Damn the Consequences: Projective Evidence and the Heterogeneity of Scientific Confirmation

P. Kyle Stanford†‡

I contrast our own evidence for the hypothesis of organic fossil origins with that available in previous centuries, suggesting that the most powerful contemporary evidence consists in a form of projective support whose distinctive features are not well captured by familiar hypothetico-deductive, abductive, or even more recent and more technically sophisticated (e.g., Bayesian) accounts of scientific confirmation. I suggest that such accounts either misrepresent or ignore something important about the heterogeneous ways in which scientific hypotheses can be supported by evidence, and I go on to suggest that the search for any single such account may be misguided in any case.

1. Introduction: Confirmation and Consequences. Karl Popper famously argued that we need not be overly concerned about our inability to answer Hume's skeptical problem of induction, because the testing and 'corroboration' of scientific hypotheses simply do not rely on the sort of enumerative induction from particulars to a universal or from known to

†To contact the author, please write to: Department of Logic and Philosophy of Science, University of California–Irvine, Irvine, CA 92697; e-mail: stanford@uci.edu.

‡It is a pleasure to acknowledge exceedingly useful discussions of this material with Jeff Barrett, Peter Lipton, Pen Maddy, Brian Skyrms, Bruce Glymour, Elliott Sober, Kevin Zollman, Gillian Barker, John Norton, Martin Thompson-Jones, Derek Turner, and others I have forgotten, as well as audiences at Oberlin, Pittsburgh's Center for History and Philosophy of Science, the Department of History and Philosophy of Science at Cambridge University, the University of Konstanz, the University of California–San Diego, the 2009 Second Conference on Integrated History and Philosophy of Science at Notre Dame University, the 2009 conference of the European Philosophy of Science Association, the 2009 Unconceived Alternatives Symposium at the Erwin Schrödinger International Institute for Mathematical Physics and the Institute Vienna Circle, the 2010 Henle Conference, and the PSA 2010. This essay is dedicated to the memory of Peter Lipton, who always asked the right questions.

Philosophy of Science, 78 (December 2011) pp. 887–899. 0031-8248/2011/7805-0015$10.00

unknown cases in the way that so troubled Hume. Instead, they rely on the very different process of conjecture and refutation.[1] Although few contemporary philosophers accept Popper's account in all its details, there is an important sense in which he nonetheless won this argument, for it remains received wisdom among most philosophers of science and scientists themselves that the fundamental pattern of confirmation for scientific hypotheses consists in showing that we can explain or predict phenomena by exhibiting them as deductive or probabilistic consequences or implications of those hypotheses, whether this is described as abduction, hypothetico-deduction, inference to the best explanation, or Popper's own method of conjecture and refutation. And, of course, the notion that the logic of scientific confirmation is fundamentally a matter of identifying and verifying a theory's further empirical consequences is often regarded as the central reason that 'naive' (supposedly Baconian) inductivism offers a quaint, charming, and hopeless picture of what scientists do.

I will argue that this 'consequentialist' view of hypothesis testing and confirmation in science is unduly restrictive, in a manner that obscures something important about the variety of ways in which evidence can bear on scientific hypotheses. I will begin with a particular example, suggesting that although the most impressive evidence we once had for the hypothesis that fossils are the remains of living organisms was indeed abductive or consequentialist in character, this is simply no longer the case. The most powerful evidence we now have in support of this hypothesis consists, instead, in a kind of inductive projection from known cases to unknown cases of just the sort that Popper suggested was no essential part of science itself. Although such projective evidence can be forced into the various molds offered by consequentialist accounts of confirmation, I will suggest that doing so comes only at the cost of obscuring the most important features of *how* that evidence provides support for the hypothesis.

Furthermore, I will argue that although the technical apparatus of the more recent and more sophisticated Bayesian approach to confirmation does not explicitly embody any such consequentialist assumptions, central attractions of that approach nonetheless arise only because it is widely presented and interpreted in a consequentialist spirit. In addition, then, the evidence examined below will illustrate why any such uniformly consequentialist understanding of Bayesianism is incomplete and why Bayesians, too, will have to recognize and incorporate the sort of heterogeneity I argue has been ignored by more traditional approaches to confirmation.

1. Of course, Popper also denied that such 'corroboration' could actually confirm scientific hypotheses, strictly speaking—but we will leave that issue for another day.

2. Fossils, History, and Taphonomy. To begin, however, it will conduce to clarity to use the shifting historical fortunes of the belief that fossils are the remains of once-living organisms to illustrate why any straightforward hypothetico-deductive or abductive account of scientific confirmation offers an impoverished picture of our best current evidence for this view. Even after the dawn of modern geology in the Renaissance the hypothesis of organic fossil origins remained highly controversial for several more centuries because it faced both substantial empirical challenges and competition from serious theoretical alternatives. Among the empirical challenges were the problems of the locations and positions in which such fossils had been found: because these (largely marine) fossils were found on hilltops and in mountains far from the sea, for example, any hypothesis of organic origin seemed to require geological change on a scale preposterously outstripping any convincing available evidence. And, of course, the fossils were *inside* rocks bearing little obvious resemblance to loose sediments: if these were organic remains, how did the original organisms get inside?

Far more important, however, existing theoretical alternatives seemed to offer more convincing explanations of the distinctive pattern or range of resemblances between fossils and living organisms. As Martin Rudwick has emphasized,

> Even when resemblances between fossils and living organisms could be clearly perceived, it did not seem to follow necessarily that the fossils were actually the remains of living organisms. This inference, so obvious to us today, was not avoided in the sixteenth century for reasons of intellectual conservatism or out of any sense of conflict with religious orthodoxy. It was usually ignored or rejected on the far more positive grounds that . . . both the renewed Aristotelianism and the synthetic Neoplatonism of the sixteenth century . . . provided the phenomenon of organic resemblance with explanations that were quite as persuasive, indeed more so, than the hypothesis of organic origin. Aristotelians could attribute organic resemblances to the growth *in situ* of objects combining the form of genuine organisms with the stony matter appropriate to all 'fossils'; objects for which the causal explanation lay in spontaneous generation or the implantation of specific 'seeds' within the Earth. Neoplatonists could attribute the same resemblances to the action of a pervasive moulding force or 'plastic virtue', which made visible the hidden web of affinities that bound all parts of the cosmos into one. In either case, the explanations successfully accounted for the fact that the resemblances varied from the striking to the barely perceptible, and they were

therefore more widely applicable and more 'successful' than the hypothesis of organic origin. (1976, 45)

It is easy to forget that theorists have not been continuously presented with fossils as they are processed, arrayed, and (perhaps most important) selected for display in modern museums. Most buried objects, of course, do not resemble living organisms at all, and some evident resemblances (say, of a stony concretion to a human fist) are obviously fortuitous. Even more important, the unfamiliarity of many fossilized organisms ensured that many of the striking shapes found in the rocks simply appeared to be animal-*like* and plant-*like*. Hermetic Neoplatonists regarded these occasional resemblances simply as another manifestation of the same web of hidden 'affinities' and 'correspondences' that they saw as underlying the coherence and intelligibility of nature as a whole and that they invoked to explain the attractive power of lodestones and pieces of amber, the medicinal value of plants, and the supposed powers and protections of nonorganic 'fossils' such as gemstones (Rudwick 1976, 26, 33–34). The opposing neo-Aristotelian tradition instead attributed the origin of fossils resembling plants and animals to the same sorts of causal processes they believed produced actual organisms: spontaneous generation, for example, might occur in the interior of the Earth (producing fossils from the stony materials there) as well as on its surface or in its seas, while the shapes of more complex organisms might be similarly generated by the form present in their characteristic seed acting on those same materials when carried there by such natural processes as the percolation of groundwater (Rudwick 1976, 33–35). The primary advantage of these theoretical alternatives over the hypothesis of organic fossil origins was their ability to explain why fossils exhibited such a wide range of variation in their degree of resemblance to living organisms: although some fossils resembled living creatures or their parts extremely closely, many other fossils were quite distant and unfamiliar. Absent the substantial ad hoc assumption that many, or even most, species of fossilized organisms were by then extinct, the range of such resemblances seemed much more plausibly explained either by the idea that ordinary generative processes produced fossils instead of organisms when operating on the quite different materials found underground or by the appeal to the sorts of affinities and correspondences already believed to hold between various parts or aspects of nature than they were by the idea that fossils were the actual physical remains of organisms that had once been alive. As the pace of fossil discoveries accumulated throughout the next several centuries, this increasingly evident lack of precise correspondences between many fossils and extant organisms, together with concerns about the positions and locations in which such fossils were found, helped ensure that the question of fossil origins remained a matter of substantial scientific controversy well into the eighteenth century.

Of course, the living forms of unfamiliar fossil organisms might simply be undiscovered, as John Ray proposed in the case of stalked crinoid fossils, whose subsequent discovery in the living state helped fuel a growing consensus in favor of organic fossil origins well before the reality of extinction became widely accepted around the turn of the nineteenth century (see Rudwick 1976, 89, 101). There were limits, however, on how far this suggestion could be plausibly pushed. Even Ray himself could not accept the same explanation for *ammonite* fossils, as he could not accept the existence of a well-populated *genus* of which not even a single species was known in the living state (Rudwick 1976, 65). And Ray ultimately found himself attracted to Edward Lywyd's 'animalculist' version of the neo-Aristotelian view that fossils were formed by the generative materials of living creatures (especially those with external fertilization) infiltrating the rocks.

Throughout this historical period, the fundamental evidential support for the hypothesis of organic fossil origins remained genuinely abductive or consequentialist in character: the battleground for competing theories of fossil origins really was a comparison of their consequences with the available evidence on such points as fossils' impressive degrees of functional detail and resemblance to living organisms, their locations and positions, and residual differences in morphology between them and extant organisms. Those who accepted the hypothesis of organic fossil origins did so because it provided the best explanation for a number of otherwise puzzling characteristics of some fossils, most important, their striking similarities of form to extant organisms. And those who rejected the hypothesis did so because many of its consequences appeared to conflict with the available evidence, especially as compared with those of any number of extant alternatives.

But I now want to contrast the evidential situation faced by our historical predecessors with our own evidential situation, for the most powerful and convincing evidence we now have for the hypothesis of organic fossil origins is no longer abductive or even consequentialist in character.[2] The most impressive contemporary evidence for the organic origin of fossils comes from what is now called actuopaleontology, or *taphonomy*, a term coined by J. A. Efremov in the 1940s to describe "the study of the transition (in all its details) of animal remains from the biosphere into the lithospere" (1940, 85). Such taphonomic research offers us direct observational and

2. This contrast between the historical situation and contemporary taphonomic evidence is developed in considerably greater detail in Stanford (2011), although the central point of that essay concerns the implications of this contrast for debates concerning scientific realism rather than confirmation.

experimental evidence concerning the various processes presently acting to sequentially transform the remains of living organisms into fossils such as those we discover in excavated natural surroundings: as a recent textbook emphasizes, taphonomic research "must rely heavily on studies of what happens to modern bone assemblages in natural environments" (Shipman 1981, 100) as well as on experimental studies of organic remains in the laboratory. Although in nature these processes are often combined sequentially over longer timescales than we can directly observe, contemporary taphonomists nonetheless quite generally conceive of their knowledge of fossilization in past environments as simply *projected* in this way from their study of the *ongoing* processes of disarticulation, scavenging, trampling, decay, weathering, hydraulic transport, mineralization, and diagenesis that they investigate in both the field and the lab: "Voorhies' work (1969) and its extensions . . . have shown that currents usually sort the skeletal elements of mammals into three groups. . . . If an assemblage contains skeletal elements from only one of these Voorhies groups, water-sorting is indicated, and it may even be possible to establish the probable velocity of the current. On the other hand, an abundance of elements from all three Voorhies groups strongly suggests that water currents were not an important force in concentrating the bones" (Shipman 1981, 131).

Beyond field observations of taphonomic processes in action, however, we can and do *make* fossils ourselves by re-creating these processes. In essays with titles like "Artificial Microfossils: Experimental Studies of Permineralization of Blue-Green Algae in Silica" (Oehler and Schopf 1971), "Experimental Taphonomy Shows the Feasibility of Fossil Embryos" (Raff et al. 2006), "Fossilization of Soft Tissue in the Laboratory" (Briggs and Kear 1993), and "Experimental Mineralization of Invertebrate Eggs and the Preservation of Neoproterozoic Embryos" (Martin, Briggs, and Parkes 2003), experimenters report their own transformation of organic remains into fossils by taphonomic processes in the laboratory: "A technique has been developed to artificially fossilize microscopic algae in crystalline silica under conditions of moderately elevated temperature and pressure. The technique is designed to simulate geochemical processes thought to have resulted in the preservation of organic microfossils in Precambrian bedded cherts. In degree of preservation and mineralogical setting, the artificially permineralized microorganisms are comparable to naturally occurring fossil algae" (Oehler and Schopf 1971, 1229).

What is most striking about this evidence for our purposes is that none of it seems particularly abductive or even consequentialist in character. That is, we do not now believe that fossils have an organic origin simply because this hypothesis would explain so many things about fossils so well, or only because the *consequences* of that hypothesis fit so well with the evidence we have,[3] but also because we have considerable independent evidence of the ongoing actual operation in nature of processes that sequentially transform organic remains into fossils like those we encounter in natural settings. Taphonomists are able to simply find or place organic remains in various settings to see how those remains are affected by taphonomic processes. And when taphonomists suspect that a puzzling sort of naturally occurring object may represent the fossil remains of a particular sort of organism, they simply fossilize those organisms (or the closest modern analogues they can find) to see what the resulting fossils are like. In the course of that inquiry, they discover which present-day objects are indeed fossilized organic remains, which organic entities we should not expect to find represented in the fossil record even if they were abundant in past environments, and so on. We might never have been able to accumulate this detailed projective evidence concerning taphonomic processes without our initial suggestive abduction concerning fossil origins, but such evidence nonetheless constitutes the most convincing present justification for the hypothesis of organic fossil origins.

The claim here is not, of course, that having any such projective support automatically renders a scientific hypothesis well confirmed or belief-worthy: as always, this will depend on the quality and quantity of the evidence in question and the breadth and depth of the parallels between the conditions under which it was accumulated and those into which it is being projected. To take a simple example, the existence of so-called rock varnish or desert varnish on rocks on the surface of Mars provides only very weak projective evidence for the existence of bacterial life on that planet. Not only does the role of bacteria in producing rock varnish on our own planet remain unclear; we also know almost nothing about the depth and breadth of the parallels between the conditions under which this distinctive pattern of coloration was formed here on Earth and those under which it would have had to have been produced in the very different environment of the Martian regolith (DiGregorio 2010). By contrast, the wide range of circumstances actually realized in nature in which we have observed various taphonomic processes or put them to work ourselves allows us to confidently *project* their operation to distant times and places, concluding that organic remains deposited long ago almost certainly underwent these same processes with the same or very similar results.

Of course, it is possible to squeeze such projective evidence into a preexisting hypothetico-deductive, abductive, or consequentialist mold. To find such a mold, we need look no further than the confirmational pattern that

3. Indeed, taphonomists sometimes reject such consequentialist reasoning quite explicitly; see Briggs and Kear (1993).

Carol Cleland has argued is "prototypical" for historical sciences such as cosmology and paleontology (2002, 480), in which

> an investigator observes puzzling traces (effects) of long-past events . . . [and] hypotheses are formulated to explain them . . . by postulating a common cause for them. . . . Traces provide evidence for past events just as successful predictions provide evidence for the generalizations examined in the lab. Instead of inferring test implications from a target hypothesis and performing a series of experiments, historical scientists focus their attention on formulating mutually exclusive hypotheses and hunting for evidentiary traces to discriminate among them. The goal is to discover a "smoking gun." A smoking gun is a trace(s) that unambiguously discriminates one hypothesis from among a set of currently available hypotheses as providing "the best explanation" of the traces thus far observed. (480–81)

We can indeed say (truly!) that postulating the operation of taphonomic processes in the distant past offers the best explanation for the existence and characteristics of fossil objects and insist that understanding the detailed operation of such processes allows them to provide a better and better explanation of the actual fossil remains that we find. But this is simply to ignore *how* this evidence makes the resulting explanation 'better', which is by giving us increasingly powerful grounds for projecting the operation of these processes into distant times and places and for making finer discriminations of legitimacy among such projections. Our explanation is 'improved' in this case only by acquiring more and more support of precisely the projective variety that Popper sought to dispense with and that abductive and consequentialist accounts of confirmation more generally assimilate to the testing of hypothetical implications. Indeed, observing and re-creating taphonomic processes ourselves seems to support the hypothesis of organic fossil origins by satisfying something much more akin to the demand for a *vera causa* so prominent in nineteenth-century science than by verifying yet another of that hypothesis's empirical implications. And it seems in no way directed toward satisfying Cleland's further description of a "smoking gun": "a trace(s) that . . . picks out one hypothesis as providing the best explanation currently available [and] need not supply direct confirming evidence for a hypothesis independently of its rivals" (2002, 490). Perhaps the hypothesis of organic fossil origins does indeed suggest or imply that we should (or might) be able to find organic remains in the process of fossilizing or to make fossils ourselves in the lab, but to see these latter findings as confirming the hypothesis *by* verifying one of its consequences is simply to cover the real confirmational story with a thin coat of consequentialist paint.

Moreover, forcing projective evidence into such a preconceived mold serves to obscure important further differences between it and more stereotypically abductive, hypothetico-deductive, or consequentialist varieties. In the case of a hypothesis like atomism or general relativity, it is not even open to us to see a theoretical claim about the constitution of nature as supported by information about events or processes we are able to observe more directly: although the behaviors of billiard balls and rubber sheets are simple physical analogies we use to help illustrate some important characteristics of atoms or space-time, they do not serve as any sort of inductive basis from which physicists actually project the causes of observed physical phenomena such as Brownian motion or gravitational lensing. By contrast, we do not think that the fossilization processes that operated in the distant past are simply analogous to or like those we study in the field and lab in various important or illuminating respects; rather, we think that these just *are* the very processes that produced fossils in the distant past and continue to do so in the present. The suggestion that contemporary physicists should study macroscopic systems of orbiting and colliding rigid objects to try to understand the atomic constitution of matter is absurd in a way that taphonomy is anything but an absurd undertaking for modern paleontology. Thus, the fact that it is even possible to understand the taphonomic evidence in support of organic fossil origins as a kind of inductive projection marks a significant difference between it and what we might think of as a more fundamentally abductive or consequentialist form of justification.

Furthermore, such projective evidence can also enjoy positive virtues that seem unavailable to its abductive counterpart. Elsewhere (Stanford 2006), I have suggested that the most daunting challenge facing abductive inferences in science is the prospect of competing explanations for a given set of phenomena that are well confirmed by the available evidence but that nonetheless remain presently unconceived—and the sort of projective support described here certainly faces an analogous kind of challenge: even if we show that some existing process *can* produce a given effect, we must still worry about whether it was actually this process rather than some (possibly unconceived) alternative that produced the effect in a particular instance. But the case for projecting the operation of a process into other times and places can be powerful enough to create an affirmative challenge for the suggestion that some alternative process (unconceived or otherwise) might actually be responsible for an observed effect. The case for projecting taphonomic processes into past environments is sufficiently strong, for example, that if we suppose some unconceived alternative process to be responsible for the fossils that we have uncovered in nature, we will then have to explain why the taphonomic processes we have investigated in such detail in the field and lab have *failed* to produce fossils over geological time or where *those* fossils have gone. We thus obtain a distinctive epistemic advantage when we can supplement abductive evidence with its projective counterpart, an advantage that is surely not equally available in every context of scientific investigation (see Stanford, 2011). Again, it seems that there are important differences between such projective and abductive forms of evidence and that such differences are simply obscured by forcing all projective evidence into a preconceived abductive or more broadly consequentialist mold.

3. Bayesian Variations. The discussion to this point will doubtless seem atavistic to the many philosophers of science who long ago concluded that hypothetico-deductive, abductive, and other classical approaches to confirmation were irredeemable and have therefore rejected all such approaches in favor of more technically sophisticated accounts. Most influentially, Bayesian confirmation theory insists that a hypothesis is tested by checking facts that are somehow deductively or probabilistically linked to it (against a background of auxiliary assumptions, including a prior probability for the hypothesis under test). In standard Bayesianism with strict conditionalization, the extent to which new evidence favors one hypothesis over its rivals is determined most fundamentally by the likelihood conferred on that evidence by the presumption that the hypothesis is indeed true in comparison to the likelihood conferred on that same evidence by each of those rivals (with background assumptions held constant). And the likelihood 'conferred on' the evidence by any given hypothesis in this formal apparatus is merely a relation of conditional probability: the evidence need not be regarded as a 'consequence' of the hypothesis in any sense of the term.

Although likelihoods can certainly be so regarded in a Bayesian framework, we must then not make the mistake of thinking that the resulting austere form of Bayesianism accomplishes or explains more about confirmation than it does. Such austere Bayesianism avoids any consequentialist presuppositions—but only by saying nothing at all about *how* and *why* it is that the truth of a given hypothesis should or would affect the likelihood we assign to any given piece or body of evidence, despite acknowledging and even insisting that the determination of such (comparative) likelihoods is what actually bakes the confirmational cookies. To be sure, this austere form of Bayesianism will automatically be able to incorporate absolutely any evidential consideration whatsoever, for these can always simply be built into the prior and conditional probabilities that the theory takes as its starting point(s), but in this form the theory quite literally does nothing more than articulate the demands imposed by considerations of synchronic and diachronic probabilistic coherence on the credences assigned to our beliefs.

This is presumably why Bayesian expositors so frequently supplement the bare formal apparatus with suggestions and examples concerning how such likelihoods might be established. It is here that consequentialism reenters the picture, for these likelihoods are typically presented as determined in familiar consequentialist ways: as a matter of how well a given hypothesis explains or accounts for the available evidence, for example, or with what probability the truth of the hypothesis under test predicts that we will find that evidence. In this way, I suggest, many of the attractions of Bayesianism as a substantive theory of confirmation arise from a widespread and implicitly consequentialist interpretation or understanding of its formal apparatus, even for many Bayesians who would not embrace the explicit claim that all confirmation must be consequentialist in character. But the taphonomic evidence for organic fossil origins that we have considered is no better captured by such Bayesian consequentialism than by simpler consequentialist accounts. The hypothesis of organic fossil origins certainly does not *explain* the existence of processes such as water sorting or mineralization of organic remains: to the contrary, it is by projecting the operation of processes such as mineralization that we explain the existence of fossils. And for all the reasons noted above it simply obscures the evidential situation to treat the existence of such processes as a probabilistic 'prediction' made by the hypothesis itself.[4]

No less than other approaches to confirmation, then, will a satisfying Bayesian account be forced to recognize and articulate the fundamental heterogeneity of the ways in which evidence can bear on hypotheses.[5] The austere Bayesian apparatus offers no substantive *theory* of confirmation at all unless and until supplemented with an account of the wide variety of ways in which hypotheses can 'confer' likelihoods on the evidence (against a set of background assumptions), and the taphonomic evidence we have considered shows why our efforts to thus meaningfully connect the Bayesian apparatus to the details of actual instances of confirmation

4. Jeff Barrett has pointed out to me that a truly austere or "pure" Bayesian will eschew any attempt to justify likelihoods or other prior probabilities in these ways, insisting instead that these likelihoods themselves come from Bayesian updating on preexisting likelihoods, ultimately tracing back to arbitrary (nondogmatic) primitive likelihoods lacking any justification whatsoever. Truly austere Bayesians put their faith in the fact that this arbitrariness is guaranteed to 'wash out' (eventually) as they condition on new evidence (just as the prior probabilities of hypotheses themselves do). But in matters epistemological as well as moral, it is hard to be pure. Moreover, even pure Bayesians must still connect the moving parts of this formal apparatus to concrete contexts of confirmation, and this is where I am suggesting the need to recognize the heterogeneity of such contexts invariably arises (see below).

5. I suggest that a parallel point also applies to other formally sophisticated accounts of confirmation such as error statistics, but I will not pursue the matter here.

will have to encompass more than just familiar variations on consequentialist themes. It simply will not do for Bayesianism to swell under a broadly consequentialist interpretation when seeking to illuminate concrete contexts of confirmation but then shrink back to an austere and wary formulation in response to the challenge that this imposes a false consequentialist homogeneity on all such contexts.

Of course, this is not the only challenge posed for contemporary Bayesianism by the taphonomic evidence for the hypothesis of organic fossil origins. After all, whether the Bayesian apparatus is conceived of in consequentialist terms or not, it treats the discovery that actual organic remains deposited in favorable circumstances mineralize as evidence on which a comparatively higher likelihood is conferred by the hypothesis of organic fossil origins than by its negation or by (serious) rival hypotheses. But contemporary thinkers are fully able to appreciate the force of this taphonomic evidence despite the fact that they (including you, dear reader, just a short time ago) cannot articulate even a single serious theoretical alternative to the hypothesis of organic fossil origins and despite the fact that it remains an utter mystery why the negation of the hypothesis of organic fossil origins should confer any particular likelihood whatsoever on processes such as water sorting or the mineralization of organic remains. The austere Bayesian apparatus does promise to allow us to formally integrate the confirmational significance of various diverse forms of evidence, but this remains a promissory note when we have no way to responsibly determine likelihoods of the sort with which taphonomic investigation seems largely unconcerned.

There is, I think, a broader moral here about the heterogeneity of forms of evidence, inference, and argument in scientific contexts. Philosophers of science have long sought the holy grail of *the* logical form of scientific confirmation, whether inductive generalization, the method of hypothesis, conjecture and refutation, Bayes's Theorem, or something else altogether. But I doubt that there is any such holy grail to be discovered. Scientific confirmation is a heterogeneous and many-splendored thing; let us count ourselves lucky to find it—in all its genuine diversity—wherever and whenever we can.

REFERENCES

Briggs, D. E. G., and A. Kear. 1993. "Fossilization of Soft Tissue in the Laboratory." *Science* 259:1439.

Cleland, C. E. 2002. "Methodological and Experimental Differences between Historical Science and Experimental Science." *Philosophy of Science* 69:447–51.

DiGregorio, B. 2010. "Martian Sheen: Life on the Rocks." *New Scientist*, no. 2747:40–43.

Efremov, J. A. 1940. "Taphonomy: A New Branch of Paleontology." *Pan-American Geologist* 74:81–93.

Martin, D., D. E. G. Briggs, and R. J. Parkes. 2003. "Experimental Mineralization of Invertebrate Eggs and the Preservation of Neoproterozoic Embryos." *Geology* 31:39.

Oehler, J. H., and J. W. Schopf. 1971. “Artificial Microfossils: Experimental Studies of Perminalization of Blue-Green Algae in Silica.” *Science* 174:1229.

Raff, E., C. J. T. Villinski, F. R. Turner, P. C. J. Donoghue, and R. A. Raff. 2006. “Experimental Taphonomy Shows the Feasibility of Fossil Embryos.” *Proceedings of the National Academy of Sciences* 103:5846–51.

Rudwick, M. J. S. 1976. *The Meaning of Fossils: Episodes in the History of Paleontology*. 2nd ed. New York: Science History.

Shipman, P. 1981. *Life History of a Fossil: An Introduction to Taphonomy and Paleoecology*. Cambridge, MA: Harvard University Press.

Stanford, P. K. 2006. *Exceeding Our Grasp: Science, History, and the Problem of Unconceived Alternatives*. New York: Oxford University Press.

———. 2011. “Getting Real: The Hypothesis of Organic Fossil Origins.” Special issue, *Modern Schoolman* 87 (March): 219–43.