The Aesthetics of Science: Beauty, Imagination, and Understanding. Steven French and Milena Ivanova (eds.). London: Routledge pp. 21-35.

Epistemic Gatekeepers: The Role of Aesthetic Factors in Science

Catherine Z. Elgin

The problem of the aesthetic

Scientific theories, models, experiments, and the like are often subject to aesthetic assessment. This is uncontroversial. What is controversial is whether such assessments have any bearing on their epistemic standing. Is there any epistemically good reason to prefer an elegant experiment to an inelegant one, a beautiful theory to an ugly one, a streamlined model to one that seems more like a Rube Goldberg machine? What are we focusing on when we make such assessments?

If we seek to understand the role of aesthetic factors in science, it would be nice to have a criterion for the aesthetic. We do not have one. The history of aesthetics is littered with failed attempts to devise such a criterion. For our purposes, however, something less may suffice. All we really need is to identify a few factors that are plausibly construed as aesthetic. Then we can try to determine what, if anything, they contribute. If we find a contribution, we can attempt to identify other factors that function analogously. We do not need an exhaustive list of aesthetic factors; nor do we need to determine what precisely makes those factors aesthetic.

Clive Bell (1913) maintained that the aesthetic response to works in the visual arts consists in the apprehension and appreciation of *significant form*. He was concerned with significant visible forms – colors, contours, configurations, and the like. I suggest that many aesthetic responses to works in the sciences consist in the apprehension and appreciation of scientifically significant forms in a logical space. That space need not be physical, and the apprehension need not be sensory. The aesthetic properties that concern us then are formal properties of scientific artifacts such as theories, models, methods, and experiments. A form, let us say, is scientifically significant to the extent that it illuminates something that bears on the scientific acceptability of the item that displays that form. Precisely which forms these are, of course, varies from one context of inquiry to the next (see McAllister 2007). Nevertheless, the extrapolation of the idea of significant form makes it plausible that features like symmetry, simplicity, systematicity and their opposites are aesthetic features of scientific artifacts.

Bell's criterion is known to be inadequate. It fails to capture features of works in the visual arts that are undeniably aesthetic. No doubt, my extrapolation is equally inadequate, and for the same reason. Still, I am not suggesting that significant form is a criterion for being aesthetic. I am suggesting that it is an aesthetic characteristic of some scientific artifacts. Just as Bell's criterion enables us to identify some aesthetically important features of works in the visual arts, my extrapolation enables us to identify some aesthetically important features of works of science. My purpose is to consider what functions such factors perform. For that, I need a few plausible candidates. I do not need (and cannot provide) a full criterion. Luckily, a sketch will do.

Truth or thereabouts

According to a familiar criterion for scientific acceptability, aesthetic factors turn out to be either irrelevant, or merely instrumental. Science, it is held, has a single overriding epistemic goal. Realists maintain that the goal is truth: science is successful when it reveals the truth or something close to the truth. Instruments and methods are scientifically valuable just because and just to the extent that they are truth conducive. Norms, standards, criteria and the like are acceptable just in case and just to the extent that they promote the discovery of truth. Constructive empiricists maintain that the goal is empirical adequacy: science is successful when it achieves empirical adequacy or something close to it. Instruments and methods are scientifically valuable just because and just to to the extent that they are conducive of the development of empirically adequate accounts. Norms, standards, criteria and the like are acceptable just in case and just to the extent that they promote the development of empirically adequate accounts. Either way, on this picture if aesthetic factors contribute to science, their value is instrumental. They are epistemically valuable because and to the extent that they are indicative of the goal's being realized or because and to the extent that they promote its realization. To streamline discussion, I will speak as though the goal is truth; a parallel argument holds if the goal is empirical adequacy.

On this view, the justification for preferring a beautiful, economical theory over an ugly, gerrymandered one with ad hoc excrescences is that the beautiful theory is more likely to be true (or approximately true) than its aesthetically unattractive rival. An elegant experiment is preferable because it is more likely to reveal the truth than one that proceeds by case hacking. If this is right, aesthetic assessments might play a useful diagnostic role. As scientists investigate a promising theory in depth, they may come to discern hidden beauties in it. The theories whose hidden beauties they appreciate are the ones that they are increasingly confident are likely to be true, or nearly so. That likelihood is why they keep working with them. And that scientists who have studied an issue in depth think that a theory is likely to be at least approximately true is reason for the rest of us to agree. If this is so, there may be a correlation between aesthetic assessments of currently credited theories and truth-related assessments. Perhaps the fact that aesthetically sensitive scientists find a theory beautiful is a good reason to consider it prima facie acceptable. If, however, beauty and truth do not run in tandem, this account holds that aesthetic judgments in science are at best epistemically idle. At worst they threaten to mislead.

Regrettably, there is no reason to believe that a scientific representation's aesthetic features correlate with the probability of its being true. Standard aesthetic assessments do not, so far as we can tell, track truth. We appreciate harmony, symmetry, elegance and simplicity, all of which can be found in works of art that make no pretense of being true. Many are not even truth-apt. Indeed, those aesthetes who advocate 'art for art's sake' would maintain that genuine aesthetic value is rightly indifferent to truth. Nor is there any a priori reason to expect the phenomena to arrange themselves so as to align with our aesthetic preferences. That scientists who have studied a theory in depth think it is true may be some reason for the rest of us to think it is true. That they think it is beautiful is not.

Many manifestly false theories are beautiful. Aristotelian biology, for example, treats species as fixed: each species has its own essence, its own proper function, and its own distinctive telos, which both determines the good for its members and explains their behavior. This seems much lovelier than contemporary Darwinian biology, which acknowledges a large role for chance in the evolution of species, indeterminacy at the boundaries between species, a non-trivial measure of adaptive opportunism. It has nothing to say about what, if anything, constitutes the distinctive good for the members of a species. Caloric theory, which takes heat to be a smoothly flowing fluid, seems at least as lovely as the kinetic theory of heat, which maintains that individual gas molecules careen randomly about. Moreover, scientists may continue to consider an account beautiful even after it has been decisively rejected. Feynman, for example, initially thought that Feynman diagrams could explain why empty space is weightless. He continued to consider his account beautiful even after he refuted it (Wilczek 2016). Scientists do not automatically lose their aesthetic appreciation for the theories they leave behind.

But if aesthetic judgments in science are epistemically idle, why do we go on making them? Perhaps scientists' aesthetic assessments supervene on, are derivative from, or even are mere expressions of truth-related judgments. Then, in calling a theory beautiful or an experiment elegant, a scientist would be simply using an aesthetic label to characterize something that she considered scientifically estimable for other, legitimate – that is, truth-related – reasons. This seems to strip the terms from their aesthetic role. The label gets exported, but its aesthetic function is left behind. 'Beautiful' seems to mean just 'good of its kind' (see McAllister 1996: 77-81). If this is so, then there is no role for genuinely aesthetic assessment in science.

Perhaps aesthetic judgments in science play a heuristic role. (I admit we are getting a bit desperate here.) They might function as fast and frugal ways to make preliminary assessments of theories, experiments, models and the like. But like the heuristics that figure in psychology's System I thinking, they are shortcuts that not infrequently lead us astray. They may be valuable for the way they enable us (with our limited minds) to come to a conclusion. But their utility is primarily practical and the heuristics are buggy. When you need a quick answer to 'Is this likely to be true?' ask yourself, 'Do you find this beautiful?' This strikes me as a dreadful idea, at least until we can find some reason to think that there is a correlation – even a loose correlation – between assessments of beauty and the probability of being true.

An Alternative

The problem with all of these proposals comes, I suggest, from thinking that if aesthetic factors are genuinely of value to science, it is because they somehow promote or sustain the justified conviction that the items that display them are true or nearly true. I will argue that aesthetic factors are integral to good science. They are not mere instruments. Nor is their utility primarily practical. But there is no reason to think that they themselves are truth-conducive. Rather, they bear on the acceptability of a scientific artifact – a theory, model, experiment or whatever. They do so, I suggest, by functioning as gatekeepers on acceptability: they play a regulative role.

The idea that factors that are not truth (or empirical adequacy) conducive are nevertheless integral to good science may seem anathema. Science, after all, seeks to understand the world. But the understanding that science delivers is embedded in models that are not, and do not purport to be, accurate representations of the phenomena they bear on (see Elgin 2017, Cartwright 1983). They simplify, streamline, augment and omit. They are not true. Some, such as Snell's law, are not even nearly true. To be epistemically acceptable, a model or theory must properly answer to the phenomena. But properly answering to the phenomena is not the same as being a true or accurate representation of the phenomena. Nor is it the same as being an empirically adequate one. Models that diverge from truth also diverge from empirical adequacy. The Lotka Volterra model of predator/prey relations, for example, construes predators as insatiably voracious. They are not. Even the hungriest shark eventually eats its fill and stops eating. Scientists already distance themselves from the idea that acceptable models must be true or empirically adequate. So the claim that the contribution of aesthetic factors to science is not a matter of their being directly or indirectly truth (or empirical adequacy) conducive is not as alarming as it might first appear.

Elsewhere, I have argued that an understanding of a topic consists of a systematically linked body of information in reflective equilibrium that is grounded in fact, is responsive to evidence and enables non-trivial inference and argument about a range of phenomena. That body of information is a network of epistemic commitments that includes beliefs about matters of fact, acceptances of models, idealizations and thought experiments that are known to diverge from truth, as well as of methods, norms and standards. Roughly, the idea is that to be justified in believing that the Higgs boson exists (or, for that matter, that the Loch Ness monster exists) requires a commitment to the methods, norms, standards and background assumptions that figure in establishing its existence. Since understanding is, in the first instance, understanding of a topic or range of phenomena – not understanding of an individual matter of fact – the various elements of a system of thought must be mutually supportive. If an aesthetic factor such as symmetry or elegance is integral to a network of scientific commitments in reflective equilibrium, and could not be eliminated from the system without threatening or undermining the system's reflective equilibrium, then if that system is, by current standards, at least as good as any available alternative, the factor is epistemically justified.

Even if beauty is not exclusively in the eye of the beholder, it is hard to characterize. Nor is it wholly a matter of significant form, however broadly construed. So I will focus on other aesthetic factors that figure in science, ones that seem to be largely if not entirely matters of form: symmetry, systematicity, simplicity, and elegance. I am not saying that these factors are aesthetic per se. The only one with any claim to that status is elegance. My point is that their function in scientific acceptance is aesthetic. A model's displaying symmetry, an experiment's being simple, a consideration's weaving seamlessly into a network of mutually supportive commitments are the sorts of things that make them epistemically attractive in science.

Symmetry

Symmetry is a matter of structural invariance. A symmetrical item – an object or a law – retains its structure under transformations. Thus, for example, a cubical block displays rotational symmetry in that it retains its shape when rotated. The members of collection that share a symmetry have the same structure and retain that structure under the same transformations. They constitute an equivalance class. With regard to the transformation in question, they are interchangeable (van Fraassen 1989: 243). Things are not symmetrical *tout court*. They are symmetrical in one respect or another. So in devising a theory or model, questions arise: What sorts of symmetries are we interested in? Under what sorts of transformations should the objects of interest retain their structure? These questions bear on the sort of theory or account we want to devise. The answers mark out our categories and shape our theorizing. The symmetries that figure in science are not like rocks; we do not just stumble over them as we go about our lives. They are products of decisions. We configure the domain, deciding what structures we want to preserve, and under what transformations we want to preserve those structures. Maybe, for example, we seek to preserve rotational symmetry, but have no interest in color invariance. Then the block would count as symmetrical even if its faces were different colors.

It is no accident that scientific models display specific symmetries. We design them to do so. In designing a model or partitioning a domain in such a way that certain symmetries are displayed, we tentatively commit ourselves to the view that those symmetries are scientifically significant forms. What is the basis of their significance? We have no reason to think that a symmetrical model is more likely to fit its target than an asymmetrical one or that it is likely to fit its target better than one that lacks the symmetry in question. A critical question is how symmetries and asymmetries affect epistemic decisions.

Buridan's ass found himself equidistant from two identical, equally nourishing and tasty bales of hay. With respect to choice worthiness, the two bales were symmetrical. Nothing favored one over the other. Symmetry paralyzed the poor beast. Rather than starve, let us hope, he would eventually choose arbitrarily – perhaps, favoring the one on the right, even though there was absolutely no reason to prefer it to the one on the left. If he did, we could not understand his decision. We could understand that he chose – indeed, understand why he had to choose. But we could not understand why he chose the pile on the right. There was no reason. This is unsatisfactory, particularly if we are in the business of studying ass psychology (which is, these days, a growth industry). We might give him a bye if this were a one-off choice. That much arbitrariness in the mental life of an ass we might be willing to tolerate. But if he found himself in the same situation on multiple occasions and choose the pile on the right significantly more often than the pile on the left, we would be apt to insist that there must be a reason. There must be something about the hay on the right or the ass's psychology, we would be apt to think, that accounts for the difference. There *must* be a tie-breaker. Then we would embark on a quest for hidden variables.

Another case: A fair coin is rotationally symmetrical. In the flip of such a coin, nothing favors heads over tails. Suppose we observed a single coin flipped seventeen times in a row. Each time, the coin came up heads. We know that in an infinite sequence of tosses of a fair coin, there is bound to be an interval where the coin comes up heads seventeen times in a row. Nevertheless, we would almost surely suppress that knowledge and judge that the coin was biased. We expect a fair coin to come up heads about half the time in even a short run of tosses. Should it not, we again seek a hidden variable – something that biases the coin.

There is, of course no guarantee that we will find the hidden variable we seek. Maybe there is none to be found. My point is that in cases like these we treat the perception of symmetry and the perception of asymmetry differently. If the bale of hay on the right were fresher, or nearer, or composed of different sorts of hay from the one on the left, we would probably consider that a symmetry-breaking difference and think we understood the ass's behavior. If the coin came up HTTH THTH TTHH TTTH H (which specific sequence is, after all, no more or less probable than a sequence of seventeen heads in a row), we would judge the coin fair. But the fact that the all heads sequence makes the coin *look* asymmetrical arouses our suspicions. We do not think we understand how it happened. Something, we believe, needs to be explained.

The conviction that a range of phenomena display a certain sort of symmetry is an initially tenable commitment (see Elgin 1996). In constructing a system in reflective equilibrium, that conviction has a slight and defeasible claim on our epistemic allegiance. We need a reason to give it up. Such reasons are often easy to find. We may discover that preserving a commitment proves too costly. Maybe the only way to preserve the conviction that the coin was biased or that the bale on the right was more desirable than the one on the left is to invoke occult forces for which we can find no independent evidence. Being unwilling to do that, we conclude that the coin was in fact fair, and that the ass's several choices were mutually independent and arbitrary – that our suspicions about the cases were unwarranted. Our quest for hidden variables in cases like these is evidence of our commitment to symmetry. We may find that we have to rethink the situation and our conviction that symmetry was broken where we thought it was. Or we may have to weaken our commitment to symmetry and recognize a measure of pure chance in nature. But we abandon the commitment with reluctance. We prefer symmetry-preserving accounts. Symmetry, I suggest, is an aesthetically pleasing feature. Our preference for it affects our behavior in accepting or rejecting certain findings.

Systematicity

Writing on the philosophy of history, Morton White (1965: 222-225), distinguishes between a chronicle and a history. A chronicle is just a list of unconnected, established facts about an historical episode. It makes no mention of dependence relations between facts; nor does it attempt to impose any order on them. A history is organized. It links the various facts to one another, displays dependency relations, imposes an explanatory framework in terms of which it makes sense that things played out as they did. Pretty clearly, the epistemic value of a history of a given event is correlated with the number and perceived importance of the facts in the corresponding chronicle that it accounts for. To be sure, the history need not mention every fact mentioned in that chronicle. But, one way or another, the history should accommodate the important facts on the list. A history of the siege of Stalingrad that left it a mystery why so many people died would be unsatisfactory.

Something similar holds in science. A science seeks systematic, integrated, mutually supportive accounts of the phenomena it investigates. It eschews danglers. A series of independently established truths about a domain might qualify as a natural chronicle. But a scientific account would be unsatisfactory if it left mysterious how those truths hang together.

Unaccommodated, isolated, but seemingly significant truths are considered problematic. Towards the end of the nineteenth century, celestial mechanics entertained a variety of increasingly strained hypotheses to accommodate the apparent anomaly in the perihelion of Mercury. But even at their most desperate, scientists did not, advocate accepting:

All planets except Mercury have regular Newtonian orbits; Mercury is just different.

Although the contention is true, was justified and confirmed by the observational evidence, and is not a Gettier case, it is unacceptable without an explanation of exactly why Mercury is different.[1]

Unexplained danglers are unacceptable. A question is considered open until putative exceptions to a proposed answer have been explained, or explained away. Unless apparent exceptions can be shown to be irrelevant, they are in one way or another expected to be woven into the scientific account. To be sure, there are anomalies – seemingly relevant issues that currently defy explanation or incorporation into our best available accounts. But anomalies are construed as outstanding debts. A scientific account is unsatisfactory to the extent that it lacks the resources to pay its debts. Moreover, once a phenomenon has been woven into an acceptable account, scientists are typically satisfied. Nothing more needs to be explained. Why Mercury's orbit is irregular was an outstanding problem for Newtonian mechanics. Why Mercury's orbit is regular is not a question for general relativity. Indeed, the question hardly makes sense. After all, given the theory, what would you expect?

Enthusiasm for systematicity runs deep. We want our fabric of scientific commitments to be tightly woven. Physicists seek a Grand Unified Theory out of a dissatisfaction with having to admit multiple fundamental forces into their ontology. It would be aesthetically more pleasing if there were only one (see Weinberg 1992). We dislike case hacking as a way of establishing an hypothesis because, it seems, the fact that the hypothesis can be shown, one by one, to apply to each individual case does not to our satisfaction show why it applies to all of them. The significant form of an acceptable scientific account consists in its being an interwoven

[1] As it turned out, of course, Mercury is not relevantly different. The Newtonian theory that construed it as different was wrong.

(preferably tightly interwoven) fabric of epistemically interdependent commitments that has few, if any danglers.

Simplicity

Simplicity is complicated. There are numerous dimensions along which a theory or model might be simple or complex. Occam's razor is traditionally formulated as a principle of ontological simplicity: Do not multiply entities beyond necessity. But, as Sober points out, there are a multiplicity of razors, each seemly worthy of adoption (2015). There is no obvious general reason to favor ontological simplicity over, for example, axiomatic simplicity, syntactic simplicity, or inferential simplicity. Ontological simplicity consists in being committed to a minimal number of primitive entities or a minimal number of primitive kinds of entities. Axiomatic simplicity consists in containing a minimal number of axioms, fundamental laws, or postulates. Syntactic simplicity concerns the number and complexity of the basic syntactical units that figure in laws or axioms. Inferential simplicity concerns length of the derivations from the fundamental laws to other commitments of the theory. No doubt this list could be extended and the distinctions sharpened. For our purposes, the point is that these are all reasonable, objective, attractive features of theories and models. Moreover, they seem to have little to do with truth or empirical adequacy. There is, on the face of it, no reason to think that a theory with a minimal number of primitives or one that admits of streamlined inferences is more likely to be true than one that has a greater number of primitives, or one that requires long, convoluted inferences. Nor is there any reason to think that the prospects of empirical adequacy are enhanced by either sort of simplicity.

It might seem then that our preference for simplicity is grounded in intelligibility or tractability. Simpler theories and models are easier to handle. This is a pragmatic asset. If so, at least one reason why a simple theory is to be preferred over a complex one is that we are less likely to make mistakes in thinking with it or acting on it. The preference then is an accommodation to our human cognitive limitations. But some simplicity comes at the cost of intelligibility. Classical propositional logic needs only one primitive – the Sheffer stroke. In terms of ontological simplicity, Sheffer stroke logic wins, hands down. But it is hard to do logic using only the Sheffer stroke. The proofs are long; many steps seem counterintuitive; mistakes are apt to be made. As a practical matter, it is preferable to formulate propositional logic using at least two, and typically three, primitives. Standard formulations of propositional logic are moreover, inferentially simpler than Scheffer-stroke logic. Proofs are shorter, more direct, and more intuitive. If intelligibility and tractability are our guides, then a somewhat complicated theory may be more attractive than the simplest theory we can devise.

In the quest for simplicity, we face trade-offs. We readily sacrifice a measure of one sort of simplicity in order to gain a different sort. Thales hypothesized that everything is water (see Aristotle, 983b27-33). Della Rocca (forthcoming) thinks that reality is one and indivisible. Albert suggests that reality consists in a single physical object – the universal wave function, or alternatively, the single universal particle (1996). All three theories achieve an admirable measure of ontological simplicity. Albert and Della Rocca contend that *au fond* there is only one thing; Thales contends that there is only one sort of thing. All face the enormous burden of accounting for the ways the one seems to configure itself into wildebeests and armadillos, starfish and galaxies, molecules and mountains, and so forth. The laws they need to invoke are apt to be exceedingly complex to account for the manifest diversity of appearances at every level, given that there is only one thing, or sort of thing to appeal to.

Still, we are not willing to jettison simplicity altogether. We reject theoretical excrescences. An electric current is standardly described as a flow of electrons. What would be wrong with elaborating this claim and saying that the electrons are propelled by tiny, undetectable gremlins nudging them along with hockey sticks? The answer is obvious: The standard account is simpler. The introduction of gremlins with hockey sticks is completely superfluous. It contributes nothing to the tenability of the account that contains it. Because it contributes nothing, the evidence for the original thesis does nothing to support the gremlin addendum. Nor do we have any independent evidence. That being so, we not only have no reason to accept the gremlin hypothesis; we have excellent reason to reject it.

This is an easy case, for the hypothesis in question is idle. It adds nothing. Things get trickier when an added hypothesis adds something, but not enough. If, to explain the behavior of Buridan's ass, we had to introduce an additional psycho-magnetic ψ-force that operated only when an organism was equidistant between two equally appetizing alternatives, we would resist, and decide on balance that we should rest satisfied with the view that the ass's choices were arbitrary. In the absence of independent evidence of the existence of such a force, we would deem the cost in added complexity too high for the payoff it promises.

Ptolemaic astronomy was committed to the view that celestial objects display a complicated pattern of circular motions involving epicycles, equants, and deferents. It was important that all the motions be circular. Later astronomers, such as Galileo, thought the pattern was unduly complicated. The price the theory had to pay to preserve the geocentric framework was too high. Still, the theory was as simple as it could be if it was to accommodate the phenomena within a geocentric system. What the Ptolemaic astronomers never did, and would have been utterly unjustified in doing, was add additional (perhaps undetectable) motions to increase the number of dimensions along which celestial objects moved in circles. Science favors a minimalist aesthetics. Rococo additions have no place.

In both the case of electron flow and the case of Ptolemaic astronomy, it might seem that concerns for simplicity are presumptively truth-conducive. The complexifying hypotheses we rejected are ones that, on the evidence available, we have no reason to consider true. But sometimes we favor simplicity over truth. We devise models that prescind from truth in order to achieve simplicity. Boyle's law – $pV = nRT$ – falsifies the behavior of gas molecules by ignoring the attractive forces between them. For certain purposes, this is reasonable. The attractive forces are too weak to make a significant difference to the phenomena we seek to accommodate. Moreover, a more realistic model, such as the virial equation, is apt to occlude patterns that the more idealized model exemplifies. Where attractive forces between gas molecules are negligible, it is not just reasonable, it is advisable to neglect them – to omit them from the representations in terms of which we understand the phenomena. This is something I have argued elsewhere (see Elgin 2017). Here the important point is that streamlined models often embody an understanding of the phenomena, precisely because they are simplified. They omit what is negligible, or in Strevens's terms not a difference-maker (2008), thereby enabling us to apprehend and appreciate the significance of what is non-negligible, the difference-makers.

Science's preference for simplicity is rather inchoate, with different sorts of simplicity trading off against one another, and different types of simplicity predominating in different contexts (see Sober 2015). Still, we reject rococo accounts. The question is: why? During his scientific realist phase, Hilary Putnam ventured the hypothesis that the laws of nature could be no more complicated than differential equations (Putnam 1975: 309). It might seem that a hard-nosed scientific realist, like Putnam at the time, should endorse such a claim only if God assured him that the challenge of figuring out the way the world is would not be too difficult for humans to meet, in much the way that a teacher might assure her students that the exam will not be too hard for them if they study assiduously. We have no such assurance. So if we have to back Putnam's hypothesis with truth-conducive reasons, it is unwarranted. We should admit that we have no clue how complicated the laws of nature are likely to be. Alternatively, however, we might take Putnam's proposal to be an aesthetic constraint on acceptance. Perhaps scientific investigators have such a strong aversion to mathematical complexity that they balk when the equations get too complicated. Being convinced that the complicated equations couldn't be right, they consider the investigation that led to them to still be open, and seek a way to either replace them by or reduce them to something simpler. Again there is no guarantee that they will succeed. But such an aesthetic preference would explain their efforts to come up with simpler laws.

Elegance

Elegance in science is a combination of effectiveness and economy. Effectiveness is an instrumental matter. There is something we seek to achieve and when we are effective our efforts pay off. The Miller-Urey experiment was effective in that it sought to demonstrate and succeeded in demonstrating that amino acids emerge from reactions of chemicals – ammonia, methane, hydrogen, and water – believed to be plentiful on Earth in prebiotic times. The four chemicals in plausible proportions were sealed in a chamber that was subjected to occasional sparks which mimicked the effect of lightning. Over several days, a series of chemical reactions occurred, eventually yielding thirteen amino acids (see Ball 2005). The elegance of the experiment lies in its simplicity which engenders a sense of inevitability. Given the experimental design, it appears, nothing but the four chemicals could account for the production of the amino acids. The result was reached with no extraneous theoretical, computational, or material factors. The background assumptions were clear and uncontroversial. The design was straightforward. Indeed, the experiment is so simple it almost looks like a nerdy high school science project.

There is no obvious reason why an elegant experiment is more likely to yield a truth, or an important truth, than an inelegant one. Still, elegance is an epistemically advantageous property. An elegant experiment makes manifest what it achieves and how it achieves what it does. It exemplifies its scientific contribution. An inelegant experiment might disclose the same truth, but we would have a harder time recognizing that or appreciating how it did so. If the inelegant experiment is sufficiently complicated, it invites the worry that unappreciated confounding factors, rather than the hypothesis being tested, account for the result. The elegant result is more illuminating. Either it readily integrates its result into a currently accepted account, or it manifestly poses a challenge to that account.

Optional Stops

So far, my claims have been largely descriptive. Scientists seek symmetry; they favor simplicity; they strive for systematicity; they appreciate elegance. I've urged that symmetry, simplicity, systematicity and elegance are aesthetic properties, but I haven't yet said much about how they contribute to scientific understanding. One contribution has already been hinted at. Deviations from the ideals of symmetry, simplicity, and systematicity demand explanation. They pose a problem that ought to be addressed. Conformity to the ideals needs no explanation. Indeed, any attempt to explain cases that conform to our desiderata is apt to look a bit weird. No one asks why things behave as expected, even when aesthetic factors figure in the expectations. There is then an imbalance in our demands for explanation.

Scientists evidently observe what Nozick calls the Optional Stop Rule (Nozick 1981 p. 2). Inquiry has no foreordained stopping point. If we reach a result we find implausible or unpalatable, it is always open to us to conclude that there must be something wrong with the investigation, the reasoning process, or the background assumptions that led to it. This is what I suggested scientists who follow Putnam's recommendation do when they arrive at a mathematically complicated scientific law. Indeed, a scientist's reason for exercising the Optional Stop rule may itself be aesthetic. Weinberg (1992) suggests that the quest for a Grand Unified Theory is at least partly motivated by a dissatisfaction with the idea that there are a multiplicity of mutually irreducible fundamental physical laws. Rather than accept a result and move on, we can decide to investigate a matter further. This may lead us to refine our methods, question our presuppositions, or look for hidden variables.

There is of course no guarantee that further investigation will lead to a conclusion we like better. It may be that the original inquiry was impeccable and the result, although unpalatable, was correct. Convinced that each individual event has a cause, a physicist might invoke the Optional Stop Rule and insist that there must be a reason why a particular radioactive particle decayed at time *t* and another, seemingly identical one did not. Rather than recognize that radioactive decay is stochastic, he might insist that scientists should keep looking. Once a no-hidden-variable theorem has been proven, it might seem, we reach a natural stopping point. What more could we want? The difficulty is that the Optional Stop Rule still applies. Rather than accept the conclusion that there is no hidden variable, we can pursue the suspicion that there is something wrong with the theorem's proof. Such resistance in not ruled out, even if there is nothing more to find.

The point is not that we are always right in our assessments. It is rather that a plausible, palatable result is, in large part because of its plausibility and palatability, apt to be deemed acceptable. Then inquiry with respect to that question ends. A sufficiently implausible or unpalatable result is apt to spark further inquiry. How could it be? The result is treated as a challenge. More work needs to be done.

The Optional Stop Rule might itself seem untenable. It seems rather unfair, even prejudicial, to subject objectionable results to greater scrutiny than attractive ones. But actually, the rule and the treatment it prescribes are reasonable. The results that strike us as plausible and palatable typically are ones that are readily integrated into accounts in reflective equilibrium. And they strengthen the accounts they are integrated into.

A constellation of epistemic commitments is in reflective equilibrium when its components are reasonable in light of one another and the constellation as a whole is at least as reasonable as any available alternative in light of our antecedent commitments. In constructing a system of thought we begin with whatever antecedent commitments (beliefs, norms, methods, goals, etc.) we take to bear on the topic we seek to understand and the sort of understanding we seek to achieve. These are a motley crew. They are apt to be mutually inconsistent, and even where consistent, non-cotenable. They are likely to be gappy, failing to cover matters we think they should cover. They frequently contain errors, omissions, and other cognitive infelicities. They are not acceptable as they stand. But they encapsulate our current understanding of the phenomena, the ways to investigate them, the norms that bear on acceptability, and so on. So we start with them. I call the starting points initially tenable commitments (see Elgin, 1996). We correct, excise, extend and amend them to bring them into accord. Although none of our initially tenable commitments is completely unrevisable, commitments have different degrees of cognitive inertia. We are more reluctant to abandon some than others. Even this reluctance is defeasible. Should it turn out that the price of retaining a commitment is too great – that is, should it turn out that we can only retain it by revising or rejecting other commitments that we think are collectively more worthy of acceptance – we will abandon even a commitment with considerable cognitive inertia (see Elgin 1996).

Aesthetic considerations may be initially tenable. At the outset we might, like Quine, simply have a fondness for desert landscapes. No matter. If we think that theories should be simple, or that laws should disclose symmetries, we can begin by favoring accounts with those features. In that case, we build into our theorizing a bias against complexity and asymmetry. This may seem question-begging, but it is not. Or anyway it is no more question-begging than a bias in favor of comprehensiveness or evidential adequacy. In any case, the bias is readily overrideable should it prove too costly. On the other hand, at the outset, we may lack aesthetic biases, having no preference for elegance, simplicity, and the rest. Then our early attempts at adjudication will be indifferent to such aesthetic considerations. But should we find that simplicity, elegance and the like are characteristics of the accounts we ultimately endorse, and discover that the closest competitors that lack those characteristics are, on the whole, less tenable, these aesthetic characteristics acquire the status of initially tenable commitments that (at least weakly) constrain future theorizing. Considerations that display these characteristics will, ceteris paribus, have an easier time gaining admission than those that lack them.

Still, to say that we would like our theory to display a certain aesthetic profile does not assure that we will get what we want. Every component of the system is subject to review. We may find that, for example, cost of simplicity is too high. If we have to sacrifice, say, a considerable measure of comprehensiveness or evidential adequacy to satisfy our current criterion for simplicity, we have a prima facie incentive to revise the criterion. We may abandon it completely; or we may restrict its scope, concluding that e.g., the payoff for recognizing a multiplicity of distinct elements has advantages that ontological monism cannot match. On the other hand, scientific developments may strengthen and reinforce an aesthetic commitment. Arguably symmetry was a rather peripheral epistemic value in classical physics. In quantum mechanics, it has moved to center stage.[2]

So we are within our epistemic rights to initially prefer theories, models, and experiments that display particular aesthetic profiles. When we do so, aesthetic factors play a regulatory role. They make no claim to track truth or empirical adequacy. Rather, their role as gatekeepers is to shape our accounts, encapsulating our evolving ways of framing our understanding so that the truths and untruths, the empirically adequate and empirically inadequate considerations that we accept are in reflective equilibrium, and therefore worthy of our reflective endorsement.

[2] I am grateful to Steven French for this point.

References

Albert, David (1996). 'Elementary Quantum Metaphysics' in *Bohmian Mechanics and QuantumTheory: An Appraisal,* ed. J. T. Cushing et al. Dordrecht: Kluwer, pp. 277-284.

Aristotle (1941). *Metaphysics. The Basic Works of Aristotle* ed. Richard McKeon. New York: Random House, pp. 689-934.

Ball, Philip (2005). *Elegant Solutions.* Cambridge: Royal Society of Chemistry.

Bell, Clive (1913). *Art.* New York: Frederick A. Stokes.

Cartwright, Nancy (1983). *How the Laws of Physics Lie.* Oxford: Clarendon.

Della Rocca, Michael (forthcoming). *The Parmenidian Ascent,* Oxford: Oxford University Press.

Elgin, Catherine (2017). *True Enough.* Cambridge: MIT Press.

Elgin, Catherine (1996). *Considered Judgment.* Princeton: Princeton University Press.

McAllister, James (1996). *Beauty and Revolution in Science.* Ithaca: Cornell University Press.

McAllister, James (2007). 'Model Selection and the Multiplicity of Patterns in Empirical Data', *Philosophy of Science* 74: 884-894.

Putnam, Hilary (1975). 'On Properties' in his *Mathematics, Matter and Method* Cambridge: Cambridge University Press, pp. 305-322.

Sober, Elliott (2015). *Ockham's Razors*. Cambridge: Cambridge University Press.

Strevens, Michael (2008). *Depth*. Cambridge MA: Harvard University Press.

Weinberg, Steven (1992). *Dreams of a Final Theory.* New York: Pantheon.

Wilczek, Frank (2016). 'How Feynman Diagrams Almost Saved Space' *Quanta Magazine*, July, 5.

Van Fraassen, Bas (1989). *Laws and Symmetry.* Oxford: Clarendon.