Chapter 11
There Is a Free Lunch After All:
William Dembski's Wrong Answers to Irrelevant Questions

Mark Perakh

The title of William Dembski's book *No Free Lunch* (Dembski 2002a) indicates that he perceives the no free lunch (NFL) theorems (Wolpert and Macready 1997) as pivotal to his thesis asserting that "specified complexity cannot be purchased without intelligence." Indeed, many statements in Dembski's book emphasize the crucial role of the NFL theorems. However, in his response (2002b) to a review rebutting his use of the NFL theorems Dembski claims that the NFL theorems are secondary for his thesis, while his principal argument is related to a "displacement problem."

I will show in this chapter that neither the NFL theorems nor the notions related to the displacement problem support Dembski's thesis.

Critiques of of Dembski's work can be found in a number of publications (Wein 2002a,b; Shallit 2002; Rosenhouse 2002; Perakh 2001a, 2002a, 2002b, 2003; Young 2002; Orr 2002; Van Till 2002). Here I will discuss only chapter 4 (*Evolutionary Algorithms*) of Dembski's book, concentrating on his use of the NFL theorems and on his "displacement problem."

Methinks It Is like a Weasel—Again

In many of his publications, including *No Free Lunch*, Dembski repeatedly discusses Richard Dawkins's METHINKS IT IS LIKE A WEASEL evolutionary algorithm, trying to prove its fallaciousness. Dawkins describes his weasel algorithm (Dawkins 1996 [1986], 47–48) as: "It . . . begins by choosing a random sequence of 28 letters. . .. It now 'breeds from' this random phrase. It duplicates it repeatedly, but with a certain chance of a random error—'mutation'—in the copying. The computer examines the mutant nonsense phrase, the 'progeny' of the original phrase, and chooses the one which, however slightly, most resembles the target phrase."

Dembski sees an inadequacy in Dawkins's algorithm: it converges on a target phrase. He says (p. 182) ". . . choosing a prespecified target sequence as Dawkins does here is deeply teleological. . .. This is a problem because evolutionary algorithms are supposed to be capable of solving complex problems without invoking teleology." But later (page 203) he says, "An evolutionary algorithm is supposed to find a target within phase space." Searching for a target *is* teleological.

Such inconsistency is Dembski's trademark. In any case, neither of his statements is correct. Evolutionary algorithms may be both targeted and targetless. Biological evolution, however, has no long-term target. Evolution is not directed toward any *specific* organism. The evolution of a species may continue indefinitely as long as the environment exerts selection pressure on that species. If a population does not show long-term change, it is not because that population has reached a target but because the environment, which coevolves with the species, acquires properties that eliminate its evolutionary pressure on the species. Dawkins's weasel algorithm, on the other hand, stops when the target phrase has been reached.

Dawkins (1996 [1986], 50) was himself the first to point out that his algorithm differs from biological evolution in that it proceeds toward a target. But then, a *model is not supposed to be a replica* of the entire modeled object or phenomenon (Perakh 2002c); models replicate only those features of the modeled objects that are crucial for analyzing a specific, usually limited, aspect of the modeled object or phenomenon and ignore all the aspects and properties which are of minor importance. Dawkins's algorithm was designed to show that a combination of *random variations* with a suitable *law* can accelerate evolution by many orders of magnitude. (The law in this case is selection). It indeed shows such an acceleration. As Dembski points out, a random search would require, on average, 10^{40} iterations of the search procedure. Dawkins's algorithm performs the task in about only forty iterations.

Dawkins's procedure is not a proof of evolution but it is a valid demonstration of a very significant acceleration of evolution if a suitable law works along with random variations.

That is why, as Dembski laments, (183) "Darwinists and even some non-Darwinists are quite taken" with Dawkins's example. Indeed, Dawkins's is a good example.

Is Specified Complexity Smuggled into Evolutionary Algorithms?

On page 193, Dembski suggests a modification of Dawkins's weasel algorithm. In his modified procedure, the algorithm will "pick a position at random in the sequence… Then randomly alter the character in that position. If the new sequence has a higher fitness function than the old, keep it and discard the old. Otherwise keep the old. Repeat the process."

I see no substantial difference between the procedure described on pages 47–48 in Dawkins's book and that suggested by Dembski. In Dembski's view, while Dawkins's algorithm compares consecutive phrases with a target, Dembski's modified algorithm "searches for the target solely on the basis of the phase space and the fitness function" (p. 194), hence "not smuggling in any obvious teleology." But the only difference between Dawkins's algorithm and its modification by Dembski is in the way they simulate mutations. Otherwise both compare intermediate phrases with the target. In Dembski's version the values of the fitness function are simply the counts of those letters in the intermediate phrases that coincide with the letters occupying the same positions in the target phrase. Indeed, on page 194, we read, "As before, fitness is determined by how close a sequence is to the target sequence." Therefore Dembski's modified algorithm is as teleological as Dawkins's original algorithm.

Continuing (194–196), Dembski insists that evolutionary algorithms cannot generate *specified complexity* (SC) but can only "smuggle" it from a "higher order phase space." This claim is irrelevant to biological evolution. In the case of the weasel algorithm, the outcome is deliberately designed. SC is injected into the algorithm via the fitness function. But since biological evolution has no long-term target, it requires no injection of SC. Natural selection is unaware of its result—the increased chance for having progeny. The advantage in proliferation occurs automatically. If Dembski thinks otherwise he needs to offer evidence that extraneous information must be injected into the natural selection algorithm *apart from that supplied by the fitness functions that arise naturally in the biosphere*. Dembski provides no such evidence.

Furthermore, in Dawkins's weasel example the evolutionary algorithm converges on a meaningful phrase—a quotation from Shakespeare. According to Dembski, the target phrase possesses SC. Michael Behe, in a foreword to (Dembski 1999), gives an example. While the *meaningful* sequence METHINKSITISLIKEAWEASEL is both *complex and specified*, a sequence NDEIRUABFDMOJHRINKE of the same length, which is gibberish, is complex but not specified. Many Dembski's statements scattered over his publications make it clear that Behe has indeed reflected correctly Dembski's position (Perakh 2001, 2003), which is that *a meaningless sequence possesses no SC.*

On the other hand, on page 195 of his NFL book Dembski indicates that Dawkins's algorithm could also be applied if the target phrase were gibberish. If, though, the target sequence is meaningless, then, according to the above quotation from Behe, it *possesses no SC*. If the target phrase possesses no SC, then obviously *no SC had to be "smuggled"* into the algorithm.

Hence, if we follow Dembski's ideas consistently, we have to conclude that the same algorithm "smuggles" SC if the target is meaningful but does not smuggle it if the target is gibberish. This notion is preposterous because algorithms are indifferent to the distinction between meaningful and gibberish targets.

This inconsistency in Dembski and Behe's approach stems from the fact that the very concept of SC is contradictory. In fact, contrary to Behe and Dembski's notions, both a meaningful phrase and a string of gibberish are "specified," if the concept of specification is given back its commonsense meaning by clearing it of the embellishments and unnecessary complications suggested by Dembski (1998a, 1999, 2002a). By having written down a gibberish sequence, Behe has *specified* it. As soon as it has been written, it becomes unequivocally distinguishable from any other sequence, which means it is specified. Whether it is meaningful or gibberish is of no consequence (Perakh 2001, 2003).

Targetless Evolutionary Algorithms

While Dembski devotes much attention to Dawkins's weasel algorithm, he ignores another procedure designed by Dawkins—the biomorphs algorithm (BA). The BA differs from the weasel algorithm in that *it has no target*.

Dawkins's BA is also a model which, though, illustrates a different aspect of evolution. The BA demonstrates how evolution, starting with a very simple "progenitor," can generate unlimited complexity, with no preselected target. Since its purpose is different, the features of evolution retained in this model are also different from the weasel algorithm.

The biomorphs algorithm has its own limitations. Dawkins points out (1996 [1986], 60 ff.) that the biomorphs algorithm "shows us the power of cumulative selection to generate an almost endless variety of quasi-biological form, but it uses artificial selection, not natural selection. The human eye does the selecting." However, in this case as well, the model, as every model, is not a replica of the reality, but is adequate to illustrate an important feature of reality—generation of complexity. Dembski's failure to discuss the targetless biomorphs algorithm undermines his critique of the weasel algorithm for its "teleological" features.

Dembski discusses, though, another model, suggested and used by Thomas Schneider (2000, 2001a,b) who claims it being targetless. Schneider maintains that his algorithm generates biologically meaningful information "from scratch," that is, without an input from intelligence. Dembski disagrees: "The No Free Lunch theorems, however, tell us this is not possible" (215).

The proper analysis of Schneider's evolutionary algorithm would require a much more complex discourse than with Dawkins's weasel algorithm and will not be attempted in this chapter which is addressed also to readers possessing no such knowledge of information theory which is necessary to digest the required mathematical arguments. I will, though, address the quoted Dembski's main argument against Schneider's algorithm. Regardless of whether or not Schneider's algorithm indeed generates information "from scratch," or information is supplied by a hidden target of Schneider's program, Dembski's argument, based on the NFL theorems, misinterprets these theorems. Dembski states repeatedly throughout his book that the *NFL theorems prohibit generation of information without intelligence.* In fact, the NFL theorems do nothing of the sort.

The No Free Lunch Theorems

I will now discuss what the NFL theorems say about biological evolution or about evolutionary algorithms such as those of Dawkins, Schneider, Altshuler and Linden (1999), Chellapilla and Fogel (1999), and others.

On page 223, Dembski states, "The No Free Lunch theorems show that for evolutionary algorithms to output CSI they had first to receive a prior input of CSI." (CSI stands for "complex specified information" which Dembski uses interchangeably with "specified complexity.")

In fact, the NFL theorems show nothing of the sort. To see why, we need a brief excursion into optimization theory.

I take the liberty of suggesting an example stemming from my life-long love for mountain climbing. Consider an expedition to a remote mountain region. If the mountain climbers are interested in finding the highest peak, they have to perform an *optimization search* over the entirety of peaks, that is over the physical relief of a mountain region. They move over the landscape, climb up a mountain at each location and note its height as measured by an altimeter, until they locate the highest peak. There are many details of the search for the highest peak which are irrelevant for the NFL theorems, including such questions as how to determine that the highest peak has indeed been found, and how to know which peak is next in height above the summit already reached. These details do not appear in the NFL theorems, which are very general and apply to a wide variety of searches and landscapes.

Now assume that before embarking on an expedition we wanted to first prepare instructions for climbing the mountains in a yet unexplored region. Let us call such sets of instructions *mountain-climbing algorithms* (MCA).

One group of mountaineers suggests to start from some peak located where the path that leads us to the region approaches its periphery. Then peak after peak is to be climbed, moving gradually toward the center of the region, regardless of whether each next peak is higher or lower than the preceding one.

We will call the MCA prepared according to such a strategy a *center-directed algorithm* (CDA).

Another group suggests a different strategy. A CDA, they say, might miss the highest peak if that is located away from the center of the region. Their preferred criterion for choosing each next peak for exploration is not its being closer to the center, but its being higher than the preceding conquered peak. Let us call that MCA a *height-oriented algorithm* (HOA). The question is: Which of the two algorithms will ensure a faster conclusion of the search for the highest peak? Which algorithm will "perform" better on the landscape of a particular mountainous region?

There is no general answer; it depends on the character of the landscape. On most landscapes, the HOA will perform better than the CDA. If, however, the highest peak is very close to the center of the region and the heights of the surrounding peaks decrease haphazardly toward the periphery, then a search that uses a CDA may outperform an HOA.

Which of the two MCAs is "better" overall? There are many mountainous countries on the globe. Some of them may have a physical relief wherein the highest point is closest to the geographical center of the region with a haphazard distribution of lower peaks around the central one. The highest peak of others may be located away from the geographical center and the heights of the lower peaks decrease gradually with the distance from that highest peak. And there are many others. Overall there may not be substantially fewer landscapes where the CDA is preferable than where the HOA is better. In other words, the performance of both the HOA and CDA *averaged over all possible landscapes* can be reasonably expected to be not much different.

This is a simplified illustration of the NFL theorems. The NFL theorems by Wolpert and Macready put the simple observation about the *equal average* *performance* of various algorithms into a rigorous mathematical form and reveal some subtle features of the algorithms' behavior which are not intuitively evident.

The NFL Theorems—Still with No Mathematics

The mountainous landscape we discussed is a particular case of what generally is called a *fitness landscape*, and the heights of the peaks in our example is a particular case of a *fitness function.*

Imagine two algorithms conducting a search on a given fitness landscape. They move from point to point over the search space (choosing the search points either at random or in a certain order). After having performed, say, *m* measurements, an algorithm produces what Wolpert and Macready call a "sample." A sample is simply a table wherein the *m* measured values of the fitness function are listed in a temporal order. Generally speaking, two arbitrarily chosen algorithms will not yield identical samples. The probability of algorithm a_1 producing a specific table which is *m* rows long is different from the probability of algorithm a_2 producing the same table after the same number of iterations.

Now enter the first NFL theorem: if the results of the two algorithms' searches are compared not for a specific fitness landscape but *averaged over all possible landscapes*, the probabilities of obtaining the same sample are equal for any pair of algorithms. The quantity which is averaged is *the probability of generating a given sample by an algorithm*.

This is an exact translation of the first NFL theorem from its mathematically symbolic form into plain words.

The NFL theorems do not restrict the value of *m*, the number of iterations. There is no condition that the search stops when a certain preselected number of iterations has been completed, or when a preselected value of the fitness function has been found.

In other words, the concept of a *target* is absent from the NFL theorems. On the other hand, the NFL theorems do not forbid the algorithms to be target-oriented. The theorems are indifferent to algorithms' having or not having a target.

The NFL theorems are often discussed in terms of algorithms' performance, though the concept of a performance measure is not part of the NFL theorems as such. As Wolpert and Macready indicate (73) "the precise way that the sample is mapped to a performance measure is unimportant." The NFL theorems allow for a wide latitude in the choice of performance measures. In particular, whereas the NFL theorems as such do not refer to any target of a search, the algorithms in question may be either target-oriented or not.

Here are examples of both targeted and non-targeted algorithms, equally subject to the NFL theorems.

Assume that the fitness function is simply the height of peaks in a specific mountainous region. If we choose a target-oriented algorithm, the target of the search can be defined as a specific peak *P* whose height is, say 6,000 meters above sea level. In this case the number *n* of iterations required to reach the predefined height of 6,000 meters may be chosen as the performance measure. Then algorithm a_1 performs better than algorithm a_2 if a_1 converges on the target in fewer steps than a_2. If two algorithms generated the *same sample* after *m* iterations, then they would have found the target—peak *P*—after the same number *n* of iterations. The first NFL theorem tells us that the *average probabilities* of reaching peak *P* in *m* steps are the same for any two algorithms. Any two algorithms will have an *equal average performance*, provided that the averaging is over all possible fitness landscapes (not all of which must in fact exist materially). Then, in the example above, the average number *n* of iterations required to locate the target is the same for any two algorithms if the averaging is done over all possible mountainous landscapes.

The NFL theorems do not say anything, however, about the relative performance of algorithms a_1 and a_2 on a *specific landscape*. On a specific landscape, either a_1 or a_2 may happen to be much better than its competitor.

Algorithms can also be compared in a targetless context. For example, rather than defining a *target* as a certain peak *P*, or even as a peak of a certain height, the algorithms may be compared by finding out which of them, a_1 or a_2, finds simply a *higher peak* after a certain number *m* of iterations. The performance measure in this case is the height of a peak reached after *m* iterations. No specific peak and no specific height is preselected as a target. An algorithm a_1 that after *m* iterations finds a *higher* peak than algorithm a_2 performs better. The first NFL theorem tells us that, if averaged over all *possible* mountainous reliefs (not all of them necessarily existing) the *probabilities* of both a_1 and a_2 generating the same sample after *m* iterations are equal. This also means that in all likelihood the height of a peak reached after *m* iterations, if *averaged over all possible landscapes*, will be the same for any two algorithms.

The NFL theorems are certainly valid for evolutionary algorithms. As Wolpert reports (2002), Wolpert and Macready (2003) have proven recently that the NFL theorems are invalid for coevolutionary algorithms, but this is a different question.

The No Free Lunch Theorems—A Little Mathematics

There are a series of NFL theorems, pertaining to various situations. The original NFL theorems (Wolpert 1996) dealt with problems of supervised learning. Later these theorems were extended to optimization problems associated with search algorithms (Wolpert and Macready 1997).

The first NFL theorem for search pertains to fixed (time-independent) fitness landscapes, while the second NFL theorem is for time-dependent landscapes. Though there is a substantial difference between the two, their principal meaning can be understood by reviewing only the first.

Imagine a finite set X called the *search space*, and a *fitness function* f that assigns a value to each point of X; the values of f are within a range denoted Y. Altogether the search points and their fitness values form the *fitness landscape*. We consider algorithms that explore X one point at a time. At each step, the algorithm decides which point to examine next, depending on the points that have been examined already and their fitness values, but it does not know the fitness of any other points. This decision might even be made at random or partly at random.

After an algorithm has iterated the search m times, we have a time-ordered set (a *sample*) denoted d_m^Y which comprises m measured values of f within the range Y. Let P be the conditional probability of having obtained a given sample d_m^Y after m iterations, for given f, Y, and m. Then, according to the first NFL theorem,

$$\sum_f P(d_m^Y \mid f, m, Y, a_1) = \sum_f P(d_m^Y \mid f, m, Y, a_2)\,, \qquad (1)$$

where a_1 and a_2 are two different algorithms. The summation is performed over *all possible fitness functions.* Equation (1) means that, in probabilistic terms, the results of a search, *if averaged over all possible fitness landscapes*, are the same for any pair of algorithms.

The equation for the second NFL theorem—for time-dependent landscapes—differs in two respects. First, it contains one more factor affecting the algorithm's behavior, an *evolution operator*, which is a rule reflecting how the landscape evolves from iteration to iteration. Second, the probabilities of obtaining a given sample d_m^Y are averaged over all possible evolution operators rather than over all possible fitness functions. For the purpose of this chapter, it is sufficient to refer to the first NFL theorem only.

Equation (1) also says that, if the performance of an algorithm a_1 is superior to that of another algorithm a_2 when applied to a specific class of fitness functions, it is necessarily inferior to the performance of a_2 on some other class of fitness functions. (In my earlier analogy the mountain climbing algorithm denoted COA performed better than HOA on one type of physical relief, but worse than HOA on another.)

Algorithms do not incorporate any prior knowledge of the properties of the fitness function. They are therefore called "black-box algorithms." They operate on a fitness landscape without prior knowledge of the landscape's relief, probing point after point, either deterministically or stochastically (according to a rule which has some random component).

Importantly, the NFL theorems do not say anything about the performance of any two algorithms on any *particular landscape.* If, in Equation (1) we remove the summation symbols, the sign of equality must be replaced with an inequality. In other words, except for rare special cases,

$$P(d_m^Y \mid f, m, Y, a_1) \neq P(d_m^Y \mid f, m, Y, a_2) \qquad (2)$$

Equation (2) means that, generally, the performance of any two arbitrarily chosen algorithms on a *specific* landscape cannot be expected to be equal.

The NFL theorems do not address a situation wherein certain algorithm a_1 *significantly outperforms* algorithm a_2 on a *few* landscapes while there are no such landscapes where a_2 is *much better* than a_1. According to the NFL theorems, in such cases a_2 outperforms a_1 on *many* landscapes but only *slightly*. This situation (a technical term for it is *head-to-head minimax asymmetry*) can be defined rigorously in quantitative terms (Wolpert and Macready 1997, 74).

As Wolpert and Macready point out (76), "… there is always a possibility of asymmetry between algorithms if one of them is stochastic." The asymmetry may be more significant than the equal average performance of algorithms established by the NFL theorems. This point is relevant to evolutionary algorithms, including Darwinian genetic algorithms and Dembski's analysis of them. Dembski (2002a, 189) states that "an evolutionary algorithm is a stochastic process." For stochastic algorithms the possibility of the "minimax" asymmetries is real, and when such asymmetries arise, they make the NFL theorems practically irrelevant.

Here is how Dembski defines the NFL theorems, "A generic NFL theorem now takes the following form: It sets up a performance measure M that characterizes how effectively an evolutionary algorithm E locates a target T within m steps using information j." According to Dembski, information j resides in an "information-resource space," which is beyond the "phase space" (his term for the search space) and usually exceeds the search space in size and complexity (200–203).

As follows from our preceding discussion, Dembski's definition misrepresents the NFL theorems. They do not "set performance measures" but only compare the generation of samples by algorithms within *m* iterations. Performance measures are introduced within the framework of corollaries to and interpretations of the NFL theorems and can be chosen in variety of ways. They should match the outcome of the search. Furthermore, there is no concept of a target (to which Demsbki offers no definition) in the NFL theorems as such; they are valid for targetless searches as well. (Each search is supposed to lead to some outcome but not necessarily to a *target*. An *outcome* is a general, qualitative concept; it may be either intended or not. It is not necessarily connected to the termination of a search; a search may be terminated for reasons unrelated to the outcome. A *target* is a specific, often quantitative concept, such as a predefined value of the fitness function that terminates the search when it is found). Moreover, there is no talk about information j in the parlance of the NFL theorems. These theorems are about black-box algorithms which start a search without a prior information about the fitness landscapes, but continue (and complete, if appropriate) the search using the information they extract gradually from the fitness function in the course of the search. This information is sufficient to continue a search at every step. Contrary to Dembski, the search algorithms do not need to go for information into a higher-order "information-resource space."

Continuing, Dembski writes (202), ". . . since blind search always constitutes a perfectly valid evolutionary algorithm, this means that the average performance of any evolutionary algorithm E is no better than blind search." This is correct, but the word *average* is crucial. Dembski forgets this word when he interprets the NFL theorems as making it impossible for evolutionary algorithms to outperform blind search on *specific* landscapes. To the contrary, the NFL theorems *do not assert* that no evolutionary algorithm performs better than a random sampling or a "blind search." Such a statement is valid only for the performance of algorithms if evaluated *on average* for all classes of problems. It is invalid when specific genetic pathways are considered.

As Wolpert and Macready emphasize, the NFL theorems do not predict performance in the real world. In fact, the uniform average is a crude tool designed to analyze the relationship between search algorithms and fitness functions.

Wolpert (2002) has pointed out that in Dembski's discourse the factors arising in the NFL theorems "are never specified in his analysis." Wolpert says also that "throughout Dembski's discourse there is a marked elision of the formal details of the biological processes under consideration. Perhaps the most glaring example of this is that neo-Darwinian evolution of ecosystems does not involve a set of genomes all searching the same, fixed fitness function, the situation considered by the NFL theorems. Rather it is a coevolutionary process. Roughly speaking, as each genome changes from one generation to the next, it modifies the surfaces that the other genomes are searching." And recent results (Wolpert and Macready 2003) indicate that "NFL results do not hold in coevolution."

The Displacement Problem

In his response to one of his critics, (Orr 2002) Dembski (2002b) says, "Given my title, it's not surprising that critics see my book *No Free Lunch* as depending crucially on the No Free Lunch theorems of Wolpert and Macready. But in fact, my key point concerns displacement, and the NFL theorems merely exemplify one instance (not the general case)."

In fact, though, Dembski introduces the "displacement problem" in the section on the NFL theorem (200–203) as a consequence of his interpretation of these theorems.

On page 202, Dembski says, "The significance of the NFL theorems is that an information-resource space J does not, and indeed cannot, privilege a target T." Dembski introduces two concepts here—a target and an "information-resource space J." In fact, the significance of the NFL theorems can hardly be seen in the quoted statement. As was discussed previously, the concept of a target as such is absent from the NFL theorems. They are equally valid for targeted and targetless searches. There is no talk about the "information-resource space" in the NFL theorems either.

On page 203 Dembski introduces the displacement problem: ". . . the problem of finding a given target has been displaced to the new problem of finding the information j capable of locating that target. Our original problem was finding a certain target within phase space. Our new problem is finding a certain j within the information-resource space J."

This quotation contains arbitrary assertions. First, the NFL theorems contain nothing about any arising "information-resource space." If Dembski wanted to introduce that concept within the framework of the NFL theorems, at the very least he should have shown what the role of an "information-resource space" is in view of the "black-box" nature of the algorithms in question. Second, the NFL theorems are indifferent to the presence or absence of a target in a search, which alone leaves Dembski's introduction of the "displacement problem," with its constant references to targets, hanging in the air.

The Irrelevance of the NFL Theorems

I submit that the real question is not whether or not the NFL theorems are valid for evolutionary algorithms (EA). Within the scope of their legitimate interpretation—when the conditions assumed for their derivation hold—the NFL theorems certainly apply to EA. The problem arises when they are applied where the assumed premises do not hold. Although Wolpert and Macready (2003) have shown that in the case of coevolution the NFL theorems may not hold, my conclusion would not change even if the NFL theorems were also valid for coevolution.

The simple fact is that the NFL theorems are *irrelevant* for the real question we face. The real (two-tiered) question is:

(a) Can an evolutionary algorithm outperform random sampling (or "blind search") in situations of interest?

(b) Can "specified complexity" be "purchased without intelligence?"

Let us discuss both (a) and (b). Dembski's answer to (a) is a categorical No. He is wrong. The correct answer is a categorical *Yes*. Let me show why.

Dembski's *No* to question (a) is partially based on the alleged mathematical certainty expressed by the NFL theorems according to which no algorithm performs better than a random search. Indeed, on page 196 we read, "The No Free Lunch theorems dash any hope of generating specified complexity via evolutionary algorithms." On page 212, "The No Free Lunch theorems show that evolutionary algorithms, apart from careful fine-tuning by a programmer, are no better than blind search and thus no better than pure chance."

What Dembski seems to ignore is the crucial point which I have stressed several times: the NFL theorems legitimately compare the performance of any two algorithms, but what they compare is performance *averaged over all possible fitness landscapes*. This is an interesting theoretical conclusion and a tool for investigating the mutual relationship between the fitness functions and search algorithms. It has no relevance for problems of practical interest encountered in real life. Practically, we are interested in finding out whether or not a given algorithm outperforms a random search *if applied to a specific class of fitness landscapes.*

There are plenty of examples showing that evolutionary algorithms indeed outperform random search when applied to fitness functions of interest. Let us recall some of them.

In Dawkins's example, as Dembski tells us himself, a random search is expected to converge on the target phrase after about 10^{40} iterations. If, though, Dawkins's evolutionary algorithm is applied to the same task, it achieves the same result after about only forty iterations. Even if Dawkins's algorithm is replaced by Dembski's version, it will reach the target, as Dembski says, after about 4,000 iterations. 4,000 *vs* 10^{40}—this *is* outperformance, and a very respectable outperformance indeed. What significance has the fact that the algorithms cannot outperform a random search if averaged over all possible fitness functions? *They outperform a random search if applied to the specific fitness functions of interest, and that is all that counts.*

In the cases of a search for an optimal shape of an antenna (Altshuler and Linden, 1999) or of a checkers playing algorithm (Chellapella and Fogel, 1999) both of which Dembski views (p. 221) more favorably than Dawkins's algorithm, again the evolutionary algorithms immensely outperform a random search. Although the NFL theorems are valid for Altshuler and Linden's and for Chellapella and Fogel's algorithms, this fact is of no consequence because what those authors are interested in is not the averaged performance over all possible fitness functions but the performance on a specific class of fitness functions, and the NFL theorems say nothing about such performance.

Rather than Equation (1), in practical situations inequality (2) is really relevant, and it says that different algorithms perform differently on specific classes of fitness functions. Hence Dembki's discussion of NFL theorems is of no consequence for the question of whether evolutionary algorithms can outperform a random search. They can and they do.

Of course, Dembski has an escape clause: He admits that evolutionary algorithms can outperform a random sampling if there is a "careful fine-tuning by a programmer." If , though, a programmer can design an evolutionary algorithm which is fine-tuned to ascend certain fitness landscapes, what can prohibit a naturally arising evolutionary algorithm to fit in with the kinds of landscape it faces? Nothing can and nothing does. If a specific evolutionary algorithm, either fine-tuned by a programmer or arising naturally, outperforms a random sampling on a specific landscape, the NFL theorems are of no consequence and Dembski's reference to these theorems is irrelevant.

This thesis can be illustrated as follows: Naturally arising fitness landscapes will frequently have a central peak topping relatively smooth slopes. If a certain property of an organism, such as its size, affects the organism's survivability, then there must be a single value of the size most favorable to the organism's fitness. If the organism is either too small or too large, its survival is at risk (Haldane 1928:20-28). If there is an optimal size that ensures the highest fitness, then the relevant fitness landscape must contain a single peak of the highest fitness surrounded by relatively smooth slopes.

The graphs in figure 11.1 (see http://www.talkreason.org/articles/fig11.cfm) schematically illustrate my thesis. The fitness function may be, for example, the average life expectancy of an animal, the average number of its surviving descendants, or some other single-valued quantity that reflects an animal's success at survival. The fitness function is represented by the solid curve, which has a well-defined peak corresponding to the optimal size, with more or less smooth slopes on both sides of the peak. We may imagine many other possible fitness functions, such as the rugged fitness function represented by the dashed curve. Such fitness functions, however, do not represent biological reality: the survivability of an animal cannot depend on its size (or on some other feature) in such a haphazard manner. It is unlikely that several different sizes will be comparably advantageous to an organism's fitness.

Obviously, in this case, the evolutionary algorithm based on natural selection is well suited to ascending the actual fitness landscape. Indeed, the closer is the organism's size to the maximum of the fitness landscape, the more it is favored by natural selection. In this case the NFL theorems, while they are correct (if we ignore coevolution), are irrelevant. Natural selection will perform well on the actual landscape, certainly better than a random search. Other algorithms can perform better on other possible fitness landscapes, such as the rugged landscape exemplified in figure 11.1, but landscapes actually encountered in the biosphere are not likely to be so rugged. (This example assumes a one-dimensional fitness function, whereas a real fitness function is multidimensional; the example is intended only to illustrate the point.)

The debate can be extended by posing the following question, as Stuart Kauffman (2000, 2003) has done: If Darwinian evolution has been indeed taking place, then obviously Darwinian evolutionary algorithms work well on the fitness landscapes which arise naturally in the biosphere. Then, according to the NFL theorems, these algorithms must perform poorly on some other possible fitness landscapes. In other words, while the natural evolutionary algorithms entailing random mutations and natural selection (plus recombination and possibly other mechanisms) do indeed outperform a random search, they should underperform a random search on different fitness landscapes which could have existed in some *alternative reality*. Then the question is: Why of all the enormous variety of possible fitness landscapes are the fitness landscapes actually observed in the biosphere exactly such that they are accessible to Darwinian evolutionary algorithms?

The answer is that because of the enormous variety of possible evolutionary algorithms and fitness functions, the probability of some fraction of algorithms being naturally "fine-tuned" to the existing landscapes is close to certainty. The above example with the fitness landscape entailing organism's optimal size may serve as an illustration of this statement.

Anyway, posing Kauffman's question shifts the discussion from the *relevance of the NFL theorems for the observed biological reality* to the realm of *anthropic coincidences*. Whatever the explanation of the anthropic coincidences may be (Drange 1998; Stenger 2002a, b; Ikeda and Jefferys 2000; Perakh 2001b), it would not alter the conclusion that the NFL theorems are irrelevant to the question of evolutionary algorithms performance as compared with a random search, *as long as we discuss existing biological reality.*

Dembski's answer to question (b), "Can specified complexity be purchased without intelligence?" is also a categorical *No*.

To my mind *Yes* is more plausible. I will now explain why.

The necessity of intelligence for generating SC is something Dembski ostensibly *sets out to prove*; in fact he instead often *uses it as a given*. He provides no evidence which would meet the requirements of scientific rigor to substantiate his thesis, but only arbitrary assumptions lacking evidence. When he attempts to apply more specific arguments, such as those based on the NFL theorems and its alleged implication in the form of the "displacement problem," his discourse is contradictory and inconsistent. On the other hand, evolutionary biologists have suggested plausible scenarios explaining how evolutionary algorithms can work without interference by an external intelligence.

ID advocates often charge that such scenarios are "just-so stories" and therefore are not convincing. First, there is often substantial empirical evidence in favor of these scenarios, and the biological literature abounds in them. Second, such a reproach sounds odd coming from ID advocates whose entire conceptual system is a "just-so story," where a blanket reference to "intelligent design" is nothing more than a *Dembskian Z-factor* (see Chapter 13) offered as a substitute for a realistic scenario.

The Displacement "Problem"

Now I will discuss some of Dembski's *specific* arguments in favor of his assertion that CSI must be necessarily "smuggled" or "front-loaded" into evolutionary algorithms.

Dembski (2002e) asserts that "the displacement problem" is in fact the core of his thesis. However, at a close inspection it becomes clear that the "displacement problem" is irrelevant to real-life situations.

Recall that Dembski defines the problem (2002a, 203) as, "…the problem of finding a given target has been displaced to the new problem of finding the information j capable of locating that target. Our original problem was finding a certain target within phase space. Our new problem is finding a certain j within the information-resource space J." As Dembski explains (p. 202) "the fitness function is of course the additional information that turns the blind search to a constrained search." Hence, the "information-resource space" J is meant by Dembski as a space of (possibly along with other sources of information) all possible fitness functions.

According to Dembski (203) the "information-resource space J" is "in practice . . . much bigger and much less tractable than the original phase space"; hence the original problem has been "displaced" to a much more intractable problem. To solve the new problem, insists Dembski, the specified complexity must be injected by intelligence. In summary, Dembski's "displacement problem" means that the space of all possible fitness functions has to be searched to determine the fitness function for the problem at hand.

Dembski provides no reason to assume that the "information-resource space" is much larger and much less tractable than the original "phase space." In fact, there seem to be no such reasons. The "information-resource space" can be larger, of about the same size, or smaller than the "phase space." Dembski provides an example (204) of a search for a treasure buried on an island. Instead of searching all over the island (whose topography constitutes the "phase space") the search may be "displaced" to a search all over the world for a map of the island wherein the location of the treasure is indicated. Now the "information-resource space" is the entire globe, which is immensely larger than the island in question. This example can easily be reversed since it could happen as well that finding the map in question is much easier than finding the treasure itself without a map. Indeed, if it is known that the map is hidden in a certain building in a certain city, the "information-resource space" becomes the specific building and will be much smaller and much more tractable than the original phase space (which was the entire island). But, regardless of which space is larger and less tractable, and regardless of the very existence or absence of the "displacement problem," it is irrelevant for a real-life optimization search. Here is why.

To start a search, an algorithm needs no information about the fitness function. That is how the "black-box" algorithms start a search. To continue the search, an algorithm needs information from the fitness function. However, no search of the space of all possible fitness function is needed. In the course of a search, the algorithm extracts the necessary information from the landscape it is exploring. The fitness landscape is always given, and automatically supplies sufficient information to continue and to complete the search.

Consider Dawkins's weasel algorithm. It explores the available phrases and selects from them using the comparison of the intermediate phrases with the target. The fitness function has in this case built-in information necessary to perform the comparison. This fitness function is *given* to the search algorithm; to provide this information to the algorithm, no search of a space of all possible fitness functions is needed and therefore is not performed.

The same is true for natural evolutionary algorithms. The evolutionary algorithms, both designed by intelligence and occurring spontaneously, deal with *given* specific fitness functions and have no need to search the "information-resource space." Dembski's "displacement problem" is a phantom.

Conclusion

Dembski (1998 d) has written, "As Christians we know naturalism is false." Obviously, if one "knows" something, this ends a discussion.

Since 1998 Dembski's attitude does not seem to have changed. Recently (2002 c) he asserted that the ID advocates would never "capitulate" to their detractors. If so, then it is a testimony to the fact noted by the critics of the ID "theory": ID is not science. Scientists normally admit that whatever theories are commonly accepted at any time, there is always a chance they may be overturned under the weight of new evidence. No genuine scientist would ever claim that he would never "capitulate" no matter what.

The points listed below encapsulate the gist of this chapter.

> Dembski's critique of Dawkins's "targeted" evolutionary algorithm fails to repudiate the illustrative value of Dawkins's example, which demonstrates how supplementing random changes with a suitable law increases the rate of evolution by many orders of magnitude.
>
> Dembski ignores Dawkins's "targetless" evolutionary algorithm, which successfully illustrates spontaneous increase of complexity in an evolutionary process.
>
> Contrary to Dembski's assertions, evolutionary algorithms routinely outperform a random search.
>
> Contrary to Dembski assertion, the NFL theorems do not make Darwinian evolution impossible. Dembski's attempt to invoke the NFL theorems to prove otherwise ignores the fact that these theorems assert the equal performance of all algorithms only if averaged over all fitness functions.
>
> Dembski's constant references to targets when he discusses optimization searches are based on his misinterpretation of the NFL theorems, which entail no concept of a target. Moreover, his discourse is irrelevant to Darwinian evolution, which is a targetless process.

The so-called displacement problem, touted by Dembski as the core of his thesis, is a phantom because evolutionary algorithms face *given specific* fitness landscapes; the landscape supplies sufficient information to continue and complete (when appropriate) a search; there is no need to search the higher-order “information-resource” space.

The question, “Why the evolutionary algorithms actually observed in the biosphere are well adjusted to the actually observed fitness functions?” belongs in the general discussion of anthropic coincidences. The arguments showing that the anthropic coincidences do not require the hypothesis of a supernatural intelligence also answer the questions about the compatibility of fitness functions and evolutionary algorithms.

Acknowledgements.

I had the privilege of obtaining advice from David H. Wolpert, the coauthor of the no free lunch theorems, who shared with me ideas pertaining both to the essence of the NFL theorems and to Dembski’s treatment of them. I also appreciate comments by Taner Edis, Gordon Elliott, Thomas D. Schneider, Jeffrey Shallit, Erik Tellgren, Matt Young, and especially Brendan McKay. Of course, the opinions and arguments in this chapter, and even more so possible errors, are all mine.

References

Altshuler, Edward E. and Derek S. Linden. 1999. “Design of Wire Antennas Using Genetic Algorithms.” *Electromagnetic Optimization by Genetic Algorithms*. Eds. Y. Rahmat-Samii and E. Michielssen. (New York: Wiley) 211–248.

Chellapilla, Kumar and David B. Fogel. 1999. "Co-Evolving Checkers Playing Programs Using Only Win, Lose or Draw." *SPIE's AeroSence '99: Applications and Science of Computational Intelligence II* (Orlando, Fla.: 5–9 April 1999).

Dawkins, Richard. 1996 [1986]. *The Blind Watchmaker. Why the Evidence Of Evolution Reveals a Universe without Design*. (NewYork: W.W. Norton & Company).

Dembski, William A. 1998d. "Introduction: Mere Creation." in *Mere Creation.* W. Dembski, ed. (Downers Grove, Illinois: InterVarsity press) p.14.

Dembski, Wlliam A. 1998 a. *The Design Inference. Eliminating Chance Through Small Probabilities*. (New York: Cambridge University Press).

Dembski, William A. 1999. *Intelligent Design. The Bridge Between Science and Theology.* (Downers Grove, Illinois: InterVarsity Press).

Dembski, William A. 2002a. *No Free Lunch. Why Specified Complexity Cannot Be Purchased Without Intelligence.* (Lanham, Maryland: Rowman and Littlefield Publishers).

Dembski, William A. 2002b. "Evolution's Logic of Credulity: An Unfettered Response to Allen Orr." www.designinference.com/documents/2002.12.Unfettered_Resp_to_Orr.htm . Accessed on December 26, 2002.

Dembski, William A. 2002c. (No title). www.arn.org/cgi-bin/ubb/ultimatebb.cgi?ubb=get_topic;f=13;t=000483 , posted 11/26/2002, accessed on December 26, 2002.

Drange, Theodore M. 1998. “The Fine-Tuning Argument.” www.infidels.org/library/modern/theodore_drange/tuning.html , accessed on December 26, 2002.

Ikeda, Michael and William Jefferys. 2001. “The Anthropic Principle Does Not Support Supernaturalism.” www.talkreason.org/articles/super.cfm , accessed on December 26, 2002.

Kauffman, Stuart A. 2000. *Investigations.* (New York: Oxford University Press).

Orr, H. Allen. 2002. “Review of William Dembski’s No Free Lunch.” *Boston Review.* Summer issue. See also http://bostonreview.mit.edu/BR27.3/orr.html , accessed on December 22,2002.

Perakh, Mark. 2001 a. “A consistent Inconsistency. How Dr. Dembski Infers Design.” www.talkreason.org/articles/dembski.cfm , accessed on December 26, 2002.

Perakh, Mark. 2001 b. “Anthropic Principles—Reasonable and Unreasonable.” www.talkreason.org/articles/anthropic.cfm , accessed on December 26, 2002.

Perakh, Mark. 2002a. “A Free Lunch In a Mousetrap.” www.talkreason.org/articles/dem_nfl.cfm , accessed on December 26, 2002.

Perakh, Mark. 2002b. “A Presentation Without Arguments.” *Skeptical Inquirer*, v 26, no 6, 2002:31. Also at www.talkreason.org/articles/presentation.cfm , accessed on December 26, 2002.

Perakh, Mark. 2002c. “Science in the Eyes of a Scientist.” www.talkreason.org/articles/good_bad_science.cfm , accessed on December 26, 2002.

Perakh, Mark. 2003. *Unintelligent Design*. (Amherst, NY: Prometheus Books), chapt.1.

Rosenhouse, Jason. 2002. “Probability, Optimization Theory, and Evolution.” *Evolution.* v.56, No 8, 1721. Also http://evol.allenpress.com/evolonline/?request=get.document&issn=0014-3820&volume=056&issue=08&page1721 , accessed on December 26, 2002.

Schneider, Thomas D. 2000. “Evolution of Biological Information.” *Nucleic Acids Research*, 28 (14) 2794–2199.

Schneider, Thomas D. 2001a. “Rebuttal to William A. Dembski Posting And To His Book ‘No Free Lunch.’” www.lecb.ncifcrf.gov/~toms/paper/ev/dembski/rebuttal.html . Accessed on December 30, 2002.

Schneider, Thomas D. 2001 b. “Effect of Ties on the Evolution of Information by the Ev Program.” www.lecb.ncifcrf.gov/~toms/paper/ev/dembski/claimtest.html . Accessed on December 30, 2002.

Shallit, Jeffrey. 2002. "Review of Dembski's NFL book." http://www.math.uwaterloo.ca/~shallit/nflr3.txt

Stenger, Victor J. 2001 a. "Anthropic Design: Does the Cosmos Show the Evidence of Purpose?" www.talkreason.org/articles/coincidence.cfm , accessed on December 26, 2002.

Stenger, Victor J. 2001 b. "Natural Explanation for the Anthropic Coincidences." www.talkreason.org/articles/anthro-philo.pdf , accessed on December 26, 2002,

Van Till, Howard. 2002. "E Coli at the No Free Lunchroom. Bacterial Flagella and Dembski's Case for Intelligent Design." , www.aaas.org/spp/evolution/perspectives/vantillecoli.pdf . Accessed on January 16, 2003.

Wein, Richard. 2002. a. "Not a Free Lunch but a Box of Chocolate." (www.talkreason.org/articles/chock_nfl.cfm , accessed on December 26, 2002.

Wein, Richard. 2002b. "Response? What Response? How Dembski Avoided Addressing My Arguments." www.talkreason.org/articles/response.cfm , accessed on December 26, 2002.

Wolpert, David H. 1996 a. "The Lack of A Priori Distinctions Between Learning Algorithms." *Neural Comp*., v. 8, no 7, 1341–1390.

Wolpert, David H. 1996 b. "The Existence Of A Priori Distinctions Between Learning Algorithms." *Neural Comp*., v. 8, no 2, 1391–1420.

Wolpert, David H. 2002. “William Dembski’s Treatment of the No Free Lunch Theorems is Written in Jello.” www.talkreason.org/articles/jello.cfm , accessed on December 26, 2002.

Wolpert, David H and William G. Macready. 1997. “The No Free Lunch Theorems For Optimization.” *IEEE Trans. Evol. Comp*. v.1, no 1, 67–82.

Wolpert, David H and William G. Macready. 2003 (to be published).

Young, Matt. 2002. “Grand Designs and Facile Analogies.” www.talkreason.org/articles/Facile.pdf , accessed on December 26, 2002.