

Why Natural Selection Can't Design Anything

William A. Dembski

Baylor University
P. O. Box 97130
Waco, Texas 76798

Discovery Institute
1402 Third Ave.
Seattle, WA 98101

1. Davies's Challenge

In *The Fifth Miracle* Paul Davies suggests that any laws capable of explaining the origin of life must be radically different from scientific laws known to date.¹ The problem, as he sees it, with currently known scientific laws, like the laws of chemistry and physics, is that they cannot explain the key feature of life that needs to be explained.² That feature is *specified complexity*. Life is both complex and specified. The basic intuition here is straightforward. A short word like the definite article "the" is specified without being complex (it conforms to an independently given pattern but is simple). A long sequence of random letters is complex without being specified (it requires a complicated instruction-set to characterize but conforms to no independently given pattern). A Shakespearean sonnet is both complex and specified.

Specified complexity is a type of information. Both complexity and specification are well-defined information-theoretic concepts. Complexity here is used in the Shannon

sense and denotes a measure of improbability.³ Take, for instance, 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. Complexity and probability therefore vary inversely—the greater the complexity, the smaller the probability.⁴ Specification here refers to the patterning of complex arrangements where the pattern is recoverable independently of the actual arrangement. In *The Design Inference* I show that specified complexity can be given a rigorous mathematical formulation.⁵

How, then, to explain specified complexity? As Davies rightly notes, *laws* (that is, necessities of nature) can explain specification but not complexity. For instance, the formation of a salt crystal follows well-defined laws, produces an independently given repetitive pattern, and is therefore specified; but that pattern will also be simple, not complex. On the other hand, as Davies also rightly notes, *contingency* (that is, chance or accidental processes of nature) can explain complexity but not specification. For instance, the exact time sequence of radioactive emissions from a piece of uranium will be contingent, complex, but not specified. The problem is to explain something like the genetic code, which is both complex and specified. As Davies puts it: “Living organisms are mysterious not for their complexity *per se*, but for their tightly specified complexity.”⁶

Nonetheless, once life (or more generally some self-replicator) has arrived, Davies thinks there is no problem accounting for specified complexity. Indeed, he thinks the Darwinian mechanism of natural selection and random variation is fully adequate to account for specified complexity once replicators are here. He writes: “Random mutations plus natural selection are one surefire way to generate biological information, extending a short random genome over time into a long random genome. Chance in the guise of mutations and law in the guise of selection form just the right combination of randomness and order needed to create ‘the impossible object.’ The necessary

information comes, as we have seen, from the environment.”⁷ In the same vein, Davies writes: “Natural selection ... acts like a ratchet, locking in the advantageous errors and discarding the bad. Starting with the DNA of some primitive ancestor microbe, bit by bit, error by error, the increasingly lengthy instructions for building more complex organisms came to be constructed”⁸

The problem with invoking the Darwinian mechanism to explain specified complexity at the origin of life is the absence of an identifiable replicator to which the mechanism might apply. Theodosius Dobzhansky was therefore right to remark that “prebiological natural selection is a contradiction in terms.”⁹ Indeed, the Darwinian mechanism of natural selection and random variation is simply not available until after the origin of life. Once life has started and self-replication has begun, however, the Darwinian mechanism is usually invoked to explain the specified complexity of living things. Though himself a dissenter from strict Darwinism, Stuart Kauffman agrees that this is the majority view within the biological community. He writes: “Biologists now tend to believe profoundly that natural selection is the invisible hand that crafts well-wrought forms. It may be an overstatement to claim that biologists view selection as the sole source of order in biology, but not by much. If current biology has a central canon, you have now heard it.”¹⁰

In this paper I will argue that the problem of explaining specified complexity is even worse than Davies makes out in *The Fifth Miracle*. Not only have we yet to explain specified complexity at the origin of life, but the Darwinian mechanism fails to explain it for the subsequent history of life as well. To see that the Darwinian mechanism is incapable of generating specified complexity, it is necessary to consider the mathematical underpinnings of that mechanism, to wit, evolutionary algorithms. Roughly speaking, an evolutionary algorithm is any well-defined mathematical procedure that generates contingency via some chance process and then sifts it via some

law-like process. The Darwinian mechanism, simulated annealing, training neural nets, and genetic algorithms all fall within this broad construal of evolutionary algorithms.¹¹

2. Blind and Constrained Search

The simplest evolutionary algorithm is blind search. Joseph Culberson illustrates blind search with the following vignette:

In the movie *UHF*, there is a marvelous scene that every computing scientist should consider. As the camera slowly pans across a small park setting, we hear a voice repeatedly asking “Is this it?” followed each time by the response “No!” As the camera continues to pan, it picks up two men on a park bench, one of them blind and holding a Rubik’s cube. He gives it a twist, holds it up to his friend and the query–response sequence is repeated. This is blind search.¹²

This scene is humorous for the same reason it is instructive. There are an enormous number of possible configurations of the Rubik’s cube. Of these only one constitutes a solution—the configuration where each face displays the same color. Within the reference class of all possible configurations, the solution therefore constitutes an instance of specified complexity.¹³ This is why the scene is instructive, because it illustrates specified complexity. This is also why it is humorous, because the two men on the park bench will long expire before reaching a solution.

To characterize blind search more precisely, let us expand on this example. Consider two interlocutors, Alice and Bob. Alice identifies a reference class of possible solutions to a problem, what in the evolutionary algorithms biz is called the *phase space*. Bob identifies not only the phase space, but also the set of actual solutions, which we will call the *target*. Provided the target has small enough probability within the phase space, the target will constitute an instance of specified complexity. For any possible solution in the phase space, Bob is able to tell Alice whether it falls within the target (that is,

whether it is in fact a solution). Alice now successively selects m possible solutions from the phase space. For each of these solutions she queries Bob, who then truthfully answers whether the candidate solution is in fact in the target. We call the number m the *sample size*. The sample size sets an upper limit on the number of candidate solutions that Alice can consider and reflects that Alice does not have infinite resources to continue the search indefinitely. Alice's search is successful if one of her candidate solutions lands in the target. Computer scientists call this blind search.

How effective is blind search at locating targets? Its effectiveness depends on two things: (1) the probability of the target and (2) the sample size m of candidate solutions to be considered. Take, for instance, a simple combination lock with three dials each numbered from zero to nine (there are thus ten possible positions of each dial). Since there are three dials, there are a thousand distinct possible combinations. This set of possible combinations constitutes the phase space. Only one of these possible combinations, however, opens the lock. The unique combination that opens the lock is the target. It follows that a single candidate solution has a one in a thousand probability of opening the lock by chance. Moreover, a sample of m possible candidate solutions has no more than m times one in a thousand probability of opening the lock by chance. If m equals ten, for instance, there is at most a one in a hundred probability of opening the lock by chance. In general, to set an upper bound on the probability that a blind search successfully locates a target, one multiplies the sample size times the probability of the target. For a sample size m and probability of target p , this number is mp .

In most cases of interest, mp is minuscule. In most cases of interest, the phase space is huge, the target is tiny, and the sample size is too small to make any headway with so tiny a target in so huge a phase space. We can think of the sample size and the probability of the target as being in competition, with the sample size needing to overcome the improbability of the target before blind search has any hope of attaining the target. For most cases of interest, the sample size is too small to overcome the

improbability of the target. And since the target will in such cases constitute an instance of specified complexity, it follows that blind search has no hope of generating specified complexity.

The crucial question now is this: Given that the target is so improbable that blind search is highly unlikely to succeed, what additional information might help Alice to make her search succeed? To answer this question, let us return to the interchange between Alice and Bob. Alice and Bob are playing a game of “*m* questions” in which Bob divulges too little information for Alice to have any hope of winning the game. Alice therefore needs some additional information from Bob. But what? Bob could just inform Alice of the exact location of the target and be done with it. But that would be too easy. If Alice is playing the role of scientist and Bob the role of nature, then Bob needs to make Alice drudge and sweat to locate the target—nature, after all, does not divulge her secrets easily. Alice and Bob are operating here with competing constraints. Bob wants to give Alice the minimum information she needs to locate the target. Alice, on the other hand, wants to make maximal use of whatever information Bob gives her to ensure that her *m* questions are as effective as possible in locating the target.

Let us therefore suppose that Bob identifies some additional information and makes it available to Alice. This information is supposed to help Alice locate the target. We have therefore a new protocol for the interchange between Alice and Bob. Before, Bob would only tell Alice whether a candidate solution belonged to the target. Now, for any candidate solution that Alice proposes, Bob will tell her what this additional information has to say about it. We therefore have a new game of “*m* questions” in which the answer to each question is not *whether* some proposed solution belongs to the target but rather *what* the additional information has to say specifically about each candidate solution. It follows that all the action in this new game of “*m* questions” centers on the additional information. Is it enough to render Alice’s *m*-step search for

the target successful? And if so, what characteristics must this additional information possess?

Evolutionary algorithms improve on blind search by exploiting such additional information.¹⁴ At the same time, evolutionary algorithms must not take undue advantage of such additional information. Suppose, for instance, the additional information takes the form of a fitness function, with the target corresponding to where in the phase space the fitness function achieves a certain optimal level of fitness. In that case it is not legitimate for the evolutionary algorithm to survey the fitness landscape induced by the fitness function, identify where in the phase space it attains a global maximum, and then head in that direction. That would be teleology. Teleology has no legitimate place in evolutionary algorithms insofar as they purport to model the Darwinian selection mechanism.

Evolutionary algorithms are sometimes called “no prior knowledge algorithms.” Accordingly, the additional information that evolutionary algorithms exploit must be confined to individual candidate solutions already proposed. In case the additional information is given by a fitness function, the evolutionary algorithm must navigate its way to the target either by randomly choosing points from the phase space or by using those as starting points and then selecting other points based solely on the local topology of the phase space and without recourse to the fitness function except to evaluate it at points already traversed by the algorithm. In other words, the algorithm must move around the phase space only on the basis of its local topology and the elevation of the fitness function at points in the phase space already traversed.¹⁵

Certainly this means that the evolutionary algorithm has to be highly constrained in its use of the fitness function. But there’s more. It means that its success in hitting the target depends crucially on the structure of the fitness function and the local topology of the phase space. If, for instance, the fitness function is totally flat and close to zero whenever it is outside the target, then it fails to discriminate between points outside the

target and so cannot be any help guiding an evolutionary algorithm into the target. For such a fitness function, the probability of the evolutionary algorithm landing in the target is no better than the probability of a blind search landing in the target (an eventuality we've dismissed out of hand—the target simply has too small a probability for blind search to stand any hope of success).

In general, then, given a target within a phase space, where the probability of the target is so small as to constitute an instance of specified complexity, the search for that target cannot be blind but must be suitably constrained if the search is to be successful. What constrains the search is additional information. Moreover, the characteristics of that information will be crucial to the success of the search. If Alice is conducting the search and Bob is providing the additional information, Bob must provide Alice with information whose characteristics and structure enable Alice with high probability to locate the target (note that this probability always depends on the sample size m of candidate solutions that Alice is able to propose). At the same time, Bob must not let Alice take undue advantage of that information. Bob therefore provides Alice with local rather than global access to that information. In particular, Bob mediates the information through the choice of candidate solutions that Alice has thus far proposed in her search. This ensures that Alice's search procedure exploits no prior knowledge of the target's location.

In limiting herself to local rather than global access to Bob's information, Alice avoids any obvious teleology and can rightly regard her search procedure as an "evolutionary algorithm" in the Darwinian sense, that is, as a faithful mathematical correlate to the Darwinian selection mechanism. The question remains whether such evolutionary algorithms smuggle in any hidden teleology and thus merely rework preexisting specified complexity rather than generate it *de novo*. In fact, evolutionary algorithms are incapable of generating specified complexity *de novo*. To see this, let us turn next to the No Free Lunch theorems.

3. The No Free Lunch Theorems

The No Free Lunch theorems are of recent vintage. David Wolpert and William Macready proved the first of these in 1996.¹⁶ The No Free Lunch theorems take our attention off the precise features of the information that transforms a blind into a constrained search and instead refocus it on the reference class of possibilities to which that information belongs. When the probability of a target is tiny, additional information is required to transform a blind into a constrained search if the search is to stand a reasonable chance of hitting the target. The No Free Lunch theorems show that the precise features of the information that guides an evolutionary algorithm to a target are irrelevant to deciding whether evolutionary algorithms can generate specified complexity. What's crucial is the reference class of possibilities to which that information belongs and from which it is drawn.

The additional information that transforms a blind into a constrained search is, after all, information. Precisely because it is information, it must by definition belong to a reference class of possibilities. Information always presupposes a reference class of multiple live possibilities, one of which is selected and the others of which are ruled out. Robert Stalnaker puts it this way: "Content requires contingency. To learn something, to acquire information, is to rule out possibilities. To understand the information conveyed in a communication is to know what possibilities would be excluded by its truth."¹⁷ Fred Dretske elaborates: "Information theory identifies the amount of information associated with, or generated by, the occurrence of an event (or the realization of a state of affairs) with the reduction in uncertainty, the elimination of possibilities, represented by that event or state of affairs."¹⁸

The additional information that constrains an otherwise blind search is drawn from a reference class of possibilities. Let us call this reference class of possibilities the

informational context. The informational context houses the additional information that transforms a blind into a constrained search. In the case of Alice and Bob, the informational context constitutes the informational resources that Bob might employ to assist Alice in locating the target. We suppose that Bob has full access to the informational context, selects some item of information from it, and makes it available to Alice to help her locate the target (though by giving her only local rather than global access to the information). The reason we speak of No Free Lunch theorems (plural) is to distinguish the different types of informational contexts from which Bob might select information to assist Alice.

What can the informational context look like? We've already seen where the informational context is the class of fitness functions on the phase space. Since the phase space is a topological space, the informational context could as well be the continuous fitness functions on the phase space. If the phase space is a differentiable manifold, the informational context could be the differentiable fitness functions on the phase space. The informational context could even be a class of temporally indexed fitness functions that identify not merely fitness, but fitness at some time t (such temporally indexed fitness functions have yet to find wide-spread use, but seem more appropriate for modeling fitness, which in any realistic environment is not likely to be static). Alternatively, the informational context need not involve any fitness functions whatsoever. The informational context could be a class of dynamical systems, describing the flow of particles through phase space.¹⁹ The possibilities for such informational contexts are limitless, and each such informational context has its own No Free Lunch theorem.

A generic No Free Lunch theorem now looks as follows: It sets up a performance measure that characterizes how effectively an evolutionary algorithm locates a given target given a sample size of at most m candidate solutions and given a particular item of information from the informational context. Next this performance measure is

averaged over all the items of information from the informational context. A generic NFL theorem then states that this averaged performance measure is independent of evolutionary algorithm—in other words, it's the same for all evolutionary algorithms. And since blind search always constitutes a perfectly valid evolutionary algorithm, this means that the average performance of any evolutionary algorithm is no better than blind search.²⁰

The significance of the No Free Lunch theorems is that an informational context does not, and indeed cannot, privilege a given target. Instead, an informational context contains information that is equally adept at guiding an evolutionary algorithm to other targets in the phase space. This was certainly the case with Richard Dawkins's well-known METHINKS IT IS LIKE A WEASEL example. In that example, Dawkins defined an evolutionary algorithm that gradually transformed a random character sequence into the target sequence METHINKS IT IS LIKE A WEASEL by randomly varying alphabetic characters and fixing them whenever they matched the corresponding character in the target sequence. Dawkins's evolutionary algorithm was a constrained search that on average attained the target sequence in about 40 steps. Unconstrained, the search would on average take 10^{40} steps.

The informational context here was a class of fitness functions, and the information rendering Dawkins's search successful was the fitness function that assigned to an arbitrary character sequence the number of characters coinciding with the target sequence. But note, in the formulation of this fitness function there was nothing special about the target sequence METHINKS IT IS LIKE A WEASEL. Any other character sequence of 28 letters and spaces would have served equally well. Given any target sequence whatsoever, we can define a fitness function that assigns the number of places where an arbitrary character sequence agrees with it. Moreover, given this fitness function, Dawkins's evolutionary algorithm will just as surely converge to the new target as previously it converged to METHINKS IT IS LIKE A WEASEL.

In general, then, there are no privileged targets, and the only thing distinguishing targets is the choice of information from the informational context. But this means that the problem of locating a target has been displaced. The new problem is locating the information needed to locate the target. The informational context thus becomes a new phase space in which we must locate a new target—the new target being the information needed to locate the original target. To say that an evolutionary algorithm has generated specified complexity within the original phase space is therefore really to say that it has borrowed specified complexity from a higher-order phase space, namely, the informational context. And since in practice this new phase space is much bigger and much less tractable than the original phase space (typically it is exponential in the original phase space), it follows that the evolutionary algorithm has in fact not generated specified complexity at all but merely shifted it around.

We have here a particularly vicious regress. For the evolutionary algorithm to generate specified complexity within the original phase space presupposes that specified complexity was first generated within the higher-order phase space, namely, the informational context. But how was this prior specified complexity generated. Clearly, it would be self-defeating to claim that some higher-order evolutionary algorithm on the higher-order phase space generated specified complexity; for then we face the even more difficult problem of generating specified complexity from a still higher-order phase space (i.e., the informational context of the informational context of the original phase space). This regress, in which evolutionary algorithms shift the problem of generating specified complexity from an original phase space to a higher-order phase space holds with perfect generality and is the take-home lesson of the No Free Lunch theorems.

4. The Displacement Problem

The essential difficulty in generating specified complexity with an evolutionary algorithm can now be stated quite simply. An evolutionary algorithm is supposed to find a target within phase space. To do this successfully, however, it needs more information than is available to a blind search. But this additional information resides within a wider informational context. And locating that additional information within the wider context is no easier than locating the original target within the original phase space. Evolutionary algorithms therefore displace the problem of generating specified complexity but do not solve it. I call this *the displacement problem*.

Think of it this way. There is an island with buried treasure. You can scour the island trying to find the buried treasure. Alternatively, you can try to find a map that tells you where the treasure is buried. Once such a map is in hand, finding the treasure is no problem. But how to find such a map? Suppose such a map exists but resides among a huge assortment of other maps. Finding the right map within that huge assortment will then be no easier than simply searching the island directly. The huge assortment of maps is the informational context. In general, an informational context is no easier to search than the original phase space. Typically the problem gets much worse since informational contexts tend to be function spaces on the original phase spaces (fitness functions being the best known case). For such function spaces, searching the informational context is exponential in the original phase space.

There is no way around the displacement problem. This is not so say that there have not been attempts to get around it. But invariably we find that when specified complexity seems to be generated for free, it has in fact been front-loaded, smuggled in, or hidden from view. I want, therefore, next to review some attempts to get around the displacement problem and uncover just where the displaced information resides once it goes underground.

First off, let us be clear about that the No Free Lunch theorems that underwrite the displacement problem apply with perfect generality – NFL applies to *any* information that might supplement a blind search, and not just to fitness functions. Usually the NFL theorems are stated in terms of fitness functions over phase spaces. Thus, in the case of biological evolution, one can try to mitigate the force of NFL by arguing that evolution is nonoptimizing. Joseph Culberson, for instance, asks, “If fitness is supposed to be increasing, then in what nontrivial way is a widespread species of today more fit than a widespread species of the middle Jurassic?”²¹ But NFL theorems can just as well be formulated for informational contexts that do not comprise fitness functions. The challenge facing biological evolution, then, is to avoid the force of NFL when evolutionary algorithms also have access to information other than fitness functions. Merely arguing that evolution is nonoptimizing is therefore not enough. Rather, one must show that finding the information that guides an evolutionary algorithm to a target is substantially easier than finding the target directly through a blind search.

Think of it this way. In trying to locate a target, you can sample no more than m points in phase space. What’s more, your problem is sufficiently complex that you will need additional information to find the target. That information resides in a broader informational context (what we have called the information-resource space). If searching through that broader informational context is no easier than searching through the original phase space, then you are no better off going with an evolutionary algorithm than going with a straight blind search. Moreover, you cannot arbitrarily truncate your informational context simply to facilitate your search, for any such truncation will itself be an act of ruling out possibilities, and that by definition means an intrusion of novel information, and in particular of specified complexity. In effect, you will be smuggling in what you are claiming to discover.²²

To resolve the displacement problem therefore requires an answer to the following question: How can the informational context be simplified sufficiently so that finding

the information needed to locate a target is easier than finding the target using blind search? There is only one way to do this without arbitrarily truncating the informational context, and that is for the phase space itself to constrain the informational context. Structures and regularities of the phase space must by themselves be enough to constrain the selection of points in the phase space and thus facilitate locating the target. The move here, then, is from contingency to necessity; from evolutionary algorithms to dynamical systems; from Darwinian evolution to complex self-organization. Stuart Kauffman's approach to biological complexity epitomizes this move, focusing on autocatalytic reactions that reliably settle into complex behaviors and patterns.²³

Nonetheless, even this proposed resolution of the displacement problem fails. Yes, structures and regularities of the phase space can simplify the informational context so that finding the information needed to locate a target is easier than finding the target using blind search. But whence those structures and regularities in the first place? Structures and regularities are constraints. And constraints, by their very specificity, could always have been otherwise. A constraint that is not specific is no constraint at all. Constraints are constraints solely in virtue of their specificity – they permit some things and rule out others. But in that case different constraints could fundamentally alter what is permitted and what is ruled out. Thus, the very structures and regularities that were supposed to eliminate contingency, information, and specified complexity merely invite them back in.

Exploiting constraints on a phase space to locate a target is therefore merely another way of displacing information. Not only does it not solve the displacement problem; its applicability is quite limited. Many phase spaces are homogeneous and provide no help in locating targets. Consider for instance a phase space comprising all possible character sequences from a fixed alphabet (such phase spaces model not only written texts but also polymers – e.g., DNA, RNA, and proteins). Such phase spaces are perfectly homogeneous, with one character string geometrically interchangeable with the next.

Whatever else the constraints on such spaces may be, they provide no help in locating targets. Rather, external semantic information (in the case of written texts) or functional information (in the case of polymers) is needed to locate a target.²⁴

To sum up, there is no getting around the displacement problem. Any output of specified complexity requires a prior input of specified complexity. In the case of evolutionary algorithms, they can yield specified complexity only if they themselves are carefully front-loaded with the right information (typically via a fitness function) and thus carefully adapted to the problem at hand. In other words, all the specified complexity we get out of an evolutionary algorithm has first to be put into its construction and into the information that guides the algorithm. Evolutionary algorithms therefore do not generate or create specified complexity but merely harness already existing specified complexity.

How, then, does one generate specified complexity? There is only one known generator of specified complexity, and that is intelligence.²⁵ In every case where we know the causal history underlying an instance of specified complexity, an intelligent agent was involved. Most human artifacts, from Shakespearean sonnets to Dürer woodcuts to Cray supercomputers, are specified and complex. For a signal from outer space to convince astronomers that extraterrestrial life is real, it too will have to be complex and specified, thus indicating that the extraterrestrial is not only alive but also intelligent (hence the search for extraterrestrial *intelligence*—with emphasis on the “intelligence”).²⁶ Thus, to claim that natural laws, even radically new ones as Paul Davies suggests, can produce specified complexity is to commit a category mistake. It is to attribute to laws something they are intrinsically incapable of delivering.²⁷ Indeed, all our evidence points to intelligence as the sole source for specified complexity.

5. Darwinian Evolution in Nature

We need now to step back and consider carefully what the displacement problem means for Darwinian evolution as it occurs in nature. Darwinists are unlikely to see the displacement problem as a serious threat to their theory. I have argued that evolutionary algorithms fail to generate specified complexity because they smuggle it in by the addition of information from an informational context. Darwinian evolution as it occurs in nature, however, seems not to smuggle in anything. Nature places organisms under environmental pressure that permits some to survive and reproduce and kills off the rest. Differential survival and reproduction is nature's criterion for optimization and induces a fitness function that guides Darwinian evolution. This fitness function is the additional information that transforms Darwinian evolution from a blind to a constrained search. But that additional information is simply a brute given. It does not appear to reside in any broader informational context. It simply is.

Richard Dawkins will therefore distinguish Darwinian evolution as it occurs in nature from evolution as it occurs in computational simulations. Commenting on his METHINKS IT IS LIKE A WEASEL example, Dawkins writes:

Although the monkey/Shakespeare model is useful for explaining the distinction between single-step selection and cumulative selection, it is misleading in important ways. One of these is that, in each generation of selective 'breeding', the mutant 'progeny' phrases were judged according to the criterion of resemblance to a distant ideal target, the phrase METHINKS IT IS LIKE A WEASEL. Life isn't like that. Evolution has no long-term goal. There is no long-distance target, no final perfection to serve as a criterion for selection.... In real life, the criterion for selection is always short-term, either simple survival or, more generally, reproductive success.... The 'watchmaker' that is cumulative natural selection is blind to the future and has no long-term goal.²⁸

The Darwinist therefore objects that “real life” Darwinian evolution can in fact generate specified complexity after all without smuggling it in. The fitness function in biological evolution follows directly from differential survival and reproduction, and this, according to the Darwinist, *can* legitimately be viewed as a “free lunch.” Organisms will sample different variants via random variation (often conceived of as genetic mutations), and then the fitness function freely bestowed by differential survival and reproduction will select those variants that constitute an improvement, which within Darwinism is defined by being better at surviving and reproducing. No specified complexity is required as input in advance.

If this objection is conceded, then the only way to show that the Darwinian mechanism cannot generate specified complexity is by demonstrating that the gradients of the fitness function induced by differential survival and reproduction are not sufficiently smooth for the Darwinian mechanism to drive large-scale biological evolution. To use another Dawkins metaphor, one must show that there is no gradual way to ascend “Mount Improbable.”²⁹ This is a separate line of argument and one Michael Behe develops in his book *Darwin's Black Box*, where he argues that irreducibly complex biochemical systems are inherently inaccessible to any gradual Darwinian pathway. This is a powerful line of argument, and one that provides empirical evidence against Darwinism. Here, however, I want to show on purely theoretical grounds that the displacement problem does indeed apply to Darwinism and that the Darwinian mechanism has no way of escaping the design problem that it raises.

Things are not nearly as simple as taking differential survival and reproduction as brute givens and from there concluding that the resulting fitness function is likewise a brute given. Differential survival and reproduction by themselves do not guarantee that any interesting evolution will occur. Consider, for instance, Sol Spiegelman's work on the evolution of polynucleotides in a replicase environment. Leaving aside that the replicase protein is supplied by the investigator (from a viral genome), as are the

activated mononucleotides needed to feed polynucleotide synthesis, the problem here and in experiments like it is the steady attenuation of information over the course of the experiment. As Brian Goodwin notes:

In a classic experiment, Spiegelman in 1967 showed what happens to a molecular replicating system in a test tube, without any cellular organization around it. The replicating molecules (the nucleic acid templates) require an energy source, building blocks (i.e., nucleotide bases), and an enzyme to help the polymerization process that is involved in self-copying of the templates. Then away it goes, making more copies of the specific nucleotide sequences that define the initial templates. But the interesting result was that these initial templates did not stay the same; they were not accurately copied. They got shorter and shorter until they reached the minimal size compatible with the sequence retaining self-copying properties. And as they got shorter, the copying process went faster. So what happened with natural selection in a test tube: the shorter templates that copied themselves faster became more numerous, while the larger ones were gradually eliminated. This looks like Darwinian evolution in a test tube. But the interesting result was that this evolution went one way: toward greater simplicity. Actual evolution tends to go toward greater complexity, species becoming more elaborate in their structure and behavior, though the process can also go in reverse, toward simplicity. But DNA on its own can go nowhere but toward greater simplicity. In order for the evolution of complexity to occur, DNA has to be within a cellular context; the whole system evolves as a reproducing unit.³⁰

My point here is not that Darwinian evolution in a test tube should be regarded as disconfirming evidence for Darwinian evolution in nature. Rather, it is that if the Darwinian mechanism of differential survival and reproduction is what in fact drives full-scale biological evolution in nature, then the fitness function induced by that

mechanism has to be very special. Indeed, many prior conditions need to be satisfied for the function to take a form consistent with the Darwinian mechanism being the principal driving force behind biological evolution. Granted, the fitness function induced by differential survival and reproduction in nature is non-arbitrary. But that doesn't make it a free lunch either.

Think of it this way. Suppose we are given a phase space of replicators that replicate according to a Darwinian mechanism of differential survival and reproduction. Suppose this mechanism induces a fitness function. Given just this information, we don't know if evolving this phase space over time will lead to anything interesting. In the case of Spiegelman's experiment, it didn't—Darwinian evolution led to increasingly simpler replicators. In real life, however, Darwinian evolution is said to lead to vast increases in the complexity of replicators, with all cellular organisms tracing their lineage back to a common unicellular ancestor. Let's grant this. The phase space then comprises a vast array of DNA-based self-replicating cellular organisms and the Darwinian mechanism of differential survival and reproduction over this phase space induces a fitness function that underwrites full-scale Darwinian evolution. In other words, the fitness function is consistent not only with the descent of all organisms from a common ancestor (i.e., common descent), but also with the Darwinian mechanism accounting for the genealogical interrelatedness of all organisms. Now suppose this is true. What prior conditions have to be satisfied for the fitness function to be the type of fitness function that allows a specifically Darwinian form of evolution to flourish?

For starters, the phase space had better be non-empty, and that presupposes raw materials like carbon, hydrogen, and oxygen. Such raw materials, however, presuppose star formation, and star formation in turn presupposes the fine-tuning of cosmological constants. Thus for the fitness function to be the type of fitness function that allows the Darwinian mechanism to flourish presupposes all the anthropic principles and cosmological fine-tuning that lead many physicists to see design in the universe. Yet

even with cosmological fine-tuning in place, many additional conditions need to be satisfied. The phase space of DNA-based self-replicating cellular organisms needs to be housed on a planet that's not too hot and not too cold. It needs a reliable light source. It needs to have a sufficient diversity of minerals and especially metals. It needs to be free from excessive bombardment by meteors. It needs not only water but enough water. Michael Denton's book *Nature's Destiny* is largely devoted to such specifically terrestrial conditions that need to be satisfied if biological evolution on earth is to stand any chance of success.³¹

But there's more. Cosmology, astrophysics, and geology fail to exhaust the conditions that a fitness function must satisfy if it is to render not just biological evolution but a specifically a Darwinian form of it the grand success we see on planet earth. Clearly, replicators needs to be robust in the sense of being able to withstand frequent and harsh environmental insults. This may seem self-evident, but computer simulations with artificial life forms tend to be quite sensitive to unexpected perturbations and thus lack the robustness we see in terrestrial biology. DNA-based replicators are indeed robust. What's more, the DNA copying mechanism of such replicators must be sufficiently reliable to avoid error catastrophes. Barring a high degree of reliability the replicators will go extinct or wallow interminably at a low state of complexity (basically just enough complexity to avoid the error catastrophe).

But perhaps most importantly, the replicators must be able to increase fitness and complexity in tandem. In particular, fitness must not be positively correlated with simplicity. This last requirement may seem easily purchased, but it is not. Stephen Jay Gould, for instance, in *Full House* argues that replication demands a certain minimal level of complexity below which things are dead (i.e., no longer replicate). Darwinian evolution is thus said to constitute a random walk off a reflecting barrier, the barrier constituting a minimal complexity threshold for which increases in complexity always permit survival but decreases below that level entail death. Enormous increases in

complexity are thus said to become not only logically possible but also highly probable.³²

The problem with this argument is that in the context of Darwinian evolution such a reflecting barrier tends also to be an absorbing barrier (i.e., there's a propensity for replicators to stay close to if not right at the minimal complexity threshold). As a consequence, such replicators will over the course of evolution remain simple and never venture into high degrees of complexity. Simplicity by definition always entails a lower cost in raw materials (be they material or computational) than increases in complexity, and so there is an inherent tendency in evolving systems for selection pressures to force such systems toward simplicity (or as it is sometimes called, *elegance*).

Fitness functions induced by differential survival and reproduction are more naturally inclined to place a premium on simplicity and regard replicators above a certain complexity threshold as too cumbersome to survive and reproduce. The Spiegelman example is a case in point. Thomas Ray's Tierra simulation gave a similar result, showing how selection acting on replicators in a computational environment also tended toward simplicity rather than complexity—unless parameters were set so that selection could favor larger sized organisms (complexity here corresponding to size).³³ This is not to say that the Darwinian mechanism automatically takes replicating systems toward a minimal level of complexity, but that if it doesn't, then some further conditions need to be satisfied, conditions reflected in the fitness function.

Vast is the catalogue of conditions that the fitness function induced by differential survival and reproduction needs to satisfy if the spectacular complexity and diversity of living forms that we see on earth is properly to be attributed to a Darwinian form of evolution. Clearly, such a catalogue is going to require a vast amount of specified complexity, and this specified complexity will be reflected in the fitness function that, as Darwinists rightly note, is non-arbitrary but, as Darwinists are reluctant to accept, is also not a free lunch. Throw together some replicators, subject them to differential

survival and reproduction, perhaps add a little game theory to the mix (à la Robert Wright),³⁴ and there's no reason to think you'll get anything interesting, and certainly not a form of Darwinian evolution that's worth spilling any ink over.

It follows that the Darwinian mechanism has no explanatory power in accounting for biological complexity. We say that X possesses explanatory power in accounting for Y insofar as X discriminates Y from non-Y. SAT scores, for instance, do a reasonable job predicting college performance and thus possess some explanatory power in accounting for success in college. But SAT scores possess no explanatory power in accounting for human genius – there are plenty of people with perfect SAT scores who will never join the company of da Vinci, Mozart, and Plato. Inflated claims by hardcore Darwinists like Daniel Dennett about the Darwinian mechanism are therefore not only false but patently seen to be false. Dennett, for instance, will enthuse:

If I were to give an award for the single best idea anyone has ever had, I'd give it to Darwin, ahead of Newton and Einstein and everyone else. In a single stroke, the idea of evolution by natural selection unifies the realm of life, meaning, and purpose with the realm of space and time, cause and effect, mechanism and physical law.³⁵

But there are a host of evolutionary contexts within which selection operates and yet no interesting evolution occurs. It follows that something other than the Darwinian selection mechanism is needed to account for biological complexity. It is remarkable how much confusion there is on this point in the biological community. Darwin's contemporary John Stuart Mill laid out the relevant logic here well before Darwin published his *Origin of Species*. In his *System of Logic*, Mill described various methods of induction, including his "method of difference."³⁶ According to the method of difference, to explain a difference in effects, a difference must be sought in the antecedent conditions for those effects. But natural selection is a common feature of evolutionary scenarios that alternately lead to increasing complexity (as with cellular

life) or wallow interminably in simplicity (as in the Spiegelman case). As a common feature of such diverse evolutionary scenarios, the Darwinian selection mechanism is therefore incapable of explaining any difference in these scenarios and thus in particular incapable of accounting for biological complexity.

What, then, besides the Darwinian selection mechanism could account for biological complexity? The answer is, of course, design. But don't take my word for it. Evolutionary algorithms researchers are themselves now realizing that for evolutionary algorithms to output design requires a prior input of design. Consider, for instance, the following remarks by Geoffrey Miller about genetic algorithms:

Genetic algorithms are rather robust search methods for [simple problems] and small design spaces. But for hard problems and very large design spaces, designing a good genetic algorithm is very, very difficult. All the expertise that human engineers would use in confronting a design problem—their knowledge base, engineering principles, analysis tools, invention heuristics and common sense—must be built into the genetic algorithm. Just as there is no general-purpose engineer, there is no general-purpose genetic algorithm.³⁷

And where exactly does design get built into an evolutionary or genetic algorithm? According to Miller, it gets built into the fitness function. He writes:

The fitness function must embody not only the engineer's conscious goals, but also her common sense. This common sense is largely intuitive and unconscious, so is hard to formalize into an explicit fitness function. Since genetic algorithm solutions are only as good as the fitness functions used to evolve them, careful development of appropriate fitness functions embodying all relevant design constraints, trade-offs and criteria is a key step in evolutionary engineering.³⁸

For some time now Darwinists have been claiming that evolutionary algorithms provide a computational justification for the Darwinian mechanism of natural selection and random variation as the primary creative force in biology. Yet if Darwinists want to

take seriously the implications of evolutionary algorithms for biology, they need to admit that the Darwinian mechanism does not resolve the problem of biological design. Indeed, evolutionary algorithms, far from eliminating the design problem, merely push it deeper. Consequently, if the Darwinian mechanism of natural selection and random variation is the means by which the complexity and diversity of life came to be, the environmental fitness that constrains biological evolution would not be a free lunch and not a brute given, but a finely crafted assemblage of smooth gradients that presupposes much prior specified complexity and therefore prior design.

In conclusion, I want to draw a pair of lessons. Both intelligent design and evolutionary algorithms have a lesson to learn from each other. The No Free Lunch theorems show that for evolutionary algorithms to output specified complexity they had first to receive a prior input of specified complexity. And since specified complexity is reliably linked to intelligence, evolutionary algorithms, insofar as they output specified complexity, do so on account of a guiding intelligence. The lesson, then, for evolutionary algorithms is that any intelligence these algorithms display is never autonomous but always derived. This is why I say that natural selection can't design anything—it can only harness preexisting designs.

Nevertheless, evolutionary algorithms do produce remarkable solutions to problems—solutions that in many cases we would never have imagined on our own. Having been given some initial input of specified complexity, evolutionary algorithms mine it to extract what value they can from it. The lesson, then, for intelligent design is that natural causes can synergize with intelligent causes to produce results far exceeding what intelligent causes left to their own abstractions might ever accomplish (this view is, of course, highly congenial to an incarnational theology). Too often design is understood in a deterministic sense in which every aspect of a designed object has to be preordained by a designing intelligence. Evolutionary algorithms underwrite a nondeterministic conception of design in which design and nature operate in tandem to

produce results that neither could produce by itself.³⁹ I close with a quote by Michael Polanyi very much in this spirit:

It is true that the teleology rejected in our day is understood as an overriding cosmic purpose necessitating all the structures and occurrences in the universe in order to accomplish itself. This form of teleology is indeed a form of determinism—perhaps even a tighter form of determinism than is provided for by a materialistic, mechanistic atomism. However, since at least the time of Charles Saunders Peirce and William James a looser view of teleology has been offered to us—one that would make it possible for us to suppose that some sort of intelligible directional tendencies may be operative in the world without our having to suppose that they *determine* all things. Actually it is possible that even Plato did not suppose that his “Good” *forced* itself upon all things. As Whitehead has pointed out, Plato tells us that the Demiurge, looking toward the Good, “persuades” an essentially free matter to structure itself, to some extent, in imitation of the Forms. Plato appeared to Whitehead to have modeled the cosmos on a struggle to achieve the Good in the somewhat recalcitrant media of space and time and matter, a struggle well known to all souls with purposes and ends and aims.⁴⁰

Notes

¹Davies claims that we are “a very long way from comprehending” how life originated. “This gulf in understanding is not merely ignorance about certain technical details, it is a major conceptual lacuna... My personal belief, for what it is worth, is that a fully satisfactory theory of the origin of life demands some radically new ideas.” Paul Davies, *The Fifth Miracle: The Search for the Origin and Meaning of Life* (New York: Simon & Schuster, 1999), 17.

²See Davies, *Fifth Miracle*, 115–122. Cf. also Michael Polanyi, “Life Transcending Physics and Chemistry,” *Chemical and Engineering News* (21 Aug. 1967): 55–66; and Michael Polanyi, “Life’s Irreducible Structure,” *Science* 113 (1968): 1308–1312.

³The (Shannon) information $I(A)$ associated with an event A is by definition $-\log_2 P(A)$, where $P(A)$ is the probability of that event and the logarithm is taken to the base 2.

⁴There are lots and lots of different complexity measures. Seth Lloyd records over thirty (see John Horgan, *The End of Science* [New York: Broadway Books, 1996], 303, n. 11). Horgan regards this abundance of complexity measures as a bad thing (194–198), but it’s not. Having many “flavors of complexity” does not subjectivize the notion. Just as we need many types of measures in daily life (volumes, densities, weights, lengths, times, etc.), so we need many different complexity measures to measure the diverse types of complication associated with diverse structures.

⁵See William A. Dembski, *The Design Inference* (New York: Cambridge University Press, 1998), chs. 5, 6, and 7; and William A. Dembski, *Intelligent Design* (Downers Grove, Ill.: InterVarsity, 1999), ch. 6.

⁶Davies, *Fifth Miracle*, 112. Compare, “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.” Quoted from Leslie Orgel, *The Origins of Life* (New York: John Wiley, 1973), 189.

⁷Davies, *Fifth Miracle*, 120.

⁸*Ibid.*, 42.

⁹Theodosius Dobzhansky, Discussion of G. Schramm’s paper, in *The Origins of Prebiological Systems and of Their Molecular Matrices*, ed. S. W. Fox (New York: Academic Press, 1965), 310.

¹⁰Stuart Kauffman, *At Home in the Universe: The Search for the Laws of Self-Organization and Complexity* (New York: Oxford University Press, 1995), 150. Note that Kauffman himself dissents from this majority view.

¹¹The definition of evolutionary algorithms given here is more general than is customary. For a popular exposition of the types of search strategies included here under evolutionary algorithms, consult Peter Coveney and Roger Highfield, *Frontiers of Complexity: The Search for Order in a Chaotic World* (New York: Fawcett Columbine, 1995). For the connection between organic evolution and evolutionary algorithms, see Thomas Bäck, *Evolutionary Algorithms in Theory and Practice: Evolution Strategies, Evolutionary Programming, Genetic Algorithms* (New York: Oxford University Press, 1996), ch. 1.

¹²Culberson, “On the Futility of Blind Search: An Algorithmic View of ‘No Free Lunch,’” *Evolutionary Computation* 6(2) (1998): 109.

¹³With this example as well as with others in this chapter I am being lax about the level of complexity needed to qualify as specified complexity. Technically, the level of complexity needs to attain at least the universal probability bound of 10^{-150} or the corresponding universal complexity bound of 500 bits. But for the purposes of illustration I am allowing less stringent bounds. See Dembski, *The Design Inference*, ch. 6.

¹⁴Though note, blind search does constitute an evolutionary algorithm. It’s just not a particularly effective one for most purposes.

¹⁵Joseph C. Culberson puts it this way: “Evolutionary algorithms (EAs) are often touted as ‘no prior knowledge’ algorithms. This means that we expect EAs to perform without special information from the environment. Similar claims are often made for other adaptive algorithms.” See Culberson, “On the Futility of Blind Search,” 109–127.

¹⁶David H. Wolpert and William G. Macready, “No Free Lunch Theorems for Optimization,” *IEEE Transactions on Evolutionary Computation* 1(1) (1997): 67–82.

¹⁷Robert Stalnaker, *Inquiry* (Cambridge, Mass.: MIT Press, 1984), 85.

¹⁸Fred Dretske, *Knowledge and the Flow of Information* (Cambridge, Mass.: MIT Press, 1981), 4.

¹⁹Cf. Morris Hirsch and Stephen Smale, *Differential Equations, Dynamical Systems, and Linear Algebra* (New York: Academic Press, 1974).

²⁰For this generic way of formulating NFL theorems, see Culberson, "On the Futility of Blind Search," 111–112.

²¹*Ibid.*, 125.

²²The essential idea behind information is the reduction of possibilities from a reference class of possibilities. That is why information theorists define information as the *reduction* or *resolution* of *uncertainty*. See John R. Pierce, *An Introduction to Information Theory: Symbols, Signals and Noise*, 2nd ed. (New York: Dover, 1980), 24.

²³Kauffman, *At Home in the Universe*, ch. 4, titled "Order for Free."

²⁴Stephen Meyer has argued this point convincingly. See his article "DNA by Design: An Inference to the Best Explanation for the Origin of Biological Information," *Rhetoric and Public Affairs* 1(4) (1998): 519–556.

²⁵See Douglas Robertson, "Algorithmic Information Theory, Free Will, and the Turing Test," *Complexity* 4(3) (1999): 25–34. Robertson argues that the defining feature of agents with free will is their ability to create (complex specified) information.

²⁶Recall the crucial signal in the movie *Contact* that convinced the radio astronomers that they had indeed established "contact" with an extraterrestrial intelligence, namely, a long sequence of prime numbers.

²⁷Davies, *The Fifth Miracle*, 17. The subtitle of Stuart Kauffman's *At Home in the Universe* demonstrates quite plainly this impulse to explain specified complexity in terms of laws: *The Search for the Laws of Self-Organization and Complexity*. Note that Kauffman refers explicitly to "the search" for such laws. At present they remain unknown. See also Roger Penrose's *The Emperor's New Mind* (Oxford: Oxford University Press, 1989) and *Shadows of the Mind* (Oxford: Oxford University Press, 1994). Penrose hopes to unravel the problem of human consciousness through unknown quantum-theoretical laws. There are no proposals for what laws that generate specified complexity might look like, much less how they might actually be formulated. The point of this paper is to argue that no such laws can exist.

²⁸Richard Dawkins, *The Blind Watchmaker* (New York: Norton, 1986), 50.

²⁹Richard Dawkins, *Climbing Mount Improbable* (New York: Norton, 1996).

³⁰Brian Goodwin, *How the Leopard Changed Its Spots: The Evolution of Complexity* (New York: Scribner's, 1994), 35–36.

³¹Michael Denton, *Nature's Destiny: How the Laws of Biology Reveal Purpose in the Universe* (New York: Free Press, 1998).

³²Stephen Jay Gould, *Full House: The Spread of Excellence from Plato to Darwin* (New York: Harmony Books, 1996), 169–173.

³³See <http://www.hip.atr.co.jp/~ray/pubs/tierra>.

³⁴Cf. Robert Wright, *Nonzero: The Logic of Human Destiny* (New York: Pantheon, 2000).

³⁵Daniel Dennett, *Darwin's Dangerous Idea* (New York: Simon & Schuster, 1995), 21.

³⁶John Stuart Mill, *A System of Logic, Ratiocinative and Inductive, Being a Connected View of the Principles of Evidence and the Methods of Scientific Investigation* (London: John W. Parker, 1843). For a contemporary look at Mill's method of agreement and method of difference, see Peter Lipton, *Inference to the Best Explanation* (London: Routledge, 1991), 20–21.

³⁷Geoffrey Miller, "Technological Evolution as Self-Fulfilling Prophecy," in *Technological Innovation as an Evolutionary Process*, ed. J. Ziman (Cambridge: Cambridge University Press, 2000), 209.

³⁸*Ibid.*, 208.

³⁹This I take to be the take-home lesson of Roger Lewin and Birute Regine's *The Soul at Work: Embracing Complexity Science for Business Success* (New York: Simon & Schuster, 2000). For a business to thrive, a framework within which the business operates must be designed. Yet once that framework is designed and in place, the business must not be micromanaged but allowed to follow its "natural course."

⁴⁰Michael Polanyi and Harry Prosch, *Meaning* (Chicago: University of Chicago Press, 1975), 162–163. Although the synergizing of intelligence and nature can be understood from the perspective of Whiteheadian process theology, it is also possible to take a more traditional theological view. Consider,

for instance, the Eastern Orthodox view on synergy as described in Timothy Ware, *The Orthodox Church* (London: Penguin, 1963), 226–227. Here Ware cites John Chrysostom and Cyril of Jerusalem in support of a synergy between a transcendent intelligence (in this case the Christian God) and nature (in particular, human nature). For a more sustained treatment of synergy from the Eastern Orthodox perspective, see Philip Sherrard, *Human Image: World Image* (Ipswich, UK: Golgonooza Press, 1992), especially chapter 7.