Divining design … again

The Design Inference, 2nd edition
William A. Dembski and Winston Ewert
Discovery Institute Press, Seattle, Washington, 2023

Shaun Doyle

What is the difference between a pile of sand (figure 1) and a sandcastle (figure 2)? The sandcastle is *designed*. We instinctively distinguish between designed and non-designed objects every day. However, it's not always as simple as telling the difference between a pile of sand and a sandcastle. So, is there a reliable way to formalize the process; to create a method that will reliably allow us to distinguish when something is designed?

Image: W.carter, Wikimedia / CC BY SA 4.0

Figure 1. A pile of sand

Image: Gaius Cornelius, Wikimedia / CC BY SA 4.0

Figure 2. A sandcastle

About 25 years ago, mathematician and philosopher Dr William Dembski proposed just such a way in his book *The Design Inference: Eliminating Chance Through Small Probabilities*. Recently, he has published a second edition of *The Design Inference* with computer scientist and researcher Dr Winston Ewert coming on board as co-author. Dr Royal Truman reviewed the first edition of *The Design Inference* in *Journal of Creation* after it came out.[1] I highly recommend readers read that review, as his comments and his general conclusion remain apt. Nonetheless, the 2nd edition has gone through a substantial revision and expansion, which warrants another review. This review will focus on those updates.

What is the design inference?

But first, it will help to recap the basics of the design inference as Dembski introduces the topic.

It's been long known that there's a problem with invoking chance to explain *some* events that have a low probability of happening. Dembski traces the intuition from Cicero to Ronald Fisher and Richard Dawkins. But why only some? For instance, any particular arrangement of sand particles in a heap of sand is vastly improbable, but we don't reject chance to explain ordinary piles of sand (figure 1). On the other hand, we do for sandcastles (figure 2). What's the difference? Dembski points out, "It's not just the sheer improbability of each of these events, but also its coincidence with or conformity to a *pattern* that compels us to look beyond chance" (p. 59).

When we see a low-probability event conform to a *prespecified* pattern (e.g., flipping a fair coin 1,000 times

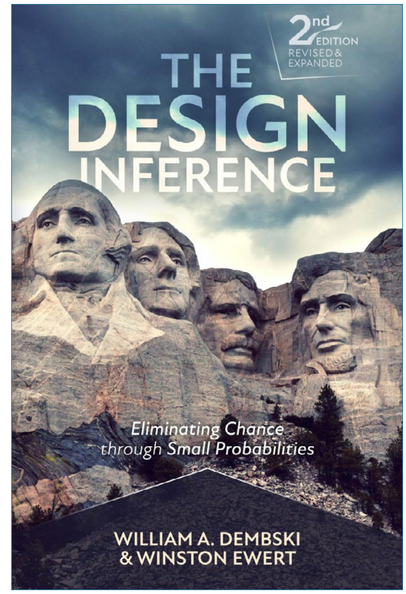

and getting 1,000 heads), eliminating chance is easy. But we don't need the pattern to be prespecified to justly reject chance. It all depends on the type of pattern. There are many patterns we can invent and impose on low-probability data (Dembski calls them 'fabrications'; e.g., describing the precise configuration of each grain in a pile of sand). However, there are very few such patterns we can target with a succinct description (e.g., when we call a pile of sand a 'sandcastle'). Dembski calls the latter 'specifications', and they are a major focus of the book.

But what counts as a low-probability event? For instance, we can describe both flipping three and 1,000 heads in a row very succinctly, and about as succinctly as each other (e.g., 'three consecutive heads' and '1,000 consecutive heads'). But we only regard the latter as an event of sufficiently low probability to reject chance. Why? We need to know two things: how many trials we can run and how many possible outcomes there are for the event in question. For the former, one measure might be the number of coincidences that all people ever could notice. Dembski calculates that to be 10^{34} (p. 64). So, any single event buried among a number of trials

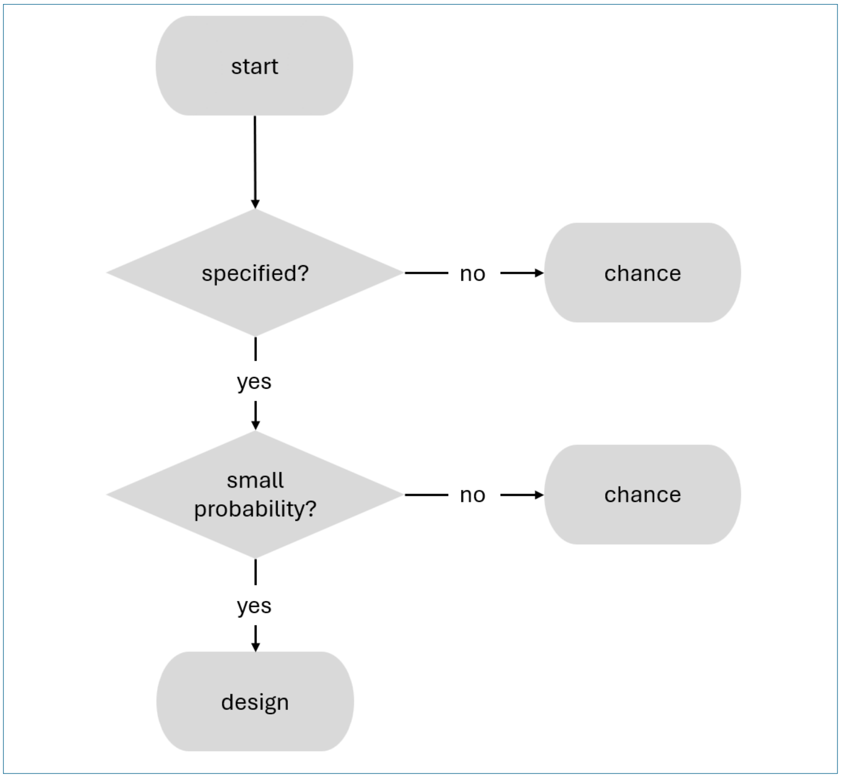

Figure 3. The updated Explanatory Filter

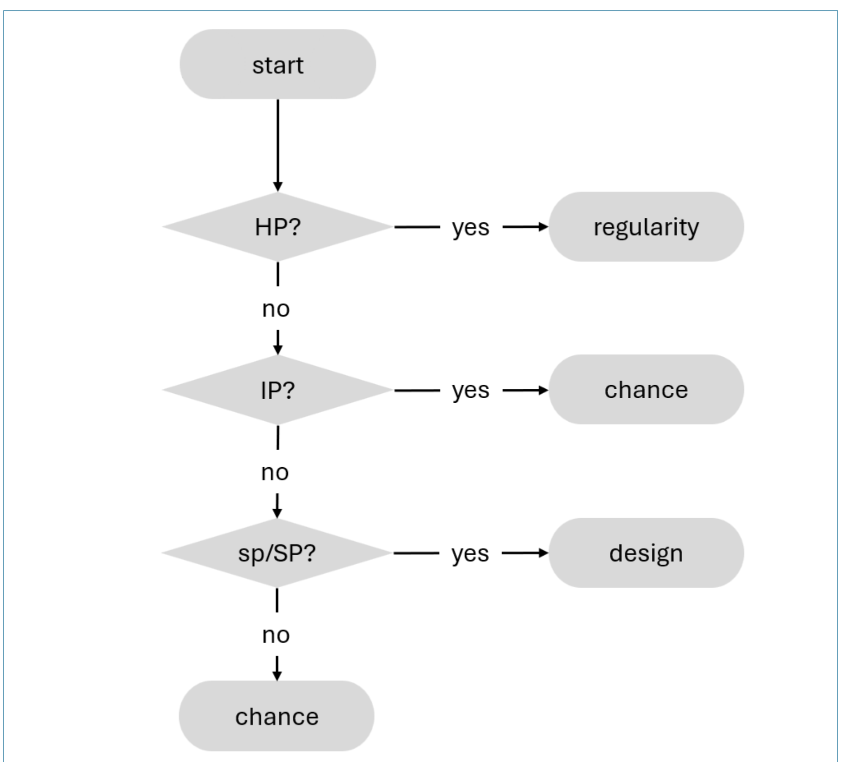

Figure 4. Dembski's original Explanatory Filter. HP means 'high probability', IP means 'intermediate probability', and sp/SP means 'specified small probability'.

higher than 10^{34} we would regard as beyond the reach of chance, i.e., a 1 in 10^{34} chance. 'Three consecutive heads' is one of eight possible outcomes, i.e., a ⅛ chance, which is much higher than 1 in 10^{34}. So, the probability is far too high to reject chance. Indeed, it would probably only take about a few minutes of flipping three coins to practically ensure getting three heads by chance.[2] But '1,000 consecutive heads' is one of about 10^{300} possible outcomes, which is well below 1 in 10^{34}. So, someone who noticed 1,000 heads flipped in a row would be justified in rejecting chance.

So, how do we get from eliminating chance to detecting design? The key is the elimination of *all relevant* chance hypotheses. There are several different ways to formulate that, but they all involve identifying a specified event of small probability. Moreover, the method Dembski has constructed aims to be as sensitive as is reasonable to false positives; i.e., ascribing design to a chance outcome. Why? While it cannot identify *all* designed outcomes, those which it does identify as designed are as secure as we can make them with probabilistic reasoning. But Dembski adds, "The central claim of this book is that design as a logical category, when ascribed by successfully drawing a design inference, correlates reliably with design as a causal category" (p. 95). In other words, Dembski's design inference reliably indicates the detection of the activity of a designing intelligence. But note that the design inference does no *more* than that. It doesn't tell us *who* the designer is or *how* the designed outcome was constructed, only that intelligence was involved.

Streamlining the design inference

How does Dembski go about achieving this? The meat of the book explains how. It centres around his 'Explanatory Filter' (EF) (figure 3).

But it is here that we see the first of the major changes from the first edition. The EF has been significantly streamlined from the first edition (figure 4). First, the number of outcomes has been reduced from three to two by eliminating the 'regularity' category. Second, the number of decision nodes in the filter has been reduced from three to two. Third, the decision nodes around the two factors that were combined in the third decision node of the original EF have been reconfigured: small probability and specification.

Removing regularity

In the first edition of *The Design Inference*, there were three possible outcomes for the design filter: 'regularity', 'chance', and 'design'. However, in the second edition, 'regularity' was removed as a category, leaving only 'chance' and 'design'. Regularity was, though, actually folded into the 'chance' category:

> "In the first edition of this book, we included a sixth predicate, namely, *nec*(*E*), which asserts that the event *E* is the result of a necessity, physical or otherwise, rendering it non-contingent and guaranteeing its occurrence. Common usage distinguishes necessary events, which are bound to happen (such as a double-headed coin landing heads) from contingent events, which could happen but are not compelled to happen (such as a fair coin landing heads). Contingent events are the sort that we typically ascribe to chance. Even so, *nec* can be conveniently assimilated to chance by including within the predicate *ch* hypotheses of the sort *H'* in which all probabilities collapse to zero and one. Thus *nec*(*E*) can be identical to *ch*(*E*;*H'*) where *H'* assigns a probability of one to *E*" (pp. 259–260).

I think this is a convenient way of streamlining options, since it effectively includes all forms of non-teleological explanation under one umbrella and creates a simple binary outcome for the EF: teleology or non-teleology.

However, the EF still seems to have a limited scope of applicability. Loke avers:

> "Dembski fails to consider the option that the event may be 'Uncaused', as has been postulated by Hawking for the Big Bang."[3]

The EF limits the explanatory options to those that can be defined probabilistically. Within the cosmos, this is uncontroversial, since even quantum events occur according to well defined probabilities. So, at the very least, this should not affect the use of the EF in *scientific* endeavours, which presuppose the existence of this universe.

It might, however, limit the scope of applicability for the EF to questions of cosmic design. Nonetheless, it will not render it useless. The EF could be used to rule out any non-design explanations that can be defined probabilistically. Moreover, Loke does a decent job defining a logically exhaustive list of options for the design of the cosmos,[4] and the only category he mentions the EF doesn't capture is 'uncaused'. But if we can't define the occurrence of uncaused events probabilistically, what can we know about them? Nothing. But that bleeds over into our knowledge of other facts. Pruss explains:

> "Starting with the observation that once we admit that some contingent states of affairs have no explanations, a completely new skeptical scenario becomes possible: no demon is deceiving you, but your perceptual states are occurring for no reason at all, with no prior causes.
>
> "Moreover, objective probabilities are tied to laws of nature or objective tendencies, and so if an objective probability attaches to some contingent fact, then that situation can be given an explanation in terms of laws of nature or objective tendencies. Hence, if the PSR [Principle of Sufficient Reason] is false of some contingent fact, no objective probability attaches to the fact.
>
> "Thus, we cannot even say that violations of the PSR are improbable if the PSR is false. Consequently, someone who does not affirm the PSR cannot say that Koons' skeptical scenario is objectively improbable. It may be taken to follow from this that if the PSR were false or maybe even not known *a priori,* we would not know any empirical truths. But we do know empirical truths. Hence, the PSR is true, and maybe even known *a priori*."[5]

So, the fact that an 'uncaused' event is outside the scope of EF actually provides us reason to believe that such an event isn't even possible. Moreover, Loke himself provides an impressive array of arguments against both the universe's beginning and its fine-tuned structure being uncaused.[6] So, it seems the EF is useful even in assessing claims of cosmic design.

Reducing the decision nodes

In the first edition, the EF also had three decision nodes that partitioned its conclusions into regularity, chance, or design (figure 4). The first decision node asked whether an event occurred with high probability. If yes, then regularity was accepted; if no, we proceed to the second node. The second node asks if an event occurred with *intermediate* probability, which Dembski defined, in the 1st edition, as "the events we reasonably expect to occur by chance in the ordinary circumstances of life."[7] If yes, then the event is ascribed to chance; if no, we proceed to the third node. The third node asks whether an event is *both* specified and has a small probability. An example of the difference is that between the pile of sand and a sandcastle (figures 1 and

2). The geometrical configuration of both 'piles of sand' has a small probability of occurring (i.e., they each demonstrate only one way of a myriad number of ways the piles could be configured), but only the sandcastle fits an independently recognizable pattern that can be characterized with a short description: 'sandcastle'. The 'generic pile', while having a small probability, is not specified, and thus we conclude chance. However, the sandcastle is both specified and low probability, and so we conclude design.

The new EF (figure 3) eliminates both the 'high probability' and 'intermediate probability' decision nodes. This would have left the EF with one decision node, but Dembski partitioned the third decision node into two: 'specification' and 'small probability'. This makes a lot more sense, since these are the two concepts that directly underwrite a successful design inference within the EF.

Specification as a short description

The first node asks whether the event we have identified can be specified by a short description. Compare, for instance, these two sequences of 100 coin tosses:

THHHTHTTHHTTTTTHTTTH-
HHTTTTHHHHHHHTTHTHH-
HHTTHTHHHTTHTTHHHHHT-
THHTHHHHTTTTTHHTHHHHH-
HHTHTTTTHTTHHHHHTHTT

HHHHHHHHHHHHHHHHHHHH-
HHHHHHHHHHHHHHHHHHHH-
HHHHHHHHHHHHHHHHHHHH-
HHHHHHHHHHHHHHHHHHHH-
HHHHHHHHHHHHHHHHHHHH

We can only generate a short description for the second sequence: '100 heads'. The first one is a random sequence I generated by tossing a fair coin 100 times. It can't really be described in any shorter way than listing the outcomes in full. Dembski delves deep into complexity theory to specify what counts as a short description. In essence, it involves the shortest description length we can find for an event using a binary code where words are individually identifiable, and the code can describe any pattern that a computer could be programmed to recognize (pp. 309–311). Still, there's no way to find a true *minimal* description length (p. 153). So, the best we can do is use the shortest description we can think of (p. 154).

Complexity as low probability

The second node concentrates defining an event with low probability. Dembski elaborates:

> "Whether an event is probable or improbable depends on the number of opportunities for the event to occur. The greater the number of opportunities, the more probable is the event" (p. 194).

Therefore, Dembski constructs an understanding of low probability by assessing *probabilistic resources*: the number of opportunities an event has to happen. Dembski elaborates:

> "To calculate this number, we need to identify the different relevant ways an event could occur and then count them" (p. 195).

This sounds simple. However, there are a lot of complicating factors for determining probabilistic resources. For instance, do we just use an assumption of uniform probability to calculate how likely an event is to happen? No. We consider *all the factors that affect how likely the event is to happen*. If there are factors that skew the probability of an event happening (e.g., natural laws), they must be taken into account. For instance, in biological systems, this *includes* the effects of 'non-random' processes like natural selection. Any analysis that fails to take such processes into account will give an incorrect assessment of the probability of the event happening.

Moreover, one must consider the difference between relative and absolute probabilistic resources. That is, those that reflect the limits of fundamental physics vs those that reflect the practical limits on the opportunities for the event to happen.

But even absolute probabilistic resources can change depending on the size of the system we're considering. For instance, the maximum resources for an event *on Earth* are smaller than they are for one that is merely in the observable universe.

One also needs to make sure we count *realizable* possibilities and not just mere possibilities. For instance, given an estimate of the informational size of the observable universe at 10^{120} bits, the number of informational states the universe could theoretically take is $2^{10^{120}}$. However, this vastly overcounts the number of *accessible* states the 10^{120} bits of the universe can take up in its history, since (assuming the big-bang timeframe of c. 13.8 billion years) there are only 10^{120} operations possible in the observable universe's history, and they change the 10^{120} bits from one state to another. Therefore, the maximum number of *achievable* states is only $10^{120} \times 10^{120} = 10^{240}$ (pp. 214–215).

But even this overestimates the probabilistic resources available to evolutionary biology on Earth. For that, Dembski calculates an upper bound of 10^{50} reproduction events in the presumed 4-billion-year evolutionary history of life on Earth.

Specified complexity as a unified measure

In the original edition of his work, Dembski maintained a conceptual distinction between 'specification' and 'small probability' and evaluated each independently to test for design. However, subsequent to the first edition of *The Design Inference* being

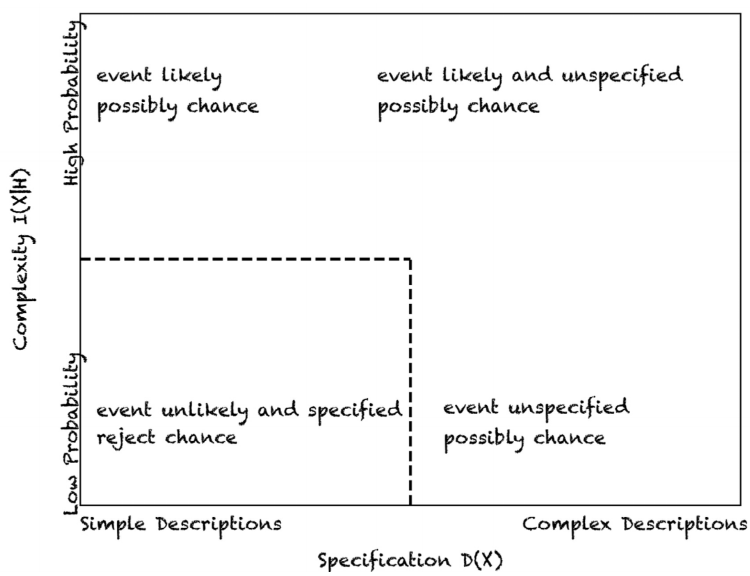

Figure 5. How specified complexity is depicted when specification and small probability constrain an event separately. It is a limited form that trades in unified 'cutoffs' in both specification (the vertical axis of the rectangle) and complexity (the horizontal axis of the rectangle). Dembski's figure 6.1 on p. 322.

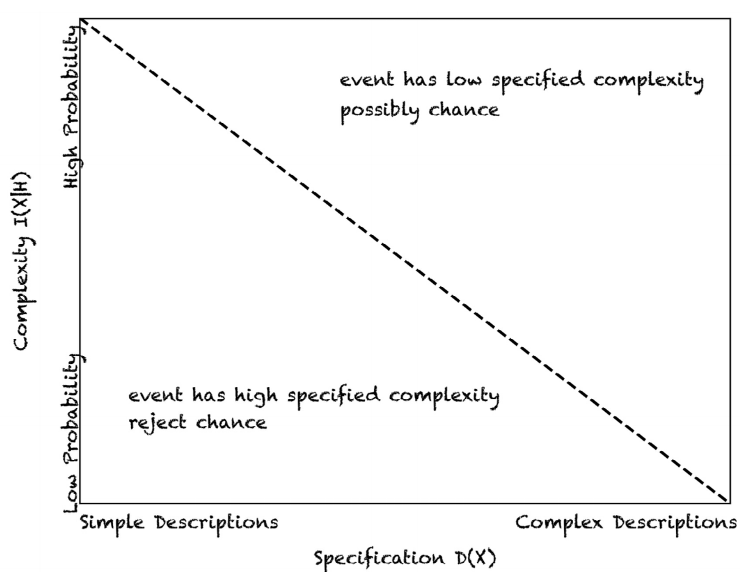

Figure 6. The resulting map of specified complexity when treated as a single measure. The diagonal line depicts the trade-off that occurs (lower probability on the one hand and longer description length on the other) in eliminating chance. Dembski's figure 6.2 on p. 324.

published, Dembski and others have engaged in much work refining their understanding of both concepts. The major fruit of this subsequent research was that Dembski and others have developed a unified 'specified complexity' measure, which he calls "algorithmic specified complexity", by combining the measures of an event's specification and probability into a single unit: bits of information. As Dembski puts it:

> "Algorithmic specified complexity is a special case of generic specified complexity. In algorithmic specified complexity, description length is fully formalized and, just as with complexity, characterized in terms of bits. Moreover, as a unified information measure, algorithmic specified complexity is the arithmetic difference of two quantities in bits: complexity in the sense of improbability (in bits) minus minimum description length (in bits)."

He contrasts the difference between treating specification and probability independently and uniting them in two helpful diagrams (figures 5 and 6). In figure 5, when they are treated independently, they set a threshold for each factor that is independent of their relation to each other. However, as Dembski points out, "This approach to chance elimination, however, does not adequately capture all the cases for which we might reject a chance hypothesis" (p. 323). Why? Because the two factors *are* related. As Dembski further points out:

> "… in eliminating a chance hypothesis to explain an event, we need to lower the probability of an event *to the degree* that we raise the length of its description. Chance elimination can live with longer descriptions provided that the probability is low enough [emphasis added]" (p. 324).

In aiming to eliminate chance explanations for an event, we strive to balance "tradeoffs between probability and description length so that a longer description length can always be offset with sufficiently small probability" (p. 326), as illustrated by figure 6. A unified specified complexity measure reflects this reasoning in ways that treating the probability and description length of an event separately cannot.

Addressing criticisms

There was a lot of reaction to the first edition of *The Design Inference*. At first, it was welcomed quite warmly. But once people started noting it could apply to evolution, the reaction soured dramatically. While a few criticisms have been aimed at the design inferential apparatus *per se,*[8] most have been aimed at its application to evolutionary biology. Dembski notes the reaction:

> "In other areas of inquiry, the design inference has proven itself useful, and indeed indispensable, for uncovering evidence of intelligent activity (recall Chapter 2). Even so, mainstream evolutionary biologists have denied that design inferential reasoning has any validity when applied to the evolution of biological systems. Whatever its merits elsewhere, the design inference is thus supposed to tell us nothing about the origin and subsequent history of life on Earth" (pp. 370–371).

Why? Did evolutionists "painstakingly show that for every biological system that exists, no design inference could legitimately be drawn?" (p. 371). Dembski rightly answers, "No". So, what happened? In essence, they simply *assumed* naturalism. To put it another way, the prior probability of Darwinian theory is regarded by mainstream biologists as essentially 1. No specific evidence for design can challenge that because, regardless of how much the specific evidence favours design, it's never enough to overcome the advantage evolution has over design in prior probability. But what justification is there for assigning such a high prior probability to Darwinism?

First, faulty reasoning props this up. Dembski's design inferential framework aims to eliminate all chance hypotheses. But evolutionists objected that, since there are chance hypotheses we haven't identified, it may be from among those that we could find one that shows how natural mechanisms could make the biological systems in question highly probable. But, as Dembski points out, no *actual* chance hypothesis has been identified. They're relying on an *unidentified, merely possible* hypothesis, which has no evidential backing. Now, maybe the evolutionist is right. But the evolutionist has a burden of proof here. They must identify such a hypothesis so it can be tested. In the search for truth in biological origins, we can't be held hostage to naturalism by unknown possibilities.

Second, what difference does Darwinism actually make? Dembski quotes John Stuart Mill on his *method of difference*:

> "If an instance in which the phenomenon under investigation occurs, and an instance in which it does not occur, have every circumstance in common save one, that one occurring only in the former; the circumstance in which alone the two instances differ is the effect, or the cause, or an indispensable part of the cause, of the phenomenon" (p. 392).[9]

Darwinian evolution operates via replication, mutation, and selection in some environment. "But that can produce wildly different types of evolution: interesting, uninteresting, complexity increasing, complexity decreasing, elegantly engineered, kludgily engineered, etc" (p. 397). Mills' method of difference tells us, though, that something *other* than these things must be what makes the difference between these different types of evolution. The key issue, it seems, is the *type* of environment that replication, mutation, and selection operate in. If they operate in an environment "capable of contributing information through the evolutionary process to its newly emerging products" (p. 393), evolution will happen. But what makes one environment capable of contributing information, and another one not? This is something that cannot be answered merely by an appeal to replication, selection, and mutation in an environment. Something else explains the difference. What is it? Design theorists say that it's design. Evolutionists say that it is …? We don't know.

Third, evolutionists must deal with the challenge of *multiple simultaneous changes*. The standard picture of evolution is that mutations required to move from structure A to structure B happen gradually, one at a time. The alternative to this, of course, is multiple simultaneous mutations. However, the more simultaneous mutations that are required to move one step on the path between A and B, the less probable that step becomes. A common evolutionary response is to say, for any given mutational step, there's no way to show that multiple simultaneous mutations are needed, but we can always reasonably assume that multiple simultaneous mutations are not needed. Once again, the evolutionist resorts to *unevidenced possibilities* as a 'refutation', shirking their own burden of proof in the process.

Nonetheless, Dembski does think the probabilities of some specified biological systems can be estimated and that some do seem to evince sufficiently low probability to underwrite a design inference. His primary example is the work of Douglas Axe on the beta-lactamase enzyme, who estimated that "the overall prevalence of sequences performing a specific function by any domain-sized fold may be as low as 1 in 10^{77}."[10] Dembski puts it through his special complexity measure, which gives us 215 bits of information, and a resultant probability bound of $2^{-215} \approx 10^{-64}$ for rejecting chance. This is well above the universal probability bound Dembski works with of 10^{-150}. However, a more realistic probability bound to use is the number of reproductive trials available to evolution in its presumed four billion years history on Earth. Dembski calculates this to be around 10^{50}, which appears to be exceedingly generous.[11] But it means that anything with less chance than c. 10^{-50} is unlikely to occur in the history of life. Since 10^{-64} is much lower than 10^{-50}, a design inference is warranted.

Conclusions

The first edition of *The Design Inference* was a landmark book. It signalled that there might actually be a way to detect design in a mathematically rigorous way. But was the second edition worth it? Definitely. I think it is a pretty substantial improvement on the first. It streamlines the design inferential scheme significantly, so that now it focuses on the most important matters: how unlikely an event is, and how easy it is to describe the pattern we see in the event. And by combining those elements into a single 'specified complexity' measure, it has made the reasoning resemble much more closely how we typically reason about design. In essence, the more unlikely and easier to describe something is, the more likely it is the result of design.

Nonetheless, none of these changes fundamentally alters the structure of Dembski's original scheme. So, one might wonder why the first edition of this should've caused so much controversy. After all, this is nothing more than a *method for detecting design*. It doesn't prejudge what is designed and what isn't. Indeed, in most areas the academic establishment seemed willing to accept its applicability. But then they realized it could also be applied to *biology*. That was a bridge too far. Darwin's ideas were too important.

And the response from the evolutionary establishment was: 'admit nothing, deny everything, and denigrate those who disagree.' There can only be evidence *for* evolution; no evidence can ever stand against it. The mere possibility that evolution could suffice to explain any biological system, *even when we have no idea how it does*, is enough reason to reject any design inference.

That doesn't change the fact that there *are* ways to apply his ideas to biological systems. And it seems that, in some cases, a design inference is warranted. But the furore was never about the particulars of Dembski's scheme. It was almost always about the very notion that anyone could suggest with a straight face that evolution might ever be wrong.

Of course, *The Design Inference* is not a biblical creationist book. Dembski himself is an old-earth creationist. But this book isn't about the biblical doctrine of creation or the relationship of Genesis to evolution and origins. It's a book that sets forth a mathematically rigorous scheme for detecting design in the world. And does it succeed at that aim? Even in biology? Yes. There never really was a question about that. Doubts only arise from those who studiously avoid a design inference, perhaps because they're trying to avoid the ultimate Designer.

References

1. Truman, R., Divining design: A review of *The Design Inference: Eliminating chance through small probabilities*, by William A. Dembski, *J. Creation* **13**(2):34–39, 1999.
2. I timed a run, and I got three heads right on the one-minute mark. Interestingly, I had already gotten *two* triple-tails results in the same run, which of course has the same probability of ⅛.
3. Loke, A., *The Teleological and Kalam Cosmological Arguments Revisited*, Palgrave Macmillan, Cham, Switzerland, p. 155, 2022.
4. Loke, ref. 3, pp. 156–157.
5. Pruss, A.R., The Leibnizian Cosmological Argument; in: Craig, W.L. and Moreland, J.P. (Eds.), *The Blackwell Companion to Natural Theology*, Blackwell Publishing, UK, Kindle Locations 1021–1030, 2009.
6. Loke, ref. 3, pp. 71–134, 297–302.
7. Dembski, W.A., *The Design Inference*, 1st edn, Cambridge University Press, Cambridge, UK, Kindle Location 960, 1998.
8. There are some, and Dembski has addressed the most significant ones this edition. See e.g., section 5.7, "Probabilistic *Modus Tollens*".
9. Mill, J.S., *A System of Logic: Ratiocinative and inductive*, 8th edn, Harper & Brothers Publishers, Franklin Square, New York, p. 280, 1881; System of Logic.djvu, en.wikisource.org, accessed 16 Aug 2024.
10. Axe, D.D., Estimating the prevalence of protein sequences adopting functional enzyme folds, *J. Mol. Biol.* **341**:1295–1315, 2004; p. 1295.
11. Dilley, S., Luskin, C., Miller, B., and Reeves, E., On the relationship between design and evolution, *Religions* **14**:850, 2023 estimated it to be around 10^{38}, based on calculations from Chatterjee, K., Pavlogiannis, A., Adlam, B., and Nowak, M.A., The time scale of evolutionary innovation, *PLoS Comp. Biol.* **10**:e1003818, 2014. Moreover, 10^{40} is the most recent estimate by evolutionists of how many cells have existed on the planet throughout their supposed 4-billion-year history: Crockford, P.W., On, Y.M.B., Ward, L.M., Milo, R., and Halevy, I., The geologic history of primary productivity, *Curr. Biol.* **33**(21):4741–4750, 2023.