

If Only Darwinists Scrutinized Their Own Work as Closely: A Response to "Erik"

By William A. Dembski

An Internet persona known as "Erik" reviewed those aspects of my book *No Free Lunch* dealing with the Law of Conservation of Information and specificational resources. Erik's review is titled "On Dembski's Law of Conservation of Information" and is available at http://www.talkreason.org/articles/dembski_LCI.pdf. I respond to the review here.

In the preface to *No Free Lunch*, I remarked that my strategy in writing this book was "to include just enough technical discussion so that experts can fill in the details as well as sufficient elaboration of the technical discussion so that nonexperts feel the force of the design inference." This approach seems not to have worked for Erik. When he fills in the technical details, he frequently botches them. And where my nontechnical discussion could have prevented his technical elaborations from going awry, he ignores them or, worse yet, attributes to me oversights that I precluded in those self-same nontechnical discussions.

Erik's review purports to fill in the technical details of chapters 2 and 3 of *No Free Lunch*. He also uses the review as a vehicle for numerous side-swipes at me and my project more generally. In this response I take up both his technical criticisms as well as the side-swipes. I shall, however, keep this response in plain English (save for one brief portion). Erik's technical elaborations, which are even more notation-heavy than the most notation-heavy parts of *No Free Lunch*, though readily comprehensible by mathematicians, will be lost on most readers. Let's get started.

Erik's first paragraph sets the tone by casting intelligent design as a usurper that tries to take the easy way out instead of doing real science. Thus he writes: "One of the characteristics of the [intelligent design] movement is a lack of interest in providing descriptive explanations and models. Instead, the ID movement sets out to establish the involvement of an intelligent designer by disproving the currently available alternatives that do not require one."

Erik offers these comments in the context of abiogenesis and evolutionary biology. But what, pray tell, are these "currently available alternatives" that purport to account for the origin and evolution of biological complexity apart from design? In fact, all evolutionists have done is describe supposedly possible mechanisms, in highly abstract and schematic terms, to which, in the case of Darwinism, no significant details have been added since the time of Darwin (and, I would urge, none has been added even since the time of Empedocles and Epicurus), and for which other naturalistic evolutionary scenarios remain even more speculative.

Critics of evolution who say it is merely a theory don't go far enough -- it doesn't even deserve to be called a theory. No Darwinist, for instance, has offered a hypothetical Darwinian production of any tightly integrated multi-part "adaptation" with enough

specificity to make the hypothesis testable (even in principle). Evolutionary biology isn't a theory -- it's a pile of promissory notes for future theories, none of which has been redeemed since the publication of Darwin's *Origin of Species* almost 150 years ago.

Let me make the same point more kindly and gently. Material mechanisms known to date offer no insight into biological complexity. Cell biologist Franklin Harold in *The Way of the Cell* (Oxford, 2001) remarks that in trying to account for biological complexity, biologists thus far have merely proposed "a variety of wishful speculations." If biologists really understood the emergence of biological complexity in purely material terms, intelligent design couldn't even get off the ground. The fact that they don't accounts for intelligent design's quick rise in public consciousness. Show us detailed, testable, mechanistic models for the origin of life, the origin of the genetic code, the origin of ubiquitous biomacromolecules and assemblages like the ribosome, and the origin of molecular machines like the bacterial flagellum, and intelligent design will die a quick death.

The mechanisms of evolutionary biology fail to specify detailed testable mechanistic pathways capable of bringing about tightly integrated multi-part complex functional biological systems. In other words, evolutionary biology trades in unspecified mechanistic causes -- indeed, that's the only currency evolutionary biology seems to know. The irony here is lost on Erik. Though Erik has no problem with unspecified mechanistic causes that gesture at (to say "account for" is far too generous) biological complexity, he objects to intelligent design introducing an "unspecified designer."

Two comments are in order here: First, scientific explanations need to be causally adequate; in other words, they need causes with sufficient power to account for the things we are trying to explain. We know that designing intelligences have the causal power to produce tightly integrated multi-part functional systems (like machines). We have no experience of undirected material processes doing the same. Thus, in introducing an "unspecified designer," intelligent design is at least identifying a cause sufficient to produce the effect in question. To be sure, intelligent design must not stop here. But it certainly must not be prevented from getting here.

Second, we can legitimately infer design even if we know no details about the designer. As Del Ratzsch points out in *Nature, Design, and Science* (SUNY Philosophy of Biology Series, 2001), if we found a bulldozer on one of Jupiter's moons, a design inference would be warranted (indeed mandated) even if we had no idea who the designer was or how the designer put the bulldozer there. Are the designers of Stonehenge specified? We know (or presume) that they were human, but we don't know much beyond that. To make a fundamental distinction between human, animal, and extraterrestrial intelligences on the one hand and unembodied intelligences on the other hand and to banish the latter for being "unspecified designers" is merely to presuppose the answer to biology's design question. This is circular reasoning, and it is endemic to evolutionary biology.

I was amused to see Erik write, "Dembski's writings are notoriously difficult to interpret." My critics and fans seem to take exactly opposite positions on every aspect of my work.

Here is a quote from a letter to me (now over a decade old) by one of the United States' top statisticians (member of the National Academy of Sciences, holding faculty positions over the years at Stanford, Harvard, and Cornell): "I'm delighted to be in touch and delighted that you're working in foundations [of probability]. I've just received your papers.... My first reaction is 'Wow, he writes well'. This is really nice to see." Indeed, most people who comment on my writing, fans and even some foes, think that I'm an exceptionally clear writer.

But it serves Erik to cast me as an obscure writer. This allows him glibly to assert: "In the cases where Dembski has responded to criticism, the response has always been that the critic has misunderstood him in some fatal way." Actually, I'd like to see some places where I've responded to critics and explicitly charged them with "misunderstanding my work" (I'm sure there are a few but not many). I frankly don't see the problem as critics misunderstanding what I'm saying but rather as critics understanding what I'm saying all too well and desperately trying to prove me wrong by distorting and redefining what I'm saying.

Erik inadvertently demonstrates this problem: "My own interpretations probably differ a little from those of other critics, but I cannot guarantee that Dembski will agree with my presentation of his ideas." Indeed, Erik's presentation of my ideas is a "re-presentation" that distorts my arguments and thereby renders them invalid. Call it misunderstanding, distortion, or whatever, Erik does have a point that my critics typically commit some fatal mistakes. But in responding to my critics, I don't leave it there -- I show how the fatal mistakes arise. Let's therefore turn to Erik's fatal mistakes.

Erik's biggest mistake, and one responsible for many of his more particular mistakes, is trying to cast my project as one of pure mathematics. It's not and never was intended to be. That's why I write about making an "in-principle mathematical argument" about the inability of material mechanisms to generate specified complexity rather than about a "strict mathematical proof." That's also why Erik constantly has to use scare quotes and qualifiers (he particularly likes the mitigating prefix "quasi-") in describing the mathematical aspects of my work. Thus Erik characterizes my project in *No Free Lunch* as providing "a subjective quasi-mathematical 'formalization'" and then applying "this 'formalization'" to biology (note the scare quotes around "formalization").

I'm not and never have been in the business of offering a strict mathematical proof for the inability of material mechanisms to generate specified complexity in the same way that no physicist is in the business of offering a strict mathematical proof for the conservation of energy. Mathematics certainly comes into the picture in both instances and is crucial in justifying these claims, but there are empirical and nonmathematical considerations that come into play as well and that make strict mathematical proof not feasible (and perhaps not even desirable).

Thus on page 3 of his review Erik will describe a chance hypothesis as inducing a probability not merely on the space of events in question but also on a space consisting of "all background knowledge," which he leaves mathematically undefined except for

assigning it the letter *K*. But as I've made clear right from the start in *The Design Inference* and throughout my subsequent work, my formalism assigns probabilities only to events. Background knowledge does not constitute an event but items of information that get factored into existing chance hypotheses and thereby update the probabilities assigned to events. Background knowledge that explicitly and univocally identifies a rejection function/region but which, under such probabilistic updating, does not alter the probabilities is, I claim, the only sort that we should be looking at in forming specifications.

Erik's reformalization of my work is therefore an exercise in irrelevance. If he really wanted to assign probabilities to knowledge claims that get used in updating probabilities, he should have tried to embed what I'm doing within a Bayesian framework (which is designed to do just that). Instead he forms a Cartesian product of an event space with a knowledge space, leaving the knowledge space undefined and concocting some ugly mathematics to boot. But Bayesianism is hardly a solution either. As I argue in sections 2.9 and 2.10 of *No Free Lunch*, there are good reasons for not assimilating my approach to design inferences to a Bayesian or likelihood scheme. What's more, when it comes to updating probabilities, the Bayesian scheme, though in some contexts useful, fails to capture our ordinary use of probabilities.

For instance, I flip a coin twice and, thinking the coin fair, assign a certain probability to the outcome. On closer inspection of the coin, I discover that it is biased and calculate that the bias favors heads with probability .6. The probability I assign to two coin flips therefore changes. This new information is now part of my background knowledge and updates the probability I assign to the coin flips. But this updating of probability doesn't require that I assign a probability to my obtaining this new information. Sometimes such probabilities can be assigned, but in many contexts they don't arise and we update probabilities without any reference to them.

The updating of probabilities does not admit a precise mathematical formalization. Such updating is a matter of human judgment, and such human judgment arises as much in everyday contexts as in the hard sciences. This is true even in the Bayesian case, where posterior probabilities get updated in terms of likelihoods and prior probabilities. Those prior probabilities are either themselves a matter of human judgment or posterior probabilities from a previous round of Bayesian decision making. But in the latter case there is a regress (i.e., priors becoming posteriors), which eventually must terminate in an act of human judgment setting some initial prior. This is not to say that our judgments in updating probabilities are subjective in the sense of floating free from all justificatory constraints. It is to say, however, that such judgments can't be shoehorned into some neat mathematical formalism, as Erik attempts to do or would like to think me as doing.

If I am belaboring this point, it is because it comes up crucially in Erik's treatment of the Law of Conservation of Information. Erik focuses principally on the deterministic form of that law. The basic idea of that law in its deterministic form is that if you have an instance of specified complexity that appears the result of some deterministic process, then the specified complexity backtracks under that process. Think of a copy machine. It

outputs some text or work of art. That text or work of art exhibits specified complexity, but is a mechanized copy of some previous text or work of art. The photocopy machine thus functions as a deterministic process that takes preexisting specified complexity to output a new instance of specified complexity. Yet in no wise was novel specified complexity created through this copying process. That's the point of the Law of Conservation of Information, and certainly in the deterministic case it is completely unobjectionable (indeed, Peter Medawar, never a fan of intelligent design, had no problem with it and called it by that very name).

What, then, is Erik's problem with it? Erik's concern seems not to be with the law as such but with my justification of it. Despite all his notation and effort, his objections are quite trivial. They are two: (1) Just because background knowledge can be used to detach a rejection function for a codomain doesn't mean that the composition of that function with the deterministic process mapping domain to codomain induces a detachable rejection function on the domain. (2) My justification of the law renders its applicability "severely limited," in Erik's words, because it presupposes that such deterministic processes have to be "known."

The picture is this: A deterministic process maps a set of antecedent possibilities (the domain) to a set of consequent possibilities (the codomain). A detachable rejection function sits over the codomain and induces an instance of specified complexity in the codomain. But the deterministic process is itself a function and can be composed with the rejection function. Erik's worry is whether this composition constitutes a detachable rejection function over the domain. I submit it does. For a deterministic process that maps domain to codomain, any probability on the codomain is a probability pushed forward from the domain (this is standard probability and ergodic theory). Any knowledge that detaches the original rejection function over the codomain, when augmented by knowledge of the deterministic process, will thus detach the composite rejection function over the domain.

This does mean that the deterministic process has to be known. It also means that the deterministic process must not incorporate any information about the original rejection function over the codomain. Erik apparently thinks these are crucial concessions that make my project founder. But he offers no justification why these conditions, which actually are quite mild, should limit the applicability of the Law of Conservation of Information.

Consider first the requirement that the deterministic process has to be known. To say that it is known doesn't mean that we have to be able to evaluate it precisely at every point or even know much about its behavior. Rather, it is to say that we've prescribed some set of conditions that specify it. This happens all the time in the study of differential equations, where solutions are proven to exist and be unique without necessarily being solvable. But the requirement for known deterministic processes in the Law of Conservation of Information is even weaker than that.

The Law of Conservation of Information in its deterministic form claims that specified complexity can't be created from scratch by deterministic processes. The very statement of the law therefore invites disproof by counterexample. What, then, would it take to disprove this law? Is it enough simply to assert that some unknown deterministic processes might disprove it (thus following evolutionary biology's example of taking refuge in unknown processes)? Obviously not. To disprove the law, one needs actually to propose some deterministic process that promises to overturn it. But any such proposal will require knowledge of the process, and that knowledge, when combined with the knowledge that detaches a rejection function from the codomain in turn detaches the rejection function from the domain.

There's one possible escape here, and that's if knowledge of the process mapping domain to codomain introduces information about the original rejection function that renders its composition with the process not detachable on the domain. Erik is sensitive to this possibility and I give him credit for it. But how could that happen? Unless the process is deliberately concocted with reference to the rejection function, there is no danger with detachability on the codomain safely backtracking to detachability on the domain. The point, after all, of the Law of Conservation of Information is to characterize physical processes happening in nature and their inability to generate specified complexity. But how do we determine what those processes in nature are? Not by focusing on individual events in the codomain (i.e., the place where the original detachable rejection function resides), but by establishing correlations between domain (antecedent circumstances) and codomain (consequent circumstances). The processes that connect domain to codomain, insofar as they have any physical significance, must be identified (by scientific investigation) without regard to the actual events in either but solely on the basis of correlations between the two. This aspect of scientific discovery guarantees that detachability on the codomain translates back to the domain. Erik's worry is therefore a tempest in a teacup with no physical significance. At best it deserves a footnote in future editions of *No Free Lunch*.

The backtracking arguments in which rejection functions on a codomain get translated back to a domain apply also in justifying the general, stochastic form of the Law of Conservation of Information. Erik takes this to imply that just as there can be no addition of specified information (information that is specified though not necessarily complex) with deterministic processes, so there can be none with the nondeterministic or stochastic processes that come up in the general form of the Law of Conservation of Information. He therefore sees as superfluous the addition of a modulus of probability that allows the generation of a certain amount of specified information by chance (albeit not actual specified complexity). Erik falters because he sees my arguments for both laws as parallel and therefore assumes the conclusions must be parallel. But there's an additional probability space thrown into the general form of the Law of Conservation of Information that validates my modulus of probability (what I refer to in the text as "mod UCB" -- the idea here being that stochastic processes cannot merely by chance contribute a specified event of probability less than the universal probability bound).

In the deterministic form of the Law of Conservation of Information, domain mapped to codomain deterministically. Thus a rejection function under composition with the deterministic process translated back to a rejection function on the domain, and there was no net gain or loss of specified information. But with a stochastic process, the rejection function on the codomain translates back to a rejection function on the Cartesian product of the domain plus an additional probability space (one that introduces the stochastic element into the stochastic process). This additional probability space induces the modulus of probability. Here is how it works (Erik can formalize this if he likes).

We've got a rejection function on the codomain that induces a rejection region, call it B , on the codomain. It induces a rejection function on the Cartesian product of domain with space of stochastic elements by composition with the stochastic process from that Cartesian product to the codomain. This composite rejection function in turn induces a rejection region that we'll call A^* (if f is the rejection function on the codomain and g is the stochastic process mapping into the codomain and d is a positive real such that $B = \{f \leq d\}$, then $A^* = \{fg \leq d\}$ -- " \leq " meaning "less than or equal" and these sets being where the functions f and fg respectively satisfy the inequality). If we left it at A^* and let the stochastic form of the Law of Conservation of Information merely connect A^* and B , then we would be back to the deterministic form of the law and Erik would be right.

But with chance in the picture, the stochastic process in question is entitled to purchase a certain amount of specified information for free. How much? I'm particularly conservative and thus allow up to 500 bits of information or equivalently an improbability no more extreme than 10^{-150} (Erik seems to have no problem with these numbers, citing with approval Seth Lloyd's work, which entails a universal probability bound 30 orders of magnitude bigger than mine). Thus, the stochastic process is entitled to identify any subset of the space of stochastic elements with marginal probability no less than 10^{-150} (the joint probability being one on the Cartesian product of the domain with the space of stochastic elements). Let X define the domain, S the space of stochastic elements, and Y the codomain. Then g maps the Cartesian product $S \times X$ into Y and B equals the subset of Y where f , the rejection function, is less than or equal to d and A^* equals the subset of $S \times X$ where the composite rejection function fg is less than or equal to d .

But since the stochastic process g is entitled to purchase any subset of S with probability as low as but not lower than 10^{-150} , g can be restricted to any subset $C \times X$ of $S \times X$ where C is a subset of S with marginal probability no less than 10^{-150} . Now C can be chosen so that the (joint) probability of the elements of $C \times X$ that map under the stochastic process g into the rejection region B is maximal. If we choose C in this way, then g 's entitlement to purchase specified information of improbability down to but no lower than 10^{-150} means that g can be restricted to this space, i.e., $C \times X$. The stochastic process g restricted to $C \times X$ then maps into Y and the probability of $C \times X$ becomes the probability on $S \times X$ conditioned on $C \times X$. A^* , the backtracked instance of specified complexity, then becomes A^* intersected with $C \times X$, which we call A . This A satisfies the general (stochastic) form of the Law of Conservation of Information in relation to B . It, as it were, gives the stochastic process the benefit of the doubt in

producing as much specified information by chance as is allowable in keeping with the universal probability bound. It also shows that the modulus of probability in the general formulation of the Law of Conservation of Information is indispensable.

Let me repeat my caution in the preface to *No Free Lunch*: My strategy in writing this book was "to include just enough technical discussion so that experts can fill in the details as well as sufficient elaboration of the technical discussion so that nonexperts feel the force of the design inference." In the deterministic form of the Law of Conservation of Information, not only is no specified complexity generated but no novel information at all is generated. Yet as soon as stochastic processes come into play, chance can produce information -- indeed, unlimited amounts of it. Chance can also produce limited amounts of specified information. Any statement of a stochastic form of the Law of Conservation of Information therefore must incorporate a modulus of probability to control for the production of limited amounts of specified information by chance. The universal probability bound sets the limit to just how much specified information can be produced by chance. In particular, specified complexity is beyond its remit. The argument sketched above shows how the modulus of probability comes into play in the stochastic formulation of the law. This argument, with its focus on the limited free play of chance over the space of stochastic elements, was always implicit. Erik has helped make it explicit.

This handles the only serious one of Erik's concerns. As for the rest, they are strictly a matter of picking at nits and straining at gnats. Moreover, these nits and gnats are of Erik's own devising. To change the metaphor, he is solicitous to remove some specks in my eye, but the specks he sees are the result of the log in his own eye. I'll therefore close with bullet point replies to a few of Erik's more notable nits and gnats:

Specificational Resources:

Specificational resources, as it were, snip away at a probability space. The more snips, the more likely that a snip will capture an event that's happened. But the snips must not be conditioned by the event itself -- that would be cheating. If the snips can be made in reference to the event, then a judiciously chosen snip can capture the event in question every time and chance can be precluded willy-nilly. The fundamental intuition behind specificational resources is therefore to quantify the number of legitimate snips wherewith chance can still be precluded.

Since precluding chance within the context of my work means sweeping the field clear of all relevant chance hypotheses, I define legitimate snips as those that are conditionally independent or detachable vis-à-vis all the relevant chance hypotheses simultaneously. Erik, by contrast, suggests that legitimate snips, rather than being detachable vis-à-vis all relevant chance hypotheses simultaneously, might be detachable only vis-à-vis individual chance hypotheses. Specifications, in that case, would need to be indexed by specific chance hypotheses and would constitute specifications for some chance hypotheses but perhaps not for others. Erik, as is common throughout his review, raises quibbles without providing concrete alternatives. What would be an example of a pattern that constitutes a specification with respect to one chance hypothesis but not with respect to another? I'm

not asking for some artificial probabilistic construction (i.e., merely an exercise in mathematical counterexample building), but one that could arise in an actual chance elimination argument.

I raise this question purely out of curiosity, not because it poses any danger to my project. The fact is that there is no reason to change my original definition, in which specifications are defined relative to the entire set of relevant chance hypotheses at once and not just relative to individual ones. The fundamental intuition behind specifications is that they can be identified, as a cognitive act, without knowledge of the event that has occurred. For a set of probability distributions that might account for the event's occurrence to contain one probability distribution with respect to which a pattern is a specification and another with respect to which it is not (i.e., conditional independence in the first case but a break down of conditional independence in the second) would indicate that the pattern depends on knowledge of the event after all and thus should not properly be regarded as a specification. <<

Universal Probability Bound:

Erik is dissatisfied with my justification of $10^{(-150)}$ as a universal probability bound. The problem for him is that I'm making the strange assumption that specifying agents must be embodied in at least one elementary particle. Well, our only experience of agents that specify events (I'm not saying that design events) is with embodied agents like ourselves who employ lots more than a single elementary particle. Erik takes this "materialistic assumption" as inappropriate, but it is not a materialistic assumption (certainly it involves no metaphysical assumption on my part about matter or energy being the only reality). It is, rather, an empirical assumption based entirely on our experience of specifying agents.

Interestingly, Erik seems quite taken with Seth Lloyd's June 2002 article in *Physical Review Letters* in which Lloyd calculates the number of bit operations possible throughout the history of the observable universe, a number that translates to a universal probability bound of $10^{(-120)}$, which is 30 orders of magnitude greater than my universal probability bound. Erik concludes that "Dembski's argument, while flawed, actually yielded a good conclusion." No, the argument is not flawed. Indeed, the reason it yielded "a good conclusion" is precisely because the physics underlying my universal probability bound is ultimately the same as the physics underlying Lloyd's. It is not a coincidence that my number and Lloyd's are in the same ballpark. Moreover, it is not an accident that my universal probability bound is smaller than Lloyd's -- I was being more conservative. That's why I continue to maintain that my universal probability bound is the most conservative in the literature. <<

Outcomes and Events:

Erik seems quite bent out of shape over my not being careful enough to distinguish between outcomes and events. Actually, I do make the distinction in *No Free Lunch* (see pp. 142-3 where I distinguish rolling a die and getting a 6 [i.e., an outcome] from rolling a die and getting an even number [i.e., an event that includes the outcome getting a 6]). There's a reason, however, why I don't distinguish too closely between outcomes and

events, and that has to do with the arbitrariness of how we individuate outcomes and events. If, for instance, all we're concerned about is the face with which a die lands, then 6 is an outcome and even numbers constitute an event (events being composites of outcomes). But if we are also concerned with whether the die lands on the floor or on the table, then the occurrence of a 6 would be an event comprised of the outcomes "6 landing on the floor" and "6 landing on the table." Similarly, even though "getting an even number tossing a die" is an event relative to the six possible outcomes of the die landing on any of its faces, relative to the partition of the probability space into "even outcomes" and "odd outcomes," the occurrence of "even outcomes" would constitute an outcome rather than an event. By suitably partitioning or refining the probability space we respectively form outcomes from events or events from outcomes.

This is why I continue to maintain that interpreting the negative logarithm to the base 2 of a probability is properly thought of as the average number of bits required to specify an event or outcome of that probability (Erik thinks that average number of bits applies to outcomes but not to events). The prototype here is coin tossing. The probability of a single head is 1/2 and corresponds to 1 bit. The probability of two heads is 1/4 and corresponds to 2 bits. And so on. But what does one do for probabilities that are not powers of 2? We have to introduce fractional bits. Such fractional bits are similar to saying that the average family has 2.4 children. Since there is no hard and fast distinction between outcomes and events, with outcomes collapsing into events and events segregating into outcomes depending on the context of inquiry, Erik has no basis for objecting to the interpretation of the negative logarithm to the base 2 of a probability as average number of bits for both outcomes and events. <<

Universal Composition Function:

And finally, we come to the most trifling of Erik's nits and gnats, which concerns my use of the universal composition function. Erik refers to my use of it as "truly bizarre" and remarks that "it is difficult [to] understand what it [i.e., my account of it] actually means." Yet the only thing bizarre here is Erik's incomprehension. Erik obviously has some mathematical training and therefore knows about universal Turing machines. Universal Turing machines are extremely simple devices (Steve Wolfram describes perhaps the simplest in his recent book *A New Kind of Science*). Now, the point of a universal Turing machine is not to incorporate any specific programs of its own but to run any program that is given to it. That, more generally, is the point of the universal composition function. Its task is essentially bookkeeping. It keeps track of how specified complexity is being shuffled around without introducing any of its own (preexisting or otherwise). At any rate, the universal composition function introduces no new mathematics, so any arguments that employ it or dispense with it are essentially equivalent. <<

Well, that about does it. The only thing I can think of that I haven't dealt with is Erik's failed counterexample in which an event involving 200 specificational resources is supposed to lead to a design inference (for which subject given which cognitive resources Erik doesn't say -- perhaps a world consisting of one specifying agent and 200 tangerines). Indeed, if such weak counterexamples could overturn the design inference, my book *The Design Inference* would never have made it as my doctoral dissertation,

much less been published through Cambridge University Press. I leave it as an exercise how Erik's counterexample runs aground against my chance elimination schema (see p. 71ff. of *No Free Lunch*).

***Acknowledgment.** I want to thank Rob Koons for some hard-hitting remarks that I've reworked in the text about the failure of Darwinism to qualify even as a theory insofar as it purports to account for biological complexity.*