determined by the specific physiological properties of the parasite and of its human hosts (as well as other environmental factors, such as the physiology and ecology of the mosquitos that transmit the disease). The core of evolutionary theory has nothing to say about this. But these facts alone are sufficient to determine the equilibrium point and so to predict the prevalence of the sickle gene and of the disease that it causes. Everything the model does, it does in virtue of its local content.

The sickle cell model belongs, in virtue of intellectual geography, in the evolutionary archipelago: it shares a certain framework, a cultural affinity, with the other islands in that vast group. But what is holding it up is the strictly local geology, the bedrock of physiological and ecological fact that pertains to its specific subject matter.

Evolutionary theory is, in a word, dappled: its models (or in some cases groups of models) are discrete islands of empirical content unified only by a formal resemblance. This is a horizontal, as opposed to a vertical, disunity, that is, a coexistence of many different micro-theories or models at a single level of description, each concerned with a different fragment of the same cross-section or plane of reality, brought together not by their hinging on the same empirical facts but by their taking a similar approach, formally speaking, to the separate fiefs over which they rule.

* * *

In her characterization of horizontal disunity, Cartwright has frequently made use of the notion of a ceteris paribus law: she suggests that the islands in the great archipelagos of science are each governed by laws with ceteris paribus clauses that give out at the shores, and so that are each informative about only a small patch of land.

Another way to say the same thing without attributing any special ceteris paribus structure to scientific generalizations is as follows: the islands of reality studied by evolutionary theory, by economics, by most of physics, are governed by *narrow-scope* laws, that is, laws or generalizations that for some reason or other cover the workings of only a minuscule portion, both horizontally and vertically speaking, of the world. Each island, then—in the evolutionary case, each venue for or episode of natural selection—is governed by its own narrow-scope laws, or if you prefer exhibits its own narrow-scope regularities, captured by the local content of the corresponding evolutionary model.

Why do the laws have only narrow scope? Why do they cease to hold off the island? Cartwright's ceteris paribus formulation suggests that offshore, there are too many interfering factors: the island provides a rarified atmosphere lacking much of the usual environmental noise. To put it another way, each narrow-scope law describes the operation of a mechanism that functions only if certain factors are absent, but such factors are almost universally present—thus, the law holds only in a few isolated systems.

This may be a good explanation of narrowness of scope in large swathes of physics, but I do not think it explains much biological narrowness. In the case of evolutionary theory, at least, the narrowness of models is to be explained in precisely the reverse way: the mechanisms described by the narrow-scope biological laws require for their operation the presence of special configurations of factors—special causal structure—that is quite rare. The physiology of the malaria parasite and the ecology of its habitat is a perfect example.

Cartwright recognizes this reason for narrowness too, and gives it a name: she calls the special configurations "nomological machines". I don't care for this term, as talk of nomological mechanisms operating side by side connotes spatiotemporally discrete pieces of hardware each doing their special thing in causal isolation. The behaviors described by evolutionary models, however, are neither spatially nor causally separate: at the same time that the malaria parasite is exacting its toll, its victims are living their rich and complicated lives, biological, social, and economic. The "islands" of evolutionary theory are not, then, spatiotemporal islands, but abstractions enjoying a more notional kind of separation.

The two explanations of narrowness are in any case complementary, a fact that Cartwright is happy to acknowledge in her characterization of ceteris paribus conditions. They specify, she says, "all the conditions necessary to [a nomological machine's] operation" (Cartwright 2003, 211); they include, therefore, both the required absences and the required presences. I think, then, that Cartwright and I are for the most part in agreement about the explanation for our worlds' regularities' narrowness of scope.

To summarize: the totality of biological reality—in its evolutionary aspect, but also perhaps you can see in its physiological and ecological aspects—is governed not by a single set of precepts encoded in evolutionary theory's core, but by a glittering mosaic of narrow-scope regularities, each dedicated to detailing, and to detailing only, a behavior caused by some particular facet of life's physical structure, such as plasmodium-human relations or the genetics of hemoglobin production.

Is that emphatically the end for the unity of science? Not exactly. The mosaic picture is logically consistent with the fundamentalist's creed, the doctrine that everything that happens is determined by the fundamental physical laws moving stuff around in accordance with the initial (or boundary) conditions.

The virulence of *P. falciparum*, the phenotypic effects of the *S* allele, and all the rest of every evolutionary model's local content, the fundamentalist will say, can be derived in principle from the configuration of the relevant physical structures and the laws of physics. Such derivations join the pieces of the mosaic into a mandala.

There are two derivations to be distinguished here. On the one hand, there is what might be called the historical derivation, which explains how the parasite came to have the virulence that it did by presenting the origin of the relevant physiological and ecological structures. On the other hand, there is the implementation derivation, which takes the structures as given and explains why those structures behave the way that they do.

Both derivations offer the opportunity for scientific unification. The implementation derivation links an evolutionary model with the underlying physical structure and law, and so promises to bring all of evolutionary theory together as a kind of applied physics—a vertical unification. The historical derivation, by contrast, promises a horizontal unification—it knits the discrete episodes represented by many evolutionary models into an encompassing biological narrative.

Let me focus on the historical derivation, thus the potential for horizontal unification. The fundamentalist claims that it is there for the taking. A dappled-world metaphysician such as Cartwright denies it, arguing: if there is the potential for historical unification, why is it not actualized? Why are there so many gaps between models, if there is both the means and the motivation to fill them?

In constructing a response, the scientific unifier has recourse to the three strategies enumerated above. The first strategy is to appeal to the complexity of the task (section 3): unification is possible in principle, but in practice it is too difficult. Too much information has been lost to history, or the computations needed to complete the story are too byzantine for our resource-constrained science. On this strategy the desirability of unification is upheld; what is lacking is the means to achieve the glorious goal.

In fact, however, evolutionary biologists experience little regret looking out over the waves lapping at their sandy shores. They are happy on their islands. There is no sense of a great unfinished task; the mosaic structure of evolutionary theory seems entirely satisfactory and natural. The complexity strategy alone cannot explain this cheerful acquiescence to the disunity of evolutionary inquiry.

The second strategy is designed precisely to explain the lack of motivation to complete the process of theoretical unification. It is the weak antireductionist move (section 4), on which the division of scientific labor lifts from the shoulders of some group of scientists any obligation to pursue the unifying derivations, even when they are eminently feasible.

The weak antireductionist strategy is well able to make sense of evolutionary biologists' lack of professional interest in implementation derivations. For what physiological reasons is *P. falciparum* so deadly? Answering that question is a job not for evolutionists but for microbiologists, epidemiologists, and medical researchers.

The strategy works less well with historical derivations and horizontal disunity more generally: whose job is it to explain the way that *P. falciparum* got the way it did, if not the evolutionary biologists'? Yet they are not doing it—and nor is anyone else. Equally, neither they nor anyone else seems very much interested in doing it. The mandala project sits idle. Such science-wide indifference is difficult for the unifier to explain; the credibility of fundamentalism suffers as a result.

And so, as foreshadowed above, we need a third explanatory strategy alongside weak antireductionism and complexity.

6. The Explanation from Contingency

Begin with an irresistible, and as far as we know quite possibly correct, story about *P. falciparum*'s virulence. As I have already noted, *P. falciparum* is the most deadly of the malaria parasites, and also the newest. Although the evidence is far from conclusive, it may have crossed over from gorillas to humans about 10,000 years ago. The sickle gene seems to have proliferated around the same time.

This hypothesis suggests, first, that *P. falciparum*'s deadliness has something to do with its recent and sudden arrival, human physiology not having had enough time to adapt to its means of infiltration and attack, and second, that the evolution of the sickle gene, or more exactly its coming to make up a significant proportion of the gene pool, was a reaction to *P. falciparum*'s lethality.[6]

There is not just a story to be told about the origins of *P. falciparum*'s virulence, then, but a really good story. In that case, why should evolutionary biologists decline to tell it? While I don't doubt that complexity, epistemic erasure, and the division of labor have some part to play in the explanation, I will suggest that the most important factor is contingency.

When a parasite crosses from one species to another, it might be far more deadly to its new host than to its old host, if it exploits some loophole that the new host has had no reason to plug. Then again, it might be far less deadly, or fail to prosper at all, if it relies on loopholes or tactics that simply do not apply in its new environment. The parasite does not know that it is making a trip, eventually, to a new species. It adapts itself to the species in which it starts out. Sometimes those adaptations make its pathological strategies especially effective in the new species, resulting in exceptional virulence, but at least as often the reverse is true. The ancestors of *P. falciparum*, then, were adapted to certain aspects of the gorilla liver, red blood cells, and immune system. As it happened, these (and perhaps other) adaptations made the microorganism especially effective in the context of the human liver, red blood cells, and immune system.[7]

Why? For what reason were gorilla and human physiology so aligned—similar in some ways, different in others—that a certain parasite that evolved in the one would prove especially deadly in the other? No reason at all, of

6. It is also of interest that the transfer seems to have occurred around the time of the Agricultural Revolution.

7. I should note that the sort of effectiveness that enhances virulence does not necessarily work to a parasite's advantage. I use the term "effectiveness", then, in a sense that is detached from biological optimality.

course. It was just a matter of chance.

By "chance", I do not mean physical probability; I have in mind, rather, what might be called happenstance, the notion that Aristotle tried to capture with his example of the person who goes to the market and "by chance" encounters a friend they were not expecting to see.[8] As Aristotle remarks, to say that such events happen by chance is not to deny that they may be causally explained—in modern terms, that they are determined by initial conditions and fundamental laws. It is rather to say that the causal explanation has a certain property, which I will gloss as contingency: the initial conditions could easily have been slightly different, and if they had been slightly different, things would have unfolded differently.

When an event's occurrence is contingent in this way, it has an explanation but it does not usually have an interesting explanation. Exceptionally important events can have explanations that are, because of the contingency of their parts, exceptionally uninteresting.

The young king Alexander of Greece was bitten, in 1920, by a monkey, and died from an ensuing infection. A tentative alliance between the monarchy and the prime minister Eleftherios Venizelos consequently fell through, and as a result Venizelos' liberal government was unexpectedly defeated in the elections shortly thereafter. Had Venizelos remained in power, it is arguable that Greece would not have pressed on with its invasion of the nascent nation of Turkey, an invasion the redoubling of which resulted in the "Great Catastrophe", costing the lives of many Turks and perhaps hundreds of thousands of Greeks. Winston Churchill conjectured famously that perhaps "a quarter of a million persons died of this monkey's bite" (Fromkin 1989, 432).

The Great Catastrophe and the deaths of the quarter million have a gripping explanation. But at least one component of that story has, in turn, an

8. *Physics*, II.4. Aristotle's chance occurrences happen only to beings capable of choice (II.6), so his notion does not apply to biological happenstance. He introduces another notion, of spontaneity, that might apply, though it has a somewhat different sense than my "chance" or "happenstance".

explanation that is devoid of significance. How did Alexander's altercation with the monkey end in his being bitten? Why did that particular bite end in sepsis and death? The initial conditions just happened to line up that way. There is a causal explanation for the conditions' chance fatal alignment, as there is a causal explanation for everything, but the explanatory question is not worth answering. At this point of utter contingency, the historical interest of the causal chain peters out.[9]

The same might be said of Ariel Sharon's stroke or the storms that shattered the Spanish Armada—and also of the virulence of *P. falciparum*. There is surely an interesting explanation of the implementation of the plasmodium's virulence—an explanation why, given its physical structure and the structure of its host, it is so liable to debilitate and kill—but the explanation why it has that particular structure is most likely not interesting at all. It is just a matter of chance that a microorganism that was adapted to certain aspects of gorilla physiology should have been, in virtue of aspects of those adaptations, so dangerous to humans.

Even if the historical derivation of *P. falciparum*'s virulence is feasible, then—even if complexity, epistemic erasure, and the norms imposed by science's division of cognitive labor do not prevent evolutionary biologists from constructing a causal explanation of how *P. falciparum* came to have those particular properties that make it especially virulent—it is likely too unrewarding to attempt. The means may be at hand, but there is no motivation for evolutionary biologists or any other breed of scientist to undertake the project.

It follows that the island that is the model for the evolution of the sickle gene, though it can be connected to the other islands of the evolutionary archipelago by historically deriving its local content, is unlikely, because it hinges on exceptionally contingent matters of fact, ever to be so connected.

9. The inverse connection between contingency and explanatory depth is investigated in Strevens (2008, §4.36); see also Hempel (1965, 345) and Hitchcock and Woodward (2003).

The historical derivation, though it exists, simply has too little to offer, explanatorily or otherwise.

The same is true, I suggest, for most or all islands in the archipelago: they are surrounded by seas of contingency. Although historical derivations link the geology of the different islands, such history will appeal too much to "matters of chance", to happenstance, to the same sort of coincidental alignments that brought about *P. falciparum*'s disastrous configuration and Alexander's inopportune death, to be worth pursuing. And so evolutionary biology will remain for the most part horizontally disunified even when science is finished.

Is this picture too pessimistic? Some historical derivations are worth having—the explanation of the proliferation of the sickle gene after *P. falciparum*'s appearance, for example. Might not the contingency of *P. falciparum*'s virulence be rather exception than rule?

There is certainly a great deal of interesting biological history. But the derivations needed to connect the islands are long and tenuous, and it takes only one contingent link to break the chain. To determine a parameter in a typical evolutionary model, such as the probability of surviving until reproductive age, so many things must come together: numerous aspects of the physiology of predators and prey, diverse features of the geography and demography of the habitat. At least a few of these are likely to be in their relevant aspects—in the aspects that make a difference to viability—matters of happenstance. But a few is all it takes to undercut the drive to bridge the islands. Even if some aspects of the story, such as diploidy in mammals (essential for the evolution of the sickle gene), are interesting to explain all the way down, the more specific features of any given evolutionary vignette, and so of the local content of an evolutionary model, are likely to trail off, in their histories, into tedious contingency. I am not saying, and I do not want to say, that no two islands will ever be bridged. I do expect, however, that the considerable horizontal disunity we see in present-day biology will be around forever.

What goes for evolutionary biology also goes for most of the other sciences. In the rest of biology, in economics, in psychology, in sociology, and even in much of physics, the primary objects of study are complexly structured objects that come in many varieties: cells, organisms, ecosystems, banking sectors, mental modules, cultures, languages, and lasers. The behavior of such things depends partly, perhaps even to a great degree, on particular features of their structure. (I think we can say this quite confidently, in spite of the exciting work on universality that we have seen over the last century or so.) Accurate models of the behavior will therefore rely to some extent, and perhaps even to a great extent, on local content—on local structural facts.

The sciences of special structures—the special sciences—will consequently reflect at first the variety of their subject matter, which is to say that they will consist of a constellation of models of particular systems, each founded on its own local content. Must things stay that way? A vertical unification will connect each model separately to the fundamental-level laws, but in many cases they can be connected also to one another, by way of a horizontal unification that subsumes many special structures within a single causal narrative explaining how each came to have its distinctive properties—so much the better to build a mandala that captures the fundamental oneness of things.

Yet it seems unlikely that science will seize this opportunity to make its representation of the world whole. One great obstacle is the difficulty of building connections, but another at least as great is the dreariness of doing so. Most structures behave the way they do in part because of features that result from chance alignments, oppositions, and misalignments. Though the behavior is interesting, contingent causes of these sorts are by and large not worth the effort to explain. We shake our heads and say: "It just happened that way". And then, if we are wise, we express our regrets and move on.

The models of the special sciences will for this reason often enough remain explanatorily isolated, all the way to the end of inquiry. Disunity is the rule in the sciences; the reason, however, has little to do with metaphysical disunity in the world outside.

* * *

Cartwright has something to say about defenses of fundamentalism according to which "the problem of unruliness... is not in nature but, rather, [is] an artifact of our postlapsarian frailties".[10] Although the explanation from contingency does not hinge on human frailty—we would decline to pursue a horizontally unified science even if we had all the power in the world—you might wonder whether the same defense nevertheless applies.

"I am an empiricist," writes Cartwright. "I know no guide to principle except successful practice" (p. 2). And on pp. 11–12: "[We should form] beliefs about laws, if we must have them at all, from the appearance of things." The thought is this: the appearance of nomological disunity gives us prima facie reason to believe that laws truly do have a mosaic nature, and this reason cannot be undercut merely by showing that fundamentalism is compatible with the appearance of unity.

I, along with other fundamentalists, reply that, though a bare demonstration of logical compatibility is indeed not enough, the argument for the metaphysical mosaic falls through if fundamentalism provides an explanation for the appearance of disunity that is equally as good as the explanation provided by the mosaic. (For the antifundamentalist argument to fall through is not, needless to say, for a pro-fundamentalist argument to succeed. Providing independent arguments in favor of fundamentalism is a matter for another day.)

Cartwright the self-identifying empiricist responds, I think, as follows. Yes; if fundamentalism provides a sufficiently good explanation for the appearances then the prima facie argument for the mosaic is overturned. But

10. Cartwright (1999, 11–12). All the following page references are to the same work.

an important component of explanatory goodness in the relevant sense is epistemic. The foundations of the fundamentalist's explanation of seeming disunity are considerably less secure than those of the mosaic explanation, appealing as they do to laws whose universal scope can be tested only very indirectly.[11]

I will not comment on the strength of this argument as an across-the-board objection to fundamentalism, but I will point out that it does not apply to my proposed explanation for horizontal disunity: the key posits of that explanation are manifold contingencies, such as the accidentalness of the virulence of *P. falciparum*, that can be uncovered through ordinary scientific investigation of exactly the sort that Cartwright considers to provide the evidence for the mosaic. Even if there is something objectionably un-empiricist about fundamentalism itself, then (which I do not concede for a moment), the proposition that *if* fundamentalism is correct, science will nevertheless exhibit considerable horizontal disunity, rests on what is by Cartwright's own lights firm empirical ground—and that, I claim, is enough to undermine the argument from horizontal disunity to the nomological mosaic.

Those antecedently attracted to fundamentalism therefore have no reason to doubt that the methodological mosaic subsists in a metaphysical mandala.

> The [Buddhist tantric] practitioner sees in the sum of his or her perceptions [not] a familiar, concrete environment, but rather… the basic pattern—purified of all misleading concepts and individual chance factors… [This is his or her] own mandala.
>
> Dagyab Rinpoche, *Buddhist Symbols in Tibetan Culture*, 80

11. Cartwright goes even further, writing that the fundamentalist foundations are based in "a priori metaphysics" rather than "hard scientific investigation" (p. 12)—a charge which most fundamentalists would strenuously deny. The better part of her empiricist argument seems to me to hinge on the indirectness of our evidence for the fundamentalist posits compared to the directness of our evidence for the mosaicists's posits.

References

Cartwright, N. (1999). *The Dappled World: A Study of the Boundaries of Science*. Cambridge University Press, Cambridge.

———. (2003). Against the completability of science. In M. W. F. Stone and J. Wolff (eds.), *The Proper Ambition of Science*. Routledge, London.

———. (2007). *Hunting Causes and Using Them: Approaches in Philosophy and Economics*. Cambridge University Press, Cambridge.

Fromkin, D. (1989). *A Peace to End All Peace: The Fall of the Ottoman Empire and the Creation of the Modern Middle East*. Henry Holt, New York.

Hempel, C. G. (1965). Aspects of scientific explanation. In *Aspects of Scientific Explanation*, chap. 12, pp. 331–496. Free Press, New York.

Hitchcock, C. and J. Woodward. (2003). Explanatory generalizations, Part II: Plumbing explanatory depth. *Noûs* 37:181–199.

Hoefer, C. (2008). For fundamentalism. In S. Hartmann, C. Hoefer, and L. Bovens (eds.), *Nancy Cartwright's Philosophy of Science*. Routledge, London.

Kitcher, P. (1999). The hegemony of molecular biology. *Biology and Philosophy* 14:195–210.

Oppenheim, P. and H. Putnam. (1958). Unity of science as a working hypothesis. In H. Feigl, M. Scriven, and G. Maxwell (eds.), *Concepts, Theories, and the Mind-Body Problem*, volume 2 of *Minnesota Studies in the Philosophy of Science*. University of Minnesota Press, Minneapolis.

Strevens, M. (2008). *Depth: An Account of Scientific Explanation*. Harvard University Press, Cambridge, MA.

———. (2016). Special-science autonomy and the division of labor. In M. Couch and J. Pfeifer (eds.), *The Philosophy of Philip Kitcher*. Oxford University Press, Oxford.