

WILLIAM A. DEMBSKI

THE THIRD MODE OF EXPLANATION:  
DETECTING EVIDENCE OF INTELLIGENT  
DESIGN IN THE SCIENCES

I. Introduction

In our workaday lives we find it important to distinguish between three modes of explanation: necessity, chance, and design. Did she fall, or was she pushed? And if she fell, was her fall accidental or unavoidable? To say she was pushed is to attribute her plunge to design. To say her fall was accidental or unavoidable is to attribute her plunge respectively to chance or necessity. More generally, given an event, object, or structure, we want to know: Did it have to happen? Did it happen by accident? Did an intelligent agent cause it to happen? In other words, did it happen by necessity, chance, or design?

At this level of analysis, necessity, chance, and design remain pretheoretical and therefore inadequate for constructing a scientific theory of design. It is therefore fair to ask whether there is a principled way to distinguish these modes of explanation. Philosophers and scientists have disagreed not only about how to distinguish these modes of explanation but also about their very legitimacy. The Epicureans, for instance, gave pride of place to chance. The Stoics, on the other hand, emphasized necessity and design but rejected chance. In the Middle Ages Moses Maimonides contended with the Islamic interpreters of Aristotle who viewed the heavens as, in Maimonides' words, "the necessary result of natural laws".<sup>1</sup> Where the Islamic philosophers saw necessity, Maimonides saw design.

In arguing for design in his *Guide for the Perplexed*, Maimonides looked to the irregular distribution of stars in the heavens. For him that irregularity demonstrated contingency (that is, an event that happened but did not have to happen and therefore was not necessary). But was that contingency the result of chance or design? Neither Maimonides nor the Islamic interpreters of Aristotle had any use for Epicurus and his views on chance. For them chance could never be fundamental but was at best a placeholder for ignorance. Thus for Maimonides and his Islamic colleagues, the question was whether a principled distinction could be drawn between necessity and design. The Islamic philosophers, intent on keeping Aristotle pure of theology, said no. Maimonides, arguing from observed contingency in nature, said yes. His argument focused on the distribution of stars in the night sky:

What determined that the one small part [of the night sky] should have ten stars, and the other portion should be without any star? . . . The answer to [this] and similar questions is very difficult and almost impossible, if we assume that all emanates from God as the necessary result of certain permanent laws, as Aristotle holds. But if we assume that all this is the result of design, there is nothing strange or improbable; the only question to be asked is this: What is the cause of this design? The answer to this question is that all this has been made for a certain purpose, though we do not know it; there is nothing that is done in vain, or by chance. . . . How, then, can any reasonable person imagine that the position, magnitude, and number of the stars, or the various courses of their spheres, are purposeless, or the result of chance? There is no doubt that every one of these things is . . . in accordance with a certain design; and it is extremely improbable that these things should be the necessary result of natural laws, and not that of design.<sup>2</sup>

Modern science has also struggled with how to distinguish between necessity, chance, and design. Newtonian mechanics, construed as a set of deterministic physical laws, seemed only to permit necessity. Nonetheless, in the General Scholium to

his *Principia*, Newton claimed that the stability of the planetary system depended not only on the regular action of the universal law of gravitation but also on the precise initial positioning of the planets and comets in relation to the sun. As he explained:

Though these bodies may, indeed, persevere in their orbits by the mere laws of gravity, yet they could by no means have at first derived the regular position of the orbits themselves from those laws. . . . [Thus] this most beautiful system of the sun, planets, and comets, could only proceed from the counsel and dominion of an intelligent and powerful being.<sup>3</sup>

Like Maimonides, Newton saw both necessity and design as legitimate explanations but gave short shrift to chance.

Newton published his *Principia* in the seventeenth century. Yet by the nineteenth century necessity was still in, chance was still out, but design had lost much of its appeal. When asked by Napoleon where God fit into his equations of celestial mechanics, Laplace famously replied, "Sire, I have no need of that hypothesis." In place of a designing intelligence that precisely positioned the heavenly bodies, Laplace proposed his nebular hypothesis, which accounted for the origin of the solar system strictly through natural gravitational forces.<sup>4</sup>

Since Laplace's day, science has largely dispensed with design. Certainly Darwin played a crucial role here by eliminating design from biology. Yet at the same time science was dispensing with design, it was also dispensing with Laplace's vision of a deterministic universe (recall Laplace's famous demon who could predict the future and retrodict the past with perfect precision provided that present positions and momenta of particles were fully known).<sup>5</sup> With the rise of statistical mechanics and then quantum mechanics, the role of chance in physics came to be regarded as ineliminable. Especially convincing here has been the failure of the Bell inequality.<sup>6</sup> Consequently, a deterministic, necessitarian universe has given way to a stochastic universe in which chance and necessity are both regarded as fundamental modes of scientific explanation, neither being re-

ducible to the other. To sum up, contemporary science allows a principled distinction between necessity and chance but repudiates design as a possible explanation for natural phenomena.

## 2. Rehabilitating Design

But was science right to repudiate design? In *The Design Inference* I argue that design is a legitimate and fundamental mode of scientific explanation, on a par with chance and necessity.<sup>7</sup> In arguing this claim, however, I want to avoid prejudging the implications of design for science. In particular, it is not my aim to guarantee creationism. Design, as I develop it, cuts both ways and might just as well be used to defeat creationism by clarifying the superfluity of design in biology. My aim is not to find design in any one place but to open up possibilities for finding design as well as for shutting it down.

My aim, then, is to rehabilitate design as a mode of scientific explanation. Given that aim, it will help to review why design was removed from science in the first place. Design, in the form of Aristotle's formal and final causes, had after all once occupied a perfectly legitimate role within natural philosophy, or what we now call science. With the rise of modern science, however, these causes fell into disrepute.

We can see how this happened by considering Francis Bacon. Bacon, a contemporary of Galileo and Kepler, though himself not a scientist, was a terrific propagandist for science. Bacon concerned himself much about the proper conduct of science, providing detailed canons for experimental observation, recording of data, and inferences from data. What interests us here, however, is what he did with Aristotle's four causes. For Aristotle, to understand any phenomenon properly, one had to understand its four causes, namely, its material, efficient, formal, and final cause.<sup>8</sup>

A standard example philosophers use to illustrate Aristotle's four causes is to consider a statue—say Michelangelo's *David*.

The material cause is what it is made of—marble. The efficient cause is the immediate activity that produced the statue—Michelangelo's actual chipping away at a marble slab with hammer and chisel. The formal cause is its structure—it is a representation of David and not some random chunk of marble. And finally, the final cause is its purpose—presumably, to beautify some Florentine palace.

Two points about Aristotle's causes are relevant to this discussion. First, Aristotle gave equal weight to all four causes. In particular, Aristotle would have regarded any inquiry that omitted one of his causes as fundamentally deficient. Second, Bacon adamantly opposed including formal and final causes within science (see his *Advancement of Learning*).<sup>9</sup> For Bacon, formal and final causes belonged to metaphysics and not to science. Science, according to Bacon, needed to limit itself to material and efficient causes, thereby freeing science from the sterility that inevitably results when science and metaphysics are conflated. This was Bacon's line, and he argued it forcefully.

We see Bacon's line championed in our own day by atheists and theists alike. In *Chance and Necessity*, biologist and Nobel laureate Jacques Monod argued that chance and necessity alone suffice to account for every aspect of the universe. Now whatever else we might want to say about chance and necessity, they provide at best a reductive account of Aristotle's formal causes and leave no room whatever for Aristotle's final causes. Indeed, Monod explicitly denies any place for purpose within science.<sup>10</sup>

Monod was an outspoken atheist. Nevertheless, as outspoken a theist as Stanley Jaki will agree with Monod about this aspect of science. Jaki is as theologically conservative a historian of science and Catholic priest as one is likely to find. Yet in his published work he explicitly states that purpose is a purely metaphysical notion and cannot legitimately be included within science. Jaki's exclusion of purpose, and more generally design, from science has practical implications. For

### 3. The Complexity-Specification Criterion

A detailed explication and justification of the complexity-specification criterion is technical and can be found in *The Design Inference*. Nevertheless, the basic idea is straightforward and easily illustrated. Consider how the radio astronomers in the movie *Contact* detected an extraterrestrial intelligence. This movie, based on a novel by Carl Sagan, was an enjoyable piece of propaganda for the SETI research program—the Search for Extraterrestrial Intelligence. To make the movie interesting, the SETI researchers in *Contact* actually did find an extraterrestrial intelligence (the *non-fictional* SETI program has yet to be so lucky).

How, then, did the SETI researchers in *Contact* convince themselves that they had found an extraterrestrial intelligence? To increase their chances of finding an extraterrestrial intelligence, SETI researchers monitor millions of radio signals from outer space. Many natural objects in space produce radio waves (for example, pulsars). Looking for signs of design among all these naturally produced radio signals is like looking for a needle in a haystack. To sift through the haystack, SETI researchers run the signals they monitor through computers programmed with pattern-matchers. So long as a signal does not match one of the preset patterns, it will pass through the pattern-matching sieve (even if it has an intelligent source). If, on the other hand, it does match one of these patterns, then, depending on the pattern matched, the SETI researchers may have cause for celebration.

The SETI researchers in *Contact* did find a signal worthy of celebration, namely, the following:

```

IIOHIOIIIIIOIIIIIIIOIIIIIIIIIIIOIIIIIIIIIIIO
IIIIIIIIIIIIIIIIIOIIIIIIIIIIIIIIIIIOIIIIIIII
IIIIIIIIIIIIIOIIIIIIIIIIIIIIIIIIIIIIIIIOII
IIIIIIIIIIIIIIIIIIIIIIIIIIIOIIIIIIIIIIIIIIII
```



be explained by chance. Finally, specification ensures that the object exhibits the type of pattern characteristic of intelligence. Let us examine these three requirements more closely.

In practice, to establish that an object, event, or structure is contingent, one must show that it is not the result of a natural law (or algorithm). For example, a crystal of salt results from forces of chemical necessity that can be described by the laws of chemistry. By contrast, a setting of silverware is not. No physical or chemical laws dictate that the fork must be on the left and the knife and spoon on the right. The place setting of silverware is therefore contingent, whereas the structure of the crystal is the result of physical necessity. Michael Polanyi and Timothy Lenoir have both described a method of establishing contingency.<sup>15</sup> 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. In the case of the radio signal in *Contact*, the pattern of o's and 1's forming a sequence of prime numbers is irreducible to the laws of physics that govern the transmission of radio signals. We therefore regard the sequence as contingent.

To see next why complexity is crucial for inferring design, consider the following sequence of bits:

110111011111

These are the first twelve bits in the previous sequence representing the prime numbers 2, 3, and 5 respectively. Now it is a sure bet that no SETI researcher, if confronted with this twelve-bit sequence, is going to contact the science editor at the *New York Times*, hold a press conference, and announce that an extraterrestrial intelligence has been discovered. No headline is going to read, "Aliens Master First Three Prime Numbers!"

The problem is that this sequence is much too short (and



thus too simple) to establish that an extraterrestrial intelligence with knowledge of prime numbers produced it. A randomly beating radio source might by chance just happen to output this sequence. A sequence of 1126 bits representing the prime numbers from 2 to 101, however, is a different story. Here the sequence is sufficiently long (and therefore sufficiently complex) that only an extraterrestrial intelligence could have produced it.

Complexity as I am describing it here is a form of probability. Later in this paper I will require a more general conception of complexity. But for now complexity as a form of probability is all we need. To see the connection between complexity and probability, consider a combination lock. The more possible combinations of the lock, the more complex the mechanism and correspondingly the more improbable that the mechanism can be opened by chance. A combination lock whose dial is numbered from 0 to 39 and which must be turned in three alternating directions will have 64,000 ( $= 40 \times 40 \times 40$ ) possible combinations and thus a  $1/64,000$  probability of being opened by chance. A more complicated combination lock whose dial is numbered from 0 to 99 and which must be turned in five alternating directions will have 10,000,000,000 ( $= 100 \times 100 \times 100 \times 100 \times 100$ ) possible combinations and thus a  $1/10,000,000,000$  probability of being opened by chance. Complexity and probability therefore vary inversely: the greater the complexity, the smaller the probability. Thus to determine whether something is sufficiently complex to warrant a design inference is to determine whether it has sufficiently small probability.

Even so, complexity (or improbability) is not enough to eliminate chance and establish design. If I flip a coin 1000 times, I will participate in a highly complex (that is, highly improbable) event. Indeed, the sequence I end up flipping will be one in a trillion trillion trillion . . . , where the ellipsis needs twenty-two more "trillions". This sequence of coin tosses will not, however, trigger a design inference. Though complex, this

sequence will not exhibit a suitable pattern. Contrast this with the previous sequence representing the prime numbers from 2 to 101. Not only is this sequence complex, but it also embodies a suitable pattern. The SETI researcher who in the movie *Contact* discovered this sequence put it this way: "This isn't noise, this has structure."

What is a *suitable* pattern for inferring design? Not just any pattern will do. Some patterns can legitimately be employed to infer design whereas others cannot. The way in which we make this distinction is easily illustrated. Consider the case of an archer. Suppose an archer stands fifty meters from a large wall with bow and arrow in hand. The wall, let us say, is sufficiently large that the archer cannot help but hit it. Now suppose each time the archer shoots an arrow at the wall, the archer paints a target around the arrow so that the arrow sits squarely in the bull's-eye. What can we conclude from this scenario? Absolutely nothing about the archer's ability as an archer. Yes, there is a pattern being matched; but it is a pattern fixed only after the arrow has been shot. Thus the pattern is contrived, or what I call "fabricated" (see below).

But suppose instead the archer paints a fixed target on the wall and then shoots at it. Suppose the archer shoots a hundred arrows and each time hits a perfect bull's-eye. What can be concluded from this second scenario? Confronted with this second scenario, we are obligated to infer that here is a world-class archer, one whose shots can legitimately be attributed, not to luck, but rather to the archer's skill and mastery. Skill and mastery are, of course, types of design.

The type of pattern where an archer fixes a target first and then shoots at it is common to statistics, where it is known as setting a rejection region prior to an experiment. In statistics, if the outcome of an experiment falls within a rejection region, the chance hypothesis supposedly responsible for the outcome is rejected. The reason for setting a rejection region prior to an experiment is to forestall what statisticians call "data snooping", or "cherry picking". Just about any data set will contain

strange and improbable patterns if we look hard enough. By forcing experimenters to set their rejection regions prior to an experiment, the statistician protects the experiment from spurious patterns that could just as well result from chance.

Now a little reflection makes clear that a pattern need not be given prior to an event to eliminate chance and implicate design. Consider the following cipher text:

nfuijolt ju jt mjlf b xfbtfn

Initially this looks like a random sequence of letters and spaces—initially one detects no pattern on the basis of which to reject chance and infer design.

But suppose next that someone comes along and tells you to treat this sequence as a Caesar cipher, moving each letter one notch down the alphabet. Now the sequence reads,

methinks it is like a weasel

Even though the pattern (in this case, the decrypted text) is given after the fact, it still is the right sort of pattern for eliminating chance and inferring design. In contrast to statistics, which always identifies its patterns before an experiment is performed, cryptanalysis must discover its patterns after the fact. In both instances, however, the patterns are suitable for inferring design.

Although in the example of the archer, the pattern (the target) is established before the event (that is, before the arrow is shot) that conforms to it, and in the example of the “methinks it is like a weasel” sequence, the pattern is only recognized after the fact, both patterns clearly indicate prior design by an intelligence. But why? What is it about these two patterns that indicates the activity of an intelligence, whereas other patterns (like the target drawn around the arrow after it is shot) do not? The key concept is that of “independence”. I define a specification as a match between an event and an independently given pattern. Events that are both highly complex and specified (that is, that match an independently given pattern) indicate design.

In the first case, where the archer hits a target that exists prior to his shooting the arrow, the pattern is clearly independent of the event. The pattern existed, and was known to exist, before the event occurred. When the arrow hits the target, an event (the arrow shot) conforms to an independently given pattern (the target). In the other case, where the archer draws the pattern around the arrow, the event does not conform to an independently existing pattern (the target). Instead, the pattern (the target) was made to conform to (or was derived from) the event in question. This type of nonindependent pattern I call a fabrication. Fabrications do not indicate anything about whether the event in question was designed.<sup>16</sup>

In the third case of the “methinks it is like a weasel” sequence, the pattern (a meaningful string of English characters) is recognized after the fact but still indicates design. Why? The answer is, again, that the pattern is independent of the event in question. In this case the event in question (the cipher text) conforms to a set of preexisting conventions of English vocabulary and grammar, indeed, to a specific sentence from a Shakespeare play. The pattern does not exist independently of the reception of the text (the event in question), even though we may only recognize the pattern later after some reflection. Indeed, upon analyzing the text we recognize that the text conforms to the independently existing conventions of English vocabulary and grammar. Thus, the pattern imbedded in the cipher text is independent of the event of our reading and analyzing it. For this reason, we have a specification, not a fabrication, and hence, evidence (in conjunction with the complexity of the sequence) for intelligent design. Technically trained readers will want to know that the distinction between a specification and a fabrication (illustrated and described above) can be justified rigorously by employing the notion of conditional independence.<sup>17</sup>

Patterns thus divide into two types, those that in the presence of complexity warrant a design inference and those that despite the presence of complexity do not warrant a design in-

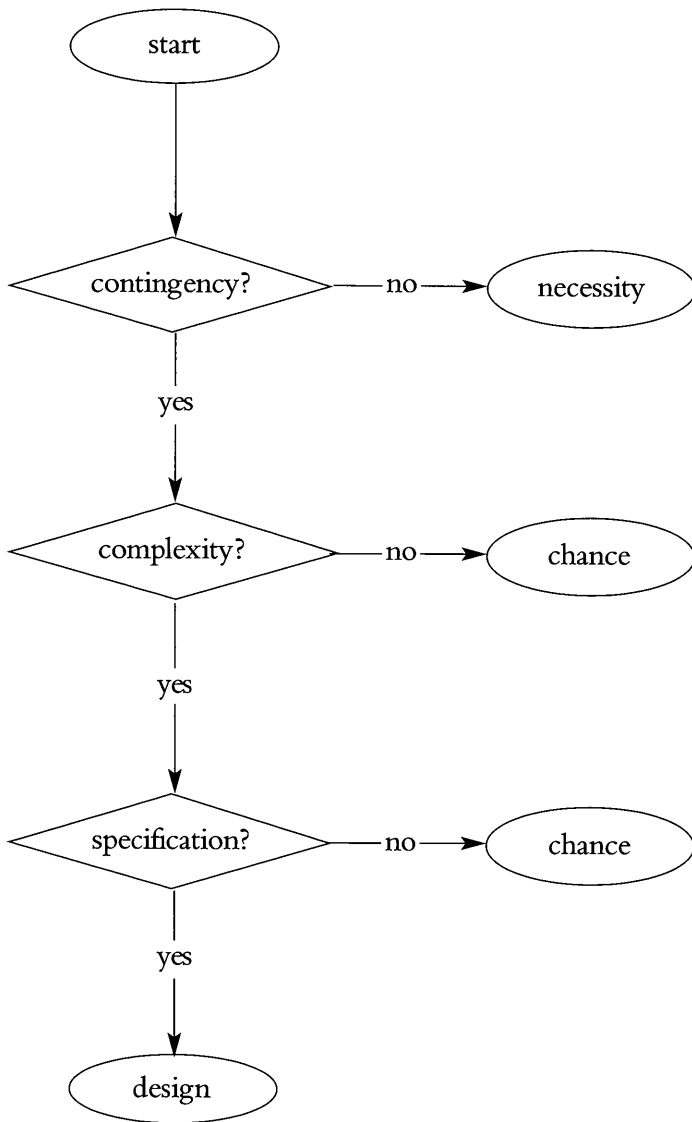
ference. The first type of pattern I call a *specification*, the second a *fabrication*. Specifications are the non-ad hoc patterns that can legitimately be used to eliminate chance and warrant a design inference. In contrast, fabrications are the ad hoc patterns that cannot legitimately be used to warrant a design inference. This distinction between specifications and fabrications can be made with full statistical rigor.<sup>18</sup>

To sum up, the complexity-specification criterion detects design by establishing three things: contingency, complexity, and specification. When called to explain an event, object, or structure, we have a decision to make—are we going to attribute it to *necessity*, *chance*, or *design*? According to the complexity-specification criterion, to answer this question is to answer three simpler questions: Is it contingent? Is it complex? Is it specified? Consequently, the complexity-specification criterion can be represented as a flow chart with three decision nodes. I call this flow chart the Explanatory Filter. [See figure on p. 32.]

### 5. False Negatives and False Positives

As with any criterion, we need to make sure that the judgments of the complexity-specification criterion agree with reality. Consider medical tests. Any medical test is a criterion. A perfectly reliable medical test would detect the presence of a disease whenever it is indeed present and fail to detect the disease whenever it is absent. Unfortunately, no medical test is perfectly reliable, and so the best we can do is keep the proportion of false positives and false negatives as low as possible.

All criteria, and not just medical tests, face the problem of false positives and false negatives. A criterion attempts to classify individuals with respect to a target group (in the case of medical tests, those who have a certain disease). When the criterion places in the target group an individual who should not be there, it commits a false positive. Alternatively, when the criterion fails to place in the target group an individual who should be there, it commits a false negative.



THE EXPLANATORY FILTER

Let us now apply these observations to the complexity-specification criterion. This criterion purports to detect design. Is it a reliable criterion? The target group for this criterion comprises all things intelligently caused. How accurate is this criterion at correctly assigning things to this target group and correctly omitting things from it? The things we are trying to explain have causal stories. In some of those causal stories intelligent causation is indispensable, whereas in others it is dispensable. An inkblot can be explained without appealing to intelligent causation; ink arranged to form meaningful text cannot. When the complexity-specification criterion assigns something to the target group, can we be confident that it actually is intelligently caused? If not, we have a problem with false positives. On the other hand, when this criterion fails to assign something to the target group, can we be confident that no intelligent cause underlies it? If not, we have a problem with false negatives.

Consider first the problem of false negatives. When the complexity-specification criterion fails to detect design in a thing, can we be sure no intelligent cause underlies it? The answer is No. For determining that something is not designed, this criterion is not reliable. False negatives are a problem for it. This problem of false negatives, however, is endemic to detecting intelligent causes.

One difficulty is that intelligent causes can mimic necessity and chance, thereby rendering their actions indistinguishable from such unintelligent causes. A bottle of ink may fall off a cupboard and spill onto a sheet of paper. Alternatively, a human agent may deliberately take a bottle of ink and pour it over a sheet of paper. The resulting inkblot may look identical in both instances but, in the one case, results by chance, in the other by design.

Another difficulty is that detecting intelligent causes requires background knowledge on our part. It takes an intelligent cause to know an intelligent cause. But if we do not know enough, we will miss it. Consider a spy listening in on a communication

channel whose messages are encrypted. Unless the spy knows how to break the cryptosystem used by the parties on whom he is eavesdropping, any messages passing the communication channel will be unintelligible and might in fact be meaningless.

The problem of false negatives therefore arises either when an intelligent agent has acted (whether consciously or unconsciously) to conceal his actions or when an intelligent agent in trying to detect design has insufficient background knowledge to determine whether design actually is present. Detectives face this problem all the time. A detective confronted with a murder needs first to determine whether a murder has indeed been committed. If the murderer was clever and made it appear that the victim died by accident, then the detective will mistake the murder for an accident. So too, if the detective is stupid and misses certain obvious clues, the detective will mistake the murder for an accident. In mistaking a murder for an accident, the detective commits a false negative. Contrast this, however, with a detective facing a murderer intent on revenge and who wants to leave no doubt that the victim was intended to die. In that case the problem of false negatives is unlikely to arise (though we can imagine an incredibly stupid detective, like Chief Inspector Clouseau, mistaking a rather obvious murder for an accident).

Intelligent causes can do things that unintelligent causes cannot and can make their actions evident. When for whatever reason an intelligent cause fails to make its actions evident, we may miss it. But when an intelligent cause succeeds in making its actions evident, we take notice. This is why false negatives do not invalidate the complexity-specification criterion. This criterion is fully capable of detecting intelligent causes intent on making their presence evident. Masters of stealth intent on concealing their actions may successfully evade the criterion. But masters of self-promotion intent on making sure their intellectual property gets properly attributed find in the complexity-specification criterion a ready friend.

And this brings us to the problem of false positives. Even



though specified complexity is not a reliable criterion for *eliminating* design, it is, I shall argue, a reliable criterion for *detecting* design. The complexity-specification criterion is a net. Things that are designed will occasionally slip past the net. We would prefer that the net catch more than it does, omitting nothing due to design. But given the ability of design to mimic unintelligent causes and the possibility that, due to ignorance, we will pass over things that are designed, this problem cannot be remedied. Nevertheless, we want to be very sure that whatever the net does catch includes only what we intend it to catch, to wit, things that are designed. If this is the case, we can have confidence that whatever the complexity-specification criterion attributes to design is indeed designed. On the other hand, if things end up in the net that are not designed, the criterion will be worthless.

I want, then, to argue that specified complexity is a reliable criterion for detecting design. Alternatively, I want to argue that the complexity-specification criterion successfully avoids false positives. Thus, whenever this criterion attributes design, it does so correctly. Let us now see why this is the case. I offer two arguments. The first is a straightforward inductive argument: in every instance where the complexity-specification criterion attributes design, and where the underlying causal story is known (that is, 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. The conclusion of this argument is a straightforward inductive generalization. It has the same logical status as concluding that all ravens are black given that all ravens observed to date have been found to be black.

Anyone with a prior commitment to naturalism is likely to object at this point, claiming that the only things we can know to be designed are artifacts manufactured by intelligent beings that are in turn the product of blind evolutionary processes (for

instance, humans). Hence to use the complexity-specification criterion to extrapolate design beyond such artifacts is illegitimate. This argument does not work. It is circular reasoning to invoke naturalism to underwrite an evolutionary account of intelligence and then, in turn, to employ this account of intelligence to insulate naturalism from critique. Naturalism is a metaphysical position, not a scientific theory based on evidence. Any account of intelligence it entails is therefore suspect and needs to be subjected to independent checks. The complexity-specification criterion provides one such check.

If we set aside the naturalist's evolutionary account of intelligence, a more serious objection remains. I am arguing inductively that the complexity-specification criterion is a reliable criterion for detecting design. The conclusion of this argument is that whenever the criterion attributes design, design actually is present. The premise of this argument is that whenever the criterion attributes design and the underlying causal story can be verified, design actually is present. Now, even though the conclusion follows as an inductive generalization from the premise, the premise itself seems false. There are a lot of coincidences out there that seem best explained without invoking design. Consider, for instance, the Shoemaker-Levy comet. The Shoemaker-Levy comet crashed into Jupiter exactly twenty-five years to the day after the Apollo 11 moon landing. What are we to make of this coincidence? Do we really want to explain it in terms of design? What if we submitted this coincidence to the complexity-specification criterion and out popped design? Our intuitions strongly suggest that the comet's trajectory and NASA's space program were operating independently and that at best this coincidence should be referred to chance—certainly not design.

This objection is readily met. The fact is that the complexity-specification criterion does not yield design all that easily, especially if the complexities are kept high (or, correspondingly, the probabilities are kept small). It is simply not the case that unusual and striking coincidences automatically yield design.

Martin Gardner is no doubt correct when he notes, "The number of events in which you participate for a month, or even a week, is so huge that the probability of noticing a startling correlation is quite high, especially if you keep a sharp outlook."<sup>19</sup> The implication he means to draw, however, is incorrect, namely, that therefore startling correlations/coincidences may uniformly be relegated to chance. Yes, the fact that the Shoemaker-Levy comet crashed into Jupiter exactly twenty-five years to the day after the Apollo 11 moon landing is a coincidence best referred to chance. But the fact that Mary Baker Eddy's writings on Christian Science bear a remarkable resemblance to Phineas Parkhurst Quimby's writings on mental healing is a coincidence that cannot be explained by chance and is properly explained by positing Quimby as a source for Eddy.<sup>20</sup>

The complexity-specification criterion is robust and easily resists counterexamples of the Shoemaker-Levy variety. Assuming, for instance, that the Apollo 11 moon landing serves as a specification for the crash of Shoemaker-Levy into Jupiter (a generous concession at that), and that the comet could have crashed at any time within a period of a year, and that the comet crashed to the very second precisely twenty-five years after the moon landing, a straightforward probability calculation indicates that the probability of this coincidence is no smaller than 1 in  $10^8$ . This simply is not all that small a probability (that is, high complexity), especially when considered in relation to all the events astronomers are observing in the solar system. Certainly this probability is nowhere near the universal probability bound of 1 in  $10^{150}$  that I propose in *The Design Inference*.<sup>21</sup> I have yet to see a convincing application of the complexity-specification criterion in which coincidences better explained by chance get attributed to design.

There is one last potential counterexample we need to consider, and that is the possibility of an evolutionary algorithm producing specified complexity. By an evolutionary algorithm I mean any clearly defined procedure that generates contingency

via some chance process and then sifts the so-generated contingency via some law-like (that is, necessitarian) process. The Darwinian mutation-selection mechanism, neural nets, and genetic algorithms all fall within this definition of evolutionary algorithms.

Now, it is widely held that evolutionary algorithms are just the means for generating specified complexity apart from design. Yet this widely held view is incorrect. The problem is that evolutionary algorithms cannot generate complexity. This may seem counterintuitive, but consider a well-known example by Richard Dawkins in which he purports to show how a cumulative selection process acting on chance can generate specified complexity.<sup>22</sup> He starts with the target sequence

METHINKS·IT·IS·LIKE·A·WEASEL

(he considers only capital Roman letters and spaces, here represented by bullets—thus 27 possibilities at each location in a symbol string).

If we tried to attain this target sequence by pure chance (for example, by randomly shaking out Scrabble pieces), the probability of getting it on the first try would be around 1 in  $10^{40}$ , and correspondingly it would take on average about  $10^{40}$  tries to stand a better than even chance of getting it. Thus, if we depended on pure chance to attain this target sequence, we would in all likelihood be unsuccessful (granted, this 1 in  $10^{40}$  improbability falls short of my universal probability bound of 1 in  $10^{150}$ , but for practical purposes 1 in  $10^{40}$  is small enough to preclude chance and, yes, implicate design). As a problem for pure chance, attaining Dawkins' target sequence is an exercise in generating specified complexity, and it becomes clear that pure chance simply is not up for the task.

But consider next Dawkins' reframing of the problem. In place of pure chance, he considers the following evolutionary algorithm: (1) Start out with a randomly selected sequence of 28 capital Roman letters and spaces, for example,

WDL·MNLT·DTJBKWIRZREZLMQCO·P

(note that the length of Dawkins' target sequence comprises a total of 28 letters and spaces—that is how many letters and spaces there are in METHINKS·IT·IS·LIKE·A·WEASEL); (2) randomly alter all the letters and spaces in this initial randomly generated sequence; (3) whenever an alteration happens to match a corresponding letter in the target sequence, leave it be and randomly alter only those remaining letters that still differ from the target sequence.

In very short order this algorithm converges to Dawkins' target sequence. In his book *The Blind Watchmaker*, Dawkins provides the following computer simulation of this algorithm:<sup>23</sup>

- (1) WDL·MNLТ·DTJBKWIRZREZLMQCO·P
- (2) WDLТMNLТ·DTJBSWIRZREZLMQCO·P
- ...
- (10) MDLDMNLS·ITJISWHRZREZ·MECS·P
- ...
- (20) MELDINLS·IT·ISWPRKE·Z·WECSEL
- ...
- (30) METHINGS·IT·ISWLIKE·B·WECSEL
- ...
- (40) METHINKS·IT·IS·LIKE·I·WEASEL
- ...
- (43) METHINKS·IT·IS·LIKE·A·WEASEL

Thus, in place of  $10^{40}$  tries on average for pure chance to generate Dawkins' target sequence, it now takes only on average 40 tries to generate it via an evolutionary algorithm. Although Dawkins has gotten a lot of mileage out of this example, exactly what it establishes is very different from what he and much of the evolutionary community think it establishes.<sup>24</sup> For one thing, choosing a target sequence is a deeply teleological move (the target is set prior to running the evolutionary algorithm, and the evolutionary algorithm is explicitly programmed to end up in the target). This is a problem because evolutionary algorithms are supposed to be devoid of teleology. But let us for the sake of argument bracket this teleological problem,

which in the case of Darwinism amounts to nature having to select its own targets.

A more serious problem then remains. We can see it by posing the following question: Given Dawkins' evolutionary algorithm, what besides the target sequence can this algorithm attain? Think of it this way. Dawkins' evolutionary algorithm is chugging along; what are the possible terminal points of this algorithm? Clearly, the algorithm is always going to converge on the target sequence (with probability 1 for that matter!). An evolutionary algorithm acts as a *probability amplifier*. Whereas it would take pure chance on average  $10^{40}$  tries to attain Dawkins' target sequence, his evolutionary algorithm on average gets it for you in the logarithm of the number of tries that it takes pure chance, that is, on average in only 40 tries (and with virtual certainty in a few hundred tries).

But a probability amplifier is also a *complexity attenuator*. Recall that the "complexity" in the complexity-specification criterion coincides with improbability. Dawkins' evolutionary algorithm vastly increases the probability of getting the target sequence but in so doing vastly decreases the complexity inherent in the target sequence. The target sequence, if it had to be obtained by randomly throwing Scrabble pieces, would be highly improbable and on average would require a vast number of iterations before it could be obtained. But with Dawkins' evolutionary algorithm, the probability of obtaining the target sequence is high given only a few iterations. In effect, Dawkins' evolutionary algorithm skews the probabilities so that what at first blush seems highly improbable or complex is nothing of the sort. It follows that evolutionary algorithms cannot generate true complexity but only the *appearance of complexity*. And since they cannot generate complexity, they cannot generate specified complexity either.

## 6. Why the Criterion Works

My second argument for showing that specified complexity reliably detects design considers the nature of intelligent agency and, specifically, what it is about intelligent agents that makes them detectable. Even though induction confirms that specified complexity is a reliable criterion for detecting design, induction does not explain why this criterion works. To see why the complexity-specification criterion is exactly the right instrument for detecting design, we need to understand what it is about intelligent agents that makes them detectable in the first place. The principal characteristic of intelligent agency is *choice*. Even the etymology of the word “intelligent” makes this clear. “Intelligent” derives from two Latin words, the preposition *inter*, meaning between, and the verb *lego*, meaning to choose or select. Thus, according to its etymology, intelligence consists in *choosing between*. For an intelligent agent to act is therefore to choose from a range of competing possibilities.

This is true not just of humans but of animals as well as of extraterrestrial intelligences. A rat navigating a maze must choose whether to go right or left at various points in the maze. When SETI researchers attempt to discover intelligence in the extraterrestrial radio transmissions they are monitoring, they assume an extraterrestrial intelligence could have chosen any number of possible radio transmissions, and then they attempt to match the transmissions they observe with certain patterns as opposed to others. Whenever a human being utters meaningful speech, a choice is made from a range of possible sound combinations that might have been uttered. Intelligent agency always entails discrimination, choosing certain things, ruling out others.

Given this characterization of intelligent agency, the crucial question is how to recognize it. Intelligent agents act by making a choice. How, then, do we recognize that an intelligent agent has made a choice? A bottle of ink spills accidentally onto a sheet of paper; someone takes a fountain pen and writes a

message on a sheet of paper. In both instances ink is applied to paper. In both instances one among an almost infinite set of possibilities is realized. In both instances a contingency is actualized and others are ruled out. Yet in one instance we ascribe agency, in the other chance.

What is the relevant difference? Not only do we need to observe that a contingency was actualized, but we ourselves need also to be able to specify that contingency. Alternatively, we need to observe the occurrence of an event that happened but did not have to happen (that is, a contingency), and we must show that this event conforms to a pattern that could be constructed independently of the event (that is, a specification). Ascribing intelligent agency therefore requires identifying both contingency and specification. A random ink blot is contingent but unspecified; a message written with ink on paper is both contingent and specified. To be sure, the exact message recorded may not be specified. But orthographic, syntactic, and semantic constraints will nonetheless specify it.

Actualizing one among several competing possibilities, ruling out the rest, and specifying the one that was actualized encapsulates how we recognize intelligent agency, or, equivalently, how we detect design. Experimental psychologists who study animal learning and behavior have known this all along. To learn a task, an animal must acquire the ability to actualize behaviors suitable for the task as well as the ability to rule out behaviors unsuitable for the task. Moreover, for a psychologist to recognize that an animal has learned a task, it is necessary not only to observe the animal making the appropriate discrimination but also to specify the discrimination.

Thus, to recognize whether a rat has successfully learned how to traverse a maze, a psychologist must first specify which sequence of right and left turns conducts the rat out of the maze. No doubt, a rat randomly wandering a maze also discriminates a sequence of right and left turns. But by randomly wandering the maze, the rat gives no indication that it can discriminate the appropriate sequence of right and left turns for exiting



the maze. Consequently, the psychologist studying the rat will have no reason to think the rat has learned how to traverse the maze.

Only if the rat executes the sequence of right and left turns specified by the psychologist will the psychologist recognize that the rat has learned how to traverse the maze. Now it is precisely the learned behaviors we regard as intelligent in animals. Hence it is no surprise that the same scheme for recognizing animal learning recurs for recognizing intelligent agency generally, to wit: actualizing one among several competing possibilities, ruling out the others, and specifying the one actualized.

Note that complexity is implicit here as well. To see this, consider again a rat traversing a maze, but now take a very simple maze in which two right turns conduct the rat out of the maze. How will a psychologist studying the rat determine whether it has learned to exit the maze? Just putting the rat in the maze will not be enough. Because the maze is so simple, the rat could by chance just happen to take two right turns and thereby exit the maze. The psychologist will therefore be uncertain whether the rat actually learned to exit this maze or whether the rat just got lucky.

But contrast this with a complicated maze in which a rat must take just the right sequence of left and right turns to exit the maze. Suppose the rat must take one hundred appropriate right and left turns and that any mistake will prevent the rat from exiting the maze. A psychologist who sees the rat take no erroneous turns and in short order exit the maze will be convinced that the rat has indeed learned how to exit the maze and that this was not dumb luck.

This general scheme for recognizing intelligent agency is but a thinly disguised form of the complexity-specification criterion. In general, to recognize intelligent agency we must observe an actualization of one among several competing possibilities, note which possibilities were ruled out, and then be able to specify the possibility that was actualized. What is more, the competing possibilities that were ruled out must be live

possibilities and sufficiently numerous so that specifying the possibility that was actualized cannot be attributed to chance. In terms of complexity, this is just another way of saying that the range of possibilities is complex. In terms of probability, this is just another way of saying that the possibility that was actualized has small probability.

All the elements in this general scheme for recognizing intelligent agency (that is, actualizing, ruling out, and specifying) find their counterpart in the complexity-specification criterion. It follows that this criterion makes precise what we have been doing right along when we recognize intelligent agency. The complexity-specification criterion pinpoints how we detect design.

## 7. Conclusion

Albert Einstein once said that in science things should be made as simple as possible but no simpler. The materialistic philosophy of science that dominated the end of the nineteenth and much of the twentieth century insists that all phenomena can be explained simply by reference to chance and/or necessity. Nevertheless, this essay has suggested, in effect, that materialistic philosophy portrays reality too simply. There are some entities and events that we cannot and, indeed, do not explain by reference to these twin modes of materialistic causation. Specifically, I have shown that when we encounter entities or events that manifest the joint properties of complexity and specification we routinely, and properly, attribute them, not to chance and/or physical/chemical necessity, but to intelligent design, that is, to mind rather than matter. Clearly, we find the complexity-specification criteria in objects that other human minds have designed. Nevertheless, this essay has not sought to answer the question of whether the criteria that reliably indicate the activity of a prior intelligent mind exist in the natural world, that is, in things that we know humans did not design, such as living organisms or the fundamental architecture of

the cosmos. In short, I have not addressed the empirical question of whether the natural world, as opposed to the world of human technology, also bears evidence of intelligent design. It is to this question that my colleagues Stephen Meyer and Michael Behe will now turn.

---

## NOTES

<sup>1</sup> Moses Maimonides, *The Guide for the Perplexed*, trans. M. Friedländer (New York: Dover, 1956), p. 188.

<sup>2</sup> *Ibid.*

<sup>3</sup> Isaac Newton, *Mathematical Principles of Natural Philosophy*, trans. A. Motte, ed. F. Cajori (Berkeley, Calif.: University of California Press, 1978), pp. 543–44.

<sup>4</sup> Pierre Simon de Laplace, *Celestial Mechanics*, 4 vols., trans. N. Bowditch (New York: Chelsea, 1966).

<sup>5</sup> See the introduction to Pierre Simon de Laplace, *A Philosophical Essay on Probabilities*, trans. F. W. Truscott and F. L. Emory (New York: Dover, 1996).

<sup>6</sup> See John S. Bell, *Speakable and Unspeakable in Quantum Mechanics* (Cambridge: Cambridge University Press, 1987).

<sup>7</sup> William A. Dembski, *The Design Inference* (Cambridge: Cambridge University Press, 1998).

<sup>8</sup> See Aristotle, *Metaphysics*, bk. 5, chap. 2, in *The Basic Works of Aristotle*, ed. R. McKeon (New York: Random House, 1941), p. 752.

<sup>9</sup> Francis Bacon, *The Advancement of Learning*, vol. 30 of *Great Books of the Western World*, ed. R. M. Hutchins (Chicago: Encyclopedia Britannica, 1952).

<sup>10</sup> Monod writes, “The cornerstone of the scientific method is the postulate that nature is objective. In other words, the *systematic* denial that ‘true’ knowledge can be got at by interpreting phenomena in terms of final causes—that is to say, of ‘purpose’.” In Jacques Monod, *Chance and Necessity* (New York: Vintage, 1972), p. 21.

<sup>11</sup> Jaki writes: “I want no part whatever with the position . . . in which science is surreptitiously taken for a means of elucidating the utterly metaphysical question of purpose.” In Stanley Jaki, *Chesterton, A Seer of Science* (Urbana, Ill.: University of Illinois Press, 1986), pp. 139–40, n. 2.

<sup>12</sup> Eliot Marshall, “Medline Searches Turn Up Cases of Suspected Plagiarism”, *Science* 279 (January 23, 1998): 473–74.

<sup>13</sup> Charles Darwin, *On the Origin of Species* (1859; reprint, Cambridge: Harvard University Press, 1964), p. 482.

<sup>14</sup> Strictly speaking, in *The Design Inference* I develop a “specification/small probability criterion”. This criterion is equivalent to the complexity-specification criterion described here.

<sup>15</sup> Michael Polanyi, “Life Transcending Physics and Chemistry”, *Chemical and Engineering News*, August 21, 1967, pp. 54–66; Michael Polanyi, “Life’s Irreducible Structure”, *Science* 113 (1968): 1308–12; Timothy Lenoir, *The Strategy of Life: Teleology and Mechanics in Nineteenth Century German Biology* (Dordrecht,

Netherlands: Reidel, 1982), pp. 7–8. See also Hubert Yockey, *Information Theory and Molecular Biology* (Cambridge: Cambridge University Press, 1992), p. 335.

<sup>16</sup> Nevertheless, fabrications may themselves derive from the design of those who observe an event after the fact, such as the archer who deliberately draws the pattern after the arrow has landed.

<sup>17</sup> Because specification is so central to inferring design, I need to elaborate on it. For a pattern to count as a specification, the important thing is not when it was identified but whether in a certain well-defined sense it is *independent* of the event it describes. Drawing a target around an arrow already embedded in a wall is not independent of the arrow's trajectory. Consequently, such a pattern cannot be used to attribute the arrow's trajectory to design. Patterns that are specifications cannot simply be read off the events whose design is in question—in other words, it is not enough to identify a pattern simply by inspecting an event and noting (i.e., “reading off”) its features. Rather, to count as specifications, patterns must be suitably independent of events. I refer to this relation of independence as *detachability* and say that a pattern is detachable if and only if it satisfies that relation.

Detachability can be understood as asking the following question: Given an event whose design is in question and a pattern describing it, would we be able to construct that pattern if we had no knowledge which event occurred? Here is the idea. An event has occurred. A pattern describing the event is given. The event is one from a range of possible events. If all we knew was the range of possible events without any specifics about which event actually occurred (e.g., we know that tomorrow's weather will be rain or shine, but we do not know which), could we still construct the pattern describing the event? If so, the pattern is detachable from the event.

For a better picture of what is at stake, I want to consider the following example. It was this example that finally clarified for me what transforms a pattern *simpliciter* into a pattern *qua* specification. Consider, therefore, the following event E, an event that to all appearances was obtained by flipping a fair coin 100 times:

THTTTHHTHHHTTTTTHTHTTHHHHTT  
 HTHHHHTHHHTTTTTTHTHTTTTHH  
 THTTTHTHTHHHTTHHHHTTHTTTHH  
 THTHTHHHHHTTHHTHHHHHTHHHHHTT

E

Is E the product of chance or not? A standard trick of statistics professors with an introductory statistics class is to divide the class in two, having students in one-half of the class each flip a coin 100 times, writing down the sequence of heads and tails on a slip of paper, and having students in the other half each generate purely with their minds a “random-looking” string of coin tosses that mimics the tossing of a coin 100 times, also writing down the sequence of heads and tails on a slip of paper. When the students then hand in their slips

of paper, it is the professor's job to sort the papers into two piles, those generated by flipping a fair coin and those concocted in the students' heads. To the amazement of the students, the statistics professor is typically able to sort the papers with 100 percent accuracy.

There is no mystery here. The statistics professor simply looks for a repetition of six or seven heads or tails in a row to distinguish the truly random from the pseudo-random sequences. In 100 coin flips, one is quite likely to see six or seven such repetitions. On the other hand, people concocting pseudo-random sequences with their minds tend to alternate between heads and tails too frequently. Whereas with a truly random sequence of coin tosses there is a 50 percent chance that one toss will differ from the next, as a matter of human psychology people expect that one toss will differ from the next around 70 percent of the time.

How, then, will our statistics professor fare when confronted with the event E described above? Will E be attributed to chance or to the musings of someone trying to mimic chance? According to the professor's crude randomness checker, E would be assigned to the pile of sequences presumed to be truly random, for E contains a repetition of seven tails in a row. Everything that at first blush would lead us to regard E as truly random checks out. There are exactly 50 alternations between heads and tails (as opposed to the 70 that would be expected from humans trying to mimic chance). What's more, the relative frequencies of heads and tails check out: there were 49 heads and 51 tails. Thus it is not as though the coin supposedly responsible for generating E was heavily biased in favor of one side versus the other.

Suppose, however, that our statistics professor suspects she is not up against a neophyte statistics student but instead a fellow statistician who is trying to put one over on her. To help organize her problem, study it more carefully, and enter it into a computer, she will find it convenient to let strings of o's and 1's represent the outcomes of coin flips, with 1 corresponding to heads and o to tails. In that case the following pattern D will correspond to the event E:

0100011011000001010011100

101110111000000100100011

0100010101100111100010011

01010111001101110111100

D

Now, the mere fact that the event E conforms to the pattern D is no reason to think that E did not occur by chance. As things stand, the pattern D has simply been read off the event E.

But D need not have been read off of E. Indeed, D could have been constructed without recourse to E. To see this, let us rewrite D as follows:

o

1

oo

o1

10  
 11  
 000  
 001  
 010  
 011  
 100  
 101  
 110  
 111  
 0000  
 0001  
 0010  
 0011  
 0100  
 0101  
 0110  
 0111  
 1000  
 1001  
 1010  
 1011  
 1100  
 1101  
 1110  
 1111  
 00

D

By viewing D this way, anyone with the least exposure to binary arithmetic immediately recognizes that D was constructed simply by writing binary numbers in ascending order, starting with the one-digit binary numbers (i.e., 0 and 1), proceeding then to the two-digit binary numbers (i.e., 00, 01, 10, and 11), and continuing on until 100 digits were recorded. It is therefore intuitively clear that D does not describe a truly random event (i.e., an event obtained by tossing a fair coin) but rather a pseudo-random event, concocted by doing a little binary arithmetic.

Although it is now intuitively clear why chance cannot properly explain E, we need to consider more closely why this is so. We started with a putative chance event E, supposedly obtained by flipping a fair coin 100 times. Since heads and tails each have probability 1/2, and since this probability gets multiplied for each flip of the coin, it follows that the probability of E is 1 in  $2^{100}$ , or approximately 1 in  $10^{30}$  (i.e., one in a thousand billion billion billion). In addition, we constructed a pattern D to which E conforms. Initially D proved insufficient to eliminate chance as the explanation of E since in its construc-

tion D was simply read off of E. Rather, to eliminate chance we had also to recognize that D exhibited a pattern independent of E (independence in this case meaning that D could have been constructed quite easily by performing some simple arithmetic operations with binary numbers). Thus, to eliminate chance we needed to employ additional *side information*, which in this case consisted of our knowledge of binary arithmetic. This side information enabled us to establish that D is independent of E (cf. the archer analogy, where the pattern functions as a fixed target that is independent of the arrow's trajectory). Consequently, this side information detaches the pattern D from the event E and thereby renders D a specification.

For side information to detach a pattern from an event, it must satisfy two conditions, a *conditional independence condition* and a *tractability condition*. According to the conditional independence condition, the side information must be conditionally independent of the event E. Conditional independence is a well-defined notion from probability theory. It means that the probability of E does not change once the side information is taken into account. Conditional independence is the standard probabilistic way of unpacking epistemic independence. Two things are epistemically independent if knowledge about one thing (in this case the side information) does not affect knowledge about the other (in this case the occurrence of E). This is certainly the case here, since our knowledge of binary arithmetic does not affect the probabilities of coin tosses.

The second condition, the tractability condition, requires that the side information enable us to construct the pattern D to which E conforms. This is evidently the case here as well, since our knowledge of binary arithmetic enables us to arrange binary numbers in ascending order and thereby construct the pattern D. But what exactly is this *ability to construct a pattern on the basis of side information*? Perhaps the most slippery words in philosophy are "can", "able", and "enable". Fortunately, just as there is a precise theory for characterizing the epistemic independence between an event and side information—namely, probability theory—so too there is a precise theory for characterizing the ability to construct a pattern on the basis of side information—namely, complexity theory.

Complexity theory, conceived now quite generally and not merely as a form of probability, assesses the difficulty of tasks given the resources available for accomplishing those tasks [see chap. 4 of my *The Design Inference*]. As a generalization of computational complexity theory, complexity theory ranks tasks according to difficulty and then determines which tasks are sufficiently manageable to be doable or tractable. For instance, given current technology we find sending a person to the moon tractable but sending a person to the nearest galaxy intractable. In the tractability condition, the task to be accomplished is the construction of a pattern and the resource for accomplishing that task is side information. Thus, for the tractability condition to be satisfied, side



information must provide the resources necessary for constructing the pattern in question. All of this admits a precise complexity-theoretic formulation and makes definite what I called "the ability to construct a pattern on the basis of side information".

Taken jointly, the tractability and conditional independence conditions mean that side information enables us to construct the pattern to which an event conforms, yet without recourse to the actual event. This is the crucial insight. Because the side information is conditionally and therefore epistemically independent of the event, any pattern constructed from this side information is obtained without recourse to the event. In this way any pattern that is constructed from such side information avoids the charge of being ad hoc. These, then, are the detachable patterns. These are the specifications.

<sup>18</sup> Dembski, *Design Inference*, chap. 5.

<sup>19</sup> Martin Gardner, "Arthur Koestler: Neoplatonism Rides Again", *World*, August 1, 1972, pp. 87-89.

<sup>20</sup> Walter Martin, *The Kingdom of the Cults*, rev. ed. (Minneapolis: Bethany House, 1985), pp. 127-30.

<sup>21</sup> Dembski, *Design Inference*, chap. 6, sec. 5.

<sup>22</sup> Richard Dawkins, *The Blind Watchmaker* (New York: Norton, 1986), pp. 47-48.

<sup>23</sup> *Ibid.*

<sup>24</sup> Cf. Bernd-Olaf Küppers, "On the Prior Probability of the Existence of Life", in *The Probabilistic Revolution*, vol. 2, ed. L. Krüger, G. Gigerenzer, and M. S. Morgan (Cambridge: MIT Press, 1987), pp. 355-69. Küppers agrees that Dawkins' METHINKS-IT-IS-LIKE-A-WEASEL example grasps an essential feature of the Darwinian mechanism.