

Intelligent Design and the NFL Theorems: Debunking Dembski

Olle Häggström*

September 8, 2005

Abstract

Another look is taken at the model assumptions involved in William Dembski's (2002a) use of the NFL theorems from optimization theory to disprove the Darwinian theory of evolution by natural selection, and his argument is shown to lack any relevance whatsoever to evolutionary biology.

1 Mathematics for good and for bad

Science, from Newton and onwards, owes much of its success to mathematics. Physics in particular is so permeated by mathematical modelling and methodology so as to become almost inconceivable without it, a situation that has even triggered some thinkers to raise the metaphysical question of why nature is so amenable to mathematical study (Wigner, 1960; Omnès, 2005). Biology is currently undergoing a mathematization which may or may not turn out as far reaching as that of physics. Looking further afield, we may note the prevalence of mathematical formalism in the social sciences (particularly in economics) where, however, its success has so far been far less overwhelming than in the natural sciences.

Unfortunately, not all uses of mathematics in other disciplines are good, and some of them are outright bad. Overall, most researchers have little or no university-level mathematical training, and are easily intimidated by those colleagues who master – or pretend to master – the language of mathematics. Members of the latter group often manage to gain acceptance for their poor

*Professor of mathematical statistics, Chalmers University of Technology, Sweden, <http://www.math.chalmers.se/~olleh/>

mathematical applications, unchallenged by a research community that is unable to penetrate their formulae.

This leads, of course, to bad research. Besides that, it has had – if I may speculate a little – another undesirable consequence: it seems likely that the extensive (ab)use of mathematics in various fields has contributed to an inferiority complex that triggered poststructuralists and postmodernists to invade much of humanities and social sciences with a jargon even more impenetrable than that of mathematics, but to a catastrophic extent lacking in meaningful content, as exposed in the Sokal Hoax (Sokal, 1996; Sokal and Bricmont, 1998).

In this paper, I will draw attention to a particularly bad application of mathematics in a context of much current interest in American politics, namely the brand of anti-Darwinian criticism known as Intelligent Design. Specifically, what I have in mind is William Dembski’s (2002a) use of the so-called NFL (No Free Lunch) theorems in optimization theory to “refute” evolution by natural selection.¹ After giving some background on the Intelligent Design movement and the NFL theorems, respectively, in Sections 2 and 3, I will outline Dembski’s use of the latter in Section 4. Then, in Sections 5, 6 and 7, I will demonstrate the error in Dembski’s argument, followed in Section 8 by some remarks on possible extensions.

Dembski’s *No Free Lunch* (2002a) has been amply debunked elsewhere; see, e.g., Orr (2002), Shallit (2002) and especially Richard Wein’s (2002a)² devastating catalogue of errors and pseudoscientific reasoning on Dembski’s part. The role of the present paper is to offer a mathematician’s account of what the NFL theorems actually say, and why Dembski’s use of them is so utterly wrong. The central argument in Section 6 concerning fitness landscapes is essentially the same as one put forth by Wein (2002a), although here I will consider Dembski’s implicit model assumption about these in more detail. Along the way, it will become clear that the NFL theorems constitute a correct but banal mathematical observation whose potential for saying anything interesting about evolutionary biology is, realistically speaking, zero.

I feel that I must warn the reader from the outset that no deep results whatsoever will be conveyed in this paper. In fact, much to the contrary, facts and results that have been described as profound by others will be

¹More recently, Dembski (2005) provides an additional smokescreen of fancy-looking but irrelevant mathematics that he claims supports his NFL-based argument against Darwinism. I will refrain from discussing that paper here, as it makes no difference to the arguments of the present paper or to the final verdict on Dembski’s approach.

²See also the subsequent exchange in Dembski (2002b, 2002c) and Wein (2002b).

shown to be fairly trivial.

2 Intelligent Design

Christian fundamentalists have long since given up their fight against heliocentrism. In contrast, Darwin's theory of natural selection is still under attack. In the United States, anti-Darwinians have gained considerable political success, with many local and state school boards having invoked regulations to the extent that Darwin's theory must be taught as just one of several theories about the origin of species – a policy that has even been endorsed by President George W. Bush.³

Opponents of Darwinism have, however, found that Biblical literalism and young-earth creationism does not always work well in courtrooms and in rational discussion. Therefore, these crude versions of creationism have in recent years to a large extent been replaced by *Intelligent Design* as the proposed main alternative to Darwinism. Proponents of Intelligent Design accept much of modern biology and natural history, insisting only that complex creatures such as ourselves cannot come about “bottom-up” in a universe governed just by natural laws, but bear unmistakable signs of being the work of an intelligent agent. As a kind of tactical retreat, they generally keep a low profile concerning the identity of this agent – God, extraterrestrial aliens, or something else.⁴

Here, I will only give a very brief recollection of some of the main features of Intelligent Design. For an excellent critical survey, see Crews (2001), or see Orr (2005) for a more recent discussion.

Intelligent Design is being promoted as a science, but as such it has some obvious difficulties. For instance, the simple question “who designed the designer?” shows that Intelligent Design cannot hope to solve the problem of emergence of complexity and life, but only move it one step away. And those of us who take seriously Popper's idea of the central role played by falsification may ask: how in the world could one ever *falsify* Intelligent Design? I will ignore these difficulties and instead focus on key arguments put forth by the two leading thinkers of the Intelligent Design movement: Michael Behe and William Dembski.

³See Dawkins and Coyne (2005) for an eloquent comment on Bush's endorsement and what is at stake.

⁴Or could it be that we live in somebody else's computer simulation, à la *The Matrix*? This mind-blowing scenario is not only a recurrent theme in fiction – it has also been defended seriously; see, e.g., Bostrom (2003).

Behe's *Darwin's Black Box* (1996) is by far the most well-known contribution to Intelligent Design. The central concept in Behe's book is that of *irreducible complexity*, neatly illustrated by a mousetrap consisting, say, of five parts (platform, spring, hammer, catch, and hold-down bar). Removing any one of these parts will not result in a device that is slightly worse at catching mice compared to a complete mousetrap, but rather something that does not catch mice at all. Likewise, Behe argues, many biological systems have the same property of indispensability of each of its components. And how could such a system have evolved along a Darwinian path of gradual improvement? This seems impossible, because up until the last of the components is formed the others are worthless, removing all selection pressure in the direction of forming the system.

The irreducible complexity argument goes back, of course, further than to Behe: classical cases that have been brought up (and accounted for by Darwinians) are the evolutions of eyes and of wings. Behe's main novelty is to focus not on the organic level but rather on the microscopic machinery inside cells.

Behe's work has, however, been sharply criticized – see, e.g., Miller (1996), Dennett (1997), and Orr (2002) – and it is clear from these critiques that Behe takes insufficient account of a variety of mechanisms including exaptation (i.e., when a system first evolves for some particular function but is later exploited by the organism for a different function) and gene duplication. His case is therefore not nearly as strong as he claims it to be in *Darwin's Black Box*.

The so-called “argument from design” is old and straightforward, but very persuasive: *Of course* the blind forces of nature cannot produce complex creatures such as ourselves – some kind of divine creator *has to* be involved. The argument suffers, however, from an obvious lack of precision. In his book *The Design Inference* (1998), Dembski sets out to remedy this (at least partially) by making precise the meaning of “complex” through his notion of *specified complexity*. It is, however, very difficult for the reader of Dembski's work to understand precisely what specified complexity means.⁵ In any case, evolutionary biologists have – using theoretical arguments supplemented by computer simulations – amply demonstrated that the blind forces of nature can in fact, via the mechanism of natural selection, produce objects that exhibit any reasonable biological notion of complexity. For Dembski and others who find themselves unable to accept this conclusion, I really have no

⁵A discussion of how inconsistently Dembski uses his own concept can be found in Wein (2002a).

better advice than to direct them to the modern classics in this field, such as Dawkins (1986, 1989) and Dennett (1995).

Of course, Dembski still has the right to stick to his view, and in his follow-up book *No Free Lunch* (2002a) he claims to give mathematical justification for “why specified complexity cannot be purchased without intelligence” by invoking the NFL theorems. These will be the topic of the next section.

3 Optimization and the NFL theorems

In combinatorial optimization, one is given a finite set V and a function $f : V \rightarrow \mathbf{R}$ which to each $x \in V$ assigns a real number.⁶ The task is to find an element $x \in V$ that maximizes $f(x)$. At first sight, this may seem like a trivial task: since V is finite, all we need to do is simply to go through all $x \in V$ systematically, calculate $f(x)$ for each of them, while keeping track of the maximum seen so far.

The reason why this “brute force” approach does not suffice is that V is usually so large that time constraints make it infeasible. Typically, the number of elements of V grows exponentially (or faster) in some parameter n that describes the size of the problem in some natural way. For instance, V could be the set of binary strings of length n , or it could be the number of permutations of n objects; this gives 2^n resp. $n!$ elements in V , in both cases making the brute force method out of the question even for moderately sized problems such as $n = 100$.

Other, less time-consuming, algorithms are therefore needed. A common approach involves so-called *local search* in V . This necessitates the introduction of some “geographic” structure in V , which can be accomplished by declaring the existence of *links* between some (but not all) pairs of elements $x, y \in V$. The set of all y that are linked to a given $x \in V$ is called the *neighborhood* of x . There is much freedom in setting up the links, but it needs to be done in such a way that, on one hand, each x has a neighborhood of manageable size, and, on the other hand, the network of links becomes “well connected” (in some sense). In specific examples, natural link structures often more or less suggest themselves: when V is the set of length- n binary strings, we may declare links precisely between those $x, y \in V$ that differ only in one bit, or when V is the set of permutations of n objects we may decide to declare a link between two permutations precisely when one of them can arise from the other by interchange of just two of the objects.

⁶We write, following convention, \mathbf{R} for the set of all real numbers.

Given the link structure, the basic local search algorithm proceeds as follows. Start at some arbitrary $x \in V$, compute f at x and at all of its neighbors, and move to the neighbor y whose f -value is the largest (unless they are all smaller than $f(x)$ in which case we stay at x). Then repeat the process, moving to the vertex z that has the largest f -value among y and *its* neighbors. This goes on until we get stuck.

This algorithm is sometimes called the *hill-climber*, as it can be pictured as a hiker in a hilly landscape, always going in the direction of the steepest climb, until the top of a hill is reached. Such hill-climbing sometimes works well, but a huge drawback is that the algorithm may get stuck on a relatively modest hill without noticing the huge mountain peak further away.

To deal with this drawback, a variety of modifications of the hill-climber algorithm have been proposed and are widely used; see, e.g., Aarts and Lenstra (1997). These modifications may for instance include randomizing the walk in such a way as to allow occasional downhill steps (as in the famous simulated annealing algorithm) or permitting occasional “long jumps” in the landscape. Many of these are quite sophisticated.

These algorithms are not only used for the pure optimization problem that we have focused on so far, but also – in fact more often – for the purpose of locating some large (but not necessarily the largest) value of f . Specifically, the goal may be to find and $x \in V$ such that $f(x)$ exceeds some given level t . The algorithm then proceeds until it encounters an element of the set T consisting of all $x \in V$ satisfying $f(x) \geq t$. The problem of finding some $x \in T$ should really be called a *search problem* rather than an optimization problem. We call T the *target set*, and it can be written in compact mathematical notation as

$$T = \{x \in V : f(x) \geq t\}. \tag{1}$$

More generally, we may not always be in a situation where “the larger value of f , the better”, so it makes sense to allow for a target set T that is not necessarily of the form (1), but may be an arbitrary subset of V . In interesting search problems, T is typically very rare, in the sense that only a very small fraction of all elements $x \in V$ are also in T .

This sets the stage for the NFL theorems of Wolpert and Macready (1997), who showed that for these optimization and search problems, no algorithm is better than any other, in a certain average sense. This may sound very surprising, so let me describe in more detail what the basic NFL theorem actually says.⁷

⁷Most of the discussion will focus on this particular NFL theorem, but see Section 8

Wolpert and Macready restrict to the setting where the function f is only allowed to take values in some prescribed finite subset S of \mathbf{R} . This is natural because in a computer implementation everything is necessarily discrete.

Given V and S , how many different ways are there to define a function $f : V \rightarrow S$? Writing $|V|$ and $|S|$ for the number of elements of V and S , respectively, there are for any particular $x \in V$ precisely $|S|$ possible choices of $f(x)$. Multiplying over all elements of V tells us that there are $|S|^{|V|}$ different choices of functions $f : V \rightarrow S$ (usually a stupendously large numbers, since already $|V|$ is typically very large). The basic NFL theorem concerns an average over all these functions.

The algorithms considered by Wolpert and Macready are of the following form. First, an $x_{(1)} \in V$ is chosen according to some rule (which, like those that follow, may or may not involve the use of random numbers), and $f(x_{(1)})$ is computed. Then $x_{(2)} \in V$ is chosen according to some rule that may take into account $x_{(1)}$ and $f(x_{(1)})$, after which $f(x_{(2)})$ is computed. And so on: given $x_{(1)}, \dots, x_{(k)}$ and $f(x_{(1)}), \dots, f(x_{(k)})$, the algorithm chooses an $x_{(k+1)}$ using a rule that may take into account all these previous values. The only other proviso that the basic NFL theorem requires is that no $x \in V$ is chosen more than once.

Imagine now that the first k f -values $f(x_{(1)}), \dots, f(x_{(k)})$ have been recorded, and define some event E_k solely in terms of these; the prototype example is to take E_k to be the event that at least one of the recorded values $f(x_{(1)}), \dots, f(x_{(k)})$ puts its corresponding $x_{(i)}$ in the target set T . The basic NFL theorem now states that

averaged over all the $|S|^{|V|}$ different possible function f , the probability of the event E_k is the same for any choice of algorithm.

Among other things, this tells us that no algorithm is better at quickly finding some x in the target set T than any other. In particular, no algorithm is better than the “blind search” algorithm that does the following: first pick $x_{(1)}$ uniformly at random from V (i.e., any element of V has the same probability $1/|V|$ of being chosen), then $x_{(2)}$ is chosen uniformly at random among the others (regardless of $f(x_{(1)})$), and so on. If, as usual, V is a very big set and the target set T is very rare, then the time taken to find some $x \in T$ will most likely be copiously large.

Thus, the basic NFL theorem seems to provide us with a disheartening message: no matter how clever we are, we cannot expect to devise algorithms

for some indication of why the plural form “theorems” is used above.

that are better than the hopelessly primitive and inefficient blind search algorithm.

In practice, however, there is no reason to despair. The key property of the basic NFL theorem that allow us to circumvent its dark message, is the averaging over all possible functions f that is involved. In almost all concrete optimization problems, we have some prior information or at least some rough idea of how f varies across V , and such information can be exploited in the construction of clever and efficient optimization algorithms, unfettered by any NFL theorem. The reason why the pessimistic message of the basic NFL theorem no longer applies in such a situation is that it averages over *all* possible f , and not just over the kinds of f that we know to be more likely.

The moral of Wolpert and Macready (1997) is, thus, that we cannot expect to construct efficient optimization or search algorithms *unless we exploit some prior knowledge of f* .⁸ Further light on their result will be shed in Section 5, but before that, I will explain how NFL is claimed to disprove Darwinian evolution.

4 Dembski's application to evolution

What, the reader may now ask, could possibly be the relevance of the theory of algorithms to evolutionary biology? Quite a lot, in fact. In the wake of the computer revolution, various search algorithms, in the spirit of those discussed in the previous section, began to flourish across the scientific literature (and still do). Some researchers turned to evolutionary biology for inspiration, and devised turbo-versions of search algorithms based on the principles of reproduction, mutation, and selection. The analogies between on one hand these algorithms, and on the other hand the Darwinian mechanisms in biology, stared researchers in the eyes with increasing intensity, and it became clear that viewing and analyzing biological evolution as an algorithm (albeit one that heavily involves randomization and parallelism – two features that computer scientists nowadays are quite good at dealing with) could be fruitful. See, e.g., Dennett (1995), who take this perspective in a very consistent manner.

The algorithmic view on Darwinian evolution is also taken up by Dembski (2002a) in his attempt to refute it. In this section, I will describe his NFL-based argument in the case of a single species evolving in a fixed environment. I will thus ignore for the moment the complications of time-dependent en-

⁸It is this observation that prompted them to use the phrase No Free Lunch.

vironments or of several species coevolving. Dembski's argument, as well as my refutation of it, extend in a straightforward manner to these situations; see Section 8 for some brief remarks in this direction.

As a preparatory lemma to his main argument, Dembski notes that the kind of blind search that was described in the previous section cannot possibly account for the occurrence of what he calls specified complexity, such as ourselves or other large animals and plants. This is absolutely correct. The human genome is about 3 000 000 000 base pairs long. Let us now take V to consist of all DNA sequences up to that length, and the target set T to be the set of all such DNA sequences giving rise to a creature exhibiting specified complexity. The number of elements of V then becomes something of the order $10^{1\ 800\ 000\ 000}$ – a truly Vast number. (Following Dennett (1995), I write Vast for “Very much larger than ASTronomical”.) The target set T is also Vast, but a more important observation is that T so much smaller than V that if we pick an element at random (uniform distribution) from V , then the odds against getting an element of T are also Vast. The precise Vast-ness of this quantity is very difficult to estimate (partly because of the difficulty in pinpointing exactly what specified complexity is), but it seems reasonably safe to state that $|V|/|T|$ is somewhere between 10^{1000} and $10^{1\ 000\ 000\ 000}$. Assuming this, the probability that a random choice from $|V|$ hits the target set $|T|$ is between 10^{-1000} and $10^{-1\ 000\ 000\ 000}$, and the number of attempts needed by the blind search algorithm before hitting T will most likely be somewhere between 10^{1000} and $10^{1\ 000\ 000\ 000}$. The age of the earth (or of the universe, for that matter) is nowhere near long enough to encompass such a search procedure – even if we take into account the massive parallelism that evolution may exploit through searching along a large number of lines of descent simultaneously. Thus, the infeasibility of the blind search algorithm is settled.

Equipped with this lemma, the basic NFL theorem does the rest, according to Dembski. Of course, no one claims that Darwinian evolution proceeds via the above blind search algorithm. The basic NFL theorem, however, tells us that no other algorithm can expect to do better, and hence Darwinian evolution cannot produce specified complexity. That is, unless either the algorithm is set up using prior knowledge of the function f (and here it is inconsequential whether this function represents some fitness quantity, or some more general phenotype aspect) to help it reach the target set T , or conversely f is set up to fit the algorithm. In either case, such prior knowledge requires (still according to Dembski) an intelligent designer.

Of course, this argument is elaborated in much more detail in *No Free Lunch*, and perhaps Dembski upon reading this will feel that the last two

sentences of the previous paragraph do not give complete justice to his line of reasoning. The rough description I have given of Dembski’s argument in this section is nevertheless sufficient to make it clear that the arguments of the next two sections refute it irreparably.

5 A probabilistic interpretation of NFL

The basic NFL theorem involves an average over all possible functions f . Whenever an average or a weighted average appears in a mathematical argument, one may stop and consider whether the averaging has some probabilistic interpretation (as it usually does), and if so, how the implicit probabilistic model might be interpreted; this can often be quite illuminating.

In the setting of Section 3, the averaging amounts to picking one of the $|S|^{|V|}$ different possible functions $f : V \rightarrow S$ at random according to uniform distribution, meaning that each one is picked with probability $1/|S|^{|V|}$. An equivalent probabilistic way of formulating the basic NFL theorem as applied to the search problem of finding some $x \in V$ belonging to the target set T , is thus as follows: the distribution of the time taken for a search algorithm A to find an element of T is – provided that the function f is generated by a random mechanism that picks one of the $|S|^{|V|}$ possible realizations with equal probability – the same regardless of the choice of A .

It is worthwhile to reflect over what it means that f is chosen according to uniform distribution on S^V . I claim that⁹

$$\begin{aligned} &\text{choosing a random function } f : V \rightarrow S \text{ according to} \\ &\text{uniform distribution on } S^V, \text{ is equivalent to choosing,} \\ &\text{for each } x \in V \text{ independently, } f(x) \text{ according to} \\ &\text{uniform distribution on } S. \end{aligned} \tag{2}$$

This is a well-known fact in probability theory, and really nothing more than a straightforward extension of the standard first-year textbook example concerning the roll of two dice: the statement that all 36 outcomes $(1, 1), (1, 2), \dots, (1, 6), (2, 1), \dots, (6, 6)$ have the same probability, is equivalent to the the statement that the two dice are independent and that the distribution for each of them is uniform on $\{1, 2, \dots, 6\}$.

For completeness and for the reader’s convenience, let me nevertheless give the explicit argument for (2): Suppose that V has m elements x_1, \dots, x_m , and that S has l elements s_1, \dots, s_l . Suppose furthermore that

⁹Here and throughout, independence means statistical independence.

for each $x \in V$ independently, we choose $f(x)$ according to uniform distribution on S . To prove the claim (2), we need to show that for any $(s_1, \dots, s_m) \in S^m$ the formula¹⁰

$$\mathbf{P}((f(x_1), \dots, f(x_m)) = (s_1, \dots, s_m)) = 1/l^m \quad (3)$$

holds. Now, the independence assumption tells us that the left-hand-side of (3) can be factorized into

$$\mathbf{P}(f(x_1) = s_1) \times \dots \times \mathbf{P}(f(x_m) = s_m). \quad (4)$$

Since each of the factors in (4) equals $1/l$, the identity (3) is verified, and the claim (2) established.

Now that we are equipped with the characterization (2), the basic NFL theorem becomes very easy to understand (and to prove). To this end, imagine an algorithm A as in Section 3, that after k steps has visited $x_{(1)}, \dots, x_{(k)} \in V$, and observed $f(x_{(1)}), \dots, f(x_{(k)})$.¹¹ Now, whichever $x_{(k+1)}$ the algorithm chooses to visit next, the conditional distribution of what it will find there (given $x_{(1)}, \dots, x_{(k)}$ and $f(x_{(1)}), \dots, f(x_{(k)})$) is, due to the independence property in (2), uniform on S . Hence, the rule for how to select $x_{(k+1)}$ does not influence what we see there, and since k was arbitrary it follows that $f(x_{(1)}), f(x_{(2)}), \dots$ form a sequence of independent and identically distributed (i.i.d.) random variables whose common distribution is uniform on S . Since this conclusion is reached regardless of the details of A , it follows that the choice of A has no influence on the distribution of the sequence $f(x_{(1)}), f(x_{(2)}), \dots$. And this is precisely what the basic NFL theorem says.

In fact, not only does the observation (2) provide us with an almost trivial proof of the basic NFL theorem – it also suggests some immediate generalizations. Indeed, the argument we just indicated uses that the $f(x)$'s are i.i.d., but not that their common distribution is uniform on S . Hence, the assertion of the basic NFL theorem holds under this weaker i.i.d. assumption. And by the same token, the assumption can be weakened even further to that of so-called *exchangeability*, which means that the joint distribution of $f(x_1), \dots, f(x_m)$ equals the joint distribution of any permutation of them (see, e.g., Kallenberg, 2005). With this latter generality in mind, the basic NFL theorem is not much more than a fancy (and more general) way of

¹⁰ \mathbf{P} is short for “the probability of”.

¹¹ The notation is worth stressing: $x_{(i)}$ denotes the i :th element visited by the algorithm, whereas x_i denotes the i :th element in some fixed but arbitrary enumeration of V .

phrasing the following fact: if we spread a well-shuffled deck of cards face-down over a table and wish to find the ace of spades by turning over as few cards as possible, then no sequential procedure for doing so is better than any other.¹²

6 Dembski's error

Let us now examine Dembski's use of NFL in the light of the probabilistic interpretation given in Section 5. For concreteness, take, as in Section 4, V to be the set of all DNA sequences of length up to 3 000 000 000. Also, take $f : V \rightarrow S$ to be some measure of fitness, so that for each $x \in V$, $f(x)$ describes the fitness of an organism with DNA sequence x . Of course, most such DNA sequences do not correspond to an organism at all, so for such x we take $f(x)$ to be the minimum of the set S of possible values – say, $f(x) = 0$.

Furthermore, let us equip V with a link structure as in Section 3. Specifically, let us declare a link between two DNA sequences $x, y \in V$ precisely when one of them can be obtained from the other either by changing a single nucleotide pair, by inserting one, or by deleting one. This choice of link structure is made in order that a move from an $x \in V$ to a neighbor $y \in V$ corresponds to a mutation of the simplest possible (single-nucleotide) kind.¹³ Thus, the reproduction-mutation-selection mechanism of Darwinian evolution can be seen as one variant or another of the local search algorithms in Section 3, with the given link structure. Although we do not know the precise details of this algorithm, let us call it A .

Dembski's (2002a) application of NFL now says that

$$\begin{aligned} & \textit{if the fitness function } f \textit{ is generated at random according} \\ & \textit{to uniform distribution among all the } |S|^{|V|} \textit{ possibilities,} \end{aligned} \tag{5}$$

then the Darwinian algorithm A cannot be expected to fare any better than

¹²This obvious card-deck example summarizes pretty much all there is to the basic NFL theorem (or any of its variants). In spite of this, Dembski is not the only one who has tried to create a hype around the result. Wolpert and Macready themselves (1997) try their best to make their result sound like some kind of breakthrough. And, with astonishing lack of perspective, Ho and Pepyne (2002) compare the basic NFL theorem to Gödel's incompleteness theorem, modestly adding that the former is "far less celebrated and much more recent".

¹³This ignores inversions, gene duplications, and other kinds of macromutations. It also ignores the recombination mechanisms of sexual reproduction. Still, it provides a good enough model of evolution to make my point clear.

blind search, and will therefore almost certainly fail to produce specified complexity (the odds against it succeeding to do so are Vast).

Phrased in this way, the result is pretty much correct.¹⁴ Its relevance to evolution depends, however, on the extent to which (5) reflects properties of the true fitness landscape. If (5) is shown (as I am about to do) to be way off in this regard, then we can conclude that NFL has nothing to offer evolutionary biology. From that, it would be tempting to conclude that Dembski's entire argument falls apart, but that is (only slightly) premature, as he has one more card up his sleeve. Namely, after having invoked NFL, Dembski's line of reasoning branches into two parts: he claims that either

- (a) assumption (5) is an accurate model the real fitness landscape, in which case Darwinian evolution by natural selection can be ruled out,

or

- (b) assumption (5) fails to be an accurate model the real fitness landscape f , in which case an intelligent designer must have been involved in setting up f .

Either way, Darwinism loses and Dembski wins.

What I will do in the remainder of this section is first to deal with (a) by showing (5) to be a totally unrealistic model for the true fitness landscape f , and then to deal with (b) by showing that the conclusion about intelligent design is unwarranted. Clearly, once those two things are taken care of, the anti-Darwinian force of Dembski's argument is reduced to zero.

Of course, the assumption (5) does not have to be a *completely accurate* description of reality in order for the conclusion in (a) to deserve taking seriously. A minimum requirement, however, is that the actual fitness landscape has to have at least some rough resemblance with what one could expect to arise from a model based on (5). Alas, it does not. I will now show that any reasonably realistic model for the actual fitness landscape will produce something that is very, very different from what (5) produces.

From the characterization (2) that we established in Section 5, we see that under assumption (5), the fitnesses of any two DNA sequences (or any collection of them, for that matter) are independent – a complete disarray.

On the other hand, any realistic model for a fitness landscape will have to exhibit a considerable amount of what I would like to call *clustering*, meaning that similar DNA sequences will tend to produce similar fitness values much

¹⁴This statement is still somewhat charitable to Dembski, as it ignores his confusion concerning what specified complexity actually means (cf. Footnote 5).

more often than could be expected under model (5). In particular, if we take the genome of a very fit creature – say, you or me, whichever you prefer – and change a single nucleotide somewhere along the DNA, then we expect with high probability that this will still produce an organism with high fitness. In contrast, under assumption (5), changing a single nucleotide is just as bad as putting together a new genome from scratch and completely at random, something that we have already noted will with overwhelming probability produce not just a slightly less fit creature, but no creature at all. (If this were true, then, given the human mutation rate, we would all be dead.)

Thus, we can safely rule out (5) in favor of fitness landscapes exhibiting clustering, and branch (a) of Dembski’s argument is thereby dismantled.

Before moving on to (b), let me just note that the distinction between clustered fitness landscapes and those produced under model assumption (5) is a very important one, for the following reason. The Darwinian algorithm A introduced above is some variant or other of local search in V with the link structure we specified. Now, what makes hill-climbers and other local search algorithms work at all is very much the amount of clustering in f . The reason why it makes sense for a local search algorithm to move to a neighbor x with a high value of $f(x)$ is not so much this high value itself, as the prospect of finding some even *higher* value among the neighbors of x , and so on. In fitness landscapes without clustering, such as those produced by (5), no such gain from moving to elements with large values of x can be expected.¹⁵

In branch (b) of his argument, Dembski claims that if the real fitness landscape f does not look like what one would expect to arise under the uniform distribution model (5), then it has to be the product of an intelligent designer. That is such an incredibly silly inference that I feel a bit embarrassed about spending ink on refuting it, but since it seems so central to Dembski’s argument, let me still try and say a few words about it.

First, there is absolutely no a priori reason to expect that the “blind forces of nature” should produce a fitness landscape distributed according to (5). Anyone reasonably experienced in probabilistic modelling in science knows that such uniform distributions have no privileged status over other models as realistic descriptions of what the laws of nature produce, and that in fact only rarely do they turn out to provide good models for physical or biological systems.

¹⁵This is not to say that clustering in itself *guarantees* that the Darwinian algorithm A can find its way up the high mountain peaks in the fitness landscape. But if Dembski wishes to argue that A cannot do this, then he needs to consider the real fitness landscape (or, more likely, realistic models thereof) rather than the utterly irrelevant model (5).

Second, let us consider specifically the phenomenon of clustering in the fitness landscape f . If we are to believe the reasoning in branch (b) of Dembski’s argument, clustering in f is such a mysterious phenomenon that it can only be explained as the work of an intelligent designer. But there is in fact nothing mysterious about it at all. Rather, it can be seen as a consequence of the very wide-spread phenomenon in science (as well as in everyday life), that “like causes often tend to have like consequences”.¹⁶ This phenomenon itself is, as every science student has seen plenty of examples of, the kind of thing that mechanistic models of nature are very good at explaining, thus removing any need to invoke an intelligent designer.

To be a little more concrete, consider one particular aspect of the strong clustering tendency of the actual fitness landscape, namely that most single-nucleotide mutations on humans lead to (as far as we can tell) *zero* change in fitness. Biologists attribute this to the fact that most parts of human DNA does not code for anything at all (something that can in turn be explained in biochemical terms, and so on down the usual reductionistic path). Dembski, instead attributes it (if we are to take his branch (b) argument seriously) to an intelligent designer. I leave it as an exercise to the reader to determine which of these two explanations belongs to the realm of science.

7 Another look at averaging

Although the above more or less concludes the argument I wanted to make in this paper, let me still consider one objection that a defender of the Dembskian argument might put forth at this stage. Namely, that the averaging that takes place in the basic NFL theorem is over *all* fitness landscapes – not just the disordered ones but also those exhibiting clustering – and that in the absence of precise knowledge of f the most fair thing to do is to take equal account of all possible fitness landscapes.

The flaw in this counterargument is that although all possible fitness landscapes are indeed taken into account in NFL, the clustered ones receive such a small fraction of the probability distribution in (5) that for all practical purposes it is zero. To illustrate just how biased the model (5) is against clustered fitness landscapes and in favor of disordered ones, consider the following thought experiment. Suppose that we have sampled 2000 elements $x_1, y_1, x_2, y_2, \dots, x_{1000}, y_{1000}$ from V , in such a way that for each i , x_i and y_i are neighbors, but each such pair is very far from all others. Suppose

¹⁶A similar remark is made by Wein (2002a). Dembski (2002b) chooses to dismiss it as “bogus”. The reader may judge.

furthermore that we have evaluated f at all these points, and found that for each pair (x_i, y_i) , the difference $|f(x_i) - f(y_i)|$ is very small, while on the other hand most of the $|f(x_i) - f(x_j)|$ differences between elements from different pairs are much larger. This would constitute strong scientific evidence that the fitness landscape generally exhibits a good deal of clustering, but Dembski’s model (5) does not allow such a conclusion: the independence property in (2) implies that the conditional distribution of the rest of the fitness landscape given these observations, is just as disordered as (5) predicted without them. This shows that model (5) takes no reasonable account of the possibility of a clustered fitness landscape.¹⁷

8 Remarks on extensions

Other than Wein, one of the most ardent public critics of Dembski’s *No Free Lunch* is the well-known evolutionary biologist H. Allen Orr (2002, 2005). And although Orr deserves kudos for these mostly pertinent contributions, his criticism does fail to identify the preponderant shortcoming of the NFL application outlined in Section 6, and some of his more mathematical concerns are unconvincing. In particular, in Orr (2002), it is claimed that the the NFL argument does not apply when the function f changes over time (corresponding to an evolving fitness landscape). But in fact, Wolpert and Macready (1997) have a variant of the basic NFL theorem for precisely such cases, and this variant can be plugged into Dembski’s argument to give a evolving-fitness-landscape analog of his constant-fitness-landscape result. Such a modified Dembski argument is vaguely hinted at in *No Free Lunch*, but the brutal truth is that it fails to be relevant to biological evolution, for very much the same reasons as those outlined in Section 6.

In Orr (2005), it is instead claimed that NFL does not apply to the situation of two or more coevolving species.¹⁸ But again, although I have not been able to find in the literature an NFL theorem adapted to this situation, it is easy to devise one¹⁹, and plug it into Dembski’s argument. But yet

¹⁷It would be an interesting idea to try to devise a prior distribution for f that takes into account “all possible models and behaviors” in a better way than does (5). Much work has been made in theoretical statistics to come up with a universal scheme for such priors (see, e.g., Rissanen, 1983), but the practical and theoretical obstacles to this are large and perhaps even unsurmountable.

¹⁸The claim seems to originate from Wolpert (2002).

¹⁹This can be achieved by applying the probabilistic reasoning in Section 5 to uniform distribution on the set of all functions $f : V \times \{0, 1\}^V \rightarrow S$, where the $\{0, 1\}^V$ -valued vector indicates which elements of V (i.e., which DNA sequences) are occupied by a species, and

again, the story is the same as in the evolving-fitness-landscape setting: the arguments in Section 6 show that also this extension is entirely devoid of relevance to evolution.

Acknowledgement. I thank Mats Rudemo, who, upon reading an earlier version of this manuscript, showed me Wein (2002a).

Bibliography

- Aarts, E. and Lenstra, J.K., eds. (1997) *Local Search in Combinatorial Optimization*, Wiley, Chichester.
- Behe, M. (1996) *Darwin's Black Box*, Free Press, New York.
- Bostrom, N. (2003) Are we living in The Matrix? The simulation argument, in *Taking the Red Pill: Science, Philosophy and Religion in The Matrix*, Yeffeth, G. (ed.), Benbella, Dallas, pp 233–242.
- Crews, F. (2001) Saving us from Darwin, *New York Review of Books*, Oct 4 and Oct 18.
- Dawkins, R. (1986) *The Blind Watchmaker: Why the Evidence of Evolution Reveals a Universe Without Design*, Penguin, London.
- Dawkins, R. (1989) *The Selfish Gene* (2nd ed.), Oxford University Press.
- Dawkins, R. and Coyne, J. (2005) One side can be wrong, *The Guardian*, Sep 1.
- Dembski, W.A. (1998) *The Design Inference: Eliminating Chance Through Small Probabilities*, Cambridge University Press.
- Dembski, W.A. (2002a) *No Free Lunch: Why Specified Complexity Cannot Be Purchased without Intelligence*, Roman & Littlefield, Lanham, MA.
- Dembski, W.A. (2002b) Obsessively criticized but scarcely refuted, http://www.designinference.com/documents/05.02.resp_to_wein.htm
- Dembski, W.A. (2002c) The fantasy life of Richard Wein: a response to a response, <http://www.designinference.com/documents/2002.06.WeinsFantasy.htm>
- Dembski, W.A. (2005) Searching large spaces: displacement and the no free lunch regress, http://www.designinference.com/documents/2005.03.Searching_Large_Spaces.pdf
- Dennett, D.C. (1995) *Darwin's Dangerous Idea*, Simon & Schuster, New York.
- Dennett, D.C. (1997) The case of the tell-tale traces: a mystery solved; a skyhook grounded, <http://ase.tufts.edu/cogstud/papers/behe.htm>.
- Ho, Y.C. and Pepyne, D.L. (2002) Simple explanation of the no-free-lunch theorem, *Journal of Optimization Theory and Applications* **115**, 549–570.
- Kallenberg, O. (2005) *Probabilistic Symmetries and Invariance Principles*, Springer, New York.

which are not. However, I wouldn't recommend readers to waste their time by pursuing this line of thought.

- Miller, K.R. (1996) Book review: Darwin's Black Box, *Creation/Evolution* **16**, 36–40.
- Omnès, R. (2005) *Converging Realities: Toward a Common Philosophy of Physics and Mathematics*, Princeton University Press.
- Orr, H.A. (2002) Book review: No Free Lunch, *Boston Review*, summer issue.
- Orr, H.A. (2005) Devolution, *The New Yorker*, May 30.
- Rissanen, J. (1983) A universal prior for integers and estimation by minimum description length, *Annals of Statistics* **11**, 416–431.
- Shallit, J. (2002) Book review: No Free Lunch, *BioSystems* **66**, 93–99.
- Sokal, A. (1996) Transgressing the boundaries: toward a transformative hermeneutics of quantum gravity, *Social Text* **46/47**, 217–252.
- Sokal, A. and Bricmont, J. (1998) *Intellectual Impostures*, Profile Books, London.
- Wein, R. (2002a) Not a free lunch but a box of chocolates,
<http://www.talkorigins.org/design/faqs/nfl/>
- Wein, R. (2002b) Response? What response?
<http://www.talkorigins.org/design/faqs/nfl/replynfl.html>
- Wigner, E. (1960) The unreasonable effectiveness of mathematics in the natural sciences, *Communications in Pure and Applied Mathematics* **13**, 1–14.
- Wolpert, D.H. (2002) William Dembski's treatment of the no free lunch theorems is written in jello, <http://www.talkreason.org/articles/jello.cfm>
- Wolpert, D.H. and Macready, W.G. (1997) No free lunch theorems for optimization, *IEEE Transactions of Evolutionary Computation* **1**, 67–82.