1

EVOLUTION AND THE METABOLISM OF ERROR

Biological Practice as Foundation for a Scientific Metaphysics

WILLIAM C. WIMSATT

1. EVOLUTIONARY FRAMEWORK

The idea that biological practice should provide a foundation for a scientific metaphysics may seem bizarre. Prior attempts to construct a scientific metaphysics have had different aims and used different methodologies, seeking deductive derivations of upper-level objects, concepts, and laws from fundamental entities and relationships, usually from the lowest-level (smallest) entities and relationships of contemporary physical theory, and sought accounts according to which the upper-level things were "nothing more" than these lowest-level posits (e.g., Steven Weinberg). The other strategy is to begin from phenomenological posits, like sense data, which were seen as immediate experience of which one had direct and certain access, out of which one could deductively construct higher-level entities (objects) and relations. The first of these was an attempt at an ontological reduction, the second an epistemological one. In either case, the aim was to justify scientific theory and practice as certain and error-free (e.g., Bertrand Russell or Bas van Frassen). The path I take is a third, which recognizes the fundamental fallibility and error tolerance of our adaptive information-gathering interactions with the world. It seeks ways to ground our natures and processes that are neither deductive nor eliminative but fundamentally related to the nature of evolutionary processes and ways of producing order in the world that are intrinsically heuristic, error-tolerant, and consistent with actual adaptive constructive processes. This ultimately necessitates not only a metaphysics based on biological practice but also a correspondingly based epistemology and methodology. I aim to provide all three.

2. AN ATTEMPTED BLACKBALL FOR ANY POSSIBLE SCIENTIFICALLY BASED METAPHYSICS

When I was an undergraduate, there was an analytically based argument against the appropriateness of any possible scientifically based metaphysics. This was to note that all scientific theories were empirically based and fallible. And not only that, given the historical record, any current or new theory was overwhelmingly likely to be false. As such, it was argued, scientific theories were categorically the wrong kind of thing on which to base a metaphysical system. But this is too quick. One can accept their fallibility and even their falsehood without being ready to accept that new theoretical revisions would totally overturn our expectations about the behavior of nature. I would argue that indeed, in main outline, our macroscopic regularities would be preserved through any possible scientific revolutions. Nothing, for example, could overthrow the main results of evolutionary theory or classical mechanics. This is true for two distinct reasons: first, most macroscopic regularities are at least approximately true “sloppy, gappy causal regularities” in the sense to be expounded later, and these are robust. This means that they are not dependent on fine details of the underlying processes: changes in these processes, or in our accounts of them, do not change the macroscopic regularities. (See Batterman 2021 on multiscale explanations, generalizing, and better founding an argument I first made in 1981.) As Cartwright (1983) argued, the ideal gas law is safe in any future scientific revolution. The second reason is that it is an adequacy condition for any new theory that it be able to capture the main successes of the older theory that it replaces, most commonly as limiting cases. I claim that the elements I draw from biological practice are both robust and sufficiently central to be preserved in any future scientific revolutions. They are thus, in those respects, appropriate bases for a scientific metaphysics.

The idea that biological practice should provide a foundation for a scientific metaphysics has at least two anchor points. The first is that evolution itself as a process provides constructive principles that are exemplified in our normal and scientific activities because of the unavoidable fact that we are biological organisms and pursue our activities in the natural world that conditions physical processes including biological evolution. As Donald

Campbell (1974) said, “Every case of fit between a system and its environment is a product of selection.” This fact alone argues for at least a basic recognition of the demands of the real world and a minimal scientific realism. It also provides a basic architecture for generative systems, including deduction, as extremely productive heuristics, as we shall see later. The second is that the investigative approaches used in science and manifested particularly clearly in biology point naturally to ontological, epistemological, and methodological approaches and conclusions. These insights may have emerged originally in biology but are no less applicable in the physical sciences once we know how and where to look.

My own work, beginning early in my career with my dissertation on functional organization (Wimsatt 1971; 1972) and paper on complexity and organization (1974) drawn from it, exemplifies this and provided a foundation for the functional analyses and complexities of scientific practice (including levels of organization and the role of mechanism) that I have pursued since (Wimsatt 1976a; 1976b; 2007; 2021b). Because practice is an end-directed activity, this analysis of function provides a framework for all of our activity and directly informs all of my other work. Thus, I take a teleological worldview of scientific and conceptual activity and return in a way—despite the substantial differences—to an Aristotelian worldview.[1] I must confess that I never (or rarely) thought of myself as a metaphysician until I was confronted with my own texts by Ken Waters and Alan Love, who reminded me that I was after all making metaphysical claims, and quite substantial ones. Thus, I join happily with the metaphysicians in this volume, although I recognize that traditional metaphysicians might find us a strange group. But it is the aim of this volume in part to convince them, through our efforts and products, that our memberships are fully paid up.

Robustness is usually viewed primarily as an epistemological tool, but it reveals as real objects other than those picked out by traditional metaphysicians, supporting a recognition of the “ontology of the tropical rainforest” rather than the traditional foundational minimalist and eliminative commitment to Quine’s “desert ontology.” So, as I argue later, it reveals an ontology as well and leads to a different way of proceeding than is common among traditional metaphysicians and epistemologists. It is also particularly interesting as a widespread architectural feature of biological organization for most more important biological functions, but the reach of this architecture extends beyond into all scientific domains, both theoretical and experimental and, as we will see, including mathematics.

3. LAWS VERSUS SLOPPY, GAPPY CAUSAL GENERALIZATIONS

Attention to evolutionary processes leads to another break with tradition in foundationalist scientific metaphysics, where exceptionless laws build the architecture of nature and the assumption is that only exceptionless regularities can be causal. The importance of "sloppy, gappy generalizations" or "sloppy, gappy causal regularities" rather than exact exceptionless laws in the compositional sciences and in engineering and technology is itself foundational[2] and due most immediately to the fact that the compositional sciences deal with mechanisms, which have tolerance levels, lifetimes, and failure rates and operate in a variety of contexts.[3] They are ultimately a product of the fact that selection in both the natural and human sciences and in the domain of technology operates most commonly by incremental differential improvement relative to the current state, so that the criteria for something to be selected are: "Is it better than what we are now using?" and, ultimately, "Is it good enough often enough?" This leads causally not only to approximate solutions but also, for material things, to the need to deal with natural populational variability that frustrates attempts at universal exceptionless generalization.

Recognizing the central role of populational variability leads us to recognize that although conceptually, causation requires that exactly the same cause (exactly the same state specification) in exactly the same circumstances will produce exactly the same effect, this formulation, while true, is irrelevant, since it is never realized in nature.[4] What is required is a formulation to recognize the kind of variation and qualifications we must deal with in the real world—a kind of statistical robustness: *quite similar causes in quite similar circumstances will sufficiently often produce quite similar effects.*

The introduction of variability within tolerances ("sloppy"), design limitations, and failure rates (exceptions or "gaps") is most obvious in technology, where design processes, confronted with the contingencies of the material world, explicitly include tolerable error rates in production and performance and designed lifetimes. But the same considerations clearly apply for biological organisms, which vary enormously in characteristic reliability and lifetimes in ways appropriate to their respective ecological niches, be it a bacterium or a baobab tree. Exceptionless laws are idealizations for real compositional systems but false or only approximately true and unrealistic for the real world. Selection goes for the satisficing solution. Like all organisms, we not only are designed to deal with error, but we utilize errors to gain information from the environment and to make more reliable structures in an ontological replay of natural selection—thus, "the metabolism of error." These "satisficing" and "fallibilist" desiderata do not articulate well with traditional foundationalist and logical paradigms, which mistake idealizations about decision making ("maximization of expected utility") and scientific inference (deduction) for analytic and foundational principles.

Another casualty of rejecting deductive foundationalism is eliminative reductionism, or "nothing-but-ism." Russell was skeptical of theoretical entities, describing them as "logical fictions," and many reductionist scientists wanted ultimately to explain everything at the lowest possible level—or, as Roger Sperry ironically put it, to explain "eventually everything in terms of essentially nothing." I have argued instead (Wimsatt 1976a; 1976b; 2007; 2021b) for a multilevel explanatory account of "articulatory reduction" in which lower- and intermediate-level interactions explain upper-level phenomena (like Brownian motion) and regularities (like the ideal gas law or the van der Waal equation of state). In this case, both the upper-level phenomena and the regularities delineating sloppy, gappy causal generalizations are explained but not eliminated by the articulation of lower-level causes. (See also Batterman 2021.)

The "Ur-principle" of the metabolism of error is profoundly heuristic and evolutionarily based and produces intrinsically heuristic robust and reliable structures that best serve the complex interactive networks of natural and human ends in our world. I will discuss the nature of heuristics after we discuss robustness.

My deepest problem with traditional philosophical views of science and of ourselves as scientific agents is thus that they adopt "in principle" idealizations of our aims, our scientific products, our data, and our methods for achieving them that are profoundly false and inappropriate for limited beings and that both flaw our attempts to truly characterize our activities and prevent us from utilizing the strengths that we can and do deploy in pursuing our aims (Wimsatt 2007). So then what resources do we have for fallible beings?

4. ROBUSTNESS

Robustness is the use of multiple independent means that may individually be of less-than-perfect reliability (often far less) to secure a net reliability higher (and often under a wider range of conditions) than possible with any single method (von Neumann 1956; Wimsatt 1981). This means that by incorporating multiple alternative and independent lines of argument to the same conclusion, we can use heuristic and inductive methods of inference as well as deductive inference in drawing reliable conclusions. Robustness can serve in the reliable and more accurate detection of objects, the construction of reliable inferences, and the production of stable results. To a deductivist, this may look like sloppy (even otiose) redundancy, but it fits our natural methods for securing reliability far better than our deductivist idealizations and allows use of our rich armamentarium of heuristic methods. Physicist Richard Feynman (1968, ch. 2) contrasted a "Babylonian" with a "Euclidian" methodology and clearly favored the former. In a "Babylonian" theory, theoretical elements were overconnected, redundant, and often approximations, and consequently results could be cross-checked and calibrated and were reliable under diverse failures. This was preferable in the real world to an elegant, sparse deductivist ("Euclidean") theory, which collapsed like a house of cards if anything failed because it had no redundancy. Our deductivism supposedly captured the structure of mature physics, though Feynman would have disagreed. Feynman is in effect arguing for a biological architecture for theories and doing so for physics as well as for biology.

But it goes deeper metaphysically. As I have argued (Wimsatt 1981), the very notion of a real object, is of something which has multiple properties, thus also has multiple means of detection and is robust. This strategy does a runaround on the traditional "argument from illusion" for the unreliability of the senses in detecting the real world by arguing that the independent means will tend to fail under different conditions and thus will reliably detect the external world when used conjointly. This is why robustness has ontological implications. (See also Eronen 2015.) And not just for objects. Within the philosophical tradition begun by Galileo and Descartes, it was the primary (multimodally detectable) qualities that were real (because they were not delimited by, and thus seen as external to, any of the senses), and the secondary qualities (detectable in only one sensory modality) were the subjective effects in us of the action of the primary qualities. That is to say, the primary qualities are robust. It is one of the ironies of the deductivist tradition from Descartes that he urges us to start with robust objects and properties in building our system but never seems to apply this lesson more broadly.

Mathematical ecologist Richard Levins (1966) urged the search for "robust theorems" that followed in multiple different independent models of a phenomenon that made diverse assumptions (Wimsatt 1981; Weisberg 2006; Odenbaugh 2011). He elsewhere spoke of this as "exploring the space of possible models" (Levins 1968). He was in effect urging the heuristic use of the criterion for logical truth as "true in all possible worlds" as a way that we could find results in the real world that we could trust. He was applying it in mathematical ecology, where all of the models used were strongly idealized approximations that were literally false. Thus, he urged that "in ecology, our truths are the intersection of independent lies" (Levins 1966). But more broadly, mathematicians in diverse areas particularly value results that show robustness through multiple independent derivational paths over those that do not (Kromer 2012; Corfield 2010; Avigad 2020), and as I mentioned earlier, Feynman (1968) argued for a similar architecture for physical theory. This is a way of preventing or massively reducing the probability of inferential failures when real people incorrectly assess the soundness or validity of an argument, and it is a principle respected and used in the most reliable of disciplines.

Robustness is the scientist's answer to the philosopher's fruitless quest for certainty. Within psychology and the social sciences, Donald Campbell (1966) sought reliability and validity through "triangulation" using multiple independent methods to calibrate and correct measurements, contrasting the heuristic that he called "multiple operationalism" with the single-linked "definitional operationalism" of philosophers and psychologists, which he regarded as fallacious. Glymour (1980) urges multiple connectedness in the localization of faults in theories (thus arguing against the Quine-Duhem hypothesis of the unfalsifiability and underdetermination of theories) and notes that a claim is untestable unless it is accessible in two or more independent ways. Robustness, though not infallible, is in effect a kind of *overdetermination* of theories. Tom Nickles gives creative extensions and qualifications to the idea of robustness in extending its use to inferential systems (Nickles 2012).

Robustness is characteristic of biological organization as well as principles of inference, as a product of selection, with redundancy, multiple means, and excess capacity to reliably accomplish all of their important functions (Wagner 2005), including their evolvability. They do so, remarkably, with stable heritability of phenotypic traits and fitness under sexual recombination (Wimsatt 1987; 2007), without which cumulative evolution would be impossible. So the same organizational design principles are found and centrally important in nature to secure reliability. So it is not surprising that scientists should follow a similar methodology in their arguments, choosing to trust entities and properties that are multiply detectable.

For most of the last century, formalist and foundationalist ideas have substantially influenced our own conceptions of what we are doing as philosophers, under the aegis of logical empiricism and analytic philosophy more broadly, but we should pay more attention to the "Babylonian" methodology of scientists. Heuristic procedures for solving problems are a species of cognitive adaptations. They are error-prone but cost-effective and applied in cross-checking ways in the complementary reduction of error. Many of the so-called "informal fallacies" of inference are prone to error but are nonetheless often quite effective and reliable. Thus, "appeal to authority" is a supposed fallacy, but it is used endemically in science, and necessarily so. No single person can know, much less validate, all of science. Campbell (1974) spoke of science as using a "95% doubt/trust ratio," in which we selectively doubted the elements under test, moving from one to another, while trusting in each case the other 95 percent—most of which we will have to take on authority or past history. This too was the lesson of Neurath's boat[5] and resonates with both biological architectures and adaptations and with real scientific methodologies. Most of the fallible inference principles that we use to increase reliability in robust arguments, such as "appeal to authority," are heuristics, so heuristics are the next topic to consider. The study of heuristics both suggests and calls for an entirely different viewpoint in constructing philosophical methodology.

5. HEURISTICS IN NATURE AND IN US

I have argued that most scientific tools of discovery, evaluation, experimental design, calibration, and theory construction are heuristic procedures.[6] I have offered a fuller analysis of these than is to be found anywhere else in the literature (Wimsatt 2007). All of the ones of which I am aware (I have primarily studied problem-solving strategies) share six properties and commonly exhibit two others. As an important marker of the connections between biological architectures and our seat-of-the pants inference procedures, with each I note an analogous or identical feature of biological adaptations. These property pairs are:

(1) By comparison with truth-preserving algorithms or other procedures for which they might be substituted, heuristics make no guarantees that they will produce a solution or the correct solution to a problem. A truth-preserving algorithm correctly applied to true premises must produce a correct conclusion. A heuristic need not.

(1′) Similarly, no adaptation guarantees success of an organism in achieving that which the adaptation is for.

(2) By comparison with procedures for which they are substituted, heuristics are "cost-effective" in terms of demands on memory, computation, or other limited resources. (This is why they are used instead of methods offering stronger guarantees.)

(2′) We take for granted that adaptations are designed in a cost-effective way—often "quick-and-dirty solutions" that work "well enough often enough."

(3) Errors produced using a heuristic are not random but systematically biased:
  (a) The heuristic will tend to break down in certain classes of cases and not others, but not at random.
  (b) With an understanding of how it works (viewing heuristics as mechanisms), it should be possible to predict the conditions under which it will fail.
  (c) Where it is meaningful to speak of a direction of error, heuristics will tend to cause errors in a certain direction, again a function of the heuristic and of the kinds of problems to which it is applied.
  (d) Heuristics may thus leave "footprints" indicating their application—results exhibiting their characteristic biases. One can work back from these to determine that a heuristic was applied and sometimes even which one.

(3′) Adaptations have a specific design that can be made to fail with a suitable choice of environmental or experimental conditions; for example, light-seeking tropism of the caterpillar of the goldtail moth is selected for because it characteristically leads the caterpillar in the spring to the top of the plant where the new leaves on which it feeds are to be found. But in a suitably designed experiment with a T-maze with the light at one end and the leaves at the other, the tropism will lead it away from rather than toward food (again, exploiting the weaknesses of the mechanism).

(4) Application of a heuristic to a problem yields a transformation of the problem into a nonequivalent but intuitively related problem that is easier to solve. Answers to the transformed problem may not be answers to the original problem, though various cognitive biases operative in learning and science may obscure this. Thus, for example, the approach to problems of heredity through transmission genetics proved much easier by ignoring the problem of how genes produced traits, and for a while transmission genetics was mistaken by some for the whole solution to the problem of heredity. (The way to completing the solution by adding an account of gene action in producing the phenotypic manifestation was begun with Jacob and Monod's elucidation of the nature and operation of the *lac* operon.)

(4′) Similarly for adaptations—for example, sensing temperature seasonality (and impending cold) by detecting changing day length.

(5) Heuristics are useful for something: they are purpose relative. Tools that are effective for one purpose may be bad for another, and increases in performance in one area are commonly accompanied by decreases elsewhere (Levins 1968). This is an instance of the "generalist" versus "specialist" trade-off. Recognizing this may help identify or predict their biases: one expects a tool to be less biased for applications it was designed for than for others it is co-opted for.

(5′) Adaptations that are clearly suboptimal because they are co-opted from a structure or procedure originally selected to serve another purpose or are subject to entrenched constraints imposed by the architecture of other adaptations are endemic in biology.

(6) Heuristics are commonly descended from other heuristics, often modified to work better in a different environment. Thus, they commonly come in related families, which may be drawn on for other resources or tools appropriate for similar tasks. Thus, Douglas Lenat (1981) exhibits some sixty heuristic rules for his theorem prover, which can also be seen as sixty instantiations of the same heuristic with slightly different antecedent conditions. Similarly one can find an array of different kinds of hammers to serve different but related purposes. On different scales of resolution, a family of heuristics may look like a single heuristic, or conversely.

(6′) Adaptations show a pattern of similarities with variations, both adaptive and not, among individuals and, at a larger scale, across species.

What does it mean that heuristics and adaptations share these six key properties? That heuristics are reducible to adaptations? No—rather that they share a common functional logic as products of evolutionary processes. But are there any other properties they might share as a result of this origin? I discuss one here, and one in the next section.

Robustness and adaptive radiations: I want to suggest that heuristics are useful in direct proportion as they are robust, even if they are both "sloppy" and "gappy" in ways indicated by the preceding six properties, so robustness is a desirable and common feature of good heuristics.

(Cartwright [1983] marks robustness as an important property of generalizations, and it is so here as well.) A heuristic that worked only under one exactly specified set of conditions and not for any of its neighbors would be of little utility, if it could be used at all. One must be careful here, since highly specialized heuristics may require very special circumstances to be useful. But when this is true, they will most commonly be a member of a larger family with a common principle that may be fruitfully specialized in different directions for different specialized uses. Thus, an adaptive radiation of heuristics is evidence of a kind of robustness of a kind of adaptation—a design principle for a class of artifacts. Thus, there are all kinds of hammers and of scissors and of threaded fasteners, each adapted to similar functions in different circumstances. (Herkimer's *Engineer's Illustrated Thesaurus* [1952] illustrates this with multiple examples in families of mechanisms.)

6. GENERATIVE ENTRENCHMENT, CONSERVATION, AND QUASI-RECAPITULATION: WHY HISTORY MATTERS

Secondly, both heuristics and adaptations, through their origin as co-opted and reselected variants of existing adaptive methods, show the marks of their history in preserved entrenched properties. This may seem merely accidental, but it is an efficient and adaptive strategy when trying to create a new tool to use what is already at hand. The co-option of existing adaptations as tools to serve new functions is central to evolutionary processes and endemic throughout, as first argued by Stephen J. Gould. It brings new selection pressures to bear on both the tool and the systematic practice or system of which it is a part. This is a major driver of new adaptive innovation and source of new complexities and opportunities.

Thus, in Wimsatt (2007), I compared nature to a "backwoods mechanic and used parts dealer," an image masterfully delineated by Douglas Harper in his *Working Knowledge: Skill and Community in a Small Shop* (1987). There the main protagonist, an able and ingenious mechanic, lived in rural upstate New York and specialized in maintaining older Saab cars. He seldom had exactly the right part in his yard of junked Saabs, but he could usually adapt (or "co-opt") and modify a similar part from another year, often in a way that improved on the intended replacement. This too uses design similarities resulting from generative entrenchment. *Generative entrenchment is an unavoidable consequence of evolution and of the evolution of increasing complexity either in biology or in cumulative culture.* In explanation, whether biological or cultural, it is the main reason that history matters.

In co-opting an element from an existing design for a new function, there may be elements, often external to the system, that aid in achieving that function. These are called scaffolds (Wimsatt and Griesemer 2007; Caporael, Griesemer, and Wimsatt 2013). A scaffold may be a behavior, an object, or a property, but in scaffolding it must exhibit an activity. A scaffold that is useful in achieving a function, if external, may subsequently be incorporated into the system for future use, increasing the organizational complexity of the design. In a complementary activity, elements of a constructed niche may be useful for other systems and become common infrastructure, tying them together into a larger ecosystem. Such processes occur with our technologies, which are highly dependent, and with ecological communities like coral reefs, where the coral backbones provide shelter and other resources for a diversity of other species.

But entrenchment alone does not determine organizational form. There may be multiple such forms, but particularly important is one whose heuristic benefits in generating easily an array of possible variations systematically have resulted in its relatively frequent occurrence and central importance in the history of evolutionary systems. These are entrenched combinatorial alphabets.

7. THE HEURISTIC ADAPTIVE ORIGINS OF COMBINATORIAL GENERATIVE SYSTEMS (INCLUDING DEDUCTION WITH FOUNDATIONAL ELEMENTS)

Elements can become fixed in their details because so much else depends on them. This is the basis for Francis Crick's famous comment that the genetic code is a "frozen accident." This I call generative entrenchment: something remains fixed or relatively stable in evolution because of its generative role; as more things come to depend upon it, the chance increases that a change in it will cause problems downstream and, if they are large enough, lethality. This decreased chance means that if other solutions exist, the waiting time to find one through random mutation is longer. This differential dependence thus leads, probabilistically, to degrees of stability with degrees of dependence and increasing stability for older elements. This feature makes cumulative evolution, both in biology and for culture and technology, possible (Wimsatt 2019). Generative entrenchment also gives biological organization a tendency to be hierarchical and fundamentally historical (Wimsatt 2015).

But another crucial property of the genetic code is its use of a small number of elements arranged in different combinations as inputs to an apparatus to assemble a correspondingly diverse array of molecular machines—linear strings of amino acids that fold into active proteins. This points to another important process and class of adaptations. As with the genetic code, entrenchment can lead to the emergence of a class of standardized parts, as with standardized thread sizes and profiles in threaded fasteners (Herkimer 1952; Wimsatt 2013). They must be standardized (within tolerances) to be interchangeable, and the tolerances commonly become more demanding as they are utilized in multiple contexts. These can themselves become "tinkertoys" *or combinatorial alphabets* in the construction of a diverse class of adaptive machines. Thus, amino acids, cells, standardized interchangeable mechanical and electronic parts, words, and program instructions are basic "alphabetical" elements capable of making large and diverse classes of machines. Some of these systems are specifically designed to facilitate reuse of parts in the evolution of complex adaptations. (Object-oriented programming, with the standardized interfaces of its objects, is also an appropriate example here. The activity there called "program maintenance" is actually more accurately called "program evolution," with the standardization of interfaces allowing the reuse of program objects in new contexts.) These are the source of major adaptive radiations and of efficient generation of "bottom-up" modular variation and increases in complexity in adaptive machines capable of occupying diverse niches.

In addition to an alphabet of stable standardized elements that can be arranged in various ways, we require for each a set of limited and constraining rules for how they may be combined—a syntax. In these, we see the basic requirements for a formal system: a set of basic elements and a set of rules for assembling them into well-formed structures, or "wffs." This is the architecture of a deductive system—and thus of all formalistic and traditional foundationalist approaches. But this arrangement is far broader. This same set of properties—basic alphabet and combination rules—is true also for atoms in making molecules, which makes "computationalism" so plausible as a form of physicalism. And it is so also for interchangeable machine parts. The stability of the standard thread specifications gives a standardized way of connecting mechanical elements. These two features are central to the generative systems mentioned previously, abstracted for the first time in the elements and propositions in Euclid's geometry and subsequently abstracted further in the development of formal symbolic logic. The possession of these two elements was the basis for Chomsky naming his theory in *Syntactic Structures* (1957) a "Generative Grammar," but if we look at Herkimer's *Engineer's Illustrated Thesaurus* (1952), we see the same thing: arrays of mechanical parts, grouped by function, mechanism, and means of articulation or connection. The last give ways of combining basic elements into a variety of diverse structures to accomplish different adaptive tasks. Function and mechanism relate to their engineering meaning, and means of articulation indicate their syntax: what can be connected to what and how.

Similar and related insights for deductive systems as adaptive structures are nicely reflected in Jeremy Avigad's (2020) analysis of the problem of reliability in the practice of the use of deductive systems in mathematics. There he notes not only the use of multiple connections or deductive paths to increase reliability and robustness but also the use of modular elements in the localization of error and the different combinatorial arrangement of these elements—an important and widespread adaptive design feature not discussed here (but see Wimsatt 2021a).

Deductive foundationalism arises from reifying the generative power and idealized error-free transmission by deductive systems of truth-like properties, starting from elements, relations, and assumptions held to be certain or true, either of basic physical properties and relations or of symbolic or logical ones, and attempting to generate or ground a much broader array of trustworthy statements, relations, and properties. These elements and the syntax used to transform them become foundations of logical inference. But taking this too seriously leads to an erroneous extrapolation of what we find to what exists in nature. This kind of constructional system is a powerful adaptation, but primarily a heuristic for dealing with the complexities in nature. In nature, nothing is certain, so we should try to ground things in robust objects, properties, or relations (like Descartes's primary qualities). As we build a system, the basic elements and transformations become successively more irreplaceable through generative entrenchment, just as the elements of the genetic code and its translation into proteins become increasingly universal and increasingly unmodifiable. But increasingly unmodifiable does not mean it is absolutely so. Any deep modifications would be very difficult, usually lethal, but when they work absolutely revolutionary. So what we need is not an absolute foundationalism but a dynamical one, one that can accomplish the ongoing repairs to Neurath's boat, however rare and difficult they may be. So a naturalistic foundationalism must be a dynamical foundationalism.

On a smaller scale, this would replicate what we see in a scientific revolution. Does this suggest a (non-progressive) Kuhnian revolution? No, it cannot be, and in fact Kuhn's views are flawed because an adequacy condition for any replacement or modification (an acceptable plank for Neurath's boat) is that it must capture the major successes of what went before. The new plank must fit (Janssen 2019; Wimsatt 2007; 2021a). It must be a functional equivalent (sometimes only approximately) for what it replaces. And now we have found the fallible, evolvable, and progressive heuristic for the generation and improvement of complex adaptive structures. The importance of deductive systems remains, but as an important heuristic methodology for problems where there is enough structure for it to be appropriate. And with this, it is seen to be no longer adequate as a traditional foundational architectonic.

8. COMBINATORIAL COMPOSITION: DEFINING, POPULATING, ANALYZING, AND USING A SPACE

Combinatorial strategies can be powerful and productive heuristics for the organized definition, population, and exploration of a space of possibilities. Definition of such a possibility space is a mark of a well-structured problem, and one of the most effective strategies in trying to bring structure to a problem that is ill-structured (Simon 1973). The Cartesian product of the defining dimensions of the space characterizes the possible elements of that space and permits systematic exploration of it and of trajectories defined by piecewise changes of the values of variables in its dimensions. Entries will have neighbors in each of these dimensions, defining their multidimensional neighborhoods, and these spaces have other useful properties. Definitions of such a space are connected with the notions of closure and completeness in logical systems. I have worked at length with two cases utilizing such spaces, one conceptual and one scientific.

The conceptual case is the characterization of types of emergence. To most scientists, unlike most philosophers, emergence and reduction are compatible, and emergence can be seen when the properties of the system are products of how the parts articulate together. But there is no direct way to characterize these different possible modes of organization. However, one can approach it negatively. Consider when the value of a system property is aggregative, so that it is nothing more than the sum of values of that property for its parts—in other words, when their organization does not matter. I delineate four requirements for aggregativity. These can be met or can fail in various ways and to various degrees in various combinations—leading to fifteen different ways in which a system property can be emergent, with aggregativity as the sixteenth case. This provides a useful combinatorial classification for different modes of dependence of a system property on the organization of its parts and tools for analysis of organization. That they are degree properties extends its usefulness by providing sensitive tools for dealing with approximations. This is discussed in Wimsatt (1997) and vastly elaborated and applied in Wimsatt (2007, ch. 12). It has rapidly become one of my most widely cited works.

Probably the best-known scientific cases are the discussions of genotype space and protein space initiated by John Maynard Smith in 1970 and widely used for a variety of conceptual and scientific arguments since. They are worth further systematic study of their use. The scientific case I have considered in detail is related: the development of the Punnett square (Wimsatt 2012) and its role in genetics through a compact and clear characterization of the possible genotypes produced in diploid matings among an array of possible gametic types. This plays a crucial role in defining matings in classical and population genetics. As discussed there, the form of algebraic expansion provided by Mendel was far less compact and intuitive, and the neat spatial organization of the Punnett square provided an easy way to reason about matings that becomes exponentially more useful and important as the number of factors considered simultaneously increases. I also discuss there why the spatial organization of the Punnett square simplified representation and problem-solving and probably was a significant factor in eliminating errors when considering multi-locus matings.

9. ROBUSTNESS, PHILOSOPHICAL METHODOLOGY, AND THE METABOLISM OF ERROR

We have elaborated the widespread use of heuristics in evolved and evolving systems in nature and in science and their foundational nature in our practice. But their reach goes further, into philosophical methodology. This "heuristic paradigm" (Tyson 1994) need not replace the current broader philosophical inspiration by various logical and more formalistic paradigms, which themselves have a heuristic origin in the nature of combinatorial generative systems, but it is surely an appropriate complement to them that should give us broader reach and more appropriate tools for a whole class of problems where variations may be familial rather than accidental. How to deal with intrinsically heterogeneous classes has been a problem for philosophy before and has been the origin of the idea of a family resemblance concept invoked by Wittgenstein in his *Philosophical Investigations* (1953). These problems should be expected to crop up particularly for products of evolutionary processes, where causally relevant variation is intrinsic to the mechanisms of change, whether due to selection or drift. This connection with evolutionary processes may sound uninterestingly narrow, but I include in this scope the three great systems of philosophical inquiry: *body, mind,* and *society*—all products of iterative design and selection. Donald Campbell (1974) argued that any case of fit between a system and its environment, including cognitive and cultural products, is a product of selection, and Herbert Simon basically urged the same thing in characterizing his *Sciences of the Artificial* (1969). It is a great irony that many philosophers who are refugees from engineering see that as a close call from which they have happily escaped and from which they have apparently learned nothing. And then they turn to the analysis of these three great designed systems without a clue about evolution or design and convinced that they do not need any (Wimsatt 2007; 2021a; 2021b).

So how do we proceed?

What is the common approach of analytic philosophy in evaluating a text? I was taught repeatedly, in introductory through graduate courses, to do something like the following in analyzing an argument:

1. Identify the intended conclusion.
2. Identify any premises used to argue for it.
3. Reformulate them as well-formed logical statements.

4. Organize them in the form of a deductive argument, leading to the conclusion.
5. If the conclusion does not follow, look for additional premises or alternative interpretations of the existing premises or of the meanings of the concepts in the premises that will yield a deductive argument to the conclusion.
6. If the premises are true, then accept the conclusion subject to the amended specifications and interpretations.
7. If the premises are false, indicate why, give counterexamples, and reject the conclusion.
8. If no set of true premises yielding the conclusion is found, reject the argument.

Here the whole aim is to construct a logically valid deduction of the conclusion from either the original set of premises or an extended or repaired or reinterpreted set. It is a powerful heuristic and can provide a systematic analysis and useful clarification. But it introduces a systematic bias. Usually in practice the aim of an analytic philosopher is critical: to find grounds to reject the argument and the conclusion. *This makes philosophy effectively predominantly negative and builds careers around the demolition of earlier attempts at systematic philosophy.*

But there are other productive ways in which we can approach a set of statements. The heuristics I propose for this class of problems would include the following, and as you can see, they provide plenty of analytical work for philosophers:

- Instead of looking for inexorable arguments, we look for robust tendencies and for conditions under which those tendencies are more likely to be realized.[7]
- Instead of looking for truths, we study errors and how and why they are made.
- Instead of looking for context-free inferences, we study commonly used but context-sensitive ones.
- Instead of classifying them as invalid because they are content or context specific, we should calibrate arguments to determine the conditions under which they work or are "locally valid."
- Look for other plausible assumptions of an inductive or abductive sort that may complement the existing argument.

- We may look for argument schemata, but if so, look for broad conditions where they are likely to work (like looking for the range of validity of a model) rather than trying to demonstrate their universal validity. In this way, we can espouse the use of formal methods but as a tool for appropriate problems, not as architectonic principles.
- Counterexamples become revealing sources of information about the limitations of a model or suggestions for probing its depths; in either case, they are a tool to refine the model, not an argument for trashing the system or something to be swept under the rug.
- It is thus often as important to try to refine, extend, and generalize counterexamples as it is to try to directly correct the original model. This may better illuminate the structure of failures of the original model and thus point to a deeper way to construct a new one. Similar suggestions were advanced by Thomas Kuhn (1962), and all of these preceding points were made, used, and powerfully elaborated by Lakatos in his *Proofs and Refutations* (1978). In this revolutionary work, Lakatos saw mathematical proofs as important means for the refinement of concepts in the light of counterexamples rather than establishment of the conclusion in its original form.
- For heuristics, we are looking at the adaptive structure of our cognition, or specific background assumptions, features of our social organization, or specific characteristics of the problem domain, for either strengths or weaknesses and the conditions under which these are realized. Thus, there is (or we can often extract) a reference context that contains more useful information about the method. This then recognizes methodologically the importance of context dependence.
- Rather than looking for universal theories or principles that are foundational to all other elements of a given domain, look for the conjoint application of robust principles that may be heterogeneous in application but complement each other to give a broader and richer fit to the details of the situation.
- Look for generative ways in which methodologies, empirical results, constraints, and conditions may have broad application to extend or support philosophical viewpoints, looking for the kinds of support that come from the preceding principles rather than entailments or similarly tight linkages. This should include studies of concept and meaning creation, change, and stabilization.

So the eight steps taken in analyzing an argument taught in analytic philosophy courses provide a useful and often powerful heuristic, but we should see them as that rather than an architectonic limitation of our methods. Heuristic methods permeate and constitute the vast majority of inferential tools that we have. It is time that we make a central place for them in our philosophy but also, fundamentally, in our meta-philosophy.

But if we adopt in philosophy all of these heuristic methods modeled more on scientific procedures, what is the difference between philosophy and science? This demands several distinct remarks.

First, remember that philosophy has been (and still should be) the midwife of the sciences, so the methodologies ought not be miles apart (and would not be, were philosophy not trying so hard to mark itself off as a distinct discipline). Remember that philosophy's break with psychology is less than a century old and that the new domain of science studies contains a great deal of philosophy or philosophically relevant material dealing more with social context and interaction than individual activities. It should be seen as complementing philosophy rather than replacing or destroying it.

Philosophy deals with concepts and with inference, both also the domains of psychology, but in a curious inversion, it is psychology that here has a theoretical interest in them, while philosophy as well as an abstract interest, often has a more applied focus in critiquing specific concepts and inferences! At times of significant theory change, the divide between philosophy and the sciences is harder to find because the theoretical revisions will generally involve both conceptual change and new experimental tools and methodological approaches. Thus, philosophers may increasingly require deeper studies of the science, and the scientists may be more open to philosophical input. (Biology has, within the last fifty years, been a productive source of new theoretical and conceptual directions, as this volume demonstrates. Most recently, the revolutionary expansion of the role of simulations within the last fifty years and the growing impact of big data are likely to have significant input to our methods of generating hypotheses and gathering and evaluating data and arguments [e.g., Evans and Rzhetsky 2010]).

Will philosophy disappear or be absorbed within the expanding methodologies of science as Quine predicted? I think not: philosophy should remain at one end of a continuum of methods in philosophy of science and philosophy of nature merging with natural philosophy and the sciences. And it must remain responsive to developments in those sciences. Thus, Quine's views are compromised by his attachment to methodological behaviorism. And even as epistemology, metaphysics, and meta-philosophy are affected by this expansion of biological perspectives, that still leaves ethics and value theory, logic, and history of philosophy. And the newly transformed aspects of philosophy will still remain as philosophical subjects. Philosophy is more robust and multidimensional than Quine supposed, and it becomes so in part by recognizing the role of biological practice in generating a scientific metaphysics—one with broader philosophical implications.

10. CONCLUSION: A MULTIPERSPECTIVAL REALIST METAPHYSICS

Robustness is based on a multiperspectival view of objects held in common. My metaphysical viewpoint is thus both multiperspectival and realist. This multiperspectivalism includes a multilevel mechanistic view that involves emergence and a non-eliminative articulatory reductionism. Because we are talking about practice, it is also intrinsically functional or teleological. We are intrinsically and objectively in the world, so we have no problems with the Kantian *Ding-an-Sich*. Because of entrenchment processes, it is also naturally historical and progressive. Finally, since our heuristics and means for metabolizing error are fundamentally knowledge-gaining processes, our epistemology must be fallibilistic, satisficing, and evolutionary rather than deductivist, maximizing, and foundationalist. And because we are products of evolution, this is an intrinsically naturalistic solution. So is this everything? One class of things not dealt with here is the scaffolding social, organizational, and technological interactions that midwife and extend our abilities to accomplish these activities, but that lead to an account of the nature and processes of cultural evolution, which I address elsewhere (Wimsatt 2019).

NOTES

1. As a satisficer, I would resist the other possibility—Leibniz—because of the optimality assumptions intrinsic to Leibniz's views.

2. It also removes a significant element of the supposed principled distance between rigorous scientific approaches and animal and human behavior and plans.

3. An important early inspiration for me here was von Neumann's classic and foundational essay on building reliable systems out of unreliable components (von Neumann 1956), which I read in Frank Rosenblatt's course

in 1964. A complement to this emphasizing what I have called "the metabolism of error" is Petroski's superb collection of essays (Petroski 1985).

4. This point plays a central role in my discussion of the role of ceteris paribus clauses and their ineliminability in functional assessments in Wimsatt (1972).

5. Otto Neurath famously compared the structure of science to a boat that needed to have potentially each of its timbers replaced (one at a time) while under sail. And for Campbell's "doubt/trust ratio," he vastly understates the trust required—I would put it at greater than 99 percent.

6. The notion of a heuristic was substantially developed by Herbert Simon in the late 1950s in the context of computer simulations of human behavior and decision making (his general problem solver or GPS program). Among psychologists, it was further elaborated in different directions by Tversky and Kahneman (1974), who emphasized the errors of heuristics, and in directions closer to that of Simon emphasizing the positive benefits by me and by Gigerenzer et al. (1999). Among philosophers, Thomas Nickles (2003; 2006) has also developed a systematic account of the uses of heuristics in science. My own account developed beginning in Wimsatt (1980) and is most fully elaborated in several chapters and appendices in Wimsatt (2007).

7. I first felt the need for this perspective when reading Sydney Shoemaker's *Self-Knowledge and Self-Identity* (1963) as an undergraduate. There he argued that it was a necessary condition for a language to work that people usually told the truth—something that we have seen sorely tested and increasingly validated under Donald Trump's regime. Philosophers scoffed at this new modality ("it is necessarily usually the case that") as having no acceptable semantics, but this was clearly a centrally important concept, and they should instead have scoffed at the semantic theories that could not deal with it. Again, they were misled by methodological idealizations.

REFERENCES

Avigad, J. 2020. "Reliability of Mathematical Inference." *Synthese* 198, no. 8: 1–23.

Batterman, R. W. 2021. "Multiscale Modeling in Active and Inactive Materials." In *Levels of Organization in the Biological Sciences*, edited by D. S. Brooks, J. DiFrisco, and W. C. Wimsatt, 215–32. Cambridge, Mass.: MIT Press.

Campbell, D. T. 1966. "Pattern Matching as an Essential in Distal Knowing." In *The Psychology of Egon Brunswik,* ed. K. R. Hammond, 81–106. New York: Holt, Rinehart and Winston.

Campbell, D. T. 1974. "Evolutionary Epistemology." In *The Philosophy of Karl Popper, vol. 2*, edited by P. Schilpp, 12–63. La Salle, Ill.: Open Court.

Caporael, L., J. Griesemer, and W. Wimsatt, eds. 2013. *Developing Scaffolding in Nature, Culture, and Cognition*. Cambridge, Mass.: MIT Press.

Cartwright, N. 1983. *How the Laws of Physics Lie.* Cambridge: Cambridge University Press.

Chomsky, N. 1957. *Syntactic Structures.* Berlin: Walter De Gruyter.

Corfield, David. 2010. "Understanding the Infinite I: Niceness, Robustness, and Realism." *Philosophia Mathematica* 18, no. 3: 253–75.

Eronen, M. 2015. "Robustness and Reality." *Synthese* 192, no. 12: 3961–77.

Evans, J., and A. Rzhetsky. 2010. "Machine Science." *Science* 329, no. 23: 399–400.

Feynman, R. 1968. *The Character of Physical Law.* Cambridge, Mass.: MIT Press.

Gigerenzer, G., P. Todd, and the ABC Group. 1999. *Simple Heuristics That Make Us Smart.* Oxford: Oxford University Press.

Glymour, C. 1980. *Theory and Evidence.* Princeton, N.J.: Princeton University Press.

Harper, D. 1987. *Working Knowledge: Skill and Community in a Small Shop.* Chicago: University of Chicago Press.

Herkimer, H. 1952. *Engineer's Illustrated Thesaurus.* New York: Chemical Publishing Company.

Janssen, M. 2019. "Arches and Scaffolds: Bridging Continuity and Discontinuity in Theory Change." In *Beyond the Meme: Development and Structure in Cultural Evolution,* edited by A. C. Love and W. C. Wimsatt, 95–199. Minnesota Studies in the Philosophy of Science, vol. 22. Minneapolis: University of Minnesota Press.

Kromer, Ralf. 2012. "Are We Still Babylonians? The Structure of the Foundations of Mathematics from a Wimsattian Perspective." In *Characterizing the Robustness of Science*, edited by L. Soler, E. Trizio, T. Nickles, and W. Wimsatt, 189–206. Boston Studies in the Philosophy of Science, vol. 292. New York: Springer.

Kuhn, T. 1962. *The Structure of Scientific Revolutions.* Chicago: University of Chicago Press.

Lakatos, I. 1978. *Proofs and Refutations.* Cambridge: Cambridge University Press.

Lenat, D. 1981. "The Nature of Heuristics." *Artificial Intelligence* 19: 189–249.

Levins, R. 1966. "The Strategy of Model-Building in Population Biology." *American Scientist* 54: 421–31.

Levins, R. 1968. *Evolution in Changing Environments.* Princeton, N.J.: Princeton University Press.

Nickles, T. 2003. "Evolutionary Models of Innovation and the Meno Problem." In *The International Handbook on Innovation*, edited by Larisa V. Shavinina, 54–78. Amsterdam: Elsevier Science Ltd.

Nickles, T. 2006. "Heuristic Appraisal: Context of Discovery or Justification?" In *Revisiting Discovery and Justification*, edited by Jutta Schickore and Friedrich Steinle, 159–82. New York: Springer.

Nickles, T. 2012. "Dynamic Robustness and Design in Nature and Artifact." In *Characterizing the Robustness of Science*, edited by Léna Soler et al., 329–60. Boston Studies in the Philosophy of Science, vol. 292. New York: Springer.

Odenbaugh, J. 2011. "True Lies: Realism, Robustness, and Models." *Philosophy of Science* 78, no. 50: 1177–88.

Petroski, H. 1985. *To Engineer Is Human: The Role of Error in Engineering Design*. New York: St. Martin's Press.

Shoemaker, S. 1963. *Self-Knowledge and Self-Identity.* Ithaca, N.Y.: Cornell University Press.

Simon, H. A. 1969. *The Sciences of the Artificial.* Cambridge, Mass.: MIT Press.

Simon, H. A. 1973. "On the Structure of Ill-Structured Problems." *Artificial Intelligence* 4: 181–201.

Tversky, A., and D. Kahneman. 1974. "Decision-Making under Uncertainty: Heuristics and Biases." *Science* 185: 1124–31.

Tyson, K. A. 1994. *New Foundations for Scientific Social and Behavioral Research: The Heuristic Paradigm*. New York: Allyn and Bacon.

von Neumann, J. 1956. "Probabilistic Logic and the Synthesis of Reliable Organisms from Unreliable Components," In *Automata Studies*, edited by C. E. Shannon and J. McCarthy, 43–98. Princeton, N.J.: Princeton University Press.

Wagner, A. 2005. *Robustness and Evolvability in Living Systems.* Princeton, N.J.: Princeton University Press.

Weisberg, M. 2006. "Robustness Analysis." *Philosophy of Science* 73, no. 5: 730–42.

Wimsatt, W. 1971. "Modern Science and the New Teleology: I—The Conceptual Foundations of Functional Analysis." PhD diss., Dept. of Philosophy, The University of Pittsburgh.

Wimsatt, W. 1972. "Teleology and the Logical Structure of Function Statements." *Studies in History and Philosophy of Science* 3: 1–80.

Wimsatt, W. 1974. "Complexity and Organization." In *PSA-1972* (Boston Studies in the Philosophy of Science, volume 20), edited by K. F. Schaffner and R. S. Cohen, 67–86. Dordrecht: Reidel.

Wimsatt, W. 1976a. “Reductionism, Levels of Organization and the Mind-Body Problem.” In *Consciousness and the Brain*, edited by G. Globus, I. Savodnik, and G. Maxwell, 199–267. New York: Plenum.

Wimsatt, W. 1976b. “Reductive Explanation: A Functional Account.” In *PSA-1974* (Boston Studies in the Philosophy of Science, volume 30), edited by A. C. Michalos, C. A. Hooker, G. Pearce, and R. S. Cohen, 671–710. Dordrecht: Reidel.

Wimsatt, W. 1980. “Reductionistic Research Strategies and Their Biases in the Units of Selection Controversy.” In *Scientific Discovery, Volume 2: Case Studies*, edited by T. Nickles, 213–59. Dordrecht: Reidel.

Wimsatt, W. 1981. “Robustness, Reliability and Overdetermination.” In *Scientific Inquiry and the Social Sciences*, edited by M. Brewer and B. Collin, 124–63. San Francisco: Jossey-Bass.

Wimsatt, W. 1987. “False Models as Means to Truer Theories.” In *Neutral Models in Biology*, edited by M. Nitecki and A. Hoffman, 23–55. London: Oxford University Press.

Wimsatt, W. 1997. “Aggregativity: Reductive Heuristics for Finding Emergence.” Philosophy of Science Association Fifteenth Biennial Meeting, 1997. *Philosophy of Science* 64, no. S4: S372–84.

Wimsatt, W. 2007. *Re-Engineering Philosophy for Limited Beings: Piecewise Approximations to Reality*. Cambridge, Mass.: Harvard University Press.

Wimsatt, W. C. 2012. “The Analytic Geometry of Genetics: The Structure, Function, and Early Evolution of Punnett Squares.” *Archive for the History of the Exact Sciences* 66: 349–96. https://doi.org/10.1007/s00407-012-0096-7.

Wimsatt, W. 2013. “Scaffolding and Entrenchment, an Architecture for a Theory of Cultural Change.” In *Developing Scaffolding in Evolution, Culture, and Cognition*, edited by L. Caporael, J. Griesemer, and W. Wimsatt, 177–206. Cambridge, Mass.: MIT Press.

Wimsatt, W. 2015. “Entrenchment as a Theoretical Tool in Evolutionary Developmental Biology.” *Conceptual Change in Biology: The Case of Evolutionary Developmental Biology*, edited by A. C. Love. Boston Studies in Philosophy of Science. Berlin: Springer.

Wimsatt, W. 2019. “Articulating Babel: A Conceptual Geography for Cultural Evolution.” *Beyond the Meme: Development and Population Structure in Cultural Evolution,* edited by A. Love and W. Wimsatt, 1–41. Minneapolis: University of Minnesota Press.

Wimsatt, W. 2021a. “Engineering Design Principles in Natural and Artificial Systems. Part I: Generative Entrenchment and Modularity.” In *Philosophy and*

Engineering: Reimagining Technology and Social Progress, edited by Zachary Pirtle, David Tomblin, and Guru Madhavan, 25–52. Berlin: Springer International Press.

Wimsatt, W. 2021b. "Levels, Robustness, Emergence, and Heterogeneous Dynamics: Finding Partial Organization in Causal Thickets." In *Levels of Organization: The Architecture of the Scientific Image*, edited by Daniel S. Brooks, James DiFrisco, and William C. Wimsatt, 21–38. Cambridge, Mass.: MIT Press.

Wimsatt, W., and J. Griesemer. 2007. "Reproducing Entrenchments to Scaffold Culture: The Central Role of Development in Cultural Evolution." In *Integrating Evolution and Development: From Theory to Practice*, edited by R. Sansome and R. Brandon, 228–323. Cambridge, Mass.: MIT Press.

Wittgenstein, L. 1953. *Philosophical Investigations.* Oxford: Blackwell.