

The Logical Underpinnings of Intelligent Design

William A. Dembski

Baylor University
Waco, TX 76798

1. Randomness

For many natural scientists, design, conceived as the action of an intelligent agent, is not a fundamental creative force in nature. Rather, material mechanisms, characterized by chance and necessity and ruled by unbroken laws, are thought sufficient to do all nature's creating. Darwin's theory epitomizes this rejection of design.

But how do we know that nature requires no help from a designing intelligence? Certainly, in special sciences ranging from forensics to archaeology to SETI (the Search for Extraterrestrial Intelligence), appeal to a designing intelligence is indispensable. What's more, within these sciences there are well-developed techniques for identifying intelligence. What if these techniques could be formalized, applied to biological systems, and registered the presence of design? Herein lies the promise of intelligent design (or ID, as it is now abbreviated).

My own work on ID began in 1988 at an interdisciplinary conference on randomness at Ohio State University. Persi Diaconis, a well-known statistician, and Harvey Friedman, a well-known logician, convened the conference. The conference came at a time when "chaos theory" or "nonlinear dynamics" were all the rage and supposed to revolutionize science. James Gleick, who had written a wildly popular book titled *Chaos*, covered the conference for the *New York Times*.

For all its promise, the conference ended on a thud. No conference proceedings were ever published. Despite a week of intense discussion,

Persi Diaconis summarized the conference with one brief concluding statement: “We know what randomness isn’t, we don’t know what it is.” For the conference participants, this was an unfortunate conclusion. The point of the conference was to provide a positive account of randomness. Instead, in discipline after discipline, randomness kept eluding our best efforts to grasp it.

That’s not to say there was a complete absence of proposals for characterizing randomness. The problem was that all such proposals approached randomness through the back door, first giving an account of what was nonrandom and then defining what was random by negating nonrandomness (complexity-theoretic approaches to randomness like that of Chaitin [1966] and Kolmogorov [1965] all shared this feature). For instance, in the case of random number generators, they were good so long as they passed a set of statistical tests. Once a statistical test was found that a random number generator no longer passed, the random number generator was discarded as no longer providing suitably random digits.

As I reflected on this asymmetry between randomness and nonrandomness, it became clear that randomness was not an intrinsic property of objects. Instead, randomness was a provisional designation for describing an absence of perceived pattern until such time as a pattern was perceived, at which time the object in question would no longer be considered random. In the case of random number generators, for instance, the statistical tests relative to which their adequacy was assessed constituted a set of patterns. So long as the random number generator passed all these tests, it was considered good and its output was considered random. But as soon as a statistical test was discovered that the random number generator no longer passed, it was no longer good and its output was considered nonrandom. George Marsaglia, a leading light in random number generation who spoke at the 1988 randomness conference, made this point beautifully, detailing one failed random number generator after another.

I wrote up these thoughts in a paper titled “Randomness by Design” (1991; see also Dembski 1998a). In that paper I argued that randomness should properly be thought of as a provisional designation that applies only so long as an object violates all of a set of patterns. Once a pattern is added to the set which the object no longer violates but rather conforms to, the object suddenly becomes nonrandom. Randomness thus becomes a relative notion, relativized to a given set of patterns. As a consequence,

2. The Design Inference

In Fisher's (1935, 13–17) approach to significance testing, a chance hypothesis is eliminated provided an event falls within a prespecified rejection region and provided that rejection region has sufficiently small probability with respect to the chance hypothesis under consideration. Fisher's rejection regions therefore constitute a type of pattern for eliminating chance. The picture here is of an arrow hitting a target. Provided the target is small enough, chance cannot plausibly explain the arrow hitting the target. Of course, the target must be given independently of the arrow's trajectory. Movable targets that can be adjusted after the arrow has landed will not do (one can't, for instance, paint a target around the arrow after it has landed).

In extending Fisher's approach to hypothesis testing, the design inference generalizes the types of rejection regions capable of eliminating chance. In Fisher's approach, to eliminate chance because an event falls within a rejection region, that rejection region must be identified prior to the occurrence of the event. This is to avoid the familiar problem known among statisticians as "data snooping" or "cherry picking," in which a pattern is imposed on an event after the fact. Requiring the rejection region to be set prior to the occurrence of an event safeguards against attributing patterns to the event that are factitious and that do not properly preclude its occurrence by chance.

This safeguard, however, is unduly restrictive. In cryptography, for instance, a pattern that breaks a cryptosystem (known as a cryptographic key) is identified after the fact (i.e., after one has listened in and recorded an enemy communication). Nonetheless, once the key is discovered, there is no doubt that the intercepted communication was not random but rather a message with semantic content and therefore designed. In contrast to statistics, which always identifies its patterns before an experiment is performed, cryptanalysis must discover its patterns after the fact. In both instances, however, the patterns are suitable for eliminating chance. Patterns suitable for eliminating chance I call specifications.

Although my work on specifications can, in hindsight, be understood as a generalization of Fisher's rejection regions, I came to this generalization without consciously attending to Fisher's theory (even though as a probabilist I was fully aware of it). Instead, having reflected on the problem of randomness and the sorts of patterns we use in practice to eliminate chance, I noticed a certain type of inference that came up

repeatedly. These were small probability arguments that, in the presence of a suitable pattern (i.e., specification), not merely eliminated a single chance hypothesis but rather swept the field clear of chance hypotheses. What's more, having swept the field of chance hypotheses, these arguments inferred to a designing intelligence.

Here is a typical example. Suppose that two parties, call them A and B, have the power to produce exactly the same artifact, call it X. Suppose further that producing X requires so much effort that it is easier to copy X once X has already been produced than to produce X from scratch. For instance, before the advent of computers, logarithmic tables had to be calculated by hand. Although there is nothing esoteric about calculating logarithms, the process is tedious if done by hand. Once the calculation has been accurately performed, however, there is no need to repeat it.

The problem, then, confronting the manufacturers of logarithmic tables was that after expending so much effort to compute logarithms, if they were to publish their results without safeguards, nothing would prevent a plagiarist from copying the logarithms directly and then simply claiming that he or she had calculated the logarithms independently. To solve this problem, manufacturers of logarithmic tables introduced occasional—but deliberate—errors into their tables, errors which they carefully noted to themselves. Thus, in a table of logarithms that was accurate to eight decimal places, errors in the seventh and eight decimal places would occasionally be introduced.

These errors then served to trap plagiarists, for even though plagiarists could always claim they computed the logarithms correctly by mechanically following a certain algorithm, they could not reasonably claim to have committed the same errors. As Aristotle remarked in his *Nicomachean Ethics* (McKeon 1941, 1106), “It is possible to fail in many ways, . . . while to succeed is possible only in one way.” Thus, when two manufacturers of logarithmic tables record identical logarithms that are correct, both receive the benefit of the doubt that they have actually done the work of calculating the logarithms. But when both record the same errors, it is perfectly legitimate to conclude that whoever published second plagiarized.

To charge whoever published second with plagiarism, of course, goes well beyond merely eliminating chance (chance in this instance being the independent origination of the same errors). To charge someone with plagiarism, copyright infringement, or cheating is to draw a design

inference. With the logarithmic table example, the crucial elements in drawing a design inference were the occurrence of a highly improbable event (in this case, getting the same incorrect digits in the seventh and eighth decimal places) and the match with an independently given pattern or specification (the same pattern of errors was repeated in different logarithmic tables).

My project, then, was to formalize and extend our commonsense understanding of design inferences so that they could be rigorously applied in scientific investigation. That my codification of design inferences happened to extend Fisher's theory of statistical significance testing was a happy, though not wholly unexpected, convergence. At the heart of my codification of design inferences was the combination of two things: improbability and specification. Improbability, as we shall see in the next section, can be conceived as a form of complexity. As a consequence, the name for this combination of improbability and specification that has now stuck is *specified complexity* or *complex specified information*.

3. Specified Complexity

The term specified complexity is about thirty years old. To my knowledge origin-of-life researcher Leslie Orgel was the first to use it. In his 1973 book *The Origins of Life* he wrote: "Living organisms are distinguished by their specified complexity. Crystals such as granite fail to qualify as living because they lack complexity; mixtures of random polymers fail to qualify because they lack specificity" (189). More recently, Paul Davies (1999, 112) identified specified complexity as the key to resolving the problem of life's origin: "Living organisms are mysterious not for their complexity *per se*, but for their tightly specified complexity." Neither Orgel nor Davies, however, provided a precise analytic account of specified complexity. I provide such an account in *The Design Inference* (1998b) and its sequel *No Free Lunch* (2002). In this section I want briefly to outline my work on specified complexity.

Orgel and Davies used specified complexity loosely. I've formalized it as a statistical criterion for identifying the effects of intelligence. Specified complexity, as I develop it, is a subtle notion that incorporates five main ingredients: (1) a probabilistic version of complexity applicable to events; (2) conditionally independent patterns; (3) probabilistic resources, which

come in two forms, replicational and specificational; (4) a specificational version of complexity applicable to patterns; and (5) a universal probability bound. Let's consider these briefly.

Probabilistic complexity. Probability can be viewed as a form of complexity. To see this, consider a combination lock. The more possible combinations of the lock, the more complex the mechanism and correspondingly the more improbable that the mechanism can be opened by chance. For instance, a combination lock whose dial is numbered from 0 to 39 and which must be turned in three alternating directions will have 64,000 ($= 40 \times 40 \times 40$) possible combinations. This number gives a measure of complexity of the combination lock but also corresponds to a $1/64,000$ probability of the lock being opened by chance. A more complicated combination lock whose dial is numbered from 0 to 99 and which must be turned in five alternating directions will have 10,000,000,000 ($= 100 \times 100 \times 100 \times 100 \times 100$) possible combinations and thus a $1/10,000,000,000$ probability of being opened by chance. Complexity and probability therefore vary inversely: the greater the complexity, the smaller the probability. The "complexity" in "specified complexity" refers to this probabilistic construal of complexity.

Conditionally independent patterns. The patterns that in the presence of complexity or improbability implicate a designing intelligence must be independent of the event whose design is in question. Crucial here is that patterns not be artificially imposed on events after the fact. For instance, if an archer shoots arrows at a wall and we then paint targets around the arrows so that they stick squarely in the bull's-eyes, we impose a pattern after the fact. Any such pattern is not independent of the arrow's trajectory. On the other hand, if the targets are set up in advance ("specified") and then the archer hits them accurately, we know it was not by chance but rather by design. The way to characterize this independence of patterns is via the probabilistic notion of conditional independence. A pattern is conditionally independent of an event if adding our knowledge of the pattern to a chance hypothesis does not alter the event's probability. The "specified" in "specified complexity" refers to such conditionally independent patterns. These are the specifications.

Probabilistic resources. Probabilistic resources refer to the number of opportunities for an event to occur or be specified. A seemingly improbable event can become quite probable once enough probabilistic resources are factored in. Alternatively, it may remain improbable even

after all the available probabilistic resources have been factored in. Probabilistic resources come in two forms: replicational and specificational. Replicational resources refer to the number of opportunities for an event to occur. Specificational resources refer to the number of opportunities to specify an event.

To see what's at stake with these two types of probabilistic resources, imagine a large wall with N identically-sized nonoverlapping targets painted on it and M arrows in your quiver. Let us say that your probability of hitting any one of these targets, taken individually, with a single arrow by chance is p . Then the probability of hitting any one of these N targets, taken collectively, with a single arrow by chance is bounded by Np , and the probability of hitting any of these N targets with at least one of your M arrows by chance is bounded by MNp . In this case, the number of replicational resources corresponds to M (the number of arrows in your quiver), the number of specificational resources corresponds to N (the number of targets on the wall), and the total number probabilistic resources corresponds to the product MN . For a specified event of probability p to be reasonably attributed to chance, the number MNp must not be too small.

Specificational complexity. The conditionally independent patterns that are specifications exhibit varying degrees of complexity. Such degrees of complexity are relativized to personal and computational agents—what I generically refer to as “subjects.” Subjects grade the complexity of patterns in light of their cognitive/computational powers and background knowledge. The degree of complexity of a specification determines the number of specificational resources that must be factored in for setting the level of improbability needed to preclude chance. The more complex the pattern, the more specificational resources must be factored in.

To see what's at stake, imagine a dictionary of 100,000 ($= 10^5$) basic concepts. There are then 10^5 1-level concepts, 10^{10} 2-level concepts, 10^{15} 3-level concepts, and so on. If “bidirectional,” “rotary,” “motor-driven,” and “propeller” are basic concepts, then the bacterial flagellum can be characterized as a 4-level concept of the form “bidirectional rotary motor-driven propeller.” Now, there are about $N = 10^{20}$ concepts of level 4 or less, which constitute the relevant specificational resources. Given p as the probability for the chance formation for the bacterial flagellum, we think of N as providing N targets for the chance formation of the bacterial

flagellum, where the probability of hitting each target is not more than p . Factoring in these N specificational resources then amounts to checking whether the probability of hitting any of these targets by chance is small, which in turn amounts to showing that the product Np is small (see last bullet point on probabilistic resources).

Universal Probability Bound. In the observable universe, probabilistic resources come in limited supplies. Within the known physical universe there are estimated around 10^{80} or so elementary particles. Moreover, the properties of matter are such that transitions from one physical state to another cannot occur at a rate faster than 10^{45} times per second. This frequency corresponds to the Planck time, which constitutes the smallest physically meaningful unit of time. Finally, the universe itself is about a billion times younger than 10^{25} seconds (assuming the universe is between ten and twenty billion years old). If we now assume that any specification of an event within the known physical universe requires at least one elementary particle to specify it and cannot be generated any faster than the Planck time, then these cosmological constraints imply that the total number of specified events throughout cosmic history cannot exceed

$$10^{80} \times 10^{45} \times 10^{25} = 10^{150}.$$

As a consequence, any specified event of probability less than 1 in 10^{150} will remain improbable even after all conceivable probabilistic resources from the observable universe have been factored in. A probability of 1 in 10^{150} is therefore a *universal probability bound* (for the details justifying this universal probability bound, see Dembski 1998b, sec. 6.5). A universal probability bound is impervious to all available probabilistic resources that may be brought against it. Indeed, all the probabilistic resources in the known physical world cannot conspire to render remotely probable an event whose probability is less than this universal probability bound.

The universal probability bound of 1 in 10^{150} is the most conservative in the literature. The French mathematician Emile Borel (1962, 28; see also Knobloch 1987, 228) proposed 1 in 10^{50} as a universal probability bound below which chance could definitively be precluded (i.e., any specified event as improbable as this could never be attributed to chance). Cryptographers assess the security of cryptosystems in terms of brute force attacks that employ as many probabilistic resources as are available in the universe to break a cryptosystem by chance. In its report on the role

of cryptography in securing the information society, the National Research Council set 1 in 10^{94} as its universal probability bound to ensure the security of cryptosystems against chance-based attacks (see Dam and Lin, 1996, 380, note 17). Theoretical computer scientist Seth Lloyd (2002) sets 10^{120} as the maximum number of bit-operations that the universe could have performed throughout its entire history. That number corresponds to a universal probability bound of 1 in 10^{120} . Stuart Kauffman (2000) in his most recent book, *Investigations*, comes up with similar numbers.

For something to exhibit specified complexity therefore means that it matches a conditionally independent pattern (i.e., specification) that corresponds to an event of probability less than the universal probability bound. Specified complexity is a widely used criterion for detecting design. For instance, when researchers in the Search for Extraterrestrial Intelligence (SETI) look for signs of intelligence from outer space, they are looking for specified complexity (recall the movie *Contact* in which contact is established when a long sequence of prime numbers comes in from outer space—such a sequence exhibits specified complexity). Let us therefore examine next the reliability of specified complexity as a criterion for detecting design.

4. Reliability of the Criterion

Specified complexity functions as a criterion for detecting design—I call it the complexity-specification criterion. In general, criteria attempt to classify individuals with respect to a target group. The target group for the complexity-specification criterion comprises all things intelligently caused. How accurate is this criterion at correctly assigning things to this target group and correctly omitting things from it?

The things we are trying to explain have causal histories. In some of those histories intelligent causation is indispensable whereas in others it is dispensable. An inkblot can be explained without appealing to intelligent causation; ink arranged to form meaningful text cannot. When the complexity-specification criterion assigns something to the target group, can we be confident that it actually is intelligently caused? If not, we have a problem with false positives. On the other hand, when this criterion fails to assign something to the target group, can we be confident that no intelligent cause underlies it? If not, we have a problem with false negatives.

Consider first the problem of false negatives. When the complexity-specification criterion fails to detect design in a thing, can we be sure that no intelligent cause underlies it? No, we cannot. To determine that something is not designed, this criterion is not reliable. False negatives are a problem for it. This problem of false negatives, however, is endemic to design detection in general. One difficulty is that intelligent causes can mimic undirected natural causes, thereby rendering their actions indistinguishable from such unintelligent causes. A bottle of ink happens to fall off a cupboard and spill onto a sheet of paper. Alternatively, a human agent deliberately takes a bottle of ink and pours it over a sheet of paper. The resulting inkblot may look identical in both instances, but in the one case results by natural causes, in the other by design.

Another difficulty is that detecting intelligent causes requires background knowledge on our part. It takes an intelligent cause to recognize an intelligent cause. But if we do not know enough, we will miss it. Consider a spy listening in on a communication channel whose messages are encrypted. Unless the spy knows how to break the cryptosystem used by the parties on whom she is eavesdropping (i.e., knows the cryptographic key), any messages traversing the communication channel will be unintelligible and might in fact be meaningless.

The problem of false negatives therefore arises either when an intelligent agent has acted (whether consciously or unconsciously) to conceal one's actions, or when an intelligent agent, in trying to detect design, has insufficient background knowledge to determine whether design actually is present. This is why false negatives do not invalidate the complexity-specification criterion. This criterion is fully capable of detecting intelligent causes intent on making their presence evident. Masters of stealth intent on concealing their actions may successfully evade the criterion. But masters of self-promotion bank on the complexity-specification criterion to make sure their intellectual property gets properly attributed. Indeed, intellectual property law would be impossible without this criterion.

And that brings us to the problem of false positives. Even though specified complexity is not a reliable criterion for *eliminating* design, it is a reliable criterion for *detecting* design. The complexity-specification criterion is a net. Things that are designed will occasionally slip past the net. We would prefer that the net catch more than it does, omitting nothing

due to design. But given the ability of design to mimic unintelligent causes and the possibility of ignorance causing us to pass over things that are designed, this problem cannot be remedied. Nevertheless, we want to be very sure that whatever the net does catch includes only what we intend it to catch—namely, things that are designed. Only things that are designed had better end up in the net. If that is the case, we can have confidence that whatever the complexity-specification criterion attributes to design is indeed designed. On the other hand, if things end up in the net that are not designed, the criterion is in trouble.

How can we see that specified complexity is a reliable criterion for detecting design? Alternatively, how can we see that the complexity-specification criterion successfully avoids false positives—that whenever it attributes design, it does so correctly? The justification for this claim is a straightforward inductive generalization: In every instance where specified complexity obtains and where the underlying causal story is known (i.e., where we are not just dealing with circumstantial evidence, but where, as it were, the video camera is running and any putative designer would be caught red-handed), it turns out design actually is present; therefore, design actually is present whenever the complexity-specification criterion attributes design.

Although this justification for the complexity-specification criterion's reliability at detecting design may seem a bit too easy, it really isn't. If something genuinely instantiates specified complexity, then it is inexplicable in terms of all material mechanism (not only those that are known but all of them). Indeed, to attribute specified complexity to something is to say that the specification to which it conforms corresponds to an event that is highly improbable with respect to all material mechanism that might give rise to the event. So take your pick—treat the item in question as inexplicable in terms of all material mechanisms or treat it as designed. But since design is uniformly associated with specified complexity when the underlying causal story is known, induction counsels attributing design in cases where the underlying causal story is not known.

To sum up, for specified complexity to eliminate chance and detect design, it is not enough that the probability be small with respect to some arbitrarily chosen probability distribution. Rather, it must be small with respect to every probability distribution that might characterize the chance occurrence of the thing in question. If that is the case, then a design

inference follows. The use of chance here is very broad and includes anything that can be captured mathematically by a stochastic process. It thus includes deterministic processes whose probabilities all collapse to zero and one (cf. necessities, regularities, and natural laws). It also includes nondeterministic processes, like evolutionary processes that combine random variation and natural selection. Indeed, chance so construed characterizes all material mechanisms.

5. Assertibility

The reliability of specified complexity as a criterion for detecting design is not a problem. Neither is there a problem with specified complexity's coherence as a meaningful concept—specified complexity is well-defined. If there's a problem, it centers on specified complexity's *assertibility*. Assertibility is a term of philosophical use that refers to the epistemic justification or warrant for a claim. Assertibility (with an "i") is distinguished from assertability (with an "a"), where the latter refers to the local factors that in the pragmatics of discourse determine whether asserting a claim is justified (see Jackson 1987, 11). For instance, as a tourist in Iraq I might be epistemically justified asserting that Saddam Hussein is a monster (in which case the claim would be assertible). Local-pragmatic considerations, however, tell against asserting this remark within Iraqi borders (the claim there would be unassertable). Unlike assertibility, assertability can depend on anything from etiquette and good manners to who happens to hold political power. Assertibility with an "i" is what interests us here.

To see what's at stake with specified complexity's assertibility, consider first a mathematical example. It's an open question in mathematics whether the number pi (the ratio of the circumference of a circle to its diameter) is regular, where by regular I mean that every number between 0 and 9 appears in the decimal expansion of pi with limiting relative frequency 1/10. Regularity is a well-defined mathematical concept. Thus, in asserting that pi is regular, we might be making a true statement. But without a mathematical proof of pi's regularity, we have no justification for asserting that pi is regular. The regularity of pi is, at least for now, unassertible (despite over 200 billion decimal digits of pi having been computed).

But what about the specified complexity of various biological systems? Are there any biological systems whose specified complexity is assertible? Critics of intelligent design argue that no attribution of specified complexity to any natural system can ever be assertible. The argument runs as follows. It starts by noting that if some natural system instantiates specified complexity, then that system must be vastly improbable with respect to all purely natural mechanisms that could be operating to produce it. But that means calculating a probability for each such mechanism. This, so the argument runs, is an impossible task. At best science could show that a given natural system is vastly improbable with respect to known mechanisms operating in known ways and for which the probability can be estimated. But that omits (1) known mechanisms operating in known ways for which the probability cannot be estimated, (2) known mechanisms operating in unknown ways, and (3) unknown mechanisms.

Thus, even if it is true that some natural system instantiates specified complexity, we could never legitimately assert its specified complexity, much less know it. Accordingly, to assert the specified complexity of any natural system constitutes an argument from ignorance. This line of reasoning against specified complexity is much like the standard agnostic line against theism—we can't prove atheism (cf. the total absence of specified complexity from nature), but we can show that theism (cf. the specified complexity of certain natural systems) cannot be justified and is therefore unassertible. This is how skeptics argue that there is no (and indeed can be no) evidence for God or design.

A little reflection, however, makes clear that this attempt by skeptics to undo specified complexity cannot be justified on the basis of scientific practice. Indeed, the skeptic imposes requirements so stringent that they are absent from every other aspect of science. If standards of scientific justification are set too high, no interesting scientific work will ever get done. Science therefore balances its standards of justification with the requirement for self-correction in light of further evidence. The possibility of self-correction in light of further evidence is absent in mathematics and accounts for mathematics' need for the highest level of justification, namely, strict logico-deductive proof. But science does not work that way.

Science must work with available evidence, and on that basis (and that basis alone) formulate the best explanation of the phenomenon in question. This means that science cannot explain a phenomenon by

appealing to the promise, prospect, or possibility of future evidence. In particular, unknown mechanisms or undiscovered ways by which those mechanisms operate cannot be invoked to explain a phenomenon. If known material mechanisms can be shown incapable of explaining a phenomenon, then it is an open question whether any mechanisms whatsoever are capable of explaining it. If, further, there are good reasons for asserting the specified complexity of certain biological systems, then design itself becomes assertible in biology. Let's now see how this could be.

6. Application to Evolutionary Biology

Evolutionary biology teaches that all biological complexity is the result of material mechanisms. These include principally the Darwinian mechanism of natural selection and random variation but also include other mechanisms (symbiogenesis, gene transfer, genetic drift, the action of regulatory genes in development, self-organizational processes, etc.). These mechanisms are just that: mindless material mechanisms that do what they do irrespective of intelligence. To be sure, mechanisms can be programmed by an intelligence. But any such intelligent programming of evolutionary mechanisms is not properly part of evolutionary biology.

Intelligent design, by contrast, teaches that biological complexity is not exclusively the result of material mechanisms but also requires intelligence, where the intelligence in question is not reducible to such mechanisms. The central issue, therefore, is not the relatedness of all organisms, or what typically is called common descent. Indeed, intelligent design is perfectly compatible with common descent. Rather, the central issue is how biological complexity emerged and whether intelligence played an indispensable (which is not to say exclusive) role in its emergence.

Suppose, therefore, for the sake of argument that intelligence—one irreducible to material mechanisms—actually did play a decisive role in the emergence of life's complexity and diversity. How could we know it? Certainly specified complexity will be required. Indeed, if specified complexity is absent or dubious, then the door is wide open for material mechanisms to explain the object of investigation. Only as specified complexity becomes assertible does the door to material mechanisms start to close.

Nevertheless, evolutionary biology teaches that within biology the door can never be closed all the way and indeed should not be closed at all. In fact, evolutionary biologists claim to have demonstrated that design is superfluous for understanding biological complexity. The only way actually to demonstrate this, however, is to exhibit material mechanisms that account for the various forms of biological complexity out there. Now, if for every instance of biological complexity some mechanism could readily be produced that accounts for it, intelligent design would drop out of scientific discussion. Occam's razor, by proscribing superfluous causes, would in this instance finish off intelligent design quite nicely.

But that hasn't happened. Why not? The reason is that there are plenty of complex biological systems for which no biologist has a clue how they emerged. I'm not talking about handwaving just-so stories. Biologists have plenty of those. I'm talking about detailed testable accounts of how such systems could have emerged. To see what's at stake, consider how biologists propose to explain the emergence of the bacterial flagellum, a molecular machine that has become the mascot of the intelligent design movement.

In public lectures Harvard biologist Howard Berg calls the bacterial flagellum "the most efficient machine in the universe." The flagellum is a nano-engineered motor-driven propeller on the backs of certain bacteria. It spins at tens of thousands of rpm, can change direction in a quarter turn, and propels a bacterium through its watery environment. According to evolutionary biology it had to emerge via some material mechanism(s). Fine, but how?

The usual story is that the flagellum is composed of parts that previously were targeted for different uses and that natural selection then co-opted to form a flagellum. This seems reasonable until we try to fill in the details. The only well-documented examples that we have of successful co-optation come from human engineering. For instance, an electrical engineer might co-opt components from a microwave oven, a radio, and a computer screen to form a working television. But in that case, we have an intelligent agent who knows all about electrical gadgets and about televisions in particular.

But natural selection doesn't know a thing about bacterial flagella. So how is natural selection going to take extant protein parts and co-opt them to form a flagellum? The problem is that natural selection can only select

for pre-existing function. It can, for instance, select for larger finch beaks when the available nuts are harder to open. Here the finch beak is already in place and natural selection merely enhances its present functionality. Natural selection might even adapt a pre-existing structure to a new function; for example, it might start with finch beaks adapted to opening nuts and end with beaks adapted to eating insects.

But for co-optation to result in a structure like the bacterial flagellum, we are not talking about enhancing the function of an existing structure or reassigning an existing structure to a different function, but reassigning multiple structures previously targeted for different functions to a novel structure exhibiting a novel function. Even the simplest bacterial flagellum requires around forty proteins for its assembly and structure. All these proteins are necessary in the sense that lacking any of them, a working flagellum does not result.

The only way for natural selection to form such a structure by co-optation, then, is for natural selection gradually to enfold existing protein parts into evolving structures whose functions co-evolve with the structures. We might, for instance, imagine a five-part mousetrap consisting of a platform, spring, hammer, holding bar, and catch evolving as follows: It starts as a doorstop (thus consisting merely of the platform), then evolves into a tie-clip (by attaching the spring and hammer to the platform), and finally becomes a full mousetrap (by also including the holding bar and catch).

Design critic Kenneth Miller finds such scenarios not only completely plausible but also deeply relevant to biology (in fact, he regularly sports a modified mousetrap cum tie-clip). Intelligent design proponents, by contrast, regard such scenarios as rubbish. Here's why. First, in such scenarios the hand of human design and intention meddles everywhere. Evolutionary biologists assure us that eventually they will discover just how the evolutionary process can take the right and needed steps without the meddling hand of design. All such assurances, however, presuppose that intelligence is dispensable in explaining biological complexity. Yet the only evidence we have of successful co-optation comes from engineering and confirms that intelligence is indispensable in explaining complex structures like the mousetrap and by implication the flagellum. Intelligence is known to have the causal power to produce such structures. We're still waiting for the promised material mechanisms.

The other reason design theorists are less than impressed with co-

optation concerns an inherent limitation of the Darwinian mechanism. The whole point of the Darwinian selection mechanism is that one can get from anywhere in biological configuration space to anywhere else provided one can take small steps. How small? Small enough that they are reasonably probable. But what guarantee is there that a sequence of baby-steps connects any two points in configuration space?

The problem is not simply one of connectivity. For the Darwinian selection mechanism to connect point A to point B in configuration space, it is not enough that there merely exist a sequence of baby-steps connecting the two. In addition, each baby-step needs in some sense to be “successful.” In biological terms, each step requires an increase in fitness as measured in terms of survival and reproduction. Natural selection, after all, is the motive force behind each baby-step, and selection only selects what is advantageous to the organism. Thus, for the Darwinian mechanism to connect two organisms, there must be a sequence of successful baby-steps connecting the two.

Richard Dawkins (1996) compares the emergence of biological complexity to climbing a mountain—Mount Improbable, as he calls it. According to him, Mount Improbable always has a gradual serpentine path leading to the top that can be traversed in baby-steps. But that’s hardly an empirical claim. Indeed, the claim is entirely gratuitous. It might be a fact about nature that Mount Improbable is sheer on all sides and getting to the top from the bottom via baby-steps is effectively impossible. A gap like that would reside in nature herself and not in our knowledge of nature (it would not, in other words, constitute a god-of-the-gaps).

Consequently, it is not enough merely to presuppose that a fitness-increasing sequence of baby steps connects two biological systems—it must be demonstrated. For instance, it is not enough to point out that some genes for the bacterial flagellum are the same as those for a type III secretory system (a type of pump) and then handwave that one was co-opted from the other. Anybody can arrange complex systems in series based on some criterion of similarity. But such series do nothing to establish whether the end evolved in Darwinian fashion from the beginning unless the probability of each step in the series can be quantified, the probability at each step turns out to be reasonably large, and each step constitutes an advantage to the evolving system.

Convinced that the Darwinian mechanism must be capable of doing such evolutionary design work, evolutionary biologists rarely ask whether

such a sequence of successful baby-steps even exists; much less do they attempt to quantify the probabilities involved. I attempt that in my book *No Free Lunch* (2002, ch. 5). There I lay out techniques for assessing the probabilistic hurdles that the Darwinian mechanism faces in trying to account for complex biological structures like the bacterial flagellum. The probabilities I calculate—and I try to be conservative—are horrendous and render natural selection utterly implausible as a mechanism for generating the flagellum and structures like it.

Is the claim that the bacterial flagellum exhibits specified complexity assertible? You bet! Science works on the basis of available evidence, not on the promise or possibility of future evidence. Our best evidence points to the specified complexity (and therefore design) of the bacterial flagellum. It is therefore incumbent on the scientific community to admit, at least provisionally, that the bacterial flagellum could be the product of design. Might there be biological examples for which the claim that they exhibit specified complexity is even more assertible? Yes there might. Unlike truth, assertibility comes in degrees, corresponding to the strength of evidence that justifies a claim. Yet even now, to say that the bacterial flagellum exhibits specified complexity is eminently assertible.

Evolutionary biology's only recourse for avoiding a design conclusion in instances like this is to look to unknown mechanisms (or known mechanisms operating in unknown ways) to overturn what our best evidence to date indicates is both complex and specified. As far as the evolutionary biologists are concerned, design theorists have failed to take into account indirect Darwinian pathways by which the bacterial flagellum might have evolved through a series of intermediate systems that changed function and structure over time in ways that we do not yet understand. But is it that we do not yet understand the indirect Darwinian evolution of the bacterial flagellum or that it never happened that way in the first place? At this point there is simply no evidence for such indirect Darwinian evolutionary pathways to account for biological systems like the bacterial flagellum.

There is further reason to be skeptical of evolutionary biology's general strategy for defeating intelligent design by looking to unknown material mechanisms. In the case of the bacterial flagellum, what keeps evolutionary biology afloat is the possibility of indirect Darwinian pathways that might account for it. Practically speaking, this means that even though no slight modification of a bacterial flagellum can continue to

serve as a motility structure, a slight modification might serve some other function. But there is now mounting evidence of biological systems for which any slight modification does not merely destroy the system's existing function but also destroys the possibility of any function of the system whatsoever (see Axe 2000). For such systems, neither direct nor indirect Darwinian pathways could account for them. In that case we would be dealing with an in-principle argument showing not merely that no known material mechanism is capable of accounting for the system but also that any unknown material mechanism is incapable of accounting for it as well. Specified complexity's assertibility in such cases would thus be even greater than in the case of the bacterial flagellum.

It is possible to rule out unknown material mechanisms once and for all provided one has independent reasons for thinking that explanations based on known material mechanisms cannot be overturned by yet-to-be-identified unknown mechanisms. Such independent reasons typically take the form of arguments from contingency that invoke numerous degrees of freedom. Thus, to establish that no material mechanism explains a phenomenon, we must establish that it is compatible with the known material mechanisms involved in its production, but that these mechanisms also permit any number of alternatives to it. By being compatible with but not required by the known material mechanisms involved in its production, a phenomenon becomes irreducible not only to the known mechanisms but also to any unknown mechanisms. How so? Because known material mechanisms can tell us conclusively that a phenomenon is contingent and allows full degrees of freedom. Any unknown mechanism would therefore have to respect that contingency and allow for the degrees of freedom already discovered.

Consider, for instance, a configuration space comprising all possible character sequences from a fixed alphabet (such spaces model not only written texts but also polymers like DNA, RNA, and proteins). Configuration spaces like this are perfectly homogeneous, with one character string geometrically interchangeable with the next. The geometry therefore precludes any underlying mechanisms from distinguishing or preferring some character strings over others. Not material mechanisms but external semantic information (in the case of written texts) or functional information (in the case of biopolymers) is needed to generate specified complexity in these instances. To argue that this semantic or functional information reduces to material mechanisms is

like arguing that Scrabble pieces have inherent in them preferential ways they like to be sequenced. They don't. Michael Polanyi (1967; 1968) made such arguments for biological design in the 1960s. Stephen Meyer (2003) has updated them for the present.

7. Eliminative Induction

To attribute specified complexity to a biological system is to engage in an eliminative induction. Eliminative inductions depend on successfully falsifying competing hypotheses (contrast this with Popperian falsification, where hypotheses are corroborated to the degree that they successfully withstand attempts to falsify them). Now, for many design skeptics, eliminative inductions are mere arguments from ignorance, that is, arguments for the truth of a proposition because it has not been shown to be false. In arguments from ignorance, the lack of evidence for a proposition is used to argue for its truth. A stereotypical argument from ignorance goes something like “ghosts and goblins exist because you haven't shown me that they don't exist.”

But that's clearly not what eliminative inductions are doing. Eliminative inductions argue that competitors to the proposition in question are false. Provided the proposition together with its competitors form a mutually exclusive and exhaustive class, eliminating all the competitors entails that the proposition is true. This the ideal case, in which eliminative inductions in fact become deductions. The problem is that in practice we don't have a neat ordering of competitors that can then all be knocked down with a few straightforward and judicious blows (like pins in a bowling alley). Philosopher of science John Earman (1992, 165) puts it this way:

The eliminative inductivist [seems to be] in a position analogous to that of Zeno's archer whose arrow can never reach the target, for faced with an infinite number of hypotheses, he can eliminate one, then two, then three, etc., but no matter how long he labors, he will never get down to just one. Indeed, it is as if the arrow never gets half way, or a quarter way, etc. to the target, since however long the eliminativist labors, he will always be faced with an infinite list [of remaining hypotheses to eliminate].

Earman offers these remarks in a chapter titled “A Plea for Eliminative Induction.” He himself thinks there is a legitimate and necessary place for

eliminative induction in scientific practice. What, then, does he make of this criticism? Here is how he handles it (Earman 1992, 165):

My response on behalf of the eliminativist has two parts. (1) Elimination need not proceed in such a plodding fashion, for the alternatives may be so ordered that an infinite number can be eliminated in one blow. (2) Even if we never get down to a single hypothesis, progress occurs if we succeed in eliminating finite or infinite chunks of the possibility space. This presupposes, of course, that we have some kind of measure, or at least topology, on the space of possibilities.

To this Earman (1992, 177) adds that eliminative inductions are typically *local inductions*, in which there is no pretense of considering all logically possible hypotheses. Rather, there is tacit agreement on the explanatory domain of the hypotheses as well as on what auxiliary hypotheses may be used in constructing explanations.

In ending this essay, I want to reflect on Earman's claim that eliminative inductions can be *progressive*. Too often critics of intelligent design charge specified complexity with underwriting a purely negative form of argumentation. But that charge is not accurate. The argument for the specified complexity of the bacterial flagellum, for instance, makes a positive contribution to our understanding of the limitations that natural mechanisms face in trying to account for it. Eliminative inductions, like all inductions and indeed all scientific claims, are fallible. But they need a place in science. To refuse them, as evolutionary biology tacitly does by rejecting specified complexity as a criterion for detecting design, does not keep science safe from disreputable influences but instead undermines scientific inquiry itself.

The way things stand now, evolutionary biology allows intelligent design only to fail but not to succeed. If evolutionary biologists can discover or construct detailed, testable, indirect Darwinian pathways that account for complex biological systems like the bacterial flagellum, then intelligent design will rightly fail. On the other hand, evolutionary biology makes it effectively impossible for intelligent design to succeed. According to evolutionary biology, intelligent design has only one way to succeed, namely, by showing that complex specified biological structures could not have evolved via any material mechanism. In other words, so long as some unknown material mechanism might have evolved the structure in question, intelligent design is proscribed.

Evolutionary theory is thereby rendered immune to disconfirmation in

principle because the universe of unknown material mechanisms can never be exhausted. Furthermore, the evolutionist has no burden of evidence. Instead, the burden of evidence is shifted entirely to the evolution skeptic. And what is required of the skeptic? The skeptic must establish a universal negative not by an eliminative induction (such inductions are invariably local and constrained) but by an exhaustive search and elimination of all conceivable possibilities—however remote, however unfounded, however unsupported by evidence. That is not how science is supposed to work.

Science is supposed to give the full range of possible explanations a fair chance to succeed. That's not to say that anything goes; but it is to say that anything might go. In particular, science may not by a priori fiat rule out logical possibilities. Evolutionary biology, by limiting itself exclusively to material mechanisms, has settled in advance which biological explanations are true apart from any consideration of empirical evidence. This is arm-chair philosophy. Intelligent design may not be correct. But the only way we could discover that is by admitting design as a real possibility, not by ruling it out a priori. Darwin (1859, 2) himself would have agreed. In the *Origin of Species* he wrote: "A fair result can be obtained only by fully stating and balancing the facts and arguments on both sides of each question."

References

- Axe, D. 2000. Extreme functional sensitivity to conservative amino acid changes on enzyme exteriors, *Journal of Molecular Biology* 301: 585–95.
- Borel, E. 1962. *Probabilities and Life*, trans. M. Baudin. New York: Dover.
- Chaitin, G. J. 1966. On the length of programs for computing finite binary sequences, *Journal of the Association for Computing Machinery* 13: 547–69.
- Dam, K. W. and H. S. Lin, eds. 1996. *Cryptography's Role in Securing the Information Society*. Washington, D.C.: National Academy Press.
- Darwin, C. [1859] 1964. *On the Origin of Species*, facsimile 1st ed. Cambridge, Mass.: Harvard University Press.
- Davies, P. 1999. *The Fifth Miracle*. New York: Simon & Schuster.
- Dawkins, R. 1996. *Climbing Mount Improbable*. New York: Norton.
- Dembski, W. A. 1991. Randomness by design. *Nous* 25(1): 75–106.
- 1998a. Randomness. In *The Routledge Encyclopedia of Philosophy*, ed. E. Craig. London: Routledge.
- 1998b. *The Design Inference: Eliminating Chance through Small Probabilities*. New York: Cambridge University Press.
2002. *No Free Lunch: Why Specified Complexity Cannot Be Purchased without Intelligence*. Lanham, Md.: Rowman and Littlefield.
- Earman, J. 1992. *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*. Cambridge, Mass.: MIT Press.
- Fisher, R. A. 1935. *The Design of Experiments*. New York: Hafner.
- Jackson, F. 1987. *Conditionals*. Oxford: Blackwell.
- Kauffman, S. 2000. *Investigations*. New York: Oxford University Press.
- Knobloch, E. 1987. Emile Borel as a probabilist. In *The Probabilistic Revolution*, vol. 1, eds. L. Krüger, L. J. Daston, and M. Heidelberger, 215–33. Cambridge, Mass.: MIT Press.

- Kolmogorov, A. 1965. Three approaches to the quantitative definition of information. *Problemy Peredachi Informatsii* (in translation) 1(1): 3–11.
- Lloyd, S. 2002. Computational capacity of the universe. *Physical Review Letters* 88(23): 7901–4.
- McKeon, R., ed. 1941. *The Basic Works of Aristotle*. New York: Random House.
- Meyer, S. C. 2003. DNA and the origin of life: Information, specification, and explanation. In J. A. Campbell and S. C. Meyer, eds., *Darwinism, Design and Public Education* (forthcoming). Lansing, Mich.: Michigan State University Press.
- Orgel, L. 1973. *The Origins of Life*. New York: Wiley.
- Polanyi, M. 1967. Life transcending physics and chemistry. *Chemical and Engineering News* 45: 54–66.
1968. Life's irreducible structure. *Science* 113: 1308–12.