

OBSESSIVELY CRITICIZED BUT SCARCELY REFUTED: A RESPONSE TO RICHARD WEIN

By William A. Dembski

Talk.origins has now officially archived Richard Wein's critique of my book *No Free Lunch* at <http://www.talkorigins.org/design/faqs/nfl>. Prior to that, the critique went through several revisions. I take it the critique is now substantially finished. In any case, I am responding to Version 1.0 last modified 04.23.02. My response here is copyright © 2002 and may be reprinted only for personal use.

1. Preamble

I have many critics. Some are measured and calm. Others are obsessive. Richard Wein is perhaps the most obsessive. His critique of my book *No Free Lunch* (hereafter NFL) weighs in at 37,000 words and purports to provide the most thorough refutation of my work to date. It certainly is long. But is it thorough and does it succeed in actually refuting my ideas? In fact, the critique fails as a refutation and skirts key issues at every opportunity. It is therefore neither thorough nor a refutation.

In this preamble I want to take up a few charges that have nothing to do with my ideas but with the context in which my ideas have been debated with Wein. I first learned of Richard Wein when he wrote me an email message two or three years back. He seemed respectful at the time and wanted clarification on some issues in my book *The Design Inference* (hereafter TDI). He also asked whether he could share our correspondence with others. I responded briefly to his questions and intended to get back to him more thoroughly. Also, I indicated that I did not want our correspondence becoming a matter of public record on the Internet. Shortly thereafter I saw a post on talk.origins in which Wein informed people that he had contacted me and that I had "fobbed him off." This public notice of our correspondence caused me to lose respect for him, and I subsequently refused to have any direct correspondence with him. After that his criticisms of my work became contemptuous and obsessive, as is evident in his most recent critique.

Wein reads his own misconceptions into my work. It is therefore no surprise that my work seems confused to him. Wein seems to make much of my citing him in the acknowledgments of NFL, as though I owed him a debt for remarkable insights. Some of the people I cite in the acknowledgments (like University of Texas philosopher Robert Koons) are perceptive critics. I do not regard Wein as one of them. Wein, in my view, is a mixed-up critic whose confusions I found instructive and forced me to clarify certain things in NFL that had been unclear in TDI. I'm grateful to him for this and acknowledged him accordingly. The acknowledgment seems to have gone to his head.

A word about credentials seems worth bringing up. Richard Wein holds a bachelor's degree in statistics -- that's it. Now, his lack of academic credentials by itself does not undercut his critique

-- any work must ultimately stand on its own merits. Yet his lack of formal training should raise a suspicion about his contempt for my knowledge of statistics and probability. I do have the requisite credentials in this field, having a Ph.D. from a good school and having published in this area (my CV is available at www.baylor.edu/~William_Dembski). NFL was thoroughly vetted by experts spanning a range of disciplines. Moreover, the endorsements appearing with the book are glowing and come from experts in computer science, machine learning, mathematics, statistics, engineering, biology, biochemistry, physics, and philosophy (see http://www3.baylor.edu/~William_Dembski/docs_books/nfl.htm). All of the endorsers are more eminent in their respective fields than Wein. How can they think NFL is the best thing since sliced bread and Wein think it is the worst thing since the Lisbon earthquake? One has to wonder if more isn't at stake than just the merit of the ideas.

One final word before we get started. Why am I responding to Wein at length? After all, won't the "big boys" weigh in soon enough (Elliott Sober, for instance, wrote a critical review of TDI)? Two reasons. First, I'm finding that the "big boys" are being skittish. For instance, when I invited a number of them to have it out on an Internet forum (I won't name names, but they include some of my most prominent critics), they all bowed out. Perhaps they are playing it safe. Better to see what sort of response Wein "the human shield" gets before weighing in. Second, groups like the NCSE and others are likely to take Wein's critique as "the definitive review" and "a thorough demolition" of my work and then tout this to their faithful.

Because Wein's critique appears only online, I shall reference it only with quotations, not page numbers. Wein's critique is long, so I shall focus only on what I regard as the most salient points where I disagree with him. There are plenty of other places of disagreement, but I forego them for now.

2. Peer Review

The first thing I want to address is a credibility issue, my own in this case. Wein complains:

The Design Inference did undergo a review process, though no details of that process are available. It is interesting to note, however, that *The Design Inference* originally constituted Dembski's thesis for his doctorate in *philosophy*, and that his doctoral supervisors were philosophers, not statisticians. The publisher (Cambridge University Press) catalogues the book under "Philosophy of Science". One suspects that the reviewers who considered the book on behalf of the publisher were philosophers who may not have had the necessary statistical background to see through Dembski's obfuscatory mathematics.

Here are the relevant facts:

- TDI is a revised and corrected version of my doctoral dissertation in philosophy. Five professional philosophers read it and had to pass on it for me to receive my Ph.D.

- TDI was voted the best dissertation in the humanities for 1995-96 at the University of Illinois at Chicago and received a \$500 cash prize. Presumably that means that it was reviewed by the humanities division of the University of Illinois at Chicago as well.
- TDI appeared in Cambridge University Press's monograph series known as *Cambridge Studies in Probability, Induction, and Decision Theory*. This series is the equivalent of a journal. It has a general editor, Brian Skyrms (who, by the way, is a member of the National Academy of Sciences). It also has an editorial board, which at the time of publication consisted of the following: Ernest Adams, Ken Binmore, Jeremy Butterfield, Persi Diaconis, William Harper, John Harsanyi, Richard Jeffrey, Wolfgang Spohn, Patrick Suppes, Amos Tversky, and Sandy Zabell. This editorial board is a literal who's who in the statistics and inductive reasoning world. Persi Diaconis (Stanford) and Sandy Zabell (Northwestern) are personal acquaintances and are housed in the statistics departments of their respective schools.
- TDI went to three anonymous referees for a grueling year-long review process. The first referee was overwhelmingly positive. The second referee was on balance negative, though s/he had some positive things to say about it. Brian Skyrms wanted the book in his series and therefore gave it to a third referee as a tie breaker. The third referee was very positive about TDI but wanted significant revisions (the referee report was about seven pages). I agreed to do the revisions, whereupon Brian Skyrms recommended the book to the Cambridge Syndicate, which in turn then issued me a contract to publish the monograph. I spent the summer of 1997 revising the monograph in accord with the third referee's suggestions.
- The review/referee process for TDI was more rigorous than anything I've experienced in the peer-reviewed journals in which I've published, and that includes math, philosophy, and theology. TDI is the equivalent of a peer-reviewed journal contribution, and I challenge anyone to argue otherwise. The reason TDI didn't appear in a journal is that the argument required a book-length treatment.
- TDI is Cambridge University Press's best selling philosophical monograph in the last several years.

That answers the history of TDI's review process, but leaves unanswered Wein's concern that it didn't appear in the actual statistics literature (like some monograph series in line with *JASA* or *Annals of Statistics*). The reason has nothing to do with the merit or lack of it of my work, but appropriateness. Journals and disciplines are these days highly specialized and segregated. Statistics journals are essentially math journals that want to see theorems and proofs or simulations. They don't want to see broader framework questions addressed in their pages, like

how do we interpret probabilities or reason with them or characterize different types of statistical inferences. That is properly handled in the philosophy of science literature (specifically under rubrics like inductive reasoning and decision theory). That's where the important work in this area is being done and that's where it gets published.

For Wein to imply that my work, if up to snuff, would have been placed elsewhere is therefore unfair and misleading. This is evident in the literature he cites to confute me. For instance, he cites with approval Colin Howson and Peter Urbach's *Scientific Reasoning: The Bayesian Approach*:

An excellent comparison of the various approaches to statistics can be found in Howson and Urbach's *Scientific Reasoning: The Bayesian Approach*. This book is strongly recommended to any reader who wishes to understand the issue in detail. The clarity of exposition makes a refreshing antidote to Dembski's muddled thinking.

And who are Howson and Urbach? They are both readers in philosophy at the London School of Economics. Moreover, their book likewise is published in a philosophy of science series and not a statistics series.

What about NFL? According to Wein, "Much of the material in *No Free Lunch*, including the application of Dembski's methods to biology, did not appear in *The Design Inference*, and so has received no review at all." Thus in particular, my claim to have contributed to the scientific understanding of the concept of information is said to be called into question. There is an obvious *tu quoque* here: presumably Wein wants his critique of my work to stand on its own merits and not depend on the process by which it was produced and disseminated. Wein has circulated his critique to trusted colleagues and feels confident that it is up to snuff. I did likewise with NFL. Neither work is peer-reviewed in Wein's preferred sense.

But there is another reason why I didn't go through the review process with NFL that I did with TDI. While I was still writing NFL, I contacted Cambridge University Press about publishing it as a sequel to TDI. In particular, I wanted on the basis of a prospectus and some sample chapters of NFL to obtain a contract for this book. I wanted this so that I didn't have to wait almost two and a half years between the time I submitted the completed manuscript and its publication as in the case of TDI. I was informed by the New York editor at CUP who oversees the production of philosophy books that even though TDI was one of their bestsellers, it was controversial, and even though the press didn't mind controversy as such, it had come to light that I was being labeled a "creationist." Thus, before CUP could issue a contract, I would have to submit the most controversial chapters of the new book. Besides this, I had some inside information that even if NFL was accepted this side of the Atlantic, it was unlikely to be accepted with the Cambridge Syndicate in England, whose several biologists were now disposed against my work.

I therefore decided to take my business elsewhere. Indeed, I decided to forego the research monograph route entirely. Given what happened with CUP, I saw the review process as working against me and not for me. My work was being widely discussed and I wanted to get NFL out in a timely fashion and not get bogged down in a biased review process. What's more, I wanted a

book priced lower than most research monographs (TDI retails for \$70; NFL retails for \$35 and can be purchased at Barnes and Noble for \$28). Rowman and Littlefield have treated me well, and I have no regrets about publishing NFL with them.

One final word about peer-review. As Frank Tipler pointed out to me, the idea of peer-review as the touchstone for truth and scientific merit is actually a post Second World War invention. In physics, peer-reviewed journals were not the norm until after 1950. In Germany, during the "Beautiful Years" -- the period when quantum mechanics was being invented in the 1920s -- one of the leading German physics journals, *Zeitschrift fur Physik*, was not peer-reviewed: any member of the German Physical Society could publish there by simply submitting the paper. So, if you had a really wild idea, all you had to do to get it published was ask a member of the GPS to submit it for you. (If you were a member, you could of course submit for yourself.) Heisenberg published his paper on the Uncertainty Principle in this journal, and Friedmann published his paper on the Friedmann universe (now the standard cosmological model) in this journal. No peer-review. Lots of brilliant physics.

Intelligent design has yet to prove itself in this way. But lack of peer-review has never barred the emergence of good science. Nor for that matter have journal articles been the sole place where groundbreaking scientific work was done. Copernicus's *De Revolutionibus*, Galileo's *On Two World Systems*, and Newton's *Principia* are cases in point. And then there's that book by Darwin whose title at the moment escapes me.

3. Argument from Ignorance

Wein's main criticism is that using specified complexity as a criterion to detect design introduces a "middleman" that serves merely to conceal what really is an argument from ignorance. Specified complexity merely eliminates all chance hypotheses that we can think of. It doesn't eliminate all of them, much less warrant a design inference.

By contrast, I contend that specified complexity be a reliable criterion for detecting design in the sense that it does not give rise to false positives (i.e., attributions of design that end up later having to be overturned). The justification for the criterion's reliability is an inductive generalization: In every instance where the complexity-specification criterion attributes design and where the underlying causal story is known (i.e., where we are not just dealing with circumstantial evidence, but where, as it were, the video camera is running and any putative designer would be caught red-handed), it turns out design actually is present; therefore, design actually is present whenever the complexity-specification criterion attributes design.

Although the complexity-specification criterion is inductively sound, its applicability depends on our ability to characterize the relevant probabilities. The "complexity" in "specified complexity" refers to the probability of an event induced by a pattern (more precisely, a family of probability values, one for each relevant chance hypothesis). Think of a target in an archery contest -- the smaller the target, the smaller the probability of hitting it by chance. Thus, for something to exhibit specified complexity, it must match a pattern (target) whose corresponding event (hitting

the target with an arrow) has small probability (less than the universal probability bound of 1 in 10^{150}).

Now, for specified complexity to eliminate chance and detect design, it is not enough that the probability be small with respect to some arbitrarily chosen probability distribution. Rather, it must be small with respect to every relevant probability distribution that might characterize the chance occurrence of the event in question. The use of chance here is very broad and includes anything that can be captured mathematically by a stochastic process. It thus includes deterministic processes whose probabilities all collapse to zero and one. It also includes nondeterministic processes, like evolutionary processes that combine random variation and natural selection. Indeed, chance so construed characterizes all material mechanisms.

But that still leaves the problem of dispensing with all relevant probability distributions. Can this be done with confidence? If by relevant probability distributions, we mean all probability distributions induced by known material mechanisms operating in known ways, then specified complexity can dispense with them and indeed must dispense with them (specified complexity by definition eliminates all probability distributions induced by known mechanisms operating in known ways). But if by relevant probability distributions we mean all probability distributions induced by any material mechanisms that might be operating, including those that are unknown, how can specified complexity dispense with them?

Specified complexity can dispense with unknown material mechanisms provided there are independent reasons for thinking that explanations based on known material mechanisms will never be overturned by yet-to-be-identified unknown mechanisms. Such independent reasons typically take the form of arguments from contingency that invoke numerous degrees of freedom. Sometimes they take the form of arguments from exhaustion: after repeatedly trying ways to achieve a result (say, transforming lead into gold), researchers become convinced that a result can't be achieved -- period. Often additional theoretical grounds help strengthen the argument from exhaustion. Alchemy, for instance, was largely discarded before chemistry provided solid theoretical grounds for its rejection. But the rise of modern chemistry put paid to alchemy.

In any event, we need to get a solid handle on the probability distributions that might apply in a given instance. How we do that must be assessed on a case-by-case basis. I pointed that out in TDI, where I contrasted the outcome of an agricultural experiment with the opening of a combination lock. With the agricultural experiment, we simply didn't have enough knowledge about the underlying mechanisms that might be operating to render one fertilizer better than another for a given crop. Combination locks, however, are a different story. Known material mechanisms (in this case, the laws of physics) prescribe two possible motions of the combination lock, namely, clockwise and counterclockwise rotations. Material mechanisms, however, cannot prescribe the exact turns that open the lock. The geometry and symmetry of the lock precludes that material mechanisms can distinguish one combination from another—one is as good as any other from the vantage of material mechanisms.

Combination locks exhibit numerous degrees of freedom in their possible combinations. In fact, it's precisely these degrees of freedom that guarantee the security of the lock. The more degrees

of freedom, the more possible combinations, and the more secure the lock. Material mechanisms are compatible with these degrees of freedom and tell us that each possible combination is physically realizable. But precisely because each possible combination is physically realizable, material mechanisms as such cannot mandate one combination to the exclusion of others. For that, we need initial and boundary conditions that, as it were, "program" the material mechanisms to actualize one combination to the exclusion of others.

More generally, to establish that no material mechanism explains a phenomenon, one typically establishes that it is compatible with the known material mechanisms involved in its production, but that these mechanisms also permit any number of alternatives to it. By being compatible with but not required by the known material mechanisms involved in its production, a phenomenon becomes irreducible not only to the known mechanisms but also to any unknown mechanisms. How so? The reason is that known material mechanisms can tell us conclusively that a phenomenon is contingent and allows full degrees of freedom. Any unknown mechanism would then have to respect that contingency and allow for the degrees of freedom already discovered.

Michael Polanyi described this method for establishing contingency via degrees of freedom in the 1960s. Moreover, he employed this method to argue for the irreducibility of biology to physics and chemistry. The method applies quite generally: the position of Scrabble pieces on a Scrabble board is irreducible to the natural laws governing the motion of Scrabble pieces; the configuration of ink on a sheet of paper is irreducible to the physics and chemistry of paper and ink; the sequencing of DNA bases is irreducible to the bonding affinities between the bases; and so on. By establishing what is possible on the basis of known material mechanisms, this method limits what is impossible on the basis of unknown material mechanisms. Scrabble pieces, for instance, can be sequenced in all possible arrangements, and no unknown material mechanism is capable of precluding some arrangement or for that matter preferring some arrangement without being suitably programmed by boundary conditions that allow at least as many degrees of freedom as the possible arrangements of Scrabble pieces. It's this regress from the output of material mechanisms to their boundary-condition input that shows the inadequacy of material mechanisms to generate specified complexity.

The reliability of specified complexity as a criterion for detecting design is therefore not merely relative to the probability distributions induced by known material mechanisms operating in known ways. In addition there have to be independent reasons why these probability distributions are secure against unknown mechanisms (in chapter 5 of NFL I argue that irreducible complexity provides such independent reasons in the case of biochemical systems). In that case, a phenomenon's specified complexity confirms the incapacity of material mechanisms to account for it. This is not to say the phenomenon is inexplicable. In cases where the underlying causal history is known, specified complexity does not occur without design. Specified complexity therefore provides inductive support not merely for inexplicability in terms of material mechanisms but also for explicability in terms of design.

When the argument from ignorance objection is raised against specified complexity and more generally against intelligent design, one is apt to think that the ignorance is on the part of design theorists who want to attribute intelligent agency to a phenomenon that, if only the design theorists understood it better, would submit to mechanistic explanation. But the fact is that it is

the scientific community as a whole that is ignorant of how irreducibly complex biological systems, for instance, emerge (molecular biologists James Shapiro and Franklin Harold, though not design theorists, readily concede this point).

But suppose for the sake of argument that intelligence -- one irreducible to material mechanisms -- actually did play a decisive role in the emergence of life's complexity and diversity. How could we know it? Certainly specified complexity will be required. Indeed, if specified complexity is absent or remains an open question, then the door is wide open for material mechanisms to explain the object of investigation. Only if specified complexity is identified does the door to material mechanisms start to close.

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

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

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

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

But natural selection doesn't know a thing about bacterial flagella. So how is natural selection going to take extant protein parts and co-opt them to form a flagellum? The problem is that natural selection can only select for pre-existing function. It can, for instance, select for larger finch beaks when the available nuts are harder to open. Here the finch beak is already in place and natural selection merely enhances its present functionality. Natural selection might even

adapt a pre-existing structure to a new function; for example, it might start with finch beaks adapted to opening nuts and end with beaks adapted to eating insects.

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

The only way for natural selection to form such a structure by co-optation, then, is for natural selection gradually to enfold existing protein parts into evolving structures whose functions co-evolve with the structures. (This is what Wein means when he writes: "The system could have evolved from a simpler system with a *different* function. In that case there could be functional intermediates after all.") We might, for instance, imagine a five-part mousetrap consisting of a platform, spring, hammer, holding bar, and catch evolving as follows: It starts as a doorstep (thus consisting merely of the platform), then evolves into a tie-clip (by attaching the spring and hammer to the platform), and finally becomes a full mousetrap (by also including the holding bar and catch).

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

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

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

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

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

Convinced that the Darwinian mechanism must be capable of doing such evolutionary design work, evolutionary biologists rarely ask whether such a sequence of successful baby-steps even exists; much less do they attempt to quantify the probabilities involved. I attempt that in chapter 5 of NFL (to which I'll return shortly). There I lay out techniques for assessing the probabilistic hurdles that the Darwinian mechanism faces in trying to account for complex biological structures like the bacterial flagellum. The probabilities I calculate -- and I try to be conservative -- are horrendous and render natural selection utterly implausible as a mechanism for generating the flagellum and structures like it.

If I'm right and the probabilities really are horrendous, then the bacterial flagellum exhibits specified complexity. Furthermore, if specified complexity is a reliable marker of intelligent agency, then systems like the bacterial flagellum bespeak intelligent design and are not solely the effect of material mechanisms.

It's here specifically that Wein raises the argument-from-ignorance objection. For something to exhibit specified complexity entails that no known material mechanism operating in known ways is able to account for it. But that leaves unknown material mechanisms. It also leaves known material mechanisms operating in unknown ways. Isn't arguing for design on the basis of specified complexity therefore merely an argument from ignorance?

Two comments to this objection: First, the great promise of Darwinian and other naturalistic accounts of evolution was precisely to show how known material mechanisms operating in known ways could produce all of biological complexity. So at the very least, specified complexity is showing that problems claimed to be solved by naturalistic means have not been solved. Some committed naturalists are actually grateful to have this pointed out. Though convinced that material mechanisms will eventually account for all of biological complexity, they see intelligent design as pressing the biological community to own up to unsolved problems. Paleontologist David Raup is a case in point.

Second, the argument from ignorance objection could in principle be raised for any design inference that employs specified complexity, including those where humans are implicated in constructing artifacts. An unknown material mechanism might explain the origin of the *Mona Lisa* in the Louvre, or the Louvre itself, or Stonehenge, or how two students wrote exactly the same essay. But no one is looking for such mechanisms. It would be madness even to try. Intelligent design caused these objects to exist, and we know that because of their specified complexity.

Specified complexity, by being defined relative to known material mechanisms operating in known ways, might always be defeated by showing that some relevant mechanism was omitted. That's always a possibility (though as with the plagiarism example and with many other cases, we don't take it seriously). As William James put it, there are live possibilities and then again there are bare possibilities. There are many design inferences which, to question or doubt, require invoking a bare possibility. Such bare possibilities, if realized, would defeat specified complexity. But defeat specified complexity in what way? Not by rendering the concept incoherent but by dissolving it.

In fact, that is how Darwinists, complexity theorists, and anyone intent on defeating specified complexity as a marker of intelligence usually attempts it, namely, by showing that it dissolves once we have a better understanding of the underlying material mechanisms that render the object in question reasonably probable. By contrast, design theorists argue that specified complexity in biology is real: that any attempt to palliate the complexities or improbabilities by invoking as yet unknown mechanisms or known mechanisms operating in unknown ways is destined to fail. This can in some cases be argued convincingly (employing a degree-of-freedom argument of the sort sketched earlier), as with Michael Behe's irreducibly complex biochemical machines and with biological structures whose geometry and symmetry allows complete freedom in possible arrangements of parts.

Consider, for instance, a configuration space comprising all possible character sequences from a fixed alphabet (such spaces model not only written texts but also polymers like DNA, RNA, and proteins). Configuration spaces like this are perfectly homogeneous, with one character string geometrically interchangeable with the next. The geometry therefore precludes any underlying mechanisms from distinguishing or preferring some character strings over others. Not material mechanisms but external semantic information (in the case of written texts) or functional information (in the case of polymers) is needed to generate specified complexity in these instances. To argue that this semantic or functional information reduces to material mechanisms is like arguing that Scrabble pieces have inherent in them preferential ways they like to be sequenced. They don't. Michael Polanyi offered such arguments for biological design in the 1960s. Steve Meyer has updated them for the present.

To sum up, evolutionary biology contends that material mechanisms are capable of accounting for all of biological complexity. Yet for biological systems that exhibit specified complexity, these mechanisms provide no explanation of how they were produced. Moreover, in contexts where the causal history is independently verifiable, specified complexity is reliably correlated with intelligence. At a minimum, biology should therefore allow the possibility of design in cases of biological specified complexity. But that's not the case.

Evolutionary biology allows only one line of criticism, namely, to show that a complex specified biological structure could not have evolved via any material mechanism. In other words, so long as some unknown material mechanism might have evolved the structure in question, intelligent design is proscribed. This renders evolutionary theory immune to disconfirmation in principle, because the universe of unknown material mechanisms can never be exhausted. Furthermore, the evolutionist has no burden of evidence. Instead, the burden of evidence is shifted entirely to the evolution skeptic. And what is required of the skeptic? The skeptic must prove nothing less than a universal negative. That is not how science is supposed to work.

Science is supposed to pursue the full range of possible explanations. Evolutionary biology, by limiting itself to material mechanisms, has settled in advance which biological explanations are true apart from any consideration of empirical evidence. This is arm-chair philosophy. Intelligent design may not be correct. But the only way we could discover that is by admitting design as a real possibility, not ruling it out a priori. Darwin himself agreed. In the *Origin of Species* he wrote: "A fair result can be obtained only by fully stating and balancing the facts and arguments on both sides of each question."

4. Tornado in a Junkyard

Richard Wein charges that the techniques I develop for calculating probabilities for irreducibly complex biochemical systems is all akin to calculating probabilities for the random formation of an airplane by a tornado whirling through a junkyard (the metaphor is Fred Hoyle's, if memory serves me correctly). Actually, no one has yet to perform this calculation, so if that's all I've accomplished, it's still a contribution to the applied probability literature. I would say, however, that I've actually accomplished quite a bit more than Wein attributes to me.

Briefly, I propose that the probability of a discrete combinatorial object (e.g., an irreducibly complex biochemical system) decomposes into three probabilities: the probability of the origination of the components of the object, the probability of all the components being localized in one place, and the probability of the components being properly configured once they are localized. These probabilities multiply since each of the three events (origination, localization, and configuration) is conditioned on the preceding event (there is no assumption of probabilistic independence being slipped in here). Also, there is nothing here that requires these probabilities to be uniform probabilities (one of Wein's concerns) or that precludes Darwinian processes from constructing a discrete combinatorial object in accord with these probabilities. A Darwinian process might with reasonable probability gradually localize the components for a discrete combinatorial object and then configure them.

So what's at issue is not the technique of decomposing the probability of a discrete combinatorial object into a set of more manageable probabilities (this technique may properly be regarded as a positive contribution on how to form some reasonable probability estimates for difficult practical problems), but its application to biology and specifically for the numbers we plug in. Now, I did stress in NFL that to settle the design question for such objects when they are irreducibly complex biochemical systems will require that scientists agree on the probabilities and be

convinced that no one is cheating. The issue, then, is whether I was cheating with the numbers. I don't believe I was, and here's why.

First off, let's be clear that Wein is blowing smoke when he claims a system like the bacterial flagellum evolved by co-optation so that natural selection gradually enfolded parts as functions co-evolve. As I pointed out in the last section, this sounds all fine and well, but Wein offers no detailed testable model for how this might actually happen in specific cases like the flagellum. We've already discussed the type III secretory system and its inadequacy as an evolutionary precursor to the flagellum. But even that supposed precursor is now thought to have arrived after the flagellum (see Nguyen L, Paulsen IT, Tchieu J, Hueck CJ, Saier MH Jr., Phylogenetic analyses of the constituents of Type III protein secretion systems, *J Mol Microbiol Biotechnol* 2000 Apr 2(2):125-44). Franklin Harold in *The Way of the Cell* (Oxford, 2001) refers to all nontelic scenarios currently proposed to explain such systems as "wishful speculations." The biological community doesn't have a clue how such systems emerge.

Next, all the biological community has to mitigate the otherwise vast improbabilities for the formation of such systems is co-optation via natural selection gradually enfolding parts as functions co-evolve. Anything other than this is going to involve saltation and therefore calculating a probability of the emergence of a multipart system by random combination. But, as Wein rightly notes, "the probability of appearance by random combination is so minuscule that this is unsatisfying as a scientific explanation." Wein therefore does not dispute my calculation of appearance by random combination, but the relevance of that calculation to systems like the flagellum. And why does he think it irrelevant? Because co-optation is supposed to be able to do it.

What we have, then, is a known material mechanism (co-optation powered by natural selection) that by unknown steps is supposed to produce a bacterial flagellum. That's the only live mechanistic option. Now, why should we believe it? What I offer in chapter 5 of NFL are reasons not to believe it. I tighten Michael Behe's notion of irreducible complexity to include a minimality condition so that no system with the same function of lower complexity can perform the same function. That was a problem with Behe's original definition because it allowed for John H. McDonald-type mousetrap examples. In such examples, however, the function in question stays put. What Wein wants to argue, however, is that irreducible complexity can never preclude complex systems arising from precursor systems whose function do not stay put but co-evolve with the systems.

I submit that there is no live possibility here but only the illusion of possibility. Wein tacitly agrees that I've correctly calculated the probability of the emergence of the flagellum apart from co-optation. Does Wein have a probability calculation that breaks the emergence of the flagellum into manageable steps each of which is reasonably probable and for each of which he actually calculates the probability? No. Yet Wein is confident that my probability calculation is in error and that the actual probability is much higher even though he cannot offer even the semblance of such a calculation. On what basis then does Wein maintain this confidence. Is he prepared to offer a Saganesque type argument of the sort "the flagellum is here so it can't have been all that improbable"?

For Wein to account for systems like the flagellum, functions of precursor systems must co-evolve. But that means the space of possible functions from which these co-evolving functions are drawn is completely unconstrained. This provides yet another recipe for insulating Darwinian theory from critique, for the space of all possible biological functions is vast and there is no way to establish the universal negative that no sequence of co-evolving functions could under co-optation have led to a given system.

Are we then at an impasse, with both Wein and me saying essentially "prove me wrong" (in my case, prove to me that my calculation doesn't hold by doing a probability calculation of your own; in Wein's case, prove to me that co-optation with co-evolving functions didn't happen)? Let me suggest that there are further reasons to be deeply skeptical of Wein's co-optation scenario. First, specified complexity is used to nail down design in cases of circumstantial evidence, so if there should happen to be design in nature, specified complexity is how we would detect it. Thus, my probability calculation for the flagellum, in the absence of a counter-calculation by Wein, is *prima facie* evidence of biological design. This may not provide sufficient reason for convinced Darwinists to abandon their paradigm, but it gives evolution skeptics reason to consider other options, including design.

Second, there is a whole field of study developed by Russian scientists and engineers known under the acronym TRIZ (Theory of Inventive Problem Solving) that details patterns of technological evolution (see, for instance, <http://www.ideationtriz.com>). As it turns out, human problem solving breaks into two types, routine problems and inventive problems. Routine problems are known to submit to Darwinian type trial-and-error solutions. Inventive problems, by contrast, are known to require an intuitive leap. This distinction has turned out to be very robust, with inventive problems not submitting to the solution strategies that work for routine problems. Since co-optation is itself an engineering metaphor, TRIZ research provides strong reason to think that systems like the flagellum, that are elegant and highly integrated, do not admit a gradual routinized decomposition of the sort required by the Darwinian mechanism. Wein disagrees and sees technological evolution in a different category entirely, and even mentions that human engineers will often engineer from the ground up whereas biological evolution won't. But in fact, engineers are stuck with existing resources as much as biology. Wein gives no evidence of knowing the TRIZ literature. If he did know it, and if he were not so biased a critic, he would perhaps more readily grant the connection between technological and biological evolution.

Third, and perhaps most telling, Wein needs fitness to vary continuously with the topology of configuration space. Small changes in configuration space need to correlate with small changes in biological function, at least some of the time. If functions are extremely isolated in the sense that small departures from a functional island in configuration space lead to complete nonfunctionality, then there is no way to evolve into or out of those islands of functionality by Darwinian means. To clear away this obstacle to the Darwinian mechanism, Wein argues that the laws of nature guarantee the continuity that permits the Darwinian mechanism to flourish. According to Wein, smooth fitness landscapes are the norm because we live in a world with regular laws of nature and these are supposed to ensure smoothness. Here is how Wein puts it:

Fitness functions are determined by rules, not generated randomly. In the real world, these rules are the physical laws of the Universe. In a computer model, they can be whatever rules the programmer chooses, but, if the model is a simulation of reality, they will be based to some degree on real physical laws. Rules inevitably give rise to patterns, so that patterned fitness functions will be favoured over totally chaotic ones. If the rules are reasonably regular, we would expect the fitness landscape to be reasonably smooth. In fact, physical laws generally are regular, in the sense that they correspond to continuous mathematical functions, like " $F = ma$ ", " $E = mc^2$ ", etc. With these functions, a small change of input leads to a small change of output. So, when fitness is determined by a combination of such laws, it's reasonable to expect that a small movement in the phase space will generally lead to a reasonably small change in the fitness value, i.e. that the fitness landscape will be smooth. On the other hand, we expect there to be exceptions, because chaos theory and catastrophe theory tell us that even smooth laws can give rise to discontinuities. But real phase spaces have many dimensions. If movement in some dimensions is blocked by discontinuities, there may still be smooth contours in other dimensions. While many potential mutations are catastrophic, many others are not.

This argument is rubbish. It invokes continuity of physical laws to underwrite continuity of fitness landscapes when the two notions of continuity are entirely different. In physical laws like $F = ma$, if you vary the quantity m slightly, the quantity F varies slightly as well. In general with physical laws, continuity is a matter of coordinating different quantities associated with physical entities. But that's not at all what's going on with fitness landscapes. Fitness landscapes coordinate a physical quantity, position in configuration space, with a teleological property, a function belonging not to a physical space but to a space of functions characterized by a functional logic. Fitness assigns a quantity to that function, but that quantity is entirely derivative from the function.

This is not to say we can't speak of continuity or smoothness of fitness landscapes. But it is to say that we can't use continuity of physical laws to underwrite it. Indeed, we have plenty of examples where there is no such continuity. Consider the case of written language and meaning. Written language is represented in character strings and there is a natural way to think of written texts as being close to each other in terms divergence of corresponding characters (the more divergence, the farther apart). How much can you randomly perturb written texts and still maintain their meaning (i.e., function)? Is it possible by small changes to the text always to maintain some meaning and thereby transform a poem by Keats into one by Ginsburg? What about computer source codes? Can gradual random changes to the code change a computer game into an accounting program all the while maintaining some function of the source code? Perhaps there is, but we have no evidence of anything like this, nor do human designers attempt anything like this. To think that continuity of physical laws underwrites continuity of fitness landscapes is without justification and another case of Darwinians taking seriously what is merely an illusion of possibility.

Finally, for all the length of Wein's critique, he ignored what I expect will be my most important contribution to assessing the probabilities of biological systems, namely, my definition of

perturbation probabilities in terms of perturbation tolerance and identify factors and their use in assessing the probability of evolving individual polypeptides and polynucleotides. Wein writes: "Since he can't calculate this directly, he uses an approximation that he calls a *perturbation probability*. We need not concern ourselves with the details." Perhaps the details are not important, but the application is. Wein contends that the only biological system to which I apply my probabilistic techniques is the bacterial flagellum. In fact, I have my eye on much different fish.

In section 5.10 of NFL, I indicated how perturbation probabilities apply to individual enzymes and how experimental evidence promises shortly to nail down the improbabilities of these systems. The beauty of work being done by ID theorists on these systems is that they are much more tractable than multiprotein molecular machines. What's more, preliminary findings of this research indicate that islands of functionality are not only extremely isolated but completely surrounded by a sea of nonfunctionality (not merely polypeptides having *different* functions but polypeptides *incapable* of function on thermodynamic grounds -- in particular, they can't fold). For such extremely isolated islands of functionality, there is no way for Wein's method of co-evolving functions to work.

Prediction: Within the next two years work on certain enzymes will demonstrate overwhelmingly that they are extremely isolated functionally, making it effectively impossible for Darwinian and other gradualistic pathways to evolve into or out of them. This will provide convincing evidence for specified complexity as a principled way to detect design and not merely as a cloak for ignorance.

5. Short Responses

I've now dealt with the substance of Wein's objections to my work. I want next in bullet-point fashion to deal with some remaining objections.

5.1 Uniform Probabilities

Wein complains that at times I focus on uniform probabilities but that at other times I consider a broader set of probabilities. The design inference requires sweeping the field of all relevant probabilities, so all of these must be considered (though Wein and I differ on what the relevant set is -- see the preceding sections). In many contexts the uniform probability is not nearly enough. As a criterion for detecting or inferring design, specified complexity must be assessed with respect to all the relevant probability distributions. Consequently, for complexity to obtain, the probability of the specified event in question must be less than the universal probability bound with respect to all the probability distributions (i.e., relevant chance hypotheses) being considered. (Note that this means the formula in the Law of Conservation of Information, $I(A\&B) = I(A) \bmod UCB$, needs to obtain for every relevant probability distribution \mathbf{P} , which gets converted to an information measure \mathbf{I} by a logarithmic transformation.)

Even so, to a first approximation, the uniform probability is often a good place to start. With polypeptides and polynucleotides, for instance, bonding affinities between amino acids or between nucleotide bases show very little preference, so from the vantage of physical laws one sequence is equivalent to any other. To be sure, it's the functional difference that is biologically significant. The question therefore is whether biology has the mechanisms to account for the evolution of such functional systems that yield probabilities different from those indicated by these systems' geometry and physics. I contend that in some instances there is no reason to take such mechanisms seriously.

There's also the question of burden of proof in probability arguments. Darwinists think that their mechanism has been "overwhelmingly confirmed," so that once natural selection is brought in, Mount Improbable can always be climbed gradually. But the presumption that probabilities can always be dissolved in this way can be turned around. Why not instead presume that the probabilities are small until a gradual Darwinian pathway up Mount Improbable is actually demonstrated (and not merely gestured at with a just-so story)? In many contexts the uniform probability distribution is a good way to see whether we're in a small probability ball park. For instance, I regularly cite the *Contact* example in which a long sequence of prime numbers represented as a bit string comes from outer space and convinces SETI researchers of the reality of extraterrestrial intelligence. What probabilities might reasonably be assigned to that sequence? What are the relevant chance hypotheses that might assign these probabilities? It's not simply going to be a uniform probability (1s vastly outnumber 0s in representing that sequence). Yet the uniform probability is much more in the right ball park than a probability that concentrates high probability on this sequence.

One clarification about uniform probabilities in the context of evolutionary computation is worth making here. When I refer to specified complexity in relation to evolutionary computation, I usually intend the complexity to correspond to the probability of the target induced by the natural topology on the configuration/search space in question. Often this is a uniform probability, but it needn't be. Placing a fitness landscape over the search space now induces a new probability distribution for the target relative to an evolutionary algorithm. This probability can be translated into a waiting time for how long (how many candidates) on average need to be checked by the evolutionary algorithm before landing in the target.

With respect to the natural probability measure (often a uniform probability) on the space, the target represents an instance of specified complexity. With respect to the new probability induced by the fitness function, however, it may not. My point in invoking the No Free Lunch theorems is that the search for the target in the original search space gets displaced to searching the space of fitness functions over that space, thus intensifying the search problem rather than resolving it. Wein disputes the relevance of this point to biology, but is left with arguing that continuity as exhibited in physical laws implies continuity/smoothness of fitness landscapes, and, as we saw in the last section, this argument is bogus.

5.2 The No Free Lunch Theorems

According to Wein the No Free Lunch theorems are irrelevant to biology. Let me indicate why I think they are relevant. The No Free Lunch theorems deal with certain families of fitness landscapes and state that when averaged with respect to those families, evolutionary algorithms do not outperform blind search. What sorts of families? Those that don't privilege any targets in configuration space over any others. The collection of all possible fitness landscapes on a configuration space fits the bill. Various subcollections do also. The details are not important. What is important is that for an evolutionary algorithm to find a target successfully, it must employ the right fitness landscape.

Regardless whether the No Free Lunch theorems are significant for biological evolution, they certainly are significant for evolutionary computation. Among other things, these theorems show that the fitness landscape does not merely represent the problem to be solved but actually constitutes its solution. As Geoffrey Miller remarks in John Ziman's (ed.) *Technological Innovation as an Evolutionary Process* (Cambridge University Press, 2000): "The fitness function [landscape] must embody not only the engineer's conscious goals, but also her common sense. This common sense is largely intuitive and unconscious, so is hard to formalize into an explicit fitness function. Since genetic algorithm solutions are only as good as the fitness functions used to evolve them, careful development of appropriate fitness functions embodying all relevant design constraints, trade-offs and criteria is a key step in evolutionary engineering."

What relevance does this have for biology? Minimally, it means that the fitness landscapes by which the Darwinian process is carried out in nature have to be smooth if the Darwinian process is going to be successful. Much of my argument until now has focused on why we should think that in fact they are not smooth (i.e., why we should think that there are islands of functionality in biological configuration space from which no gradual paths that maintain functional advantage either enter or exit) -- Michael Behe's irreducibly complex biochemical systems being a case in point. But what if they are smooth? Does that indicate design in the choice of fitness landscapes that operate in nature?

Wein thinks not. But to argue that he employs a bogus argument saying that continuity in the laws of nature underwrites continuity/smoothness of fitness landscapes in biology. We've seen that this argument doesn't work. Wein's other tack is to say that looking for design in the fitness landscapes merely reduces to a standard cosmological fine-tuning argument -- the laws of physics must be fine-tuned to allow for stable configurations of matter from which to construct biological systems. But there's more than Wein lets on. Given, for the sake of argument, that Darwinism is the means by which biological complexity emerged, why is nature such that the fitness landscapes operating in biology are smooth and allow the Darwinian mechanism to be successful? This question, under the supposition of Darwinism's truth, requires an answer, and it is not enough to say that nature has given us a free lunch.

We can imagine nature being rule governed, fine-tuned, allowing all sorts of interesting physics, but not having any smooth Darwinian pathways connecting complex replicating systems. What's more, we know of rule governed, fine-tuned systems where the disconnect between physical embodiment and functional logic is extremely sharp (e.g., natural language and computer source

codes). Why is ours a world where the Darwinian mechanism works (if indeed it works)? In NFL I contend that design looms here as well. Granted, I don't think this is the strongest place to argue for biological design. The strongest place is to argue head-on that the fitness landscapes are not smooth. But design at the level of fitness landscapes (under the supposition that Darwinism is correct) remains to my mind a powerful argument in light of all the cases we know where the Darwinian mechanism operates but fails to produce anything interesting.

5.3 Specification

Wein remains unhappy with my account of specification. Even granting Fisher's theory of significance testing, Wein does not think that my design inference makes good on generalizing Fisher's theory. The problem, according to Wein, is with my generalization of Fisher's rejection regions in terms of specifications. Fisher's rejection regions are identified prior to an experiment, and thus are objectively identified without any tailoring of rejection regions to events. But with specifications that are formulated after events occur, the subject doing the specifying first observes the event and then formulates the specification on the basis of the event. How can this be done without illegitimately tailoring a pattern to an event, or alternatively without reading onto the event a pattern that merely strikes one's fancy? Wein considers a dice tossing example and summarizes his concern as follows: "I suggest that, in this example, no matter what outcome we observed, Dembski's approach could be used to justify a rejection region consisting of just the specific observed outcome."

As is so often the case with Wein's critique, he conveniently omits the key that unlocks the door. Nowhere in his review does he mention the tractability condition, and yet it is precisely that condition that circumvents his worry of artificially tailoring patterns to events. He's right to stress this concern, but I stress it too in NFL and I'm at pains to show how this concern can be met. The way I meet it is through what I call the tractability condition (why no mention in Wein's critique?). Essentially two things are going on with specified complexity. There is high improbability or probabilistic complexity of the event signified by a specification. On the other hand, there is low "patterned" complexity of the specification taken simply as a pattern. In this way specified complexity combines order (low patterned complexity) with chaos (high probabilistic complexity), which sits nicely with the Whiteheadian notion of novelty and the complexity theorist's notion of the edge of chaos.

To repeat, what one needs for specified complexity (among other things) is a combination of high probabilistic complexity or improbability of the event associated with the specification but low patterned complexity of the specification treated simply as a pattern. But how does one make sense of that patterned complexity? Think of it this way. Imagine a dictionary of 100,000 (= 10^5) basic concepts. There are then 10^5 1-level concepts, 10^{10} 2-level concepts, 10^{15} 3-level concepts, and so on. If "outboard," "rotary," "propeller," and "motor" are basic concepts, then the bacterial flagellum can be characterized as a 4-level concept of the form "ouboard propeller-driven rotary motor." Now, there are about 10^{20} concepts of level 4 or less. Thus for concepts that specify an event of probability less than the universal probability bound of 1 in 10^{150} , all concepts of level 4 or less, even when their probabilities are combined, will have minuscule probability.

That's the point of the tractability condition, namely, to balance the patterned complexity of a specification with the improbability of the event signified by the specification. By factoring in specificational resources whose patterned complexity is no greater than the specification in question, one avoids the problem of artificially tailoring a specification to an event. The tractability condition says, as it were, "Okay, you've witnessed the event already; now let's make sure the pattern you see in the event is really there."

To see that the tractability condition is just what the doctor ordered, let's consider a case where it blocks a design inference. Consider a long sequence of coin tosses of length 1000 (the improbability is 1 in 10^{300} , well below the universal probability bound). Now let's set up a 1000-level concept. We imagine that our basic concepts include the numbers 0 and 1. Let our 1000-level concept be the sequence of 0s and 1s in the sequence of coin tosses we observed (1 for heads, 0 for tails). In that case, there are on the order of 10^{5000} concepts of level 1000 or less. That number corresponds to the specificational resources we would need to exhaust before we could effectively draw a design inference for this sequence. The level of improbability at which such an inference could be drawn would therefore have to be on the order of 1 in 10^{5000} . But this won't work because bit strings of length 1000 each have probability only around 1 in 10^{300} , which is far greater.

In fact, the 1000-level concept considered here is problematic not just for the vast number of specificational resources it requires, but also because the pattern it denotes is not conditionally independent from the coin-tossing event we observed. Consequently, it is not even a specification. It's the combination of the conditional independence condition and the tractability condition that, in light of the physical limits of the observable universe, enables us to set the universal probability bound at 1 in 10^{150} (see chapter 6 of TDI).

The short of it is that Wein, by ignoring the tractability condition, has missed precisely the key insight that allows my project to succeed.

5.4 Algorithmic Information Theory

Wein attributes confusion to me about algorithmic information theory, but the confusion is his. The reason Wein is confused is because within algorithmic information theory, the highly incompressible bit strings are the ones with "high complexity." Thus when I attribute specified complexity to highly compressible bit strings, he thinks I'm confused. But that's because he is reading in his own preconceptions rather than applying the framework I lay out in NFL.

Let me therefore try one more go at it. Specified complexity is a property of patterns that signify events. For something to exhibit specified complexity, it must conform to a specification that signifies an event that has small probability (i.e., is probabilistically complex) but also is simple as far as patterns go (i.e., has low patterned complexity). In algorithmic information theory, we think of the bit strings as coin tosses and therefore as elementary events or outcomes. Specified complexity therefore does not apply to any of these outcomes individually but rather to collections of such outcomes (composite events) that exhibit a suitable pattern, namely, a

specification where the signified event has high improbability (i.e., high probabilistic complexity) and low patterned complexity. Specifications with low patterned complexity in this case include bit strings that are highly compressible as well as bit strings that are highly incompressible. But these specifications are very different in terms of the probabilities of the events they describe.

It is a combinatorial fact that most bit strings are highly incompressible. As a consequence, the specification of highly incompressible bit strings does not signify a high improbability event and thus cannot exhibit specified complexity. On the other hand, bit strings that are highly compressible constitute (on combinatorial grounds) but a minuscule fraction of all bit strings. As a consequence, the specification of highly compressible bit strings does signify a high improbability event and thus exhibits specified complexity. We thus find that highly compressible strings taken collectively exhibit specified complexity even though taken individually they exhibit low algorithmic information and thus low complexity in the Kolmogorov-Chaitin sense. That may seem counterintuitive, but actually it makes good intuitive sense. It is the highly compressible strings that are nonrandom and that lead to design inferences, not the highly incompressible ones.

5.5 Predictive Power of Darwinian Evolution

Citing talk.origins documents, Wein claims: "Evolutionary theory certainly does make many predictions all the same. For example, evolutionary theory predicts that there will be a high degree of congruence between phylogenetic trees derived from morphological studies and from independent molecular studies. This prediction has been confirmed, and continues to be confirmed as more species are tested."

This is nonsense. Wein needs to consult the primary literature. Let me suggest Simon Conway Morris's review article in *Cell* (100 [Jan. 7, 2000]:1-11) titled "Evolution: Bringing Molecules into the Fold." Conway Morris writes: "Constructing phylogenies is central to the evolutionary enterprise, yet rival schemes are often strongly contradictory. Can we really recover the true history of life?" Among the rival schemes are morphological and molecular studies.

5.6 Explanatory Power

Criticizing my account of explanatory power, Wein writes: "The term *explanatory power* is widely used but difficult to define. I will not attempt a definition, but will note that, in part, it is another face of *predictive power*, referring to the ability to 'retrodict' past observations."

Actually, the term explanatory power has a well-defined sense in the philosophy of science and epistemology literature, which Wein would do well to consult. Consilience rather than retrodiction is the key. Here's a standard definition of explanatory power: "Explanations are also sometimes taken to be more plausible the more explanatory 'power' they have. This power is usually defined in terms of the number of things or more likely, the number of kinds of things, the theory can explain. Thus Newtonian mechanics was so attractive, the argument goes, partly

because of the range of phenomena the theory could explain." (Jonathan Dancy and Ernest Sosa, *A Companion to Epistemology* [Oxford: Blackwell, 1992], p. 208, s.v. "Inference to the Best Explanation.")

5.7 Wein's Acknowledgment

Wein offers the following acknowledgment for help on his critique of NFL: "I am grateful for the assistance of Wesley Elsberry, Jeffrey Shallit, Erik Tellgren and others who have shared their ideas with me." Am I to assume that Wein speaks for Elsberry, Shallit, Tellgren, and others and that in his critique of my work he is relating their best arguments as well as his own? Or can they do better? For Darwinism's continued health and vigor, let's hope others can do better.