

Book Review. William A. Dembski, *No Free Lunch: Why Specified Complexity Cannot Be Purchased without Intelligence* (Rowman & Littlefield, Lanham, Md., 2002).

In this book, Dembski bridges the gap between the highly theoretical discussion of design detection in *The Design Inference* and the nuts-and-bolts arguments for the design of biological systems in Michael Behe's *Darwin's Black Box*. Although there is still further work to be done in closing this gap, Dembski's new book puts the design-theoretic challenge to Darwinian accounts of evolution on a new level of clarity and logical rigor. As new data about the specified complexity pours in, especially from computer-driven analyses of the functional regions in the configurational space of proteins, the informational challenge to Darwinism should soon reach a critical level. At that point, even the best efforts by the biological establishment to squash this challenge through name-calling, bullying, and stonewalling (which have already been so much in evidence) will be doomed to failure.

In the first three chapters of *No Free Lunch*, Dembski explains the formal theory of design inference that he began developing in *The Design Inference*. However, this is no mere recapitulation: Dembski has effected several quite important improvements in the theory. Although *The Design Inference* is still useful as a reference source on certain points, I would recommend reading these chapters in *No Free Lunch* first, aware of the fact that it represents a new and improved version of the theory. The theory of design inference in *No Free Lunch* is cleaner and more elegant than that of *The Design Inference*. Dembski has been moving in exactly the direction I have hoped he would, although I believe that still further movement in that same direction would be desirable.

The most important difference between the two versions lies in the greater use Dembski makes of the concept of *specificational resources*. An event triggers an inference to design if the total available probabilistic resources are not great enough to make the event more likely than not. Probabilistic resources come in two kinds: replicational resources and specificational resources. Replicational resources consist of our best estimate as to the number of opportunities for this sort of event to occur within the universe of discourse under consideration. For example, if a lottery has been won, the replicational resources consist of the number of tickets to the lottery that were in fact issued. In the most general case, we look at the total number of distinct events that could possibly occur in the entire history of the recorded universe, which Dembski puts at 10 to the 150th power.

However, consideration of replicational resources is not sufficient, since any event is, if specified with enough detail, very unlikely to have occurred even once in the history of the world. An unlikely event triggers a design inference only if it fits some relatively simple pattern, like that of a hundred consecutive royal flushes dealt from a single continuously reshuffled deck. Thus, we must also take into account the relevant specificational resources. When we specify an event by giving a unique description or identifier of it, we exert a certain amount of linguistic or computational resources. We can count the number of words in a most succinct description of the event, or the number

of lines of code in a standard programming language in the shortest program capable of recognizing the event in question.

Suppose we observe an event E fitting a description D , which involves computational resources of measure μ . Let's suppose that the probability of an event's occurring that meets description D is r . To see if r is low enough to trigger a design inference, we must consult all of the relevant probabilistic resources. Let's say that there is the potentiality for M events like E to occur in the history of the world. Even if $M \cdot r$ (our estimate of the probability that an event will occur meeting description D will occur at least once) is extremely small, we cannot conclude that E was the result of design, since it may be that we have packed so much information into D ex post facto that it really isn't surprising that the probability $M \cdot r$ is so low. Instead, we must also take into account the relevant specificational resources. The specificational resources are identified by assembling a class C of event-descriptions meeting two conditions: (1) each description D' in C has a probability that is lower than or equal to r , and (2) each description D' in C can be specified with computational resources of measure μ or lower. Ignoring for the moment the possibility of overlap between the events in class C , we can estimate the specificational resources by simply counting the number of descriptions in C . Let's say that C contains N descriptions. Now, we simply look at the new probability $M \cdot N \cdot r$. This represents our estimate of the probability that an event will occur at least once fitting some description or other that can be constructed with no more resources than are needed to construct the description D . If this estimate is lower than one-half (50%), then we are justified in inferring design.

Taking into account the total specificational resources excludes the possibility that the pattern we claim to have found is a mere post hoc fabrication, like drawing a bull's eye target around an arrow stuck in a fence. It is always very unlikely that the arrow hit the exact point on the fence that it did, but not every such unlikely event triggers the inference that the arrow struck where it did by design. We have to consider all the events that are just as unlikely and that can be specified with the same or fewer computational resources. In the case of a fabrication, these specificational resources will include nearly every possible event (no matter where the arrow landed, we could, with the same effort, draw a target around it). When we multiply the low probability of the observed event by the large number of the events in the specificational resources, we end up with a probability close to 100%, and no design inference is triggered. In contrast, if the arrow strikes the bull's-eye of a pre-existing target, then the specificational resources would consist only of bull's-eyes on such pre-existing targets, a much smaller number, since it takes additional computational resources to specify points outside such pre-existing bull's-eyes.¹

¹ As Dembski recognizes on page 77, the model I've sketched is a little too simple, since by simply counting the number of events in the specificational resources, we can reach too high a number by ignoring that many of the events in that set may overlap (be such that one or more of the events could occur simultaneously). However, Dembski's repair of this problem (condition 5 on page 77) is itself somewhat defective: Dembski requires that each member of the class be a set-theoretically maximal event meeting the two conditions (i.e., a set of outcomes that is not a proper subset of another set meeting the two conditions). This won't exclude the possibility of two such maximal events overlapping, leading to continued double counting of probabilities. Instead, we should try to estimate the probability of the

Students of Dembski's work may have noticed that so far I have said nothing about the detachability of the pattern, a condition that figures largely in *The Design Inference* and that continues to be mentioned and defended in *No Free Lunch*.

Detachability is supposed to eliminate fabrications by ensuring that the pattern we detect is, in some relevant sense, *independent* of our knowledge of what actually happened. This is the right idea, and this is in fact what the introduction of specificational resources accomplishes by forcing us to consider every event of comparable probability and comparable descriptive complexity, regardless of whether we have actually observed it or not. Specificational resources are well defined and fully up to the task. In contrast, Dembski's detachability condition, introduced for the same purpose, is very difficult to make sense of.

In introducing the concept of *detachability*, Dembski tries to use probabilistic or *conditional* independence as a way of securing the proper sort of independence of the pattern from our knowledge of the event. However, conditional independence is a relation that holds between propositions or items of knowledge, not between patterns and propositions. This forces Dembski to introduce a proposition K that is supposed to correspond to the pattern π . K is an item of knowledge that enables us to "explicitly and unambiguously identify" the pattern π , or its characteristic function f . Since π and f are abstract, mathematical objects, it is not at all clear what it means for a proposition to enable us to "identify" one of these objects. K is supposed to be a "description of f that leaves no doubt about its identity" (p. 64). But a description of a mathematical object is not a proposition. The description 'the positive square root of 9' is a description that leaves no doubt about the identity of its bearer (viz., the number 3), but this description is not a proposition. It cannot be true or false, and so we cannot determine whether another proposition is "conditionally independent" of this description. Perhaps K is the proposition that there exists one (and only one?) mathematical object meeting this

disjunction of all of the events in the class of specificational resources, and then let the specificational measure N be the multiplicative inverse of this probability.

There is another minor technical problem with the definition of specificational resources on page 77. There Dembski takes into account the fact that we may not have a single estimate of the probabilities of the event-description D and of the descriptions in the class C : instead, there may be a number of "chance hypotheses" $H_1 \dots H_n$, each of which assigns distinct probabilities to the various descriptions. Dembski requires, for including D' in the class C of specificational resources, that the maximum probability $P(D'/H_i)$, i between 1 and n , be no greater than the maximum probability $P(D/H_j)$, j between 1 and n . However, this is too stringent: we should require only that $P(D'/H_i)$ be no greater than $P(D/H_i)$ for *every* chance hypothesis H_i , since the point of the design inference is to eliminate each chance hypothesis individually. This stronger condition can easily be met by simply building one of the H_i 's into each description D' . Then, the $P(D'/H_i)$, for $j \neq i$, will be zero, and when we take the disjunction of all of the events in C for each hypothesis H_i , we will simply be looking at the descriptions in C that do include H_i .

It should be noted that in both cases, these minor repairs make the condition for inferring design less stringent than Dembski's version – so any phenomenon that passes Dembski's more stringent filter will also pass my alternative as well. To the extent that Dembski is worried only about producing a condition that is sufficient to warrant design, these repairs are unneeded. However, I understand the project to be one that ideally will give us a condition both necessary and sufficient for the warranted inference of design.

description: e.g., ‘there exists a unique positive square root of 9’. Such mathematical propositions are standardly all given the probability 1 (or 100%), but perhaps a non-standard probability measure could be devised that gives mathematical truths and falsehoods probabilities intermediate between 0 and 1. However, even though this seems like a worthy project, it won’t help Dembski, since when I become aware of the unique existence of such a mathematical object, my knowledge will be based on some mathematical intuition or proof, and so it will always be probabilistically or conditionally independent of any empirical knowledge, including my knowledge of the occurrence of event E.

Sometimes, Dembski talks as if the proposition K had something to do with our awareness or acquaintance with the mathematical object (p. 65). It enables us to focus only on certain mathematical objects thanks to the fact that, as finite thinkers, we are capable of being aware of only a limited number of mathematical objects. So, perhaps we should interpret K as the proposition that I am presently aware of (or acquainted with) the object f. However, such a blatantly autobiographical proposition seems out of place here, and in any case, it could have conditional independence from the proposition that E occurred, even if f represents a post hoc fabrication. Suppose I draw a target around the arrow. I am now aware of a mathematical object (one that picks out the coordinates of a small region on the fence). My knowledge that I am aware of this region is based on simple introspection and is probabilistically independent of my knowledge that the arrow landed where it did (I know by introspection that I’m aware of these coordinates, and my estimate of the probability that I am so aware is not affected by supposing that I’m wrong about where the arrow actually landed.) My knowledge of this awareness is not *causally* independent of my knowledge of where the arrow landed, but Dembski’s detachability condition does not demand that kind of independence. I have to conclude that detachability was an interesting idea that simply hasn’t panned out. (Perhaps it could be salvaged by requiring that the saliency of the pattern π be causally independent of our knowledge of E.)

Dembski introduced the detachability condition as a device intended, like the device of specificational resources, to exclude fabrications. In fact, it is a separate and independent device, raising the question of whether Dembski really needs both, and if not, which is superior. In my opinion, the role of specificational resources renders the detachability condition redundant. Moreover, as I argued above, the detachability condition remains problematic at best, and so I would recommend jettisoning it entirely. The resulting theory will be even more elegant than the version sketched in *No Free Lunch*.

Consequently, I would simplify the summary of the design inference on pages 72 and 73 by simply dropping step 4 (the detachability condition), which I believe is redundant in any case, given steps 5 and 6 (the introduction of specificational resources).

In *No Free Lunch*, Dembski helps to clarify the fact that, although he makes use of mathematical models, the design inference itself is a posteriori, not a purely a priori or mathematical inference. As Dembski points out on page 68, the design inference involves “sweeping the field clear of all chance hypotheses”, but this cannot mean excluding all possible chance hypotheses. It is always possible to introduce ad hoc a

chance hypothesis that assigns high probability to the observed event's occurring as a result of natural (unintelligent) causes. However, the burden of proof is on the design skeptic at that point to find some independent grounds for considering such a hypothesis. The fact that it saves the observed phenomenon *ex post facto* gives us no grounds for taking the hypothesis seriously. The bare possibility that some unknown natural process might account for the phenomenon no more provides grounds for doubting the inference to design, than the bare possibility that I am a brain in a vat provides grounds for doubting the deliverances of my senses.

It is important to bear in mind, as many of Dembski's critics do not, the nature of Dembski's project. Dembski is attempting to provide a mathematically rigorous model of an ability that all of us share: the ability to recognize certain reliable indicators of design. Dembski is engaged in a project of theoretical epistemology, taking our considered judgments about which inferences are reasonable as data. He is not engaged in refuting a Pyrrhonian skeptic, who denies that we are ever reasonable in any of our inferences. This is why Dembski rejects out of hand any model of design inference that would make it impossible for us ever to be reasonable in inferring the presence of design agency. Such a model clearly would not fit the fact that we are often reasonable in so inferring. It is for this reason that Dembski labels as fallacious any attempt to avoid a design conclusion by arbitrarily increasing the world's replicational resources. I could avoid the conclusion that Shakespeare wrote his plays by hypothesizing that the universe is large and old enough (far larger and older than anything we have actually observed) that the chance occurrence of Shakespeare's plays at least once in the history of the universe was quite probable. Obviously, such inflation of replicational resources would render every attempt to infer design unsuccessful, a blatantly unsatisfactory result. It is special pleading of the worst kind for cosmologists to resort to *ad hoc* inflation of the universe in order to avoid the inference that the fine-tuning of observed constants was designed, when they refuse to resort to such inflation to avoid design inferences that they find more congenial.

As Dembski points out, in their review of *The Design Inference in Philosophy of Science* in 1999, Fitelson, Stephens and Sober reject Dembski's model simply because it is not identical to the likelihood account of probabilistic inference that they favor. Such a likelihood account is simply an extension of philosopher David Hume's project, in which all empirical conclusions are based on extrapolation from past experience. This should mean that we can infer design only when we have personal experience of similar things having been designed in the past. As Dembski points out, this runs into two difficulties. First, archeologists are often able to recognize the existence of artifacts, even when the intended use and means of production are almost entirely unknown, and SETI research is predicated on the assumption that a similar kind of recognition of design is possible in the total absence of prior knowledge of the producers. Second, the Humean likelihood approach is unable to explain how we were able to recognize design in the first place. If all we can use as data are regularities in observed behavior, how do we ever come to know that some of this behavior is intelligent and purpose-driven? Sober et al. might appeal to introspection, but even there, as Hume *himself* was the first to point out, all we

observe are successions of psychic events. We never have direct access to the intelligent self and its causal efficacy in ordering events to ends.

In addition, as Dembski argues, the likelihood model relies on a myth of false precision. In many cases, we have no way of judging how likely it is that an intelligent agent should act precisely as he does. The human and social sciences are notoriously poor predictors for just this reason. Finally, the likelihood model ignores the central role that causation plays in scientific inference. Recent work on causation, especially by British philosopher Nancy Cartwright (see her *Nature's Capacities and Their Measurement*), has brought back into focus the crucial importance of the kind of eliminative inference championed by Dembski and opposed so dogmatically by Sober et al.

There is one minor point that could be further clarified. In both *The Design Inference* and *No Free Lunch*, Dembski discusses the case of Nicholas Caputo, an election judge who, while claiming to use a randomizing device, put Democrats first on the ballot 40 out of 41 times. On pages 80-82, Dembski says that the Caputo case is a model of a correct and successful design inference. He concludes, "The New Jersey Supreme Court is warranted in concluding that E did not occur according to the chance hypothesis." However, on page 107, Sober et al. are criticized for jumping to the erroneous conclusion that Caputo was found guilty of fraud. Dembski says that the likelihood method that Sober et al. used to reach the conclusion that Caputo was guilty is rightly viewed with skepticism by the legal community. This seems an inconsistency: if both Dembski's inference machine and the likelihood model reach the conclusion that design was involved, and if only Caputo had the means and opportunity, then how can Dembski fault the likelihood method for being too incautious in assigning guilt to Caputo?

In Chapter 3, Dembski discusses the relationship between various kinds of information: statistical, syntactic, semantic, and complex specified information. He rightly concludes that syntactic and semantic information are special cases of complex specified information (CSI). When CSI is present, we can conclude that the event has been constructed for some purpose or other. Semantic information consists of CSI that has been introduced for one very special kind of purpose: that of storing or conveying meaning. In this chapter, Dembski constructs an elegant proof of his most important formal result: the Law of Conservation of CSI. I have already compared this in importance to Newton's laws of motion, and I think it is worthy of being considered a fourth law of thermodynamics. Dembski concludes the chapter with a fascinating discussion of the paradox of Maxwell's demon, a thought-experiment in which a nano-equipped intelligent agent is able (apparently) to produce a clear violation of the Second Law of Thermodynamics. I think that further work needs to be done on the relationship between entropy and CSI. I conjecture that Dembski's theory of CSI could be used to put all of thermodynamics on a new and firmer footing: in fact, I suspect that thermodynamic entropy is nothing other than the inverse of CSI. In thermodynamics, some states are supposed to be "more probable" than others, a claim which I have found puzzling. Surely, every state is, if specified with sufficient detail, as improbable as any other. A state of low entropy must consist of a state that is, not only improbable, but also highly

specified (in Dembski's senses). If this is correct, then Dembski's law of conservation of CSI may turn out to be a new first law (or perhaps, zeroeth law) of thermodynamics.

Chapter 4, on evolutionary algorithms, is in some ways the crux of the book. It is here that Dembski discusses the No Free Lunch theorems that give the book its title. Dembski demonstrates that so-called "evolutionary" (or trial-and-error) algorithms cannot generate CSI. Instead the problem of generating CSI is simply pushed back from a high CSI result to a very high CSI fitness function. A fitness function is a function that differentially "rewards" or "selects" different trials, whether these trials be algorithms, behavior patterns, neural-net configurations, or biological genotypes. To solve a problem using an evolutionary algorithm, the designer must find a fitness function that can enable the sort of evolution likely to lead to an acceptable solution to the problem. Where there is a large amount of CSI required to solve a problem directly, the NFL theorems entail that an equally large quantity of CSI is required to find a suitable fitness function.

I hope that committed Darwinists will accept this brilliant and illuminating analysis. Even if you are convinced that Dembski is completely wrong about Darwinism, he has made an indispensable contribution to the development of the information-theoretic analysis of evolution. I hope that relatively few commit the genetic fallacy of rejecting Dembski's work on the grounds that he holds "scientifically incorrect" opinions on other matters.

As I see it, Dembski's work in chapter 4 by no means excludes the possibility that Darwinian mechanisms can explain complex adaptations. Dembski's analysis leaves the Darwinist with a two-pronged response to the informational challenges:

1. The Darwinist can claim that there is a large store of CSI implicit in the natural world:
 - (a) There is CSI implicit in the laws of nature and in the anthropic values of the constants.
 - (b) The stability of the laws of nature adds CSI (as compared with a world in which the laws of nature vary randomly from one moment or region to the next).
 - (c) There is CSI implicit in the selection of "Darwinian" fitness functions: i.e., fitness functions that reward organisms in proportion to their rate of reproductive fecundity. (as compared with fitness functions that select organisms for other, arbitrarily selected features).
2. The Darwinist can claim that the complexity of adaptations is only apparent, where "complexity" is understood in probabilistic terms. He can argue that we are prone to exaggerate the amount of CSI contained in adaptations, by failing to take into account how likely it is that adaptations of that kind should arise in a world like ours after a billion years of evolution.

What the Darwinist should not dispute is Dembski's demonstration that natural causes (including mutations and natural selection) cannot increase CSI. That thesis should be taken as common ground, so the discussion can instead focus on the two claims above.

In particular, I think much more attention should be given to point 1(c), which might be called logical or conceptual or investigator selection. By focusing on living things, on things that replicate themselves, considered as self-replicators, the biologist is considering only a very small subset of the class of all possible fitness functions. The focus of biological investigation itself introduces a significant amount of CSI into the biological domain. This needs to be quantified and compared with the amount of CSI actually found there. My uninformed guess is that the quantity of information provided by such logical or investigator selection will prove to be quite limited, far less than is needed to account for the complex adaptations we find.

The importance of this point emerges clearly in Dembski's discussion (on pages 221-223) of the development by Chellapilla and Fogel of a checkers-playing neural net algorithm that was able to attain expert status in 840 generations. According to Dembski, the CSI was inserted through the "coordination of local fitness functions":

It is important to understand that there is nothing requiring one local fitness function defined for 30 neural nets to match up with another local fitness function defined for another 30 neural nets. A local fitness function, as it were, hands off winning neural nets satisfying its criterion of success to another local fitness function *imposing the same criterion of success* [emphasis mine]. Chellapilla and Fogel kept the criterion of winning constant from one set of neural nets to the next. But this was a choice on their part. To be sure, it was the appropriate choice given that they were trying to optimize checker playing. But it was a choice nonetheless. Indeed, it was a choice that inserted an enormous amount of specified complexity (the space of all possible combinations of local fitness functions from which they chose their coordinated set of local fitness functions is enormous). [pages 222-3]

This source of CSI seems closely related to my point 1(b) and 1(c) above – the stability of the laws of nature, and the general stability of the conditions on the earth's surface, ensure that each succeeding generation of organisms confronts a Darwinian fitness function very similar to the one faced by its ancestors. Yet, surprisingly, Dembski goes on to say that Chellapilla and Fogel's choice of inter-generational stability is "without a natural analogue". Perhaps I'm missing something, but this just seems clearly false. Darwinian selection does involve the repeated imposition of essentially the same fitness function on a succession of generations. Dembski says that in nature the "criterion for 'tournament victory' will vary considerably depending on who is playing the tournament". I guess what Dembski means is that as the biological world evolves, the relevant fitness function for each organism changes. However, this evolution is, ex hypothesi, a very gradual one, and so organisms do normally face very stable environments.

In my view, more attention should be paid to the question of how much CSI is really embodied in attaining the expert level in checkers. We shouldn't assume that everything that human beings find difficult to do involves a large amount of CSI. We find games like

checkers and chess interesting precisely because they require us to do things that we are not naturally adept at. It is much easier to program a computer to play chess well than it is to program a robot to find its way down a hallway, despite the fact that any human child finds the latter much easier than the former. Most of us are not very good at multiplying three-digit numbers in our head, an operation involving a trivial amount of information, and children can spend years playing tic-tac-toe, despite the fact that the finding an optimal strategy for the game is laughably easy.

In Chapter 5, Dembski brings his analysis to bear on the problem of irreducibly complex systems in biology. This will certainly be the most controversial chapters of the book, aiming at it does a shot across the bow of the Darwinian establishment, but I found it entirely convincing. Here Dembski clarifies that his target is not evolution in the broad sense, the hypothesis of common descent, but only the adequacy of non-intelligent mechanisms (such as that of Darwinian natural selection) to account for the complex adaptations we find in so many places in the world of living things. This is critical, since many of the supposed “objections” to Behe and Dembski on this point (including appeals to cooptation and bricolage) are in fact merely defenses of the thesis of common descent, and have nothing to say about how new CSI becomes incorporated in living things.

I won't repeat here the substance of Behe's argument that Darwinism cannot account for irreducibly complex molecular machines – I take it that this should already be familiar to my readers. The core of the chapter is taken up with Dembski's masterful refutation of the extant objections on the part of biologists and philosophers to Behe's book. There are four objections that deserve the most attention: the scaffolding theory, the co-optation theory, the reducible complexity argument, and the redundancy argument. The scaffolding theory points out that it is possible to reach an irreducibly complex system not by simple addition, but by a combination of addition and subtraction. That is, it could be that there was once a reducibly complex flagellum (one that could be reached in a stepwise Darwinian fashion), and that at some point some features or parts of that flagellum were lost, resulting in the present, irreducibly complex system. Eliminating any of the present parts of the flagellum results in a complete loss of function, but that might not have been true about the still more complex intermediate version.

This theory is the product of some real ingenuity, exploiting a loophole in Behe's original argument. However, it is hard to believe that it is a promising strategy for explaining the kind of examples Behe adduces. As Dembski points out, the very nature of the function of the flagellum guarantees that there must be an irreducibly complex core. You simply cannot gain rotary propulsion in a stepwise, Darwinian fashion. The function requires, not just in fact, but in principle, a large number of mutually adjusted parts. Second, no one has produced any specific, testable hypotheses about these missing intermediate forms that are reducibly complex. A mere conceptual possibility that such a thing might exist does nothing toward weakening Behe's design inference. Are commercial power tools the product of intelligent design? It is logically possible that all such tools are actually waste products extruded by heretofore unobserved organisms, but this logical possibility gives no basis for real doubt about the actual design of those tools. Similarly, the bare possibility that there might be unimagined flagella that are producible by

Darwinian methods, from which present-day flagella are descended by drastic reductions, does not count as a cogent rebuttal of Behe's argument.

The second theory is the co-optation idea, first proposed by Stephen Jay Gould and Elizabeth S. Vrba. On this view, structures that are irreducibly complex in relation to their present function F actually evolved because they served some other, entirely distinct function G, in relation to which they were reducibly complex. Here we have to make a distinction between two possible kinds of co-optation: total co-optation, and *modular co-optation*. Total co-optation would mean that the entire structure, with all of its present features, evolved because it served function G, but is now preserved because it serves a new and entirely disjoint function F. Total co-optation, like scaffolding, is a bit of science fantasy, since no one has discovered a distinct function served by present-day flagella (or the blood-clotting mechanism, or intracellular transportation systems, and so on). Moreover, when the mechanisms are complex, it is astronomically improbable that the very same structure that serves function G would be simultaneously an irreducibly complex mechanism capable of serving unrelated function F. If such miracles of coincidence really were common in the history of life, this would surely demand an intelligent designer who had worked out in advance how to build an optimal G-er that could also serve as an F-er.

Modular co-optation is much more plausible, and we can even find evidence of such modular co-optation in similarities between existing parts of distinct molecular machines. In modular co-optation, an irreducibly complex structure with function F can be broken down into a number of separate modules. Each module serves some general-purpose function (like forming an aperture, or supplying energy, or constituting a rigid rod). At some point, pre-existing modules are borrowed or "co-opted" from a variety of mechanisms and re-assembled to form a new mechanism with a new function.

Taking into account the possibility of modular co-optation may force us to re-evaluate the degree of complexity of a molecular machine. The internal complexity of each module is no longer a problem, since that complexity is now seen to be possibly reducible. However, once we have identified the relevant modules, there is still a large amount of CSI that needs to be explained: namely, the CSI involved in selecting the right number of the right kinds of modules, and in putting them together in the right spatial arrangement. In a complex cell, like a bacterium, the number of available modules is quite high, posing the serious problem of excluding those modules that would, if incorporated into the new structure, interfere with the performance of the new function. In other words, as Dembski puts it, structures like the flagellum require "multiple coordinated co-optations".

In fact, the problem for the Darwinists is even greater than this. We must take into account the dynamical and developmental aspects of organic systems. An organism is not really a set of three-dimensional structures, but rather a set of four-dimensional processes. Darwinists assume that it is simple matter to introduce a small, incremental change in the adult form of an organism (for example, to move a module from one location to another). In fact, however, small changes in one time-slice of a complex system like an organism (say, in the embryonic stage) tend to produce very large changes further downstream.

Chaos theorists call this the Butterfly Effect (in principle, the fluttering of a butterfly's wing now could produce a hurricane in China later). To re-organize a large number of pre-existing modules into a new structure involves a carefully fine-tuned re-ordering of the organic stream of active life, so as to bring about new construction and maintenance processes, processes that are themselves irreducibly complex. The introduction of new, internally coherent processes of this kind involves fantastically large amounts of CSI, as is well known by those who study these intracellular construction projects.

John H. McDonald has attacked Behe's illustration – the irreducibly complex mousetrap. By ingeniously folding a single wire, McDonald has shown how to build a mousetrap consisting of a single piece, in place of Behe's irreducibly complex, seven-piece mousetrap. I think it's obvious that a large amount of CSI went into the twisting and retwisting of McDonald's supposedly "simple" mousetrap, at least as much CSI as was present in Behe's original model. There is a serious point, however, that McDonald's stunt illustrates: we must be careful about identifying what counts as "parts" of a system. The fact that the various parts of McDonald's wire are physically connected to one another does not make them a single part. I think that instead of talking about the various "parts" of a system, we should instead focus on the causally relevant features of the system. We can clearly identify a large number of features of McDonald's mousetrap that are indispensable to its functioning.

Finally, Niall Shanks and Karl Joplin have made much of the fact that many biological systems are multiply redundant. If one system fails, there are often others that can step in and pick up the slack. These redundant systems seem, at first glance, to be incompatible with Behe's definition of irreducible complexity, since we can delete any part of any of the redundant systems without any loss of function. However, this involves a simple mistake. In fact, the redundant systems that Shanks and Joplin are pointing to are especially good examples of Behe's irreducible complexity, once we identify the function correctly. Suppose that we have a triply redundant system for some effect E (such as temperature maintenance). There are three independent sub-systems, each of which is by itself sufficient to produce E. In this case, there is no irreducibly complex system with producing E as its function, as Shanks and Joplin observe. However, consider instead the function of *the triply redundant production of E*. The three sub-systems, taken together, is an irreducibly complex system in relation to this function. If we eliminate any essential part of any of the three sub-systems, the remainder can perform, at best, the function of the doubly redundant production of E. The function of triple redundancy has been lost.

Dembski's Chapter 5 ends with an estimate of the number of bits of CSI incorporated in the bacterial flagellum. Dembski offers the quite conservative estimate of 10 to the minus 234th power as the probability of the formation of the flagellum without intelligent agency, which corresponds to about 780 bits of CSI. Dembski's procedure for reaching this estimate is an excellent model to follow. He recognizes that one cannot simply estimate the probability of finding exactly the structure that actually occurs in the extant organism. Instead, one must estimate the probability of finding some structure that fulfills the same function. At first glance, this seems to be an intractable problem, since

we will never be able to discover exactly how many alternative structures could fulfill the same minimal function. However, there is a tractable solution. Instead of trying to search the entire space of organic structures, for example, one can focus instead upon the perturbation neighborhood of the existing structure. If we find that the probability of reaching a functional alternative inside that neighborhood is sufficiently low, a design inference can be triggered.

In effect, what we do is to increase somewhat the computational resources involved in specifying our target. Instead of looking for structures that match the pattern of, say, *enabling propulsion through rotary motion*, we match the slightly more complicated pattern of *being a structure in the neighborhood of the flagellum that enables propulsion through rotary motion*. So long as this neighborhood is large enough to enable a sufficiently low probability (below the universal bound of 10 to the minus 150^{th} power), and so long as the neighborhood is defined in a simple and natural way (and not by means of a computationally-expensive gerrymandering), then this neighborhood-specific target is both tractable and capable of triggering a design inference.

As Dembski notes, this procedure is inspired by an analogy constructed by John Leslie in his book, *Universes: the Fly on the Wall* analogy. Suppose that a gunshot strikes and kills a fly on a wall. Do we look for a design explanation? Initially, we would want to know how much of the wall was covered by flies. However, even if it turns out that much of the wall is inaccessible to us (so we can't say how much of those parts were covered with flies), we may still be able to infer design. In fact, we may be able to infer design even if we know that the great majority of the wall was covered in flies. Suppose, for example, that the wall is largely covered with flies, except for one, isolated spot, about one mile wide. On this stretch, there is a single fly, located roughly in the center of the otherwise bare stretch. If it is the isolated fly that is shot, we would take a design explanation seriously. Similarly, even if we can't estimate just how many organic structures could function as a rotary-motion propulsion system, we may infer design if we find that the proportion of structures in the neighborhood of the bacterial flagellum that can do so is fantastically small.

In the final chapter, Dembski turns to the question of the future of intelligent design as a scientific research program. Dembski's own project, that of formalizing our commonsense ability to detect design, has already proved fruitful. I would compare what Dembski has done for the design inference with Alonzo Church's formulation of our intuitive idea of computability in terms of the mathematically rigorous definition of recursibility. Dembski points out a large number of additional questions that a design-theoretic approach raises: which structures have been designed? What is their function? Are they optimally designed, and if so, for what constraints? Has the design been perturbed or degraded, and if so, how? At what point in evolutionary history did the design take shape? As Dembski suggests, we can study the history of evolution, looking for patterns in the introduction and development of design, including instances of intentional preadaptation (of the kinds hypothesized by Michael Denton in *Nature's Destiny*).

Dembski points out that the design-theoretic program is fully compatible with the front-loading of design, despite the claims of critics like Howard van Till. Taking design seriously does not require the postulation of miraculous interventions.

Some philosophers and scientists have objected to the intelligent design program on the grounds that at some point, a disembodied intelligent agent will be required, and that disembodied agent must somehow move the particles making up the physical and biological worlds. Dembski argues that quantum indeterminism opens up the possibility of such disembodied action without any violation of fundamental physical laws.

I have several reasons for thinking that this is not the tack that design theorists should take. First, there are deterministic versions of quantum mechanics, and not just the many-worlds version Dembski criticizes. For example, David Bohm's mechanics can be given a deterministic interpretation. (By the way, I think there are more powerful objections to the many-worlds interpretation of QM that Dembski overlooks – see D. H. Hodgson's *The Mind matters: Consciousness and Choice in a Quantum World*, pages 335-341, for a good overview of these.)

Second, it is not at all clear that a disembodied agent can take advantage of quantum indeterminism without violating quantum laws. Suppose that a particle has (according to its wave function) a 50% of passing through a given membrane. If the agent makes the particle pass through the membrane, he makes the probability of its being there 100%, which contradicts the probability assigned to the event by the fundamental quantum law. Such a deviation in objective probabilities is just as much a violation of quantum law as the spontaneous generation of energy would be a violation of Newton's laws.

Thirdly, there is something odd about a powerful disembodied agent (especially if that agent is omnipotent) being forced to find obscure loopholes in the laws of physics before he can act. The laws of nature describe what happens when the physical things are left to their own devices. It is hard to see how they could act as constraints on a supernatural agent.

I think a much better approach is to recognize that teleological principles, such as least time and least action principles, have a long and distinguished history in physical theory.² In fact, it is at least arguable that teleological explanations (couched in terms of integral equations) are the most fundamental, and that the dynamic laws, involving forces and conserved energies, and expressed as differential equations, are in fact mere by-products or epiphenomena.³ Forces and energies should be thought of as mere bookkeeping entries, taking the form they do as is dictated by the underlying teleological laws. In this sense, we can expect, in due course, to find evolutionary forces and energies, but only

² Wolfgang Yourgrau and Stanley Mandelstam, *Variational Principles in Dynamics and Quantum Theory* (Dover Publications, New York, 1979), pp. 19-23, 164-167; Cornelius Lanczos, *The Variational Principles of Mechanics* (4th edition, Dover Publications, New York, 1986), xxvii, 345-6; Robert Bruce Lindsay and Henry Morgenaw, *Foundations of Physics* (Dover Publications, New York, 1957), pp. 133-6.

³ Jim Hall, "Least Action Hero," *Lingua Franca* 9 (October 1999): 68.

after we have discovered the higher-order, biological analogues of the physical principles of minimum action (perhaps something like “best function” principles).

Efficient explanation by means of differential equations would be a heuristically useful but ultimately fictional. Such a teleological recasting of modern physics (advocated most energetically by Max Planck⁴) would provide a physical model far friendlier to intelligent agency than the usual efficient-causality model. The introduction of new design into nature would be merely an intermittent and large-scale version of a universal phenomenon.

Finally, Dembski considers the metaphysical objections to intelligent design offered by doctrinaire naturalists like Richard Dawkins and Daniel Dennett. The most potent of these is the “Who designed the designer?” retort, attempting to show that ID theory solves one mystery only at the cost of postulating a still greater mystery. This is a venerable objection – Hume’s Philo poses it in his *Dialogues Concerning Natural Religion*. Dembski argues that the ID theorist can “readily decline the regress”. I think this is too casual a dismissal. Although neither Dawkins nor Dennett would put it quite this way, I think that they are both implicitly recognizing that canons of scientific acceptability are answerable to metaphysical theory (just as metaphysical theory is answerable to scientific data). So long as our dominant metaphysical theory is some version of materialistic naturalism, inferences to design as an explanatory category will always be seen as unacceptable, unless design itself can be reduced (in part, by Darwinian means) to naturalistic categories.

However, materialism has come and is increasingly coming under very heavy fire within the philosophical community. Powerful epistemological objections to materialistic naturalism have been raised by Alvin Plantinga (in *Warrant and Proper Function*) and by Michael Rea (in *World without Design*). Other relevant works are recent collections edited by William Lane Craig and J. P. Moreland (*Naturalism: A Critical Analysis*) and by Paul Moser and Paul Copan (*Theism: A Philosophical Defense*), as well as my own *Realism Regained*, and the revival of Thomistic metaphysics by Barry Miller and Norman Kretzmann.

Of course, the intelligent design movement is itself a powerful challenge to the hegemony of materialism, and William Dembski continues to demonstrate that he is one of the leading innovators in that movement. I hope this book will be widely read and profoundly influential.

⁴ Max Planck, “The Principle of Least Action,” *A Survey of Physical Theory*, R. Jones and D. H. Williams, trans. (Dover Publications, New York, 1960), pp. 69-81; “Science and Faith,” in *Scientific Autobiography and Other Papers*, W. H. Johnson, trans. (W. W. Norton & Co., New York, 1936), pp. 119-126.