REVIVING THE ARGUMENT FROM DESIGN

DETECTING DESIGN THROUGH SMALL PROBABILITIES†

William A. Dembski

Northwestern University
Department of Mathematics
Evanston, IL 60208

Portions of this work were presented at the biannual meeting of the ***Association of Christians in the Mathematical Sciences*** *(30 May 1991) and at the 50th anniversary meeting of the* ***American Scientific Affiliation*** *(29 July 1991).*

---

†This work was completed while I was a postdoctoral fellow in mathematics at Northwestern University under the direction of probabilist Mark Pinsky.

Reviving the Argument from Design
Detecting Design through Small Probabilities

William A. Dembski

Institute Apollos-Leonidas, Box 1216, Evanston, IL 60204

and

Northwestern University, Department of Mathematics, Evanston, IL 60208

1. The Old Design Argument

In its most general form the argument from design contends that the order and complexity exhibited in nature could not have sprung from nature, but instead required the skill and power of an intelligent designer.[1] The argument is one act in the philosophical drama known as metaphysics and ontology. In recent times metaphysics and ontology have fallen into disrepute. The reason generally cited is that past metaphysicians made a habit of introducing specious entities which on the one hand hindered the progress of science, and which on the other hand were rendered obsolete by science. Yet despite scientific naturalism's desire to be rid of metaphysics, there is a sense in which metaphysics is unavoidable, and this concerns the question of ontological commitment, i.e., what do we believe exists?[2]

The question about what exists and what constitutes the fundamental principle of existence has a long and pugnacious history. Two views have been notable in the West. One takes the fundamental principle of existence as supremely intelligent and rational; the other takes the fundamental principle of existence as unintelligent and material. These views are not only incompatible, but they have been slugging it out since the dawn of philosophical speculation. Thus even among the ancients we find Plato and Aristotle championing the primacy of the rational against atomists like Democritus and Epicurus who affirm the preeminence of matter in motion. The argument from design arises when these opposing views collide in the marketplace of ideas.

So long as the fundamental principle of reality is regarded as intelligent and rational, an argument from design is strictly speaking unnecessary. Within a framework that makes intelligence fundamental, order and complexity in nature are accounted as the work of an intelligence having the power and skill to manipulate nature with total precision and competence. When, however, nature is regarded as sufficient in itself, not requiring the services of an intervening intelligence, anyone who believes intelligence is fundamental faces a challenge. The challenge is to show that intelligence is not superfluous to nature, but necessary to account for the order we observe in nature.

Depending on one's sociological context, the challenge can be easy to meet. Thus David was blithely able to say, "The fool hath said in his heart there is no God."[3] Where theism dominates, no argument is needed. But when naturalism starts to be taken seriously, David's facile response becomes inappropriate. Now this is where the argument from design comes in. For an argument from design to be effective it must demonstrate not only that a designer explains the order we witness in nature, but also that explanations which fail to incorporate a designer are in some way inadequate. Against the argument from design the naturalist will argue that an explanation within the bounds of nature is not only possible but eminently plausible, thereby rendering the designer superfluous. Occam's razor is the naturalist's primary weapon for debunking an argument from design (e.g., "Why do we need God when we can explain everything without him?").

Christian theology has historically been the dominant shareholder in the argument from design. The apostle Paul himself gave the imprimatur to the design argument, certifying that it constitutes a valid theological enterprise, when he wrote:

> That which may be known of God is manifest in them; for God hath showed it unto them. For the invisible things of him from the creation of the world are clearly seen, being understood by the things that are made, even his eternal power and Godhead; so that they are without excuse.[4]

If those who deny God in spite of the natural order are without excuse, then surely they are without rational excuse. The invitation is clear and theologians have been ready to accept it: employ reason to demonstrate the existence of God from "the things that are made" (i.e., nature).

The first full-fledged argument from design that I'm aware of stems from the 4th century Cappadocian father Gregory of Nazianzus. His argument compares God not only to a lutemaker, but also to a luteplayer. In this respect, Gregory's argument extends William Paley's more recent (18th century) watchmaker argument—to parallel Gregory of Nazianzus, Paley would have had to introduce a watchwinder in addition to his famous watchmaker. Here is Gregory's argument:

> Now our very eyes and the Law of Nature teach us that God exists and that He is the Efficient and Maintaining Cause of all things: our eyes, because they fall on visible objects, and see them in beautiful stability and progress, immovably moving and revolving if I may so say; natural Law, because through these visible things and their order, it reasons back to their Author. For how could this Universe have come into being or been put together, unless God had called it into existence, and held it together? For

> every one who sees a beautifully made lute, and considers the skill with which it has been fitted together and arranged, or who hears its melody, would think of none but the lutemaker, or the luteplayer, and would recur to him in mind, though he might not know him by sight. And thus to us also is manifested That which made and moves and preserves all created things, even though He be not comprehended by the mind. And very wanting in sense is he who will not willingly go thus far in following natural proofs. . . .[5]

As the centuries passed, the argument from design was refined and tailored to the apologetic needs of the time. Aquinas, for instance, working within an Aristotelian context held that the following formulation of the argument was cogent:

> The fifth [argument for the existence of God] is taken from the governance of things. We see that things which lack knowledge, such as natural bodies, act for an end, and this is evident from their acting always, or nearly always, in the same way, so as to obtain the best result. Hence it is plain that they achieve their end not by chance, but by design. Now whatever lacks knowledge cannot move towards an end, unless it be directed by some being endowed with knowledge and intelligence, as the arrow is directed by the archer. Therefore some intelligent being exists by whom all natural things are ordered to their end; and this being we call God.[6]

So long as Aristotelian philosophy and science are presupposed, this argument works. But once it is denied that "nature works for a determinate end,"[7] the argument loses its punch.

It's worth comparing the arguments of Thomas Aquinas and Gregory of Nazianzus. Despite Gregory's reference to archaic musical instruments like the lute, his argument is more relevant to our age. Gregory concentrates on efficient causes, the stuff of modern science. Aquinas, on the other hand, looks to Aristotle and his final causes. Not only is Aristotelian science dead, but its scientific sterility is acknowledged even among 20th century Thomists.[8] To say that modern science ignores purpose in nature is an understatement—it actively avoids it. Nobel laureate Jacques Monod in his book *Chance and Necessity* goes so far as to define the scientific enterprise in terms of a systematic repudiation of purpose, goals, and final causes in nature. He 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."[9]

Monod is of course saying too much, and in saying too much conflates the scientific method with scientism. Aside from this philosophical point (which I don't mean to minimize), however, Monod does capture the spirit of the modern scientific enterprise. Insofar as causality is a scientifically reputable notion at all, science seeks to understand nature through efficient, not final causes.

With the advent of modern science in the 16th and 17th centuries, the argument from design returned to the form outlined by Gregory of Nazianzus, with its emphasis on specific objects and their efficient causes. This type of design argument became an intellectual growth industry during 17th, 18th, and early 19th centuries in England, and fell under the generic title of physicotheology. Names associated with this movement included Robert Boyle, Robert Hooke, John Ray, and William Paley. It is significant that despite the repudiation of Aristotle and his final causes in nature, the notion of purpose did not disappear from the writings of the physicotheologians. Granted, physical laws of nature à la Newtonian mechanics had no room for final causes. But within the physical universe governed by those laws purposes could still be discovered.

An analogy might be helpful here. Think of the cosmos as an oil painting on canvas. The canvas represents the fabric of the universe, i.e., its physical laws. The canvas is a precondition for the work of art that will appear on it. Aquinas's approach to the design argument was to find purpose in the canvas itself, something he could easily do within Aristotelian science. The physicotheologians, on the other hand, looked for purpose in the painting rather than canvas. This difference between design arguments is significant because it shows how design arguments can be accommodated to the science of the day, Aristotelian in the one case, Newtonian in the other.

Paley's watchmaker argument is worth reviewing at this point. Not only does it present the design argument in its heyday, evincing a confidence and vigor it has yet to recover, but it does so within a framework whose scientific commitments are not too foreign to ours. Paley's argument begins as follows:

> In crossing a heath, suppose I pitched my foot against a *stone*, and were asked how the stone came to be there; I might possibly answer, that, for any thing I knew to the contrary, it had lain there for ever: nor would it perhaps be very easy to show the absurdity of this answer. But suppose I had found a *watch* upon the ground, and it should be inquired how the watch happened to be in that place: I should hardly think of the answer which I had before given, that, for any thing I knew, the watch might have always been there. Yet why should not this answer serve for the watch as well as for the stone? why is it not as admissible in the second case, as in the first? For this reason and for no other, viz. that, when we come to inspect the watch, we perceive (what we could not discover in the stone) that its several parts are framed and put together for a purpose. . . . The inference, we think, is inevitable, that the watch must have had a maker: that there must have existed, at some time, and at some place or other, an artificer or artificers who formed it for the purpose which we find it actually to answer: who comprehended its construction, and designed its use.[10]

The form of Paley's argument is clear: watches presuppose watchmakers. More generally, contrivances presuppose contrivers. But there are instances of contrivance in the world too intricate and marvelous to have a human author. For such contrivances we must go outside nature and look to God himself as author. Paley therefore sets himself the task of enumerating putative instances of divine contrivance and recounting them in their most winsome light, arguing that these contrivances have no explanation apart from divine craftsmanship. Thus he examines the human eye, demonstrates what a marvelous instrument it is, and concludes that only a divine being could be responsible for it.

Since the criticisms of Hume, Kant, and Darwin, Paley's argument has fallen on hard times. This is not to say that Paley's argument has lost its popular appeal. But in secular academic circles it no longer carries any weight. Thus when a contemporary philosopher of religion like Richard Swinburne tailors his design argument to the 20th century intellectual, he is constrained by past criticisms of the design argument as well as his own scientific picture of the world. Since Swinburne, for instance, accepts the neo-Darwinian account that "simple animals and plants can be produced by natural processes from inorganic matter" and that "complex animals and plants can be produced through generation by less complex animals and plants,"[11] he cannot argue from contrivance in the plant and animal kingdoms to a divine creator.

The tendency in philosophy of religion these days is therefore to move the design argument to deeper levels. The laws of nature are supposed to be adequate to account for life (the key example of divine contrivance in the past). Nature's laws henceforward assume the role Paley previously assigned to God. But what about the very laws that make life possible? What about the order which these laws demonstrate? Current design arguments reason their way back to God in this more subtle way, starting from the laws of nature and the order they impart to the universe.[12] Note that there is a cost. Watches and eyeballs are readily understood by the man on the street. Scientific laws are not. Hence a design argument that looks for order at deeper levels has the practical problem of going unappreciated.

There is another problem with this approach to the design argument, however, and this concerns the question of explanation: What needs to be explained and what constitutes a valid explanation? Animals and their sensory organs do need to be explained. Animals haven't been on this planet forever. Did therefore God place them here or did nature place them here apart from God? This is a definite question of fact: trace back the genealogy of any organism; does the genealogy terminate in a primeval soup or in a divine act?[13] In the history of ideas attention is usually focused on arguments and counterarguments for choosing sides. Thus as experts in intellectual history we would analyze the reasons for Paley and Darwin coming down on opposite sides of the previous question. The point remains, however, that Paley and Darwin were addressing a question that required explanation.

On the other hand, it's not so clear whether natural laws themselves require explanation. Recall our analogy of the oil painting on canvas. Usually an artist is known by what appears on the canvas, not by the canvas itself. Nevertheless, might it be possible to infer an artist from a blank canvas? This type of question confronts any design argument that wants to make the laws of nature its point of departure. If physical laws correspond to a blank canvas, serving solely as a framework within which the cosmos unfolds, it's not clear what characteristics physical laws must have to enable us legitimately to infer a designer. Does the very existence of physical laws suffice to conclude design? Surely such a claim is question begging since design cannot be ascertained except within a cosmos already conditioned by those laws. The argument from design is supposed to hinge on an analogy, with human contrivance the point of departure. Once design is raised to so abstract a level as natural law, it's not clear any analogy is left.

This concludes our brief history of the design argument within Christian theology. Before turning to more detailed criticisms of the design argument, I need to clarify a point which up to now I have left deliberately vague. I have and shall continue to refer to the *argument from design* or the *design argument*. By now it's clear that there are many distinct philosophical arguments which go by that name. Thus when I refer to the design argument, I really mean a class of arguments. Given this usage it's natural to ask what is common to the class. Richard Swinburne offers a helpful summary of the features common to design arguments:

> [The design argument] is basically an argument by analogy, an analogy between the order in the natural world . . . and the patterns of order which men often produce. . . . It argues from similarity between phenomena of two kinds B and B* to similarity between their causes A and A*. In view of the similarities between the two kinds of order B and B*, the theist postulates a cause (A*) in some respects similar to A (men); yet in view of the dissimilarities the theist must postulate a cause in other respects different. All arguments by analogy do and must proceed in this way. They cannot postulate a cause in all respects similar. They postulate a cause who is such that one would expect him to produce phenomena similar to B in the respects in which B* [is] similar to B and different from B in the respects in which B* [is] different from B. All argument from analogy works like this.[14]

Let me stress Swinburne's point that the argument from analogy does not demand that the cause of **B*** match the cause of **B** in every aspect, but only in those aspects where **B** and **B*** are themselves similar. This observation meets several of Hume's celebrated objections to the design argument, viz., that we should expect the designer of the universe to comprise several individuals, fashion the universe out of pre-existing materials, assume embodied forms, etc., by analogy with human designers. The argument from analogy is also an argument from disanalogy, attempting to give a rational account of those differences which keep the analogy from constituting an identity.[15]

Schematically Swinburne's framework for characterizing the argument from design can be represented as follows, with arrows representing causal relations and dashed lines similarity relations:

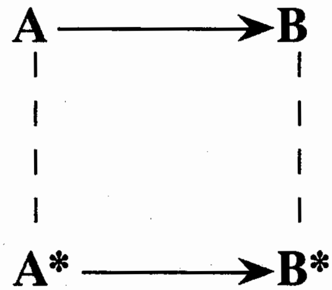

(1.1)

It will be instructive to see how critics of the design argument challenge schema (1.1). The criticisms of the design argument we take up next are formulated largely in reference to this schema.

2. Criticisms of the Old Design Argument

The most persuasive critics of the design argument historically have been David Hume, Immanuel Kant, and Charles Darwin. They are chiefly responsible for the state of disrepair in which the argument presently finds itself. We may distinguish the philosophical critique of Hume and Kant from the scientific critique of Darwin. If this distinction seems artificial, it probably is. Nevertheless, for as brief a discussion of the subject as I intend to give, the distinction is worth making. Moreover, the distinction underscores a significant historical fact, namely that until science offered an explanation of apparent design apart from a designer—as in the theory of natural selection—the philosophical critique of the design argument was not all that compelling. Richard Dawkins, though defending Darwinism throughout his book *The Blind Watchmaker*, offers the following remarkable concession:

> I feel more in common with the Reverend William Paley than I do with the distinguished modern philosopher, a well-known atheist, with whom I once discussed the matter [of apparent design in nature] at dinner. I said that I could not imagine being an atheist at any time before 1859, when Darwin's *Origin of Species* was published. "What about Hume?" replied the philosopher. "How did Hume explain the organized complexity of the living world?" I asked. "He didn't," said the philosopher. "Why does it need any special explanation?"
>
> Paley knew that it needed a special explanation; Darwin knew it, and I suspect that in his heart of hearts my philosopher companion knew it too.... It is sometimes said that [Hume] disposed of the Argument from Design a century before Darwin. But what Hume did was criticize the logic of using apparent design in nature as *positive* evidence for the existence of a God. He did not offer any *alternative* explanation for apparent design, but left the question open. An atheist before Darwin could have said, following Hume: "I have no explanation for complex biological design. All I know is that God isn't a good explanation, so we must wait and hope that somebody comes up with a better one...." Darwin made it possible to be an intellectually fulfilled atheist.[16]

Following Dawkins we might say that philosophical criticisms of the design argument challenge schema (1.1) by replacing it with the following schema:

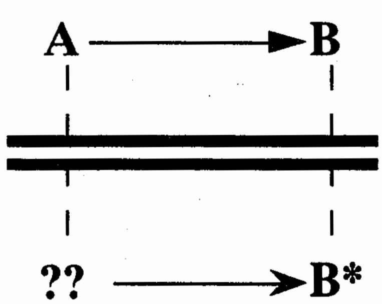

(2.1)

The two thick horizontal lines indicate a breakdown of the analogies represented by the dashed vertical lines; the question marks indicate the resulting ignorance of the cause of B*. The question marks render this critique unsatisfying, for the critique offers no counter-explanation.

Against an unsatisfying philosophical criticism, scientific criticisms of the design argument challenge schema (1.1) by replacing it with the following schema:

$$\mathbf{A} \longrightarrow \mathbf{B}$$

(2.2)

$$\mathbf{Nat} \longrightarrow \mathbf{B^*}$$

The vertical lines which had depicted a fuzzy similarity relation are gone. **B*** is to be understood on its own terms and its cause is to be identified with a natural process. In fact, the scientific critique of the design argument splits the top level **A** → **B** from the bottom level **Nat** → **B*** so that any analogy between the two levels becomes superfluous. Indeed, a scientist would prefer to replace schema (2.2) with the simpler

$$\mathbf{Nat} \longrightarrow \mathbf{B^*} \qquad (2.2')$$

There is a third, probabilistic criticism of the design argument which has its roots in the philosophies that postulate eternally recurrent cycles. Hume states it in his *Dialogues Concerning Natural Religion*. Nevertheless, in a suitably rationalized or mysticized form, Hume's criticism can be attributed to anyone from Spinoza to Epicurus to an eastern pantheist. I choose Hume's version of the criticism because it is suitably naturalized to modern tastes, having been stripped of any inherent mysticism. Schematically, the probabilistic criticism looks as follows:

$$\mathbf{Chance} \longrightarrow \mathbf{B^*} \qquad (2.3)$$

This may appear as no criticism at all since to say that chance caused **B*** appears to say that **B*** requires no explanation—blind matter in motion suffices to explain **B***.

The probabilistic criticism, however, is more subtle than this caricature suggests. Here is how Hume puts it:

> A finite number of particles is only susceptible of finite transpositions: and it must happen, in an eternal duration, that every possible order or position must be tried an infinite number of times. This world, therefore, with all its events, even the most minute, has before been produced and destroyed, and will again be produced and destroyed, without any bounds and limitations. No one, who has a conception of the power of infinite, in comparison of finite, will ever scruple this determination.[17]

This criticism hinges on what modern probabilists call the strong law of large numbers. One consequence of this law is that any event with positive probability, however small, will definitely occur if given sufficiently many (independent) opportunities to occur. In general, for events of probability p, sufficiently many corresponds to 1/p trials (for instance, the event of dealing a royal flush in poker has probability p around one in a million; hence, one expects to deal such a hand about every million deals). Since an event of probability p can be expected to occur every 1/p trials, and since 1/p is strictly finite, it follows that with an infinite number of trials this event will occur infinitely often.

Hume's argument presupposes that any configuration of the universe is only one of finitely many possibilities, each with an extremely small positive probability, but also with infinitely many opportunities to occur. Since infinite opportunity always swamps finite probability, all possible configurations of the universe are not only guaranteed to occur, but to recur infinitely often. Although Hume's probabilistic intuitions are impeccable, his cosmological assumptions are not.[18] Not only does modern cosmology support a universe finite in extent and number of particles, but it also supports a universe finite in duration (cf. the big bang). If these strong finiteness conditions obtain,[19] then Hume's premises are simply wrong and his argument founders. Nevertheless, we shall return to Hume's probabilistic criticism inasmuch as there are more subtle versions which will prove relevant our eventual reformulation of the design argument.

Having outlined the three generic criticisms of the design argument schematically, I want to turn to **B***, which appears in all the schemas recorded so far. The argument from design starts with a complex ordered system **B** produced by a human agent **A**, next finds a complex ordered system **B*** in nature which could not have been produced by a human agent, and finally attributes **B***'s cause to a unique divine agent **A***. To ascribe divine causality to **B*** is philosophically problematic, as was pointed out by Kant.[20] Although **B*** will typically exceed **B** in grandeur and complexity, that **B*** should require God for a cause does not follow. Since any given **B*** is finite, inference to an infinite God is not strictly speaking warranted. Thus according to Kant the rational agent **A*** responsible for **B*** need merely be an architect capable of producing **B***—albeit an architect whose powers and skill far exceed ours. Kant was willing to concede that natural reason could legitimately infer such an architect. Thus

Kant did not dismiss the argument from design as a total waste; indeed, inference to an architect short of God is still progress in the face of scientific naturalism. Moreover, once an architect is admitted into the discussion, an argument on the grounds of simplicity can be made that the most straightforward, least problematic candidate for the architect responsible for **B*** is none other than the "one of infinite power, knowledge, and freedom, i.e. God."[21]

B* is the pivot on which the argument from design turns. Whether the design argument is cogent depends on our analysis of the pivot. Two questions therefore immediately arise: What properties should a pivot **B*** possess to promote a successful design argument? Are there any actual systems **B*** which possess these properties? A science fiction example might be helpful here. Imagine a fantastic Rube Goldberg device. Let us say that the purpose of the device is to pour a cup of tea, but that to accomplish this simple task, a causal chain of ten trillion events needs —in typical Rube Goldberg fashion—to be set in motion: thus the chain starts with a million dominoes tumbling down in succession which in turn causes steel balls to run down inclined planes which in turn causes a gun to fired, etc., etc., which finally ends in the cup of tea being poured. This ultra Rube Goldberg device, let us say, requires as much space as the entire state to Texas. Suppose Christopher Columbus in 1492 had happened to land in Texas and found this device. What would he have thought? How would he have explained it?

Columbus's options are limited. First he might have looked to human agency. But this seems impossible (we assume human technology in the alternate world is identical to our's). The device could have arisen by chance. Again this seems insupportable: the device does serve a purpose, albeit a silly one, and is fantastically complex. The laws of nature might be so constituted as to play an occasional joke. The workings of nature may require us to assimilate a Law of the Cosmic Joke to our physics so that fantastic Rube Goldberg devices will on rare, though not implausibly rare, occasions arise according to this law. But without a physical theory which describes how the Law of the Cosmic Joke works and how it coheres with the rest of our physics, this last option seems again implausible. Finally, some alien intelligence qua architect might have visited this planet and placed it here. This seems the least implausible explanation so far, although Columbus would be loath to attribute the device to an all-wise God—surely there are more intelligent ways to pour a cup of tea.

The last example may seem silly, but it underscores an important point, namely that inference to a superhuman architect cannot be rejected on a priori grounds. Indeed, there are ways the world could be which would make it difficult to reject a superhuman interloper. Moreover, we can imagine devices which serve more serious purposes than the preceding Rube Goldberg device, and whose authors are worthy of greater respect than the architect of that Rube Goldberg device. For instance, a concrete device which answers all mathematical questions humans might pose—either by giving a proof or by showing that the question is undecidable—exceeds the capability of human engineers (Gödel's incompleteness theorem is relevant here), and would give the superhuman designer no little respect within the mathematical community.

It's time to return to reality. As I see it, a given **B*** has four possible explanations: a superhuman architect, a human architect, chance, or nature operating according to natural laws. Schematically we can lay out these possibilities as follows:

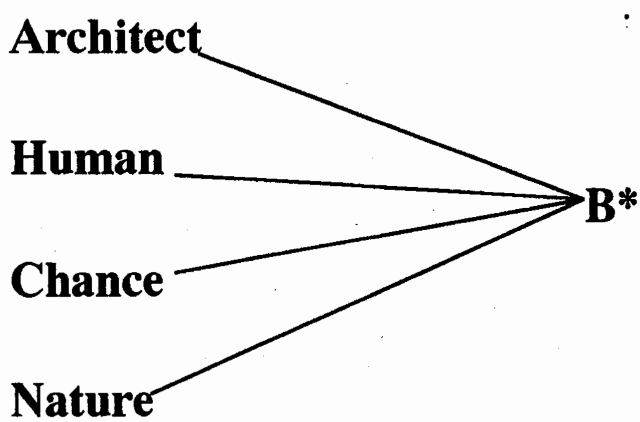

(2.4)

Usually human agency and chance can be excluded from the start. Indeed, let us agree to ignore pivots **B*** where human agency and chance are serious possibilities—the design argument has enough problems without introducing weak pivots. Hence the design argument can be understood as an attempt to make superhuman, intelligent causation of **B*** more plausible than natural causation. Rebuttals to the design argument are therefore intended to reverse this order of plausibility. Thus if **B*** is a virus or a living organism or a planet or a solar system or any physical object up to and including the cosmos, arguments and counter-arguments will be made to credit respectively a superhuman architect or nature.

Although we permit the cosmos to serve as a pivot **B*** (and indeed it has been a much favored pivot among advocates of the design argument), there is a way of taking **B*** as the cosmos which is relatively recent to the design argument. The method philosophy of religion currently favors in the design argument (insofar as it favors the design argument at all) is to let **B*** comprise the cosmos in the sense of the natural laws governing the cosmos. There is a reason for this move. With the advance of scientific knowledge, the laws of nature have come to appear increasingly powerful, giving any would-be architect less and less to do. By letting **B*** comprise the laws which govern the cosmos, one short-circuits the god-of-the-gaps problem inherent in design arguments where the pivot **B*** is a physical object; for if **B*** is an object, there is always the possibility that **B*** can be explained as the effect of some natural cause. But if **B*** comprises the laws which govern the cosmos, then the question is no longer about what causes lie behind this or that object, but about the very nature of physical causality: Do the laws themselves need explaining, and if so do they need to be explained via a designer? In the past this question has usually been answered by the cosmological argument. Nevertheless, recent formulations of the design argument attempt to infer an architect when the pivot **B*** comprises the laws of nature.

Observe what becomes of schema (2.4) when the laws of nature are substituted for **B***. When this substitution is made, **B*** becomes identified with one of its possible explanations, viz. **Nature**. An obvious question now arises: Can an architect explain nature better than nature can account for herself? Why does nature need an explanation? Any design argument which addresses this question will clearly appeal to order and regularities in nature. A design argument which therefore takes **B*** as the laws governing the cosmos really ends up being a meta-design argument, or to use Kantian terminology, a "transcendental" design argument, looking for conditions of the possibility of (designed) order in the cosmos.

Meta-design arguments like this begin with specific instances of order in the cosmos, agree that immediate causes of these instances are natural and in accord with natural laws, but then contend that an architect is needed to account for the very laws. The starting point for the meta-design argument is therefore a cluster of phenomena each of which has a naturalistic explanation, but when taken jointly evince an order which requires an explanation for the possibility of naturalistic explanation. The meta-design argument can be represented by altering schema (2.4) as follows:

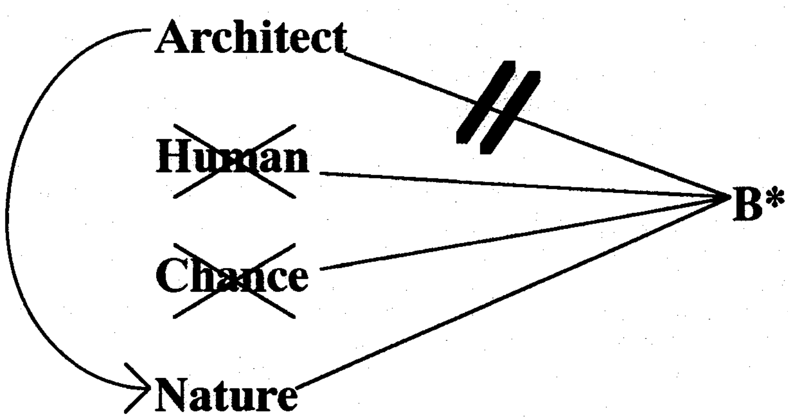

(2.5)

Chance and human agency are excluded, the architect is no longer directly responsible for **B***, but instead is responsible for natural laws which are in turn responsible for **B***.

Even at so general a level of abstraction as schemas (2.4) and (2.5), the design and meta-design arguments have their problems. The design argument tries to make the existence and causal power of an architect an empirical question, but in so doing depends on the science of the day. Hence a design argument based on the pivot **B*** fails to the degree that current scientific theory provides a naturalistic explanation for **B***. Lately the tendency in science has been to augment the power of natural law to account for phenomena formerly attributed to divine agency. It would seem, therefore, that design arguments should become less and less compelling as science progresses. But this doesn't follow.

By way of analogy, consider the state of mathematics before Gödel proved his famous incompleteness theorem. Hilbert had declaimed, "In mathematics there is no *ignorabimus*."[22] Hilbert thereby denied the existence of undecidable problems in mathematics. This was a supreme optimism, that by finitary methods all problems in mathematics could be decided. It would seem that by augmenting our mathematical knowledge, we could only confirm Hilbert's optimism. But instead, by augmenting our mathematical knowledge with an impossibility result, and thus placing limits on what can be accomplished in mathematics generally, Gödel decisively undermined Hilbert's optimism.

So too with the argument from design, an impossibility result can conceivably resuscitate a design argument whose stock has dwindled under a previous scientific paradigm. Impossibility results exist in physics. For instance, the speed of light is considered an upper bound on the communication speed of signals. Physical theory would therefore keep AT&T from investing in a communications system between earth and Mars (once it's been colonized) that required interplanetary telephone conversations to occur without time lags (as they do here on earth). Similarly, molecular biology may some day discover a physical property of living things which no amount of tinkering with nonliving things can duplicate. Should such a physical property be discovered, a design argument in the manner of Paley with pivot **B*** = {living things} might again flourish.

Note that in the last instance I am not advocating vitalism, nor am I unaware that current molecular biology utterly rejects a chasm between living and nonliving things. My point is simply that when the design argument is wedded to science, its fortunes depend on what naturalistic explanations of the pivots **B*** are currently available; and since science and the resulting naturalistic explanations are not fixed, the stability of a design argument over time is always in danger. Given a pivot **B***, science may plead ignorance about it, offer a naturalistic explanation for it, or say that no naturalistic explanation is possible. None of these responses to **B*** is etched in stone. It may happen that science settles on one response and never changes. But there are no guarantees. Science may well execute a random walk among its possible responses to **B*** as long as there are scientists.

Although the meta-design argument begins with concrete regularities in the universe, it is not tied to any specific scientific paradigm. This gives it an advantage over design arguments which depend on what naturalistic explanations are currently available or fashionable. The meta-design argument therefore sidesteps the scientific seesaw on which the design argument oscillates. The meta-design argument starts with instances of order in the universe, and then freely admits that the cause of such order is consistent with the laws of nature. But what about these very laws—don't they need a designer? The meta-design argument wants to answer this question affirmatively. How can it do this?

Several strategies are available. One strategy examines anthropic coincidences in the laws of nature—e.g., "if certain fundamental constants of physics had been altered ever so slightly, then life would have been impossible." The standard countermove is to assert the anthropic principle—viz., "if the fundamental constants had been altered as you say, then you wouldn't be here to know the difference." The anthropic principle asserts that anthropic coincidences don't require explanation for the simple reason that there would be no one to do any explaining if the coincidences hadn't occurred. The meta-design argument must now argue that anthropic coincidences cannot be dismissed so easily, that they really do need an explanation, and that the most plausible explanation is a superhuman architect. Other strategies for handling the meta-design argument include aesthetic, value-laden appeals (e.g., the laws of nature evince beauty and simplicity) and considerations of subjective probability (e.g., using Bayes' theorem to examine the probability of a superhuman architect given the laws of nature).[23]

This is only the merest sketch of how the meta-design argument might be applied in practice. Whether such an argument can be made compelling and what is required to make it compelling are questions I leave for another occasion. The point I want to make, however, is that what the meta-design argument gains by insulating itself against god-of-the-gaps assaults, it loses at the gut level of raw, common sense appeal. On the other hand, the standard design argument retains its gut level appeal[24]—the objects of nature are after all quite marvelous—but is vulnerable to the naturalistic explanations science has currently to offer. Unfortunately, the standard design argument is tied too closely to the fortunes of the prevailing science, and is therefore always suspect.

What we want is a design argument which is empirically grounded, but cannot be confounded by the winds of scientific change. I claim such an argument is possible, but requires we probe more deeply into what constitutes purpose and design. An analysis of design is therefore the next topic we consider.

3. Design and Designatum

Instead of explicating the entire cluster of concepts relevant to the design argument (e.g., *purpose*, *goal*, *end*, *telos*, *function*, *design*, *regularity*, *order*, *complexity*, and *pattern*), I shall concentrate on the one I consider central and analyze it, viz., design. Designs are blueprints. Designs are therefore incomplete, always pointing to realities beyond themselves. What realities? Why the things to be constructed from the designs. Designs always imply a duality—the design as well as the object(s) built from the design. Let us call any object constructed according to the specifications of a design a *designatum*. If there are multiple copies of an object sharing the same design, we refer to these as *designata*. In general, however, one instance of a designatum will suffice for our purposes. Thus our attention will usually be restricted to pairs of the form (design, designatum). We shall refer to design-designatum pairs as instances of design. Thus *design* is the generic term. The following picture represents the relation between design and designatum:

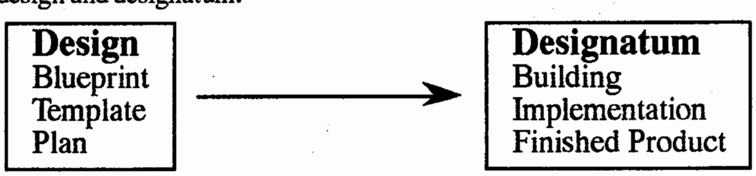

(3.1)

The arrow indicates not so much causal connection as conceptual priority. We shall return to this point later.

To clarify the design-designatum relation, let us apply it to Paley's watchmaker example. The watch discovered in the heath is the designatum. The watchmaker's mechanical picture of the watch, whether in his mind or on paper is the design. Observe the the design-designatum relation does not presuppose a designer. In the watchmaker example, if the design is reified on paper (i.e., if precise plans for the construction of the watch are written down), then together with the designatum (the watch), this design-designatum pair can be analyzed without reference to a human or otherwise intelligent agent. The source of both the design and the designatum are left entirely open within this framework. The old argument from design invariably embeds the design-designatum framework (3.1) within the following more elaborate framework:

Designer → Design → Builder → Designatum (3.2)

When it comes to artifacts, this picture is perfectly accurate: a designer conceives a design which is handed to a builder who constructs the designatum according to the specifications of the design. Thus an architect draws the plans which a construction company follows to erect a building. Thus engineers in Detroit make blueprints for auto plants to produce cars we eventually drive. The same pattern recurs over and over. (3.2) is a perfectly adequate picture of human contrivance. Sometimes the picture can be simplified, as when the designer and builder are one and the same. But this is the picture which holds generally—for human contrivance.

This picture, however, cannot be the starting point for a design argument, for in this case the designer is presupposed. In (3.2) the designer is prior to the design, the builder, and the designatum. But this is precisely what an argument from design is supposed to establish. Hence to make (3.2) the basis of a design argument is to beg the question. We therefore stick with (3.1). (3.1) does not presuppose a designer. The way I intend to approach the design argument is to start with a design-designatum pair, analyze it on its own terms, and from this analysis—without presupposing a designer—conclude that such a pair can best be understood within an explanatory framework that includes a designer. (3.2) therefore becomes the conclusion, not the premise of my design argument.

Let us examine the relation between design and designatum more closely. Since this relation does not presuppose a designer, it does not fall prey to the teleological problems that beset design when design is automatically referred to a designer. Consider the case of Paley and what he had to say about purpose in the design of a watch:

> When we come to inspect the watch, we perceive . . . that its several parts are framed and put together for a purpose, e.g., that they are formed and adjusted as to produce motion, and that motion so regulated as to point out the hour of the day.[25]

By itself there is nothing wrong with Paley's functional understanding of the watch. Indeed, watches are made to tell time. The problem with this functional understanding of the watch, however, is that it is inextricably tied to human agency and intentionality—timekeeping is after all a purely human preoccupation. Intrinsic to Paley's notion of purpose is an intelligent agent—the watchmaker who gave it purpose. Paley presupposes the designer in the design.

For Paley design is evidenced solely in the purposes of a rational agent. Contrast this with design as it arises within a design-designatum framework. Within such a framework design is evident in Paley's example because there is a watch (the designatum) and a blueprint (the design) from which the watch can be constructed. The design and designatum are physical objects (assume the design has been written down). All that is required of the design is that it give us an unambiguous specification of the designatum. The blueprint need not explain anything about the time-keeping or motion of the watch. What it must do is list all the components of the watch and show how they are to be assembled. A design-designatum pair is evident whenever an idiot with no greater capacity than obeying simple instructions can execute the design and thereby produce the designatum.[26]

Earlier we said that design is conceptually prior to designatum. This is certainly true of designs that spring from human ingenuity. The question, however, arises whether there is an intrinsic way of distinguishing design from designatum, particularly when the role of intelligence (i.e., a designer) as the source of the design is unclear. In fact, a subtle duality exists between design and designatum. Often it is evident that a design-designatum pair is present without it being clear what plays the role of design and designatum. Consider for instance the latest American jet fighter and its blueprints. This pair constitutes a design-designatum pair. Which is the design and which is the designatum seems clear: any Lockheed engineer will tell you that the blueprints constitute the design and the jet the designatum (given our manner of speaking). But now change the context. Suppose this jet was stolen by a hostile power and taken to an airbase in that hostile nation. The scientists of that nation decide they want to mass produce our jet. Painstakingly do they disassemble our jet, noting each component (its tensile strength, its composition, its relative position, etc.). After this elaborate process they produce a blueprint from which a jet identical to the one stolen can be mass produced. Given our original plane and the blueprint produced in this way, what is the design and what is the designatum? I think it's clear that the design is the original plane and the designatum is the blueprint. In general, for cases of plagiarism and patent infringement the role of design and designatum are reversed.

Although the theory of design I'm developing is designer-independent (no designer is presupposed), the examples I've given up till now have all involved human designers. Let us therefore consider an example of design free of human designers. The example I have in mind involves cell division and the genetic code. When a cell divides, it provides its daughter cells with exact copies (modulo highly infrequent copying errors) of its genetic makeup. Let us for now ignore the molecular biology which underlies the encoding of genetic information, and simply assume that the parent cell encodes its genetic makeup in a long string of 0's and 1's—call it σ—which it then passes on to its daughter cells upon cell division. Consider now the parent cell C_p and one of its daughter cells C_d. Both C_p and C_d possess the same genetic makeup which we denote by σ. C_p therefore possess an instance of σ, which we can call σ_p, as does C_d, which we can call σ_d. Whereas σ is an abstract representation of the genetic information inherent in C_p and its successors, σ_p and σ_d are concrete instantiations of this abstract σ.[27] When I say concrete, I mean physical—σ_p and σ_d will in fact be polynucleotides.

Now I claim that the relation of σ_p to σ_d is that of design to designatum. Although σ_p and σ_d are chemically indistinguishable, σ_p serves as the template from which σ_d is constructed upon cell division. Our picture is that of an unthinking chemical machine which inputs σ_p and outputs σ_d. The point I want to stress is that no designer is immediately apparent for this design-designatum pair (σ_p,σ_d). Granted, an 18th century physicotheologian would now argue that the cell itself shows evidence of design, and thus what I am calling unintelligent design in the pair (σ_p,σ_d) is really part of a grander, overall design. My point, however, remains, namely that the pair (σ_p,σ_d) is an instance of design which can be understood without recourse to a designer.

Is it necessary that a causal connection exist between design and designatum with the direction of causality proceeding from design to designatum? Alternatively, in a design-designatum instance, does the designatum presuppose a cause which depends fundamentally on the design. By now it's clear that the design-designatum relation does not presuppose a designer. Hence we needn't embed (3.1) in (3.2). If, however, we eliminate the designer from (3.2), do we obtain an appropriate representation of the design-designatum relationship? In the following modification of (3.2), is it necessary to retain the builder?

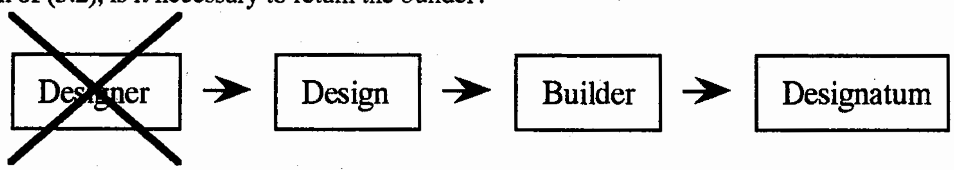

(3.3)

In the examples we've considered so far, there has always been a builder—someone or something to implement the design and thereby produce the designatum. Is this necessary to our conception of design? Earlier I said that designs are blueprints and designata are objects constructed according to the specifications of these blueprints. It would appear therefore that a builder is unavoidable in any design-designatum instance.

Is the builder in fact unavoidable? Yes and No. It will be important to my subsequent argument that design-designatum instances not require an actual, known builder. The reason for foregoing this requirement is to avoid prejudging the causal relationship between design and designatum. Hence it will be important that design-designatum instances be perspicuous without there being a causal story of how this particular design relates to that particular designatum. We've already seen an example of a design-designatum instance where the causal connection was unclear: in the fighter jet example, is the blueprint causally prior to the aircraft or vice versa? The answer depends on one's perspective, whether that of Lockheed's engineers or that of the hostile power. This is the problem of duality, where design and designatum can interchange roles.

There is a deeper problem, however, about the way design and designatum are causally connected. In general to say that A and B are causally related means that either A causes B or B causes A or some third event C causes both A and B. Now duality addresses only the first two possibilities. What about design-designatum instances in which we know for a fact that no builder was involved in constructing one from the other? Recall the previous example of cell division. We considered a parent cell C_p and one of its daughter cells C_d. Both cells share the same genetic makeup σ, instantiating it as σ_p in the one case, σ_d in the other. Moreover, the causal connection is clear: the parent passes the full complement of its genetic makeup to its offspring. The relevant design-designatum pair is therefore (σ_p,σ_d).

In this example we considered only one daughter cell. Nevertheless, when a cell divides, it inevitably forms two daughter cells. Thus C_p had two daughter cells, call them C_d and $C_{d'}$, having genetic endowment respectively σ_d and $\sigma_{d'}$, both instances of σ. What about the pair $(\sigma_d,\sigma_{d'})$? Can we understand it as a design-designatum pair even though the causal connection between σ_d and $\sigma_{d'}$ is mediated through a common cause, σ_p? I definitely want to include pairs like $(\sigma_d,\sigma_{d'})$ within the design-designatum framework. To do this, however, it must be possible to pick out design-designatum instances independent of any causal history that might connect the design and designatum. Here is how we do it.

What I said about designs being blueprints and designata being objects constructed according to the specifications of these blueprints now needs a slight modification. What I really mean to say is that designs are blueprints and designata are objects *that can be* constructed according to the specifications of these blueprints. The role of the builder is therefore not eliminated but left open ended. The question is not how the design and the designatum are in fact related, but *how they could be related*. We still have to tell a story about how a builder could conceivably use the design to construct the designatum, but we permit ourselves a counterfactual story. The story could coincide with reality, as in the case the parent cell C_p giving birth to the daughter cell C_d and thereby eliciting a causal connection between σ_p and σ_d. Nevertheless, the same account serves to relate σ_d and $\sigma_{d'}$ as a design-designatum pair: $\sigma_{d'}$ could have been produced from σ_d upon cell division. Another way of making $(\sigma_d,\sigma_{d'})$ a design-designatum pair is to say that $\sigma_{d'}$ could have been produced by an incredibly sophisticated genetic engineer who simply copied σ_d.

If this counterfactual approach seems outlandish and smacks of science fiction, think again. What if you were simply presented with the genetic materials σ_d and $\sigma_{d'}$ and given no knowledge of their causal history? You would say that they were identical (better yet, isomorphic). Fine. And how do you know they are identical? Well, by observing that one could be copied from the other. I propose to view identity as the copying of a prototype: take the prototype (the design), give it to a copy machine, and get out a copy (the designatum). In this way the machine's input becomes the design, its output the designatum. Does this mean that photocopy machines produce design-designatum instances whenever they make copies? The answer is yes.

There is a reason for this apparent silliness. As you travel around the universe, what enables you to recognize design? What are the conditions for the possibility of recognizing design? Do you need to presuppose a designer? Do you need to know the builder who took the design and therewith fashioned the designatum? If you do, then you unduly limit what you can discern as design. Consider again Paley's watch in the heath example, but this time assume it's not a human who stumbles across the watch, but a Plutonian (i.e., an inhabitant of the planet Pluto). To the Plutonian, what distinguishes the watch from a rock? The Plutonian does not know the purpose of the watch. He does not know how to tell time. If, however, he stays on earth long enough, he will discover other watches identical to the watch he discovered in the heath. He may even come across the watchmaker's blueprints for the watch. In either case he will have a design-designatum match which should alert him to design. In contrast, no "rock in the heath" will ever alert our Plutonian friend to design. The reason is that no two rocks are identical, nor is it likely that a specification for the rock exists which would enable a sculptor to copy that rock. In short, an arbitrary rock will never serve as a design (resp. designatum) since its complementary designatum (resp. design) doesn't exist. Watches always have such a complement, either because they were mass produced or because there is a blueprint.

Identity is a special case of design that will command much of our attention. In fact, identity in the sense of specification lies at the heart of the probabilistic design argument I have yet to formulate. Design is, however, a broader notion than identity. In general, blueprints differ substantially from their implementation. Indeed, the difference may be so substantial that without special knowledge of how the designatum can be constructed from the design, the presence of design may go undetected. The question therefore arises what constructions of designata from designs shall count as instances of design.

Not all constructions are permissible. Consider for instance a large slab of marble, call it S. Consider next the pair (S,D) where D is Michelangelo's David. Does (S,D) constitute a design-designatum instance? This seems utterly silly. Yet one can argue that a statue isomorphic to Michelangelo's David is inherent in S (modulo some excess marble which needs to be removed[28]). Now find a machine which removes this excess marble from just the right places. In this way S becomes a design and D a designatum for the design-designatum pair (S,D).

What's wrong with this example? To appreciate the difficulty, contrast it with the following example—an example we do want to count as a design-designatum instance. Consider the following two strings of characters:

```
A:      TO BE OR NOT TO BE, THAT IS THE QUESTION?
B:      UP CF PS OPU UP CF, UIBU JT UIF RWFTUJPO?
```

String B can be derived from string A by moving each letter of A up a notch in the alphabet. Similarly, A can be derived from B by moving each letter down a notch. A and B are related cryptographically, and the way to get from one to the other is by applying the appropriate encoding or decoding function.

Now why should (B,A) (or for that matter (A,B)) count as a design-designatum instance, but (S,D) fail? Thusfar all we've required of a putative design-designatum pair is that it be possible to construct the second component from the first. We placed no limitation on the possible builders which could transform a putative design into a putative designatum. It's now clear, however, that some restrictions must apply. In fact, there is one overriding restriction. To see it, consider what must be true of a builder that transforms S into D (this is the problem case). How could a builder possibly get from a slab of marble S to a copy of Michelangelo's David D? Let us assume that builders function reliably, not transforming designs randomly into designata. Then the only way I can imagine a builder transforming S into D is by already possessing a copy or blueprint of D before the actual transformation from S into D is effected. A builder cannot start from S alone to obtain D. Rather, the builder must already possess D in some form and therewith construct the actual D from S.

Compare this with the pair (B,A), which epitomizes the problem of cryptography. Cryptography's problem is twofold, transforming plaintext into ciphertext (encoding) and transforming ciphertext back into plaintext (decoding). Encoding and decoding must be quick and painless if a cryptographic system is to find widespread use. For this reason the actual method of encoding and decoding must not depend on the messages handed to the encoding and decoding functions. Thus when decoding B, we don't need a copy of A on hand. Rather we have a simple rule which does not presuppose A, but simply tells us to move every letter of the alphabet down one notch. In this way we recover A from B.[29]

The difference between (S,D) and (B,A) can be characterized this way: any builder/machine which transforms S into D anticipates D; on the other hand, there is some builder/machine which transforms B into A without anticipating A. We speak respectively of *anticipating* and *nonanticipating* builders/machines. An anticipating machine is a question begging machine. It is supposed to construct the designatum from the design, but in fact the designatum is already available to it (either concretely or abstractly). A nonanticipating machine, on the other hand, has only the design at its disposal and must construct the designatum without, if you will, any foreknowledge of it. We therefore require of any design-designatum instance not only that it be possible to construct the designatum from the design, but also that this construction can be accomplished by a machine which does not anticipate the designatum, i.e., by a nonanticipating machine. Note that the pair (U,V) where U and V are identical (isomorphic) therefore always constitutes a design-designatum instance inasmuch as a copy machine—which is always nonanticipating—suffices to produce V from U.

One further point about design-designatum pairs needs to be clarified. We saw how (S,D) failed as a design-designatum instance once we required that the transformation procedure taking S into D be nonanticipating. Consider now a pair similar to (S,D), but which reverses the order of design and designatum: (D,s). D is still Michelangelo's David, but s is a little slab of marble, say a one inch cube. It's certainly possible to derive s from D by suitably vandalizing D: simply slice off a one inch cube from D (D is after all a big statue). Whether there exist nonanticipating machines which derive s from D can be debated. My sense is that a machine which makes cubes from big pieces of marble (not just D) can be considered nonanticipating, if only because its output is so simple. I don't, however, want to get into a controversy over complexity, simplicity, and information, and how these notions might elucidate the idea of nonanticipating machines. As I've already noted, I'm primarily interested in instances of design where the design and designatum are identical/isomorphic.

The question remains, however, whether (D,s) should be considered a legitimate design-designatum pair. Let us for the sake of argument suppose that a nonanticipating machine can suitably vandalize D to produce s. According to our characterization of design up to this point, we would then have to admit (D,s) among the legitimate design-designatum instances. But this seems absurd. D is no more a design for s than the mountain of marble from which D was itself taken. There is a simple way to avoid such silly instances of design, and this is by appealing to duality. Even if it's possible to get from D to s via a nonanticipating machine, the reverse is not possible. As we've seen, no nonanticipating machine transforms a slab of marble into D (or a suitably miniaturized version of D).

The difference between design and designatum is a matter of context and conceptual point of view—this was the point of duality. What serves as design in one instance can serve as designatum in another (cf. the jet fighter example). Hence for any pair (X,Y) to be a design-designatum pair, it must be the case that (Y,X) is also a design-designatum pair (cf. the cryptography example). This is not to say that we should ignore the order of X and Y and focus instead on the binary set {X,Y} where order doesn't matter—the conceptual priority of X over Y in the design-designatum pair (X,Y) is important. Loosely speaking, duality requires there be no information loss in going from design to designatum. Note that this is never a problem when X and Y are identical/isomorphic.

With these observations in place, we can now offer the following, final characterization of design: a pair (D,Δ) is a design-designatum pair if there exist two nonanticipating machines, one capable of producing Δ from D, the other capable of producing D form Δ. In this way, whenever (D,Δ) is a design-designatum instance, so is (Δ,D). The difference between (D,Δ) and (Δ,D) is therefore a matter of which nonanticipating machine we care to focus on. In the case of (D,Δ) we focus on the nonanticipating machine which uses D to construct Δ (mutatis mutandis in the case of (Δ,D)). As we have repeatedly noted, design can always accommodate identity/isomorphism.

With a precise definition of design in hand, let us now shift gears and consider the problem of interpreting small probabilities. Once we understand what is at stake with small probabilities, the connection between small probabilities and design will be evident. The next section therefore introduces the problem of small probabilities. Our aim is twofold. First, to develop mathematical prerequisites and illustrative examples necessary to appreciate that there really is a problem interpreting small probabilities; second, to propose an explanatory framework for interpreting small probabilities. Once these aims are accomplished, we shall be able to connect small probabilities with design and thereby reformulate the design argument.

4. The Problem of Small Probabilities

It used to be said that if you owe a million dollars, you are in trouble; if you owe a billion dollars, your bank is in trouble; and if you owe 100 billion dollars, the world is in trouble. Although this claim needs to be adjusted for inflation, it makes the point that debt can be divided into roughly three categories: that which primarily affects the individual, that which is limited to a given community, and that which threatens the entire world. Now I claim that probabilities come in three types corresponding to this economic division. These I designate as pedestrian (type π), rare (type ρ), and surreal (type σ) probabilities.

Events with pedestrian probabilities are those an individual is apt to experience personally. For example, the probability of rolling double sixes with a pair of dice is a pedestrian probability since an individual given a few minutes is virtually assured of rolling double sixes. Events with rare probabilities are those no one individual is apt to experience, but which a community as a whole is apt to experience. For example, no single individual is likely to win a state lottery, but it is assured that the state will have a winner. Finally, events with surreal probabilities are those which the entire physical universe cannot render plausible. For example, the actual world is extremely unlikely to witness 1000 heads in a row from tossing a coin.

I shall refer to surreal probabilities as *probabilities of type* σ or *σ-probabilities*. Similarly, for pedestrian and rare probabilities I use the Greek letters π and ρ. σ-probabilities are the problem. The average individual has a pretty good sense for π-probabilities, having to deal with π-events (events having π-probabilities) all the time. Although ρ-probabilities aren't as well understood by the man on the street, actuaries and statisticians earn their living understanding these probabilities. σ-probabilities, however, constitute a probabilistic wasteland.

An analogy from statistics is useful for understanding π, ρ, and σ. In statistics one defines concrete error bounds of α and β corresponding to type 1 and type 2 errors respectively. To commit a type 1 error is to reject the null hypothesis when it in fact holds. To commit a type 2 error is to fail to reject the null hypothesis when it is false. α and β prescribe real numbers between 0 and 1 which measure the probability of committing a type 1 or a type 2 error. For instance, a commonly used α level in the social sciences is .05, implying that if the observed outcome is so extreme as to occur only 5 percent of the time when the null hypothesis is in fact true, then we can feel free to reject the null hypothesis. Suffice it to say that intuitions about α and β levels are well established.[30]

As with α and β, let us think of π, ρ, and σ as prescribing real numbers between 0 and 1 which represent pedestrian, rare, and surreal probabilities. Pedestrian probabilities are those we expect to see regularly. Thus if π is a pedestrian probability, then so is π' for $\pi' \geq \pi$. If on the other hand ρ is a rare probability, then any probability still smaller than ρ will be rare (or surreal if it is too much smaller). Similarly, a probability smaller than a surreal probability is surreal. There is a problem of demarcation between π and ρ, and between ρ and σ, but we won't worry about it. Just where the cutoffs occur is a problem of vagueness and is far less a concern than determining rational grounds for assigning probabilities to types π, ρ, and σ in ideal cases where vagueness near the cutoffs plays no role. Besides, our primary concern is with σ-probabilities.

Let us therefore massage our intuitions about surreality. We noted that our intuitions about pedestrian and rare probabilities aren't bad: the man on the street is fairly comfortable with π-probabilities, while the actuary has a pretty good handle on ρ-probabilities. In fact, we suggested that $\pi = 2.8 \times 10^{-2}$ is a nice pedestrian probability (this is the probability of rolling double sixes with a pair of dice), while $\rho = 10^{-6}$ is a nice rare probability (this is a typical order of magnitude for state lotteries). What about surreality? Can we find a value for σ which is in some way prototypical, that satisfies the predicate "σ is surreal" without being too, too small?

Before answering this question, we need a clearer picture of what we mean by surreality. What does it mean for a probability to satisfy "x is surreal"? Our naive intuition is that a σ-event is so unlikely that we can exclude it from all rational consideration and discussion. To take an extreme example, consider the possibility of a thermodynamic accident whereby a loaded gun (say a perfect replica of a Colt .45) materializes in your hand, gets aimed at your favorite enemy, fires, and kills him. The laws of physics don't bar this event from happening by chance. Nevertheless, a court will be compelled to convict you of willful homicide. Why does a court refuse to attribute such an event to chance (thereby exonerating you)? How would a jury respond to a defence attorney who tries to argue the gun just materialized?[31]

Our intuition then is that σ-events don't happen and can be safely ignored. To see that this can't quite be right, consider the following naive formulation of our intuition:

> Emile Borel . . . formulated a basic law of probability. It states that the occurrence of any event where the chances are beyond one in 10^{50} . . . is an event which we can state with certainty will *never* happen—no matter how much time is allotted, no matter how many conceivable opportunities could exist for the event to take place.[32]

This formulation seems to capture our intuition while at the same time fixing a concrete surreality level of $\sigma = 10^{-50}$. Nevertheless, this formulation of surreality is entirely inadequate, as we can see by considering the following simple counterexample. Suppose I flip a fair coin 200 times (by "fair" I mean the coin has distinguishable sides and is evenly balanced). As I flip the coin I note the outcomes. There are 2^{200}, or approximately 10^{60}, equiprobable sequences of possible coin flips. I will therefore participate in an event having probability 1 in 10^{60}. Because this probability is "beyond one in 10^{50}," the preceding passage appears to say that the event I experience by flipping the coin 200 times can't happen.[33] Yet given that I flip a coin 200 times, some exceedingly unlikely event must happen—a σ-event, if we accept Borel's surreality level.

We want to maintain that σ-events are so improbable that they can't happen. Yet we cannot deny that highly improbable events—σ-events—do happen (keep tossing a coin or rolling a die, and the sequence you observe will constitute a σ-event). In order to resolve the paradox we must introduce an extra-probabilistic factor I call *specification*. If a probabilistic set-up, like tossing a coin 200 times, demands that some σ-event occur, then necessarily some extremely improbable event will occur. If, however, we specify a σ-event, and thereafter observe it, we have cause for surprise and alarm. It is the match between specification and event that probability theory can't explain. Unspecified σ-events happen all the time. Specified σ-events don't. In fact, if σ is small enough, specified σ-events must not be attributed to chance.

Several examples should clarify the difference between specified and unspecified σ-events. Imagine that before you is a large, grassy field. You have 100 stones and 100 flags each marked from 1 to 100. With a helicopter you fly over the field, releasing the stones indiscriminately. After you have dropped your last stone, you land the helicopter safely away from the field, leave the helicopter on foot, and examine where your stones have landed, placing next to each stone a flag with the corresponding number. There are an exceedingly large number of ways the stones could have landed. They had to land in some particular way. You are looking at it. You are not surprised or shocked. You don't think a miracle has occurred because you are witnessing an event of exceedingly small probability. Some improbable event had to occur. Placing the flags next to the stones *after* the stones have fallen does not alter these conclusions.

Now modify the procedure. As before you have a field, stones, flags, and a helicopter. As before you take your helicopter and stones, and fly over the field, dropping the stones indiscriminately. But this time before you take off, you first walk around your field and stick the flags in the ground at will. Having dropped the stones, you land the helicopter and now examine the field. Lo and behold, all the stones are next to their matching flags. Do you have a right to be surprised? Absolutely. When an extremely unlikely event matches a preset pattern, there is cause for surprise. In fact when such an event becomes too unlikely, non-probabilistic factors must be taken into account. Placing the flags after the stones have fallen (or for that matter not placing any flags whatsoever) leaves the event of falling stones unspecified. But placing the flags before the stones have fallen specifies the event.

Another example that distinguishes specified from unspecified events concerns an archer who stands 70 meters from a large wall with bow and arrow in hand. The wall is sufficiently large that the archer cannot help hitting it. Suppose every time he shoots an arrow at the wall, he paints a target around the arrow, so that the arrow is squarely in the bull's-eye. What can be concluded? Absolutely nothing about the archer's ability as an archer. But suppose now he paints a fixed target on the wall and then shoots at it. Behold, 100 times in a row he hits a perfect bull's-eye. Nobody in his right mind would attribute this performance to beginner's luck. In fact, one is obliged to conclude this is a world-class archer. Fixing the target specifies the event of hitting a bull's-eye. Adjusting the target after every shot leaves hitting the bull's-eye unspecified.

In his book *The Blind Watchmaker* Richard Dawkins uses the following example to illustrate specified σ-events:

> Hitting upon the lucky number that opens the bank's safe is the equivalent, in our analogy, of hurling scrap metal around at random and happening to assemble a Boeing 747. Of all the millions of unique and, with hindsight equally improbable, positions of the combination lock, only one opens the lock. Similarly, of all the millions of unique and, with hindsight equally improbable, arrangements of a heap of junk, only one (or very few) will fly. The uniqueness of the arrangement that flies, or that opens the safe, [has] nothing to do with hindsight. It is *specified in advance*.[34]

Dawkins' purpose in *The Blind Watchmaker* is to explore the origin and development of complex living systems—systems whose appearance marks a σ-event. Now as he sees it complexity and specification are intimately related; in fact, complexity presupposes specification. Thus he defines a complex system as having "some quality, *specifiable in advance*, that is *highly unlikely* to have been acquired by random chance alone."[35] What's more, Dawkins demands that the occurrence of a specified σ-event be explained apart from chance: "What I do care about is that, whatever we choose to *call* the quality of being statistically-improbable-in-a-direction-specified-without-hindsight, it is an important quality that needs a special effort of explanation."[36] We shall return to the question of explaining specified σ-events.

To understand why specified σ-events "need a special effort of explanation," we must understand how specification contrasts with the strong law of large numbers (SLLN). Specification can be viewed as the antithesis of SLLN. SLLN asserts that an event with positive probability of occurring, no matter how small, will definitely occur if the probabilistic occasion for producing the event is repeated *often enough*.[37] Just how often is often enough depends on the probability of the event. For σ-events the expected number of repetitions before success is $1/\sigma$. Hence about 10^{50} repetitions will on average be needed to attain a desired σ-event if σ equals Borel's 10^{-50}. SLLN is specification of a very weak sort. It says that *eventually* we are guaranteed to observe a desired (i.e., specified) outcome. It says nothing, however, about what the *next* outcome shall be. Immediacy is inherent in specification. Specification focuses on what will happen next, not on the distant future after myriads of trials.

To clarify this point, imagine a chance mechanism **C** which emits only 0's and 1's. Suppose 1 has the surreal probability σ, so that 0 has the pedestrian probability $1-\sigma$. We assume the chance mechanism operates independently of prior and future outcomes. Thus we observe outcomes X_1, X_2, X_3, etc. which constitute an independent and identically distributed sequence of 0–1 random variables taking probability σ at 1. Now SLLN guarantees that if we could observe this sequence indefinitely, we will eventually come across an X_n that equals 1 (on average the first n for which this happens is around $1/\sigma$). In opposition to SLLN, specification spurns the sequence X_1, X_2, X_3, etc. Specification dismisses the long run entirely and concentrates exclusively on the single event. Given the chance mechanism **C**, we ask it to give us but one chance outcome—no more. Independently we specify one σ-event—in this case a 1. Our attitude is that a match between observed outcome and specified σ-event results from tampering with **C**, i.e., the chance mechanism **C** did not legitimately deliver 1 on the first trial.

In distinguishing SLLN and specification, we must address the question of probabilistic resources. To understand what I mean by probabilistic resources, let us recall the naive formulation of surreality quoted earlier:

> The occurrence of any event where the chances are beyond one in 10^{50} ... is an event which we can state with certainty will *never* happen—no matter how much time is allotted, no matter how many conceivable opportunities could exist for the event to take place.[38]

Even if we ignore the failure of this passage to address specification, we still have a problem with interpreting, "no matter how much time is allotted, no matter how many conceivable opportunities could exist for the event to take place." Clearly, if either time or space is boundless, then SLLN has free rein to accomplish anything. Recall Hume's point in his *Dialogues Concerning Natural Religion* quoted earlier:

> A finite number of particles is only susceptible of finite transpositions, and it must happen in an *eternal duration* that every possible order or position must be tried an infinite number of times.[39]

When we introduce a proviso like "no matter how much time or space is allotted," we clearly intend to restrict our quantifiers to the actual world, taking into account whatever limitations physical law imposes on our relentless attempts to replicate a desired result. The problem with such a proviso is that we conflate specification, which essentially hinges on the single event, with SLLN, which essentially depends on the long run. If we are not perfectly clear what constitutes an upper bound on the number of attempts we can make to attain a specified result, i.e., what our probabilistic resources comprise, then any fixed surreality index σ will be vague and misleading. An event with probability σ is likely to occur if it is given $1/\sigma$ opportunities to occur. The crucial question is how many opportunities obtain. If enough opportunities can be packed into the universe, then SLLN takes over and gives us the desired σ-event—even if σ is miniscule.

We therefore adopt the convention of referring σ levels solely to one-time probabilities. Given multiple opportunities for observing an event, we follow a straightforward procedure for collapsing multiple opportunities into a single opportunity: suppose an event E has positive probability p, i.e., P(E) = p. Suppose, furthermore, that the probabilistic resources for observing E are limited to at most n opportunities. We assume these opportunities are stochastically independent, and that the underlying chance mechanism is stable in the sense that probabilities don't change over time. Let E(n) denote the event that E happens at least once in n trials. From elementary probability theory it follows that

$$\begin{aligned} P(E(n)) &= P(\text{E happens at least once in n trials}) \\ &= \textstyle\sum_{1\le k\le n} P(\text{E happens for the first time at trial k}) \\ &= \textstyle\sum_{1\le k\le n} (1-p)^{k-1}p \\ &= 1-(1-p)^n . \end{aligned}$$ [40]

Given E and given n opportunities to observe E, the important question is whether E(n)—not whether E—is a σ-event. In terms of probabilities, the question is whether $P(E(n)) = 1-(1-p)^n < \sigma$, not whether $P(E) = p < \sigma$. Multiple sampling to attain E has to be collapsed into one-time sampling to attain E(n). Only if E(n) is a specified σ-event, will we want to say that the original E didn't happen by chance.

Let us summarize our findings. We began with a fundamental, though fuzzy, intuition about events with incredibly small probabilities (σ-probabilities) not happening. On closer examination, however, we found that events with incredibly small probabilities happen all the time. Was our intuition wrong? No. In naively formulating our intuition, we failed to mention that the σ-events we want to exclude from happening must also be specified. Is this enough? No. In addition to specification we found it necessary to control for the number of trials available to attain a σ-event. Only by taking into account the available probabilistic resources are we able to short-circuit the strong law of large numbers. The simplest course is to recast all events so that they have exactly one opportunity to occur. Our discussion of specification and probabilistic resources leads to the following formulation of our fundamental intuition: a specified σ-event with exactly one opportunity to occur never occurs by chance. Even this isn't good enough. The questions that remain include, (1) What exactly is a specification? (2) What is meant by chance? and (3) Just how small does σ have to be before we can confidently reject the chance occurrence of a specified σ-event? The next three sections will be devoted to answering these questions and thereby further clarifying our fundamental intuition.

5. Specification as Design

Specification is an instance of design. Just what sort of instance is the subject of this section. Our starting point is the chance mechanism **C**. For simplicity we assume that **C** has only finitely many possible outcomes: $\{\varphi_1,\varphi_2,\ldots,\varphi_n\}$. The outcomes, or equivalently elementary events, are mutually exclusive and have associated with them positive probabilities p_i: the probability of φ_i is p_i and the sum $p_1 + p_2 + \cdots + p_n$ equals 1. Our view is that the probabilities p_i are genuine properties of the elementary events φ_i. Thus the probabilities p_i are not limits of relative frequencies computed over a long number of trials, nor are they subjective probabilities.[41] Rather they spring from the dynamics and constitution of the chance mechanism **C**.

Whether probability assignments based on long run relative frequencies confirm or disconfirm our a priori assignment of probabilities p_i to φ_i will prove irrelevant to our discussion. The remarkable agreement between observed relative frequencies and a priori probability assignments across a wide assortment of chance mechanisms is of course empirical confirmation for our a priori assignments. But this agreement is logically speaking unnecessary, and in the cases we consider the a priori probabilities are always more secure than any probabilities derived from observed relative frequencies. In the words of Emile Borel, our chance mechanism **C** involves

> phenomena, for which the probability can be calculated from the very nature of the phenomena themselves; this is true of the throwing of a die, the observation of an honestly constructed roulette wheel and certain physical and biological phenomena. . . .[42]

The physical phenomenon I have in mind to serve as a prototype for **C** is quantum mechanics: the possible outcomes $\{\varphi_1,\varphi_2,\ldots,\varphi_n\}$ are supposed to suggest eigenstates for a quantum mechanical observable such that the probability p_i is the square of the absolute value of the φ_i-coefficient in the eigenstate decomposition of an arbitrary state vector ψ.[43]

In offering this conception of the chance mechanism **C**, I dispense with frequentist and subjective approaches to the probabilities p_i.[44] In games of chance like coin tossing, dice, cards, and roulette, the probabilities follow from symmetry considerations. With urn models where an urn's content comprises different colored balls of identical size and weight, the probabilities for drawing various colored balls follow from considerations of symmetry (the balls are isomorphic) and mixing (the balls are stirred after each replacement). In quantum mechanics, we have a theory which is inherently probabilistic, not only in its mathematical formalism, but also in the randomness that gets attributed to the particles of nature, randomness which the theory claims cannot be eliminated by refining our measurement apparatus.[45] In each of these instances the probabilities are intrinsic to the chance set-up. Thus in each case we can speak unequivocally about the probability of the single event.

The constitution of the chance mechanism **C** is what it is by virtue of its construction/ experimental preparation. Its dynamics, however, arises in but one way: *random sampling*, or more briefly, *sampling*. To sample from **C** is to allow chance to operate and produce an outcome φ_i. Inherent in sampling is the idea of independence: sampling from **C** on one occasion is unaffected by past sampling from **C** and does not affect future sampling from **C**. This type of stochastic independence is represented mathematically by taking products: the joint probability of two stochastically independent events is computed by multiplying the probabilities of the individual events. This picture of sampling holds for games of chance, urn models, and quantum mechanics. We shall return to these ideas in Section 6.

Although sampling always produces an outcome φ_i, we extend the idea of sampling to permit events generally as possible outputs. An event is a subset S of the collection of possible outcomes $\{\varphi_1,\varphi_2,\ldots,\varphi_n\}$. We say that sampling from **C** produced S if the actual output of sampling—some outcome φ_i—is a member of S (i.e., $\varphi_i \in S$). With this characterization of event, it is possible to assign probabilities to events: $\mathbf{P}(S)$ equals the sum of all the p_i's such that φ_i is in S. Note that if S is empty (i.e., $S = \varnothing$), then $\mathbf{P}(S) = \mathbf{P}(\varnothing) = 0$; if S equals the entire set of possible outcomes $\{\varphi_1,\varphi_2,\ldots,\varphi_n\}$ (which we denote by Ω), then $\mathbf{P}(S) = \mathbf{P}(\Omega) = 1$; and if S equals a singleton $\{\varphi_i\}$, then $\mathbf{P}(S) = \mathbf{P}(\{\varphi_i\}) = p_i$. Therefore the events S sampled from **C** are subsets of $\Omega = \{\varphi_1,\varphi_2,\ldots,\varphi_n\}$ and have probability $\mathbf{P}(S)$.

The distinction between events and descriptions must now be introduced. Events are what happen in the world. If, however, we want to discourse and reason about events, we need descriptions for these events. A philosopher of language will no doubt find this last remark trite. Nevertheless, my aim is straightforward and confined to experimental observation and measurement. In measurement nature presents itself to a measuring device whereupon a measurement is made and a datum recorded. Data constitute the descriptions of events. That's all I'm saying. Given a collection of events we want a language to describe those events.

The chance mechanism **C** can be represented mathematically by the standard probabilistic triple $(\Omega,\Sigma,\mathbf{P})$ where $\Omega = \{\varphi_1,\varphi_2,\ldots,\varphi_n\}$, Σ = the subsets of Ω, and $\mathbf{P}$ is the probability measure on Σ which assigns to each $\{\varphi_i\}$ the probability p_i. A descriptive language for **C** is now a pair $\mathbf{D} = (\Lambda,\tau)$ where Λ is a collection of expressions and τ is a function from Λ into Σ. Because Ω is finite we make the simplifying assumption that τ is a one-to-one correspondence. Thus for each event S in Σ, there is precisely one expression S^* in Λ such that $\tau(S^*) = S$. τ, if you will, translates or interprets the expressions of Λ into the events of Σ. We refer to the expressions of Λ as descriptions (descriptions of the events in Σ) and to the function τ as a translator.[46]

A *specification* for the chance mechanism **C** given the language **D** can now be defined as a description-event pair (R*,S) such that $R^* \in \Lambda$, $S \in \Sigma$, and $\tau(R^*) \supset S$. The description R* therefore describes an event which includes the event S. In particular the pair (S*,S) where $\tau(S^*) = S$ is always a specification. A concrete example might help. Let **C** model rolling a pair of dice, i.e., craps. Let **D** be a language which describes the possible events of **C**. Then the following pairs constitute specifications: ("the sum of the faces is greater than 10", double sixes were rolled), ("the product of the faces is 1", snake-eyes were rolled), and ("the faces sum to 12", double sixes were rolled). In the first instance the event expressed by the description properly includes the actual event. In the other instances the events described coincide with the actual events. Note that the description-event distinction parallels the use-mention distinction in the philosophy of language, with quotes distinguishing description and mention from event and use respectively.[47]

How then is specification an instance of design? More particularly, given a chance mechanism $\mathbf{C} = (\Omega,\Sigma,\mathbf{P})$ and a descriptive language $\mathbf{D} = (\Lambda,\tau)$, how is the specification (R*,S) an instance of design? For (R*,S) to be a specification we said the event described by R* must include the event S, i.e., $\tau(R^*) \supset S$. Now we also said that for the event S to occur means that one of the outcomes φ_i contained in S occurred. This can be generalized: the event S occurs only if the event R occurs for any R that includes S (i.e., $R \supset S$). In other words, to say an event occurs is to say all events containing that event occur (one consequence of this remark is that Ω—the necessary event—always occurs). Thus for any specification (R*,S), it is not only true that $(R^*,\tau(R^*))$ is always a specification, but also true that any explanation of (R*,S) is equally an explanation of $(R^*,\tau(R^*))$ as well. Associated with any specification (R*,S) therefore is a privileged specification (R*,R) for which $\tau(R^*) = R$. It is this privileged specification which properly speaking is an instance of design. By extension we include specifications generally as instances of design.

When $\tau(R^*) = R$, the specification (R*,R) is a design-designatum pair. Specifications like this are essentially instances of identity (recall that identity is always an instance of design). In fact, if we fail to distinguish events from their descriptions (as is common practice), then the identity is exact. Thus we have a quick and dirty way of making (R*,R) into a design instance—simply conflate events and descriptions. The distinction between events and descriptions is, however, worth preserving. Thus to argue that (R*,R) is a design instance, we need to find nonanticipating procedures for converting R* into R and vice versa. The obvious place to look is the translator τ. Recall that τ is a one-to-one correspondence of the language Λ with the event space Σ. Thus τ and its inverse τ^{-1} are well-defined procedures taking us from R* to R and back again ($\tau(R^*) = R$ and $\tau^{-1}(R) = R^*$).

Are the procedures τ and τ^{-1} which relate description and event nonanticipating? I submit that unless they are, the very relation between description and event becomes unintelligible. Scientific measurement never presupposes the event that shall be measured—insofar as there is any objectivity to science, this is it. On the other hand, descriptions are always formulated so that it is possible to decide whether an event which occurs satisfies the description. For example, in coin tossing there is a straightforward way to record a sequence of tosses—represent tails by 0, heads by 1, and record the appropriate 0-1 string. On the other hand, given a 0-1 string of length n, there is a straightforward way to decide whether a given outcome of n coin tosses matches the description. In neither direction is the description or event anticipated. It follows therefore that specifications (R*,R) for which $\tau(R^*) = R$ are design instances. By extension we say that arbitrary specifications (R*,S) are design instances as well.

Before leaving this section, let us be clear about terminology. In Section 3 we examined the design-designatum relation. A typical design-designatum pair took the form (D,Δ). We called D the design and Δ the designatum. The pair itself we called an instance of design. Moreover, the term "design" was used generically to signal the presence of a design-designatum instance. In this section we examined the description-event relation, specifically those which constitute specifications. A typical description-event pair has the form (R*,S). We call R* the description and S the event. Given the chance mechanism $\mathbf{C} = (\Omega,\Sigma,\mathbf{P})$, the descriptive language $\mathbf{D} = (\Lambda,\tau)$, and the inclusion $\tau(R^*) \supset S$, we say (R*,S) is a specification. Until now we have used specification solely in referring to pairs (R*,S). We now extend this usage. In case (R*,S) is a specification, we call (R*,S) a specification-event pair, referring to R* as the specification and S as the event. Specifications are therefore pairs of the form (R*,S) satisfying $\tau(R^*) \supset S$ as well as the first components of such pairs (i.e., the descriptions). The *likelihood* of a specification (R*,S) can now be unambiguously defined as the probability of the specified event $\tau(R^*)$, i.e., $\mathbf{P}(\tau(R^*))$.

6. Chance and Coincidence

Chance is used in so many different and incompatible ways that I would prefer to avoid the issue entirely. Yet because I've introduced the chance mechanism $\mathbf{C} = (\Omega,\Sigma,\mathbf{P})$ and the notion of random sampling from **C**, it will be necessary to say a few words about chance. The problem with chance is the philosophical baggage that invariably seems to come along for the ride. Someone like Laplace will equate chance with ignorance, arguing that for a super-intelligence (the Laplacian demon) the world is deterministic in strict Newtonian terms. A scholastic philosopher like Thomas Aquinas, though not a determinist in the mechanical sense like Laplace, will hold a similar view about chance since God is the first cause, causality is universal, no aspect of reality is uncaused, and God's omniscience includes perfect knowledge of the entire causal structure of the universe. On this view, the universe holds no surprises for God. Hence chance is significant only within our finite perspective and only insofar as we use it to steer an optimal course through life's uncertainties. Whereas Laplace reduces chance to ignorance via a mechanical determinism, Thomas does so via the principle of sufficient reason (a principle which among other things requires that every event have a cause).[48]

At the other extreme are views of chance arising from a radical indeterminism. The atomism of Epicurus is a case in point. For Epicurus the world was a buzz of particles swerving off in various directions where the direction of swerve could not be assessed even in principle. Irreducible or pure chance events have been defended more recently by the 19th century philosophers Charles Peirce and William James. With the advent of quantum mechanics a whole morass of questions about chance has been raised: What is the nature of causality at the micro-level? How does causality at the micro-level relate to the macro-level (cf. Schrödinger's cat)? Although the mathematical theory of quantum mechanics is thoroughly probabilistic, does the mathematics accurately mirror randomness as it exists in nature (Einstein thought no)?

Causality seems to lie at the root of most debates about chance. Perhaps if we could explicate causality, chance would take care of itself. But causality is itself philosophically loaded, with many problems of its own. Hume's criticism of causality whereby he reduces it to a habit of mind relating events which are contiguous, temporally ordered, and in constant conjunction is still potent. 20th century philosophers actively avoid metaphysical investments. Hence to equate chance events with uncaused events may be saying more than some philosophers care to admit. On the other hand, philosophers who view philosophy as continuous with science are apt to conclude that nature is, at least in part, acausal and therefore subject to irreducible chance influences. Thus when such a philosopher investigates radioactivity and finds what he takes to be identical radioactive particles decaying at different rates, he concludes that nature operates indeterministically since identical antecedent conditions (the identical particles) lead to different results (varying rates of decay). At this point the philosopher can be challenged to analyze his understanding of identity, particularly whether the "identical radioactive particles" of physics are identical in fact. This proliferation of questions is the stuff of philosophy, but little use in establishing a broad consensus.

I therefore propose we take a minimalist view of chance, a view we might equally well call the gambler's view of chance. What's important to a professional gambler when he enters a casino, sits himself down at a gaming table, and starts to play? I claim that philosophical and theological considerations don't enter his mind. The first thing that interests him is the chance mechanism $\mathbf{C} = (\Omega,\Sigma,\mathbf{P})$: What are the possible outcomes $\varphi_1,\varphi_2,\ldots,\varphi_n$? What are the probabilities p_i associated with each outcome φ_i? What are the cash payoffs associated with the events $E \subset \Omega = \{\varphi_1,\varphi_2,\ldots,\varphi_n\}$? Once these basic questions are answered, his next step is to formulate a strategy in line with rational decision theory as well as the risks he willing or eager to incur. The first consideration is largely a matter of arithmetic and accounting; the second is more psychological and subjective. There is a third point, however, which is usually presupposed and therefore left unstated, but which is vitally important to the gambler. It lies at the heart of our minimalist conception of chance; namely, the gambler wants the chance mechanism employed properly—he wants to preclude cheating.

The gambler is concerned that magnets aren't influencing the roulette ball, that cards are thoroughly shuffled and unmarked, and that dice aren't loaded. Although cheating is always a possibility, the gambler is not paranoid. He wants to play the game. He likes the thrill of the gamble. But he isn't neurotic in his validation of the chance mechanism. If he is betting on the toss of a coin, he does not think about an underlying determinism which Newtonian mechanics could exploit by calculating the coin's trajectory. It is sufficient that the coin be evenly balanced and that the person flipping the coin give it a good jolt. In the case of dice it is enough that they be reasonably precise cubes constructed out of a homogeneous material, and that the dice be vigorously tossed and not merely placed on the craps table.

Instead of saying that the chance mechanism was properly or fairly employed, we usually say that it was randomly sampled. We therefore distinguish between random and crooked sampling. Random sampling derives its outcomes from the chance mechanism by properly using the mechanism, crooked sampling by improperly using the mechanism. It's important to realize that this characterization of chance isn't circular. A chance mechanism $\mathbf{C}$ is not simply a static object with fixed outcomes and fixed associated probabilities. It also has a dynamics which is properly accessed through random sampling and short-circuited through crooked sampling. What's more, whether a random or a crooked sample was taken is a definite question of fact which can be judged empirically on the basis of evidence.

What constitutes random sampling for a given chance mechanism is specific to that mechanism. For a fair coin random sampling means having someone toss the coin by giving it a good jolt. For geometrically regular dice random sampling means having the gambler throw the dice vigorously. For an honestly constructed roulette wheel random sampling means having the croupier simultaneously throw the ball onto the wheel and spin the wheel. For a quantum mechanical spin 1/2 particle, random sampling means measuring spin in a given direction. Granted, cheating is always a possibility and the gambler must always keep his wits about him. But inherent in the chance mechanism is a sampling procedure which delineates random sampling and excludes crooked sampling. It is this sampling procedure which in any instance of sampling can be empirically validated, and which makes random sampling a well-defined operation.

No professional gambler is content to place a single bet on a single random sample from a single chance mechanism. In fact, any gambling strategy requires placing multiple bets over an extended series of plays. The gambler therefore needs a way of unifying diverse random samples into a composite random sample. The key to this unification is *stochastic independence*. Inherent in random sampling is that a random sample from one chance mechanism $\mathbf{C} = (\Omega,\Sigma,\mathbf{P})$ and a random sample from another chance mechanism $\mathbf{C}' = (\Omega',\Sigma',\mathbf{P}')$ is equivalent to a random sample from the composite chance mechanism $\mathbf{C} \otimes \mathbf{C}' = (\Omega \otimes \Omega', \Sigma \otimes \Sigma', \mathbf{P} \otimes \mathbf{P}')$ where $\Omega \otimes \Omega'$ is the cartesian product of Ω and Ω', $\Sigma \otimes \Sigma'$ is the product σ-algebra of Σ and Σ', and $\mathbf{P} \otimes \mathbf{P}'$ is the product measure of $\mathbf{P}$ and $\mathbf{P}'$ on $\Sigma \otimes \Sigma'$. The sampling procedure for this composite, or better yet, product chance mechanism is the mereological sum of the sampling procedures for the component chance mechanisms. In plain language, to sample randomly from the composite, sample randomly from each component. In this context stochastic independence simply means that distinct random samples combine to form a composite random sample according to the probability distribution given by the product chance mechanism $\mathbf{C} \otimes \mathbf{C}' = (\Omega \otimes \Omega', \Sigma \otimes \Sigma', \mathbf{P} \otimes \mathbf{P}')$. This method of combining random samples extends to multiple sampling from the same as well as from diverse chance mechanisms. Notice that what we are calling stochastic independence is inherent in random sampling and need not be referred, as is commonly the case, to any sort of causal independence among the samples. Distinct random samples are stochastically independent, as it were, by definition—by virtue of what it is to sample randomly.

Our minimalist picture of chance is therefore this. Associated with a chance mechanism $\mathbf{C}$ is a sampling procedure which delineates precisely what it is for $\mathbf{C}$ to be randomly sampled. (Often this procedure involves some sort of vigorous mixing.) It is this sampling procedure which makes $\mathbf{C}$ into a *chance* mechanism, and not merely a mechanism simpliciter. Inherent in random sampling is that distinct random samples, whether from the same or from different chance mechanisms, are stochastically independent. It is an empirical question whether in sampling from a chance mechanism the associated sampling procedure was followed. If it was, we say that a random sample was taken and that chance was operative. Otherwise, the sample was crooked and cheating likely. Notice that this picture refuses to commit itself to any metaphysical doctrine of chance—all the deep philosophical questions remain unanswered. This is exactly what we want: a common core to the concept of chance which is capable of supporting divergent philosophical positions about chance.

When last we saw our fundamental intuition about events with incredibly small (i.e., surreal or σ-) probabilities, we formulated it as follows: a specified σ-event with exactly one opportunity to occur never occurs by chance. Our convention about probabilistic resources constrains us to consider only those events with precisely one opportunity to occur, collapsing events with multiple opportunities into events with a single opportunity. Our fundamental intuition could therefore be stated more simply as, No specified σ-event ever occurs by chance. With the preceding minimalist account of chance as well as the account of specification given in the last section, we are now able further to refine this intuition. The vehicle for this further refinement is coincidence, the next topic for consideration.

The most important class of specifications for our purposes are the coincidences. Before giving a formal definition of coincidence, however, I want to consider a series of coin tossing examples which should clarify what is at stake with coincidence as well as what distinguishes it from specification generally. Suppose $\mathbf{C}$ is a chance mechanism for tossing a fair coin while $\mathbf{D}$ is a language for describing sequences of coin tosses. Let us represent such sequences as strings of 0's and 1's (0 for tails, 1 for heads). Consider the following specification:

("01000010111010001010", 01000010111010001010) (6.1)

The first component is the description of an event comprising 20 coin tosses. The second component is that very event.

How could (6.1) fail to be a coincidence? By itself nothing precludes (6.1) from constituting a coincidence. Associated with any specification, however, is a story about how that particular specification came about. With (6.1) I tossed a coin 20 times and observed the sequence 01000010111010001010. Having observed the sequence, I then recorded it. Thus the description "01000010111010001010" is my observation of the chance event 01000010111010001010. The point to recognize is the causal link that exists between the first and second component of (6.1): only after duly noting the event did I record the description.[49] For coincidences, descriptions must not be a mere record of the observed chance events.

The preceding account of how (6.1) arose is in fact what happened. Nevertheless, another story could be told about (6.1)'s genesis which would still keep it from constituting a coincidence. Suppose I happened to write down the sequence "01000010111010001010"—for whatever reason. After writing it down I took my fair coin and instead of tossing it, placed it 20 times on my desk so that the first face up was tails, the second heads, the next four tails, etc. The causal link between description and event in (6.1) is obvious, only this time the direction of causality is reversed—from description to event and not vice versa as in the last example. Essentially I used the description to force the event. Clearly the chance mechanism $\mathbf{C}$ was being used improperly. Although the coin which is so central to $\mathbf{C}$ was employed, $\mathbf{C}$ itself was not randomly sampled. In fact, by deliberately manipulating the coin, I short-circuited the chance mechanism. Hence according to this second story (6.1) again fails to be a coincidence.

How then might (6.1) have arisen as a coincidence? Prediction is perhaps the simplest way: first record the sequence "01000010111010001010" and afterwards observe 01000010111010001010 as the random flips of a fair coin. In this way the description is formulated independent of the event and the event results from random sampling. Prediction, however, implies a temporal priority which is not strictly speaking necessary to keep description and event independent. Suppose I sample from **C**, observe event 01000010111010001010, and record it as a datum. Having done this, I refuse to divulge any information about the occurrence of this event (assume I place all evidence of the event in a safe). Suppose next that a week later I bump into somebody who proceeds to write down the string "01000010111010001010." In this case independence between description and event is again preserved even though the description postdates the actual event. On this last account (6.1) is again a coincidence.

With these examples in mind we can now formulate a precise definition of coincidence: a *coincidence* (R*,S) for the chance-mechanism/descriptive-language pair (**C**,**D**) is a specification for (**C**,**D**) whose genesis satisfies the following conditions:

(Ch) The event S resulted from randomly sampling **C**.

(ZInf) The description R* was formulated without utilizing any information about the sampled event S.

Condition (Ch) is the *chance condition*. It says that a chance mechanism was randomly sampled to produce S (more briefly, that S is a chance or random event). Condition (ZInf) is the *zero-information condition*. It says that no information about S was exploited in the formulation of R*. It is intuitively obvious how these conditions were violated in the preceding examples: (Ch) prevents us from forcing a coin to come up in certain ways by reading from an already existing description while (ZInf) prevents us from reading a description off a duly noted event (in this case a sequence of 20 coin tosses).

Condition (Ch) was explicated earlier in this section when we analyzed random sampling and chance. Condition (ZInf), however, still requires explanation. Given a specification (R*,S), what is it for R* to arise without exploiting any information about S? I propose to understand information as physical evidence about S which might influence the formulation of R*. Notice that I restrict information to physical evidence—angels whispering hints about S in your ear don't count. The question remains what sort of information about S might influence the formulation of R*. This question is more conveniently rephrased negatively: what must be true for no information about S to influence the formulation of R*? Since information is physical evidence, any answer to this last question depends on our understanding of the physical world.

There are, however, certain invariants which hold universally regardless what physical theory of the world is currently in vogue. One such invariant is temporal. Thus if R* is prior to or simultaneous with S, then no information about S was available to influence R*. Another invariant is spatial. Thus if no evidence about S has by time t propagated outside a certain radius and if at time t R* arises outside this radius, then again no information about S was available to influence the formulation of R*. Still another invariant is aleatory. Thus if R* is the description of an event R which is itself the result of random sampling from a chance mechanism, then any event S distinct from R cannot be said to influence R*.

The important thing to recognize about conditions (Ch) and (ZInf) is that they involve definite questions of fact open to empirical verification. Given a specification (R*,S) from the chance-mechanism/descriptive-language pair (**C**,**D**), condition (Ch) asks whether S resulted by randomly sampling **C** according to the prescribed sampling procedure associated with this chance mechanism. Condition (ZInf), on the other hand, asks whether any information about S was available in the formulation of R*. Given certain physical assumptions about the way the world works (assumptions which, by the way, are always presupposed in scientific explanation), (ZInf) is as much a question of fact and equally open to empirical scrutiny as (Ch).

With this characterization of coincidence, the problem of small probabilities now takes the form of a question: Do specifications below a certain likelihood invariably fail to be coincidences? Think of it this way. There is a huge set of specifications $\mathit{Spec} = \{ (R_i^*,S_i) \mid i \in I \}$ where as i runs through the index set I, (R_i^*,S_i) runs through all possible specifications. Associated with each i in I is also a chance-mechanism/descriptive-language pair ($\mathbf{C}_i = (\Omega_i,\Sigma_i,\mathbf{P}_i)$, $\mathbf{D}_i = (\Lambda_i,\tau_i)$) with respect to which (R_i^*,S_i) is a specification. Similarly, there is a slightly less huge set of coincidences $\mathit{Coincid} = \{ (R_j^*,S_j) \mid j \in J \} \subset \mathit{Spec}$ where J is a subset of the original index set I and each coincidence (R_j^*,S_j) is associated with ($\mathbf{C}_j = (\Omega_j,\Sigma_j,\mathbf{P}_j)$, $\mathbf{D}_j = (\Lambda_j,\tau_j)$). The members of *Spec* are specifications simpliciter whereas the members of *Coincid* are specifications which additionally satisfy (Ch) and (ZInf). Given our discussion of surreality, we now concentrate on specifications and coincidences whose likelihoods fall below a fixed surreality level σ. We therefore define $\mathit{Spec}_\sigma = \{ (R_i^*,S_i) \mid i \in I \text{ such that } \mathbf{P}_i(\tau_i(R_i^*)) \le \sigma \}$ and $\mathit{Coincid}_\sigma = \{ (R_j^*,S_j) \mid j \in J \text{ such that } \mathbf{P}_j(\tau_j(R_j^*)) \le \sigma \}$, the σ-specifications and σ-coincidences respectively. The problem of small probabilities therefore reduces to finding a positive σ for which $\mathit{Coincid}_\sigma = \varnothing$. The search for this σ is our next task.

7. How small σ?

Our fundamental intuition about events of sufficiently small probability not happening can now be formulated as a law. We call it the Law of Small Probability (LSP) and state four equivalent versions of it:

(LSP1) Specifications of sufficiently small likelihood never occur by chance.
(LSP2) Specifications of sufficiently small likelihood are never coincidences.
(LSP3) There is a positive σ for which $Coincid_{\sigma} = \varnothing$.
(LSP4) There is a positive σ such that no σ-specification is also a coincidence.

For LSP to have any practical significance, it will be necessary to characterize the "sufficiently small likelihoods" referred to in (LSP1) and (LSP2) via concrete numbers. The obvious candidate for such concrete numbers is of course the surreal probabilities of Section 4. (LSP3) and (LSP4) therefore substitute the surreal probability σ for the "sufficiently small likelihood" of (LSP1) and (LSP2). Our task henceforth is to find a concrete σ which without exception obviates all σ-coincidences. Borel proposes one such σ for his Single Law of Chance[50]: $\sigma = 10^{-50}$. The problem with Borel's σ, however, is that it derives from a loose argument based on the size of the universe—an argument which fails adequately to address the question of probabilistic resources and fails utterly to address the problem of specification. Both issues must now be squarely faced.

To discover a rational basis for surreal probabilities and thereby give substance to LSP we need to consider two types of resources—probabilistic and specificational. Probabilistic resources were addressed in Section 4. The point there was to bypass problems that might arise on account of the strong law of large numbers when events have an indeterminate number of opportunities to occur. To do this we limited our attention to events with a single opportunity to occur. Given therefore an event E with up to n opportunities to occur, we adopted the convention of analyzing not E but E(n)—the event that E occurs in at least one of n independent trials. Just how many opportunities n are available for a given E will typically depend on physical constraints and the physical theory underlying those constraints. But because we always collapse events with multiple opportunities into events with a single opportunity, the question of exactly what physical theory governs our probabilistic resources is bypassed. In LSP, physical constraints are mediated through the collapse of probabilistic resources into one-time events.

While probabilistic resources are easily understood and widely appreciated, specificational resources are not. This is not to say that the concept is wholly lacking in the probabilistic literature. Its use, however, has at best been intermittent and dimly conceived. The clearest formulation I know of what's at stake with a specificational resource is due to Charles Peirce. Peirce begins by examining the claim that

> the idea of probability essentially belongs to a kind of inference which is repeated indefinitely. An individual inference must be either true or false, and can show no effect of probability; and, therefore, in reference to a single case considered in itself, probability can have no meaning.[51]

There is a problem with this frequentist conception of probability, however. As Peirce continues,

> Yet if a man had to choose between drawing a card from a pack containing twenty-five red cards and a black one, or from a pack containing twenty-five black cards and a red one, and if the drawing of a red card were destined to transport him to eternal felicity, and that of a black one to consign him to everlasting woe, it would be folly to deny that he ought to prefer the pack containing the larger portion of red cards, although, from the nature of the risk, it could not be repeated. It is not easy to reconcile this with our analysis of the [frequentist] conception of chance. But suppose he should choose the red pack, and should draw the wrong card, what consolation would he have? He might say that he had acted in accordance with reason, but that would only show that his reason was absolutely worthless [consigned as he now is to "everlasting woe"]. And if he should choose the right card, how could he regard it as anything but a happy accident? He could not say that if he had drawn from the other pack, he might have drawn the wrong one, because [such] an hypothetical proposition . . . means nothing with reference to a single case.[52]

Peirce's solution isn't obvious:

> The man whom we have supposed as having to draw from the two packs . . . cannot be logical so long as he is concerned only with his own fate.[53]

With whose fate then should this man be concerned? Again Peirce:

> It seems to me that we are driven to this, that logicality inexorably requires that our interests shall *not* be limited. They must not stop at our own fate, but must embrace the whole community. This community, again, must not be limited, but must extend to all races of beings with whom we can come into immediate or mediate intellectual relation. I must reach . . . beyond this geological epoch, beyond all bounds. He who would not sacrifice his own soul to save the whole world, is, as it seems to me, illogical in all his inferences, collectively. Logic is rooted in the social principle.[54]

Peirce concludes,

> It may seem strange that I should put forward three sentiments, namely, interest in an indefinite community, recognition of the possibility of this interest being made supreme, and hope in the unlimited continuance of intellectual activity, as indispensable requirements of logic. Yet, when we consider that logic depends on a mere struggle to escape doubt, which, as it terminates in action, must begin in emotion, and that, furthermore, the only cause of our planting ourselves on reason is that other methods of escaping doubt fail on account of the social impulse, why should we wonder to find social sentiment presupposed in reasoning? As for the other two sentiments which I find necessary, they are so only as supports and accessories of that. It interests me to notice that these three sentiments seem to be pretty much the same as

that famous trio of Charity, Faith, and Hope, which, in the estimation of St. Paul, are the finest and greatest of spiritual gifts. Neither Old nor New Testament is a textbook of the logic of science, but the latter is certainly the highest existing authority in regard to the dispositions of heart which a man ought to have.[55]

I've quote Peirce at length not because I agree with his conception of logic and probability, but because he clarifies the community's role in probabilistic reasoning. Let me recast Peirce's example of the red and black card decks into the language of specification. Recall that the red deck has 25 red cards and one black card; for the black deck the numbers are reversed. The specified event is the drawing of a red card. If the specified event occurs, bliss; otherwise, torment. Given a choice whether to sample from the red or the black deck, which do I choose? If I look only to myself, I find it difficult to justify picking the red deck: whichever choice I make, if red turns up, heaven; if black, hell. I may reproach myself for picking the black deck if black comes up, but what consolation have I in picking the red deck if this choice leads to black as well? If, however, I refer my choice to the community, then I have a rationale for picking the red deck. Granted, black may still turn up, but the community as a whole will fare better by consistently choosing the red deck. My specification of the red outcome by itself provides no rationale for picking the red deck. But when I consider all the specifications of this outcome advanced by the members of my community, all of whom presumably want to escape hell and enter heaven, then according to Peirce's social conception of logic, reason mandates picking the red deck.

To make his point Peirce need not have invoked heaven and hell. Consider a national election in a country with one-hundred million citizens. Citizen X claims that his vote will not affect the outcome of the election. X goes so far as to claim that the probability of his being killed in an accident on way to the election booth is incomparably greater than the probability that his refusal to vote will affect the election's outcome. X therefore on principle refuses to vote. Is X justified in his refusal? The conscientious citizen has but one rebuttal: if everyone assumed your attitude, there would be no election at all. X can be refuted only by appealing to the community. Kant's categorical imperative comes to mind: "Act only on that maxim whereby thou canst at the same time will that it should become a universal law." Such universality requires community.

If we adopt the categorical imperative in probabilistic reasoning as not only morally obligatory (à la Kant), but also constitutive of rationality (à la Peirce), then two prescriptions for probabilistic decision theory become axiomatic. We state them in the language of chance mechanisms and descriptive languages.

(Bet1) Given a chance mechanism $\mathbf{C} = (\Omega,\Sigma,\mathbf{P})$ and a descriptive language $\mathbf{D} = (\Lambda,\tau)$, suppose that $\mathbf{C}$ is to be randomly sampled exactly once, and that from a finite collection of descriptions $\{E_1{}^*,E_2{}^*,\ldots,E_n{}^*\} \subset \Lambda$ exactly one must be selected with the object of describing the sampled event. Then choose any $E_i{}^*$ for which $\mathbf{P}(\tau(E_i{}^*)) = \max_{1\leq j\leq n}\mathbf{P}(\tau(E_j{}^*))$.

(Bet2) Given a finite family of chance mechanisms $\mathbf{C}_j = (\Omega,\Sigma,\mathbf{P}_j)$ $(j = 1,\ldots,n)$ based on the same underlying outcome/event space (Ω,Σ) and a descriptive language $\mathbf{D} = (\Lambda,\tau)$, suppose that we fix a description $E^* \in \Lambda$ and that exactly one of the $\mathbf{C}_j$'s is to be randomly sampled a single time with the object of producing the event described by E^*. Then choose any chance mechanism $\mathbf{C}_i$ such that the associated probability $\mathbf{P}_i$ satisfies $\mathbf{P}_i(\tau(E^*)) = \max_{1\leq j\leq n}\mathbf{P}_j(\tau(E^*))$.

(Bet1) and (Bet2) are, if you like, betting axioms. In fact they are betting axioms of the most primitive sort, instructing the gambler to bet where probabilities are maximal. Not only do these axioms receive the gambler's seal of approval, but according to Peirce they are constitutive of rationality. Both the gambler's and Peirce's approbation are significant: because our minimalist conception of chance is modeled on gambling (see Section 6), the gambler's approval is indispensable; because our aim is philosophical, it's nice to know these axioms don't fly in the face of reason. Now the case I shall make argues that LSP, the centerpiece of this discussion, is a consequence of these axioms once we take specificational resources into account.

What then is a specificational resource? Think of specificational resources as follows. Given a chance mechanism $\mathbf{C} = (\Omega,\Sigma,\mathbf{P})$ and a descriptive language $\mathbf{D} = (\Lambda,\tau)$, suppose some event $E \in \Sigma$ is going to be randomly sampled from $\mathbf{C}$. We don't know what E is, but we assume it has small probability, say $\mathbf{P}(E) \leq p$ for some small positive p. (This is an information-theoretic assumption: the information of an event E is by definition $-\log_2\mathbf{P}(E)$ which is large only to the degree that $\mathbf{P}(E)$ is small. E is the event we want to comprehend. E is informative only insofar as it excludes possibilities, and this is the case only when E has small probability. Compare this, say, with $E = \Omega$ which tells us that indeed something happened, but nothing about what didn't happen.) Suppose next that each member in the community has one opportunity to guess what E is going to be. Thus if J indexes the community, each community member $j \in J$ formulates one description/guess $E_j{}^*$. A guess $E_j{}^*$ is successful or unsuccessful depending on whether $\tau(E_j{}^*)$ includes of fails to include E. Without further restriction, however, the community can always successfully guess E simply by selecting a description for the entire space of possibilities Ω ($\Omega \supset E$ is always true). The guesses therefore must attempt not only to describe E, but also to describe it efficiently. We therefore add the proviso that each $E_j{}^*$ describe an event of probability no greater than p (i.e., $\mathbf{P}(\tau(E_j{}^*)) \leq p$)—this is in keeping with E having probability no greater than p. When E is finally sampled, we ask which E_j*'s, if any, describe E (i.e., for which E_j*'s is it true that $\tau(E_j^*) \supset E$). The relevant specificational resource here is simply the number of specifications E_j* available for describing E, in this case the cardinality of the community J.

Probabilistic resources answer the question how many independent opportunities exist for observing a given event. Specificational resources answer the question how many specifications are available for describing a yet unobserved event. Although the two resources complement each other, they are distinct. To see this, consider an event E from a chance mechanism $\mathbf{C} = (\Omega,\Sigma,\mathbf{P})$. Suppose E has up to m independent opportunities to occur. Then as we saw in Section 4, E(m) (the event that E occurs at least once in these m independent trials) has probability $1-(1-\mathbf{P}(E))^m$. In accordance with the convention of Section 4 we collapse events having multiple opportunities into events having a single opportunity. We therefore disregard E and henceforward concentrate on E(m). Observe that the chance mechanism which gives rise to E(m) is not the original C, but the m-fold product of C, $\mathbf{C}^m = (\Omega^m,\Sigma^m,\mathbf{P}^m)$, with Ω^m representing the standard cartesian product, Σ^m the product σ-algebra, and $\mathbf{P}^m$ the product measure.

To keep the notation simple, however, let us rename the collapsed event E(m) as E and the corresponding product mechanism $\mathbf{C}^m$ as $\mathbf{C} = (\Omega,\Sigma,\mathbf{P})$ (given our convention about collapsing multiple into one-time events, we might just as well have assumed m = 1). E we suppose was randomly sampled from the chance mechanism C. Suppose moreover that C has descriptive language $\mathbf{D} = (\Lambda,\tau)$ and that we've fixed a positive σ with $\mathbf{P}(E) \le \sigma$. Assume that to specify E the relevant community can generate at most n descriptions from D, each description referring to an event with probability no greater than σ. Then n is a specificational resource, setting an upper bound on the number of ways the community can specify E within a tolerance σ. Now let $\{E_1^*,E_2^*,\ldots,E_n^*\} \subset \Lambda$ be any n descriptions satisfying $\max_{1\le j\le n}\mathbf{P}(\tau(E_j^*)) \le \sigma$ (i.e., each description specifies an event within the prescribed tolerance). Then the probability that one of these E_j*'s describes E (i.e., $\tau(E_j^*) \supset E$) cannot exceed the probability of $\tau(E_1^*) \cup \tau(E_2^*) \cup \cdots \cup \tau(E_n^*)$—i.e., the event that at least one of the events described in $\{E_1^*,E_2^*,\ldots,E_n^*\}$ happens. But the probability of this event is bounded by nσ:

$$\begin{aligned}\mathbf{P}(\tau(E_1^*) \cup \tau(E_2^*) \cup \cdots \cup \tau(E_n^*)) &\le \textstyle\sum_{1\le j\le n} \mathbf{P}(\tau(E_j^*)) \\ &\le n\left[\max_{1\le j\le n}\mathbf{P}(\tau(E_j^*))\right] \\ &\le n\sigma.\end{aligned}$$

Suppose now that σ is chosen so that $\sigma < 1/(2n)$ (alternatively, so that $n\sigma < 1/2$). Consider the two complementary events Φ and Φ^c. Φ is the event that at least one of the descriptions $\{E_1^*,E_2^*,\ldots,E_n^*\}$ produced by the community in question specified E; Φ^c is the complementary event that none of the proposed descriptions specified E. Since E is a random event, if the descriptions $\{E_1^*,E_2^*,\ldots,E_n^*\}$ are framed with zero-information about the event E (as they must in the case of coincidences), then Φ and Φ^c have definite probabilities. Moreover, by the preceding argument we know that the probability of Φ is bounded by nσ, which because we chose $\sigma < 1/(2n)$ is in turn less than 1/2. It follows that Φ^c has probability greater than 1/2. Given a choice between events Φ and Φ^c, axioms (Bet1) and (Bet2) tell us to go with event of greater probability, in this case Φ^c. But Φ^c is the event that for an arbitrary collection of n descriptions $\{E_1^*,E_2^*,\ldots,E_n^*\}$ such that $\max_{1\le j\le n}\mathbf{P}(\tau(E_j^*)) \le \sigma$, each E_j* failed to describe E. In particular, Φ^c denies that any of the pairs (E_1^*,E), (E_2^*,E), ..., (E_n^*,E) is a specification, much less a coincidence. To bet on Φ^c is therefore to bet against the community generating a σ-coincidence.

Let us reconsider this argument. Confronted with the claim that the σ-specification (E*,E) is a coincidence, our first question is how many specifications E_j* of size σ (i.e., $\mathbf{P}(\tau(E_j^*)) \le \sigma$) can conceivably be generated by the community with a stake in (E*,E). An upper bound on this number is the relevant specificational resource. Call this number n. The community can therefore propose at most n descriptions of size σ, say $\{E_1^*,E_2^*,\ldots,E_n^*\}$. Is (E*,E) a coincidence? If so, E is the random output of a chance mechanism. Given our specificational resources, (E*,E) has no more probability of being a coincidence than any of the pairs (E_1^*,E), (E_2^*,E), ..., (E_n^*,E) has of being a coincidence (note that for an arbitrary E_j*, $\tau(E_j^*)$ need not include E and therefore (E_j^*,E) need not in general be a specification/coincidence). This probabilistic equivalence is crucial to my argument. A specification (E*,E) is likely to be a coincidence only if one of the would-be specifications (E_j^*,E) generated by the community can with reasonable probability be expected to be a coincidence.

What then is the probability that one of these n would-be specifications (E_j^*,E) $(j = 1,\ldots,n)$ is in fact a specification? E is a random event; moreover, each of the E_j*'s is proposed without any information about E. The probability in question is therefore identical with the probability that the random event E is included in at least one of the events $\tau(E_1^*)$, $\tau(E_2^*)$, ..., $\tau(E_n^*)$. But this cannot exceed $\mathbf{P}(\tau(E_1^*) \cup \tau(E_2^*) \cup \cdots \cup \tau(E_n^*))$, which in turn is bounded by nσ. If therefore $\sigma < 1/(2n)$, the event Φ (that at least one of the E_j*'s describes E) has probability less than 1/2, and the complementary event Φ^c (that none of the E_j*'s describes E) has probability greater than 1/2. Given axioms (Bet1) and (Bet2) we are obliged to prefer Φ^c over Φ, and therefore to reject (E*,E) as a σ-coincidence.

If the preceding argument seems difficult, it might help to contrast it with the lottery paradox. In the lottery paradox a large number of tickets each have a small (but not too small) probability of being selected. Although any one ticket holder is unlikely to win and therefore can justly be pessimistic about his chances of winning the lottery, some ticket holder is sure to be a winner. For each ticket holder therefore to follow his pessimist inclinations and baldly deny he will win leads to a paradox since at least one ticket holder will be wrong, viz., the winner.[56] More formally, a community J of lottery players take part in a lottery. Associated with player j $\in$ J is a specification (E_j*,E_j) where E_j is the event that j wins the lottery and E_j* is a description of E_j. Each player j, however, is pessimistic and therefore denies he will win the lottery—in our framework, each j denies that (E_j*,E_j) is a coincidence.

Although the details of the lottery need not concern us, one fact stands out: real-world lotteries typically have winners, thereby ensuring that at least one of the specifications (E_j*,E_j) is a coincidence. Why do real-world lotteries work this way? If a real-world lottery has n players, the probability of any player winning is usually 1/n. This probability of winning, however, isn't etched in stone. Any probability σ could serve as a player's probability for winning the lottery. Lottery managers would like to keep σ as small as possible. The problem, however, with taking σ too small is that no one would be willing to play the lottery. If σ is too small, each player's pessimism about his chances of winning would not only be well-founded, but also extend to his fellow players' prospects of winning. In our framework, each player j will be right to deny not only that (E_j*,E_j) is a coincidence, but also that all such specifications are coincidences.

Now this is exactly what we did by choosing $\sigma < 1/(2n)$. We can think of the specificational resource n as the number of lottery players relevant to a putative coincidence. If σ, the likelihood of the putative coincidence, is less than 1/(2n), then the probability of this "lottery" having a winner is less than 1/2 (this is the event we called Φ), whereas the probability of this "lottery" failing to have a winner is greater than 1/2 (this is the event we called Φ^c). We can think of the Law of Small Probability (LSP) as describing a lottery in which the probability of any player winning is so small that no player is likely to win. In this way we cash out the hitherto vague notion of "sufficiently small." Indeed, σ is sufficiently small if it is small enough to preclude a lottery paradox.

One question remains: What exactly is the role of the community in LSP? LSP as formulated makes no mention of a community. Yet the claim that specifications of sufficiently small likelihood are never coincidences must necessarily be relativized to a community. We've argued that σ is sufficiently small if $\sigma < 1/(2n)$ where n is the specificational resourcebelonging to the relevant community. What then is the relevant community? It's certainly possible to specify σ on a case by case basis, in each case determining the relevant community and specificational resources. But this approach is unsatisfying since any given community can be conceived as a sub-community of a bigger community whose specificational resources exceed those of the sub-community. "Sufficiently small" for this bigger community will be smaller than "sufficiently small" for the sub-community.

Given a community with specificational resources n, I can compute a probability σ which embodies the notion of "sufficiently small" for that community (according to our previous argument $\sigma < 1/(2n)$ works). Given a specification (E*,E) with likelihood less than σ, LSP has me conclude that (E*,E) is not a coincidence. But suppose my adversary now comes along and informs me that mine was not the relevant community; that the community I considered is embedded in a much larger community which alone is the relevant community; that the relevant specificational resources are in fact much larger than n; and finally, that "sufficiently small" is captured by a σ' which is much smaller than σ. Having referred me to a larger community, he finds that (E*,E) has likelihood greater than σ' (though still less than σ). According to my adversary LSP is in this case inconclusive: (E*,E) could conceivably be a coincidence.

The way out of these ever expanding communities is to find a super-community which includes as sub-communities all the communities that might arise in practice. This is not as far-fetched as it seems. Indeed, Peirce refers explicitly to such a super-community:

> Logicality inexorably requires that our interests shall *not* be limited. They must not stop at our own fate, but must embrace the whole community. This community, again, must not be limited, but must extend to all races of beings with whom we can come into immediate or mediate intellectual relation. I must reach . . . beyond this geological epoch, beyond all bounds.[57]

This super-community is not infinite, nor does it possess infinite specificational resources. (With no limit on specificational resources LSP becomes vacuous—"sufficiently small" in this case would have to be identified with probability zero.) The super-community is of course the universe, which we define as the totality of things in immediate or mediate relation with planet earth. The super-community is therefore not an inflationary universe: an inflationary universe comprises isolated sub-universes which, to recall Peirce's phrase, have no "immediate or mediate intellectual relations" to ourselves. The super-community is the actual world, the world we inhabit.

Now the amazing thing about the world we inhabit is that it is very finite. In fact, specificational resources for the universe are readily computed and easily understood. To see just how limited our specificational resources are, imagine tossing a fair coin a thousand times. There are approximately 10^{300} possible outcomes. It is perfectly feasible to record any bit string a thousand bits long—a single sheet of paper will do. Indeed, any conceivable outcome from a thousand coin tosses can be recorded. In the actual world, however, it is impossible effectively to record any but a small number of these strings. There is a distinction worth making between theoretical and effective possibility. Theoretical possibility asks what can conceivably be put down on paper; effective possibility asks how much paper is available. Now specification is fundamentally a question of effective possibility, not theoretical possibility. To specify an event is to identify it, not necessarily on paper, but in some physical medium. Hence to specify an event is to consume a limited physical resource from the actual world.

The physical resources relevant to specification are data storage and transmission. How many bits of information can the universe store? How fast can this information be moved? How long can this information be stored before it decays? This is Shannon's information theory.[58] The number of elementary particles in the universe constrains the first question, the speed of light the second, thermodynamics the third. Quantum mechanical considerations indicate that information storage below the atomic level is infeasible. Hence with under 10^{80} elementary particles in the actual world, the number of bits available at any moment is less than 10^{80}. The Planck time ($> 10^{-45}$ s) certainly bounds the speed with which bits in the actual world can be switched on and off.[59] Let us assume no more than a thousand trillion years ($< 10^{25}$ s) from the big bang to the heat death of the universe. Then an upper bound on the total number of bits available for specification in the actual world is $10^{150} = 10^{80} \times (10^{25}/10^{-45})$ (i.e., the number of bits available at any moment multiplied by the number of times these bits can switch on and off). Since any specification will require at least one bit, 10^{150} supplies a conservative bound for the total number of specifications effectively possible in the actual world.

Because I don't expect physics to overturn this bound, I dub 10^{150} the *universal specification bound* (USB)and denote it by Π.[60] Π limits the total number of events (pedestrian, rare, or surreal) that can be specified in the actual world. Given a lottery with Π players, how small does σ—the probability of a player winning the lottery—have to be before it is a safe bet that the lottery has no winner at all? Each player specifies the σ-event that he will win the lottery. The probability that at least one player correctly specifies a win is therefore bounded by $\Pi\sigma$. If in turn $\Pi\sigma < 1/2$, then the probability of the lottery having no winner exceeds the probability of it having a winner. Thus if $\sigma < 1/(2\Pi)$, we are safer betting that the lottery has no winner than betting that it has a winner. We've seen this argument before. Confronted with a σ-specification (E*,E) where $\sigma < 1/(2\Pi)$, this line of reasoning leads us to reject (E*,E) as a coincidence. Note that σ has the same order of magnitude as the reciprocal of Π, the universal specification bound. This is convenient since 10^{-150}, though smaller than Borel's 10^{-50}, is still readily accessible.[61]

Given a universal specification bound Π which constitutes an upper bound on the totality of specificational resources available in the universe, we can formulate a final version of the Law of Small Probability:

LSP For $\sigma < 1/(2\Pi)$ no σ-specification is a coincidence.

Stated in this way, LSP is not a recommendation we adopt on subjective grounds, but rather one of many rational bets we place as rational agents operating in the actual world. I say "rational bets" because betting on LSP is consonant with preferring Peirce's red deck over his black deck. In line with Peirce I've argued that axioms (Bet1) and (Bet2) are constitutive of rationality. LSP is a consequence of these axioms once we have in hand a universal specification bound. Other things being equal, the rational agent prefers to bet where probabilities are greatest. This is not only true of gamblers, but also of scientists whose inductions are invariably made against a backdrop of uncertainties. To attain a red card Peirce's rational agent prefers the red deck over the black. To be correct about which specifications are coincidences, LSP's rational agent excludes from the space of coincidences those specifications whose likelihood is less than $1/(2\Pi)$. The rationale in both cases is the same.

8. The Dilemma of Scientific Naturalism

The formal definition of coincidence introduced in Section 6 agrees with the popular conception of coincidence. Coincidence in whatever guise, however, is the butt of a common prejudice that needs to be addressed at the outset. Coincidence frequently connotes brute concurrence of events for which explanation is not only superfluous, but downright unwelcome. Galen Pletcher describes this sense of coincidence as follows:

> To call something a "coincidence" explains nothing. It does exactly the opposite: it asserts that the fact that two events are closely related—in time, and in other ways—does not need to be explained. It says more than that the relation between them cannot (at present) be explained.... To call something a coincidence is to express (even if only implicitly or perhaps even unwittingly) the opinion that it is misguided to search for an explanation (in the proper sense) of the coinciding of the phenomena at issue.[62]

To see that Pletcher has accurately assessed the attitude of scientific naturalism toward coincidence, consider how two notable skeptics play down coincidence. Martin Gardner writes,

> 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.[63]

B. F. Skinner expands on this criticism:

> The world in which we live is an extremely complex sample space, in which it is doubtful whether there are any "laws of chance" which apply to many of the single events occurring in it. Coincidences are certainly to be expected, and the sheer number may be felt to build up a case for a force or agent which is metaphysical, supernatural, or at least not part of the current corpus of science. But the mere accumulation of instances has less to do with probability than with the striking force of coincidence.... Science has not ignored some underlying order; [rather,] it has not yet devised ways of protecting us against spurious evidences of order.[64]

Gardner and Skinner offer the standard skeptical reply against using coincidence to infer an "agent which is metaphysical, supernatural, or at least not part of the current corpus of science." First, we tend for psychological reasons to read into coincidence more than is actually there. Thus if by a premonition I delay stepping outside and thereby avoid having a safe fall on my head, my naive reaction is likely to be that God or a guardian angel spared my life. Second, in many cases of coincidence no straightforward assignment of (im)probability to the coincidence is possible. This is Skinner's comment about the inapplicability of the laws of chance to many coincidences. Thus in the example where the safe misses my head, it's unclear what chance process alternatively protects or mangles me. Finally, when plausible probability assignments can be made, comparing the numerous opportunities to observe coincidences with the actual number of coincidences observed, we find that the relative frequency of coincidences agrees with the assigned probabilities. This is Gardner's argument, and it is supposed to justify assimilating coincidences to "spurious evidences of order."

Now I agree that any rigorous analysis of coincidence must repudiate the psychological significance of this or that coincidence.[65] Moreover, I am content to forego analyzing coincidences which cannot be grounded in a plausible chance mechanism. Nevertheless, I wholly reject the claim that even when probabilities can be coherently assigned, coincidences can be dismissed because, in the words of Gardner, "the probability of noticing a startling correlation is quite high, especially if you keep a sharp outlook." Stated in this way, the claim seems to assert a necessary, analytic truth. But this is false. The claim is in fact empirical. If "the number of events in which you participate for a month"—to use Gardner's phrase—is some number N, and if the coincidence in which you participate has probability p which is much smaller than 1/N, then the probability of noticing that startling correlation is quite low, even if you keep a sharp lookout. Without realizing it, by equating coincidence with a highly probable event once the vast number of opportunities for observing that event are factored in, Gardner issues a challenge. For if the analysis of coincidence is a probabilistic and therefore empirical question, as Gardner seems to admit, then we can legitimately ask how to explain highly improbable coincidences which remain improbable even with a vast number of opportunities to observe them. We are back to the question of probabilistic and specificational resources: how many opportunities are there to observe and specify an event and how do these numbers relate to the probability of the event itself? This is a question we've spared no effort elucidating and answering.

The problem therefore is not to explain coincidence generally, but σ-coincidence. Gardner and Skinner do not formulate the problem of σ-coincidence, much less address it. The reason for this failure, however, cannot be credited to ignorance or oversight on their part. Unlike unicorns and fire-breathing dragons, σ-coincidences are credible to a large assortment of reasonable people whose grip on reality seems unimpaired (e.g., Newton, Boyle, and Maxwell all held to the Biblical signs, wonders, miracles, and prophecies). Gardner and Skinner are therefore regularly confronted with supposed σ-coincidences, i.e., coincidences which if actual would be immune to their criticisms of coincidence. This is particularly true of Martin Gardner who as the author of *Science: Good, Bad, and Bogus* as well as a frequent contributor to *The Skeptical Inquirer* (a quarterly for debunking the supernatural) knows all too well the claims made in support of σ-coincidence. Why then do they dismiss the problem of coincidence generally and sweep the problem of σ-coincidence under a rug?

To answer this question we need to understand the dilemma that confronts scientific naturalism whenever it is presented with a putative σ-coincidence. Cosmologist Edwin Hubble perhaps said it best: "Not until the empirical resources are exhausted, need we pass on to the dreamy realms of speculation."[66] I refer to this statement as Hubble's Canon. Hubble's Canon epitomizes the scientific naturalist's commitment to naturalistic explanation as the sole avenue to reliable knowledge. Indeed, to have one's empirical resources exhausted is to be at a loss for naturalistic explanation. Any other kind of explanation is therefore wholly suspect since it resides in the dreamy realms of speculation. Now there is a way to exhaust Hubble's empirical resources without going astray in Hubble's "dreamy realms of speculation." The vehicle which travels this fine line between empirical reliability and metaphysical intangibility is coincidence. Confronted with a well accredited σ-coincidence, the scientific naturalist experiences an inner conflict psychologists call cognitive dissonance. Hubble's Canon lays bare this inner conflict: on the one hand the scientific naturalist is confronted with a breakdown of empirical resources, on the other he must face his aversion to alternative explanatory frameworks which remove him from the realm of science and translate him into that other realm of dreamy speculation.

The scientific naturalist's dilemma is this. On the one hand, he is committed to letting nature speak for herself. He will not legislate what nature can do; nature alone must decide. His task therefore is to describe and understand nature, not to prescribe and confine nature according to some artificial set of presuppositions. This commitment is not a philosophical position, but an empirical attitude. I call it the Normal Empirical Attitude (NEA).[67] On the other hand, the scientific naturalist cannot escape the conclusion that his cognitive faculties condition his understanding of the world. These faculties are subject to limitations and work in certain prescribed ways. Philosophers are fond of finding such limitations and prescribing norms for rationality. Aristotle's three laws of thought is one recommendation. Kant's *Critique of Pure Reason* proposes an entire conceptual system. In line with Peirce, we argued that reason is obliged to respect the Law of Small Probability (LSP). Confronted with a σ-coincidence NEA and LSP pull in opposite directions. Coincidence as defined in Section 6 is a question of fact open to empirical scrutiny. NEA therefore has no quarrel with σ-coincidence—if it happens, it happens. Against this laissez-faire attitude, LSP asserts as a principle of reason that σ-coincidences don't happen. This is the dilemma.

An example might help. In the interest of eliminating the national deficit, the federal government agrees to hold a national lottery in which the grand prize is to be dictator of the United States for a single day—i.e., for 24 hours the winner will have full power over every aspect of government. If a white supremacist wins, he can order the wholesale execution of non-whites. If a porn king wins, he can order this country turned into a giant debauch. If a pacifist wins, he can order the destruction of all our weapons. . . . The more moderate elements of the society will clearly want to prevent the looney fringe from winning, and will therefore be inclined to invest heavily in this lottery. This natural inclination, however, is mitigated by the following consideration: the probability of any one ticket winning is on the order of 10^{-300}. To buy a ticket, the lottery player pays a fixed price and then records a 0–1 string of length 1000—whichever string he chooses. He is permitted to purchase as many tickets as he wishes, subject only to his financial resources and the time it takes to record the 0–1 strings of length 1000. The lottery is drawn at a joint session of congress with each member of congress flipping a coin one or two times (an adjustment needs to be made to get exactly 1000 flips) and recording the resulting coin tosses in alphabetical order by name of congressman or senator.

Suppose now that the fateful day has arrived. Billions of tickets have been sold at ten dollars a piece. To prevent cheating Congress has enlisted the services of the National Academy of Sciences. Following the NAS's recommendation each ticket holder's name is duly entered onto a secure data base, together with the tickets purchased and the ticket numbers (i.e., the bit strings relevant to deciding the winner). All this information is now in place. After much fanfare the congressmen and senators start flipping their coins. As soon as Congressman Zybrowski has announced his final toss, the data base is consulted to determine whether the lottery had a winner. Lo and behold, the lottery did indeed have a winner—Joe "Killdozer" Skinhead, leader of the White Trash Nation. Joe's first act as dictator is to raise a swastika over the Capitol.

The scientific naturalist's dilemma is this. The number of lottery tickets sold is so miniscule compared to the total number of tickets possible that the Law of Small Probability bars the lottery from having a winner (for the lottery to have a winner would constitute a σ-coincidence). The Law of Small Probability therefore refuses to attribute Joe Skinhead's win to a genuine σ-coincidence. On the other hand, with the National Academy of Sciences monitoring the lottery, the empirical evidence is overwhelming that the lottery was properly conducted, and therefore that Joe Skinhead's win did in fact constitute a genuine σ-coincidence. The Normal Empirical Attitude therefore supports the conclusion that Joe's win was a genuine σ-coincidence.

How then does the scientific naturalist resolve this dilemma? In practice the scientific naturalist opts for one of two courses, which correspond to siding with either LSP or NEA. The one course is to charge fraud, the other to invoke chance. Charging fraud is consistent with LSP, for if fraud was involved, then the putative σ-coincidence was no coincidence at all. On the other hand, invoking chance is consistent with NEA, for if the empirical evidence confirms that a σ-coincidence occurred, then no rationalistic principle like LSP must be allowed to interfere with this conclusion. In practice the choice is between LSP and NEA, with no middle road between the two.

The hard-boiled scientific naturalist prefers LSP over NEA, at least when push comes to shove. Consider for instance David Hume's celebrated critique of miracles: no testimony can establish the truth of a miracle (a special type of σ-coincidence) because uniformities in the natural order render the supposed miracle less credible than fraud and collusion among the witnesses.[68] According to Hume it is always more probable that someone is lying than that a reputed miracle happened. Fraud is the method of choice for dismissing σ-coincidence. If fraud can be consistently maintained, it is unnecessary to entertain other modes of explanation. If there really is fraud, the hard-boiled scientific naturalist will be quick to spot and expose it. But if after a diligent search he finds not only that evidence of fraud is lacking, but also that solid empirical evidence precludes the possibility of fraud, he will claim it occurred nonetheless. This attitude underlies Martin Gardner's book *Science: Good, Bad, and Bogus* and characterizes his colleagues at the *Skeptical Inquirer*.[69]

The hard-boiled scientific naturalist as a matter of principle debunks any putative σ-coincidence. He concurs with the Law of Small Probability and therefore automatically charges fraud whenever a σ-coincidence is advertised. He invariably demands a naturalistic explanation for why a σ-specification (E*,E) isn't a coincidence. Thus he is committed to finding a scientific explanation which establishes that at least one of the defining conditions for coincidence ((Ch) or (ZInf)) was violated. But what if the most thorough scientific investigation into (E*,E) fails to produce any evidence that either (Ch) or (ZInf) was violated? What if this same investigation in fact confirms that both (Ch) and (ZInf) were satisfied? To persist in charging fraud, or alternatively, to continue to advocate LSP is therefore to commit oneself to a crypto-causality which maintains that a causal explanation acceptable to science does indeed exist—one demonstrating how either (Ch) or (ZInf) were violated—but is simply unknown to us at present. This is a naturalistic fideism, a belief that the events occurring in nature operate solely according to the laws of nature, regardless what happens.

Suppose now that a thorough scientific investigation of the σ-specification (E*,E) confirms that conditions (Ch) and (ZInf) were satisfied. If sufficiently many such specifications become confirmed, NEA will start to press on the scientific naturalist's tendency to maintain LSP. LSP is sound gambling strategy. Nevertheless, the soundest gambling strategy combined with the most favorable odds do not guarantee success if by guarantee we mean absolute and total certainty. LSP is not an analytic truth. Given sufficiently many well-attested σ-coincidences, the scientific naturalist will feel the pull of reality. If he finds the evidence for σ-coincidence convincing, and if he is committed to his scientific naturalism, he appears to have but one option: invoke chance. Now there are crude and sophisticated ways of invoking chance. The crude way is simply to shrug one's shoulders and mumble in amazement, "Gee, what dumb luck," and leave it at that. If, however, a good deal is riding on the interpretation of a σ-coincidence—for example, whether God was responsible for it—then a more subtle tack is required. Indeed, the "dumb luck" argument allows anyone with an alternate explanation to come along and foist it on his listeners.

More sophisticated ways to invoke chance are, however, available. Most congenial to the scientific naturalist is a strategy that systematically augments probabilistic and specificational resources. For instance, given the standard big bang cosmology, I computed a universal specification bound $\Pi = 10^{150}$ so that $\sigma < 1/(2\Pi)$ captures the idea of "sufficiently small" for the Law of Small Probability. The scientific naturalist may therefore argue that my standard universe is too small and that the true cosmology is inflationary, thereby justifying a much bigger universal specification bound Π and therefore a much smaller surreality index σ. The problem with an inflationary universe is that its justification resides solely in increasing the probabilistic and specificational resources needed to make coincidences plausible which with fewer resources would have seemed implausible. In fact, the only empirical support for an inflationary universe is this increase of resources.[70] Think of it this way: to convince me you flipped heads 30 times in a row with a fair coin (the probability is around one in a billion), you have to convince me you flipped the coin sufficiently often to make this otherwise rare event likely. The conservative cosmologist claims you lacked opportunity to attain this sequence. The inflationary cosmologist in turn says you had plenty of opportunity—just look at all your counterparts in parallel worlds who were also flipping coins.[71]

In times past it used to be much easier to "inflate" probabilistic and specificational resources than it is now. The question whether the universe is finite or infinite used to be a philosophical, not an empirical question. Thomas Aquinas claimed it was only by revelation that we could know that the universe was finite. Reason, according to him, left the possibility of an infinite universe open. Spinoza's monism required an infinite universe, but this was again philosophically conditioned. In a passage already quoted, Hume noted the benefits that accrue to scientific naturalism when a universe of infinite duration is presupposed:

> A finite number of particles is only susceptible of finite transpositions: and it must happen, in an eternal duration, that every possible order or position must be tried an infinite number of times. This world, therefore, with all its events, even the most minute, has before been produced and destroyed, and will again be produced and destroyed, without any bounds and limitations. No one, who has a conception of the power of infinite, in comparison of finite, will ever scruple this determination.[72]

In his younger days Einstein had been committed to Spinoza's God (= Spinoza's infinite universe). This commitment led him to formulate his field equations to accord with an infinite universe. But "when in 1927 the Abbé Lemaître derived from Einstein's cosmological equations the expansion of the universe and correlated that rate with data on galactic red-shifts already available,"[73] the infinity of the universe became an empirical question. The "data on galactic red-shifts already available" was that of Hubble and Humason. When in the early 1930's Einstein visited Hubble in California and inspected his data, Einstein came away convinced that the universe was indeed finite.[74] The inflationary universe of Alan Guth and his successors, much like the steady state theory of the 1950's, attempts to recapture Spinoza's lost infinity. In my view, these theories arise solely out of a need to preserve scientific naturalism, in this case by increasing probabilistic and specificational resources and thereby making the appeal to chance plausible.

Still another way to invoke chance looks like the shoulder shrug ("Gee, what dumb luck"), but adds a theory of explanation, or, if you will, a theory of non-explanation. In his travels through the universe, the naive inquirer attaches question marks to whatever arouses his curiosity. Each question mark remains in force until a satisfactory explanation is found for the phenomenon in question. Once such an explanation is found the question mark can be erased. The philosopher can therefore pose a meta-question: Are all such questions valid? Is it legitimate to attach question marks and search for solutions without restriction, or are there limits to the questions we can legitimately ask and the answers we can in good conscience propose?[75] Recall Dawkins' philosopher. When asked by Dawkins "How did Hume explain the organized complexity of the living world?" the philosopher responded, "He didn't—why does it need any special explanation?"[76] A theory of non-explanation removes coincidence from the realm of scientific discourse. If a coincidence can be dismissed as fraud, fine. But if the best scientific evidence supports its occurrence, no problem. On this view, coincidences that are intractable to scientific explanation must simply be accepted as brute facts without recourse to non-scientific explanatory frameworks. This approach preserves scientific naturalism by refusing to address questions that science finds intractable.

A final sophisticated way to invoke chance is to cloak our ignorance in terminology. This is accomplished by transmuting the object of our ignorance into a first principle of the understanding. Carl Jung does precisely this with his notion of *synchronicity*.[77] Coincidences which are incredibly improbable, subjectively meaningful, and beyond the explanatory capabilities of science according to Jung happen, and happen often. Jung writes,

> As a psychiatrist and psychotherapist I have often come up against the phenomena in question [i.e., σ-coincidences] and could convince myself how much these inner experiences meant to my patients. In most cases they were things which people do not talk about for fear of exposing themselves to thoughtless ridicule. I was amazed to see how many people have had experiences of this kind and how carefully the secret was guarded. So my interest in this problem has a human as well as a scientific foundation.[78]

Synchronicity, however, has nothing to do with non-physical agents causally interfering with the natural order. Nor is it to be identified with divine providence. Synchronicity is thoroughly non-causal. In fact Jung defines it as an "acausal connecting principle."[79] Jung contrasts synchronicity with causality as follows:

> Causality is the way we explain the link between two successive events. Synchronicity designates the parallelism of time and meaning between psychic and psychophysical events, which scientific knowledge so far has been unable to reduce to a common principle. The term [synchronicity] explains nothing, it simply formulates the occurrence of meaningful coincidences which, in themselves, are chance happenings, but are so improbable that we must assume them to be based on some kind of principle. . . .[80]

Jung as a psychologist is interested in the psychological impact of coincidence. Coincidence for Jung is not coincidence simpliciter, but "meaningful coincidence." Jung is therefore not interested in coincidence as brute fact, but rather coincidence as it affects the psyche. There is nothing intellectually disreputable about investigating the psychological significance of coincidence, especially since it is a historical fact that humans through the ages have been non-trivially influenced by coincidence. Having said this, however, I must stress that Jung's work on synchronicity does nothing to disturb the scientific picture of the world. Synchronicity, insofar as it is an acausal connecting principle and therefore has anything to do with causality at all, is thoroughly Kantian in spirit. The causality Jung refers to is in the mind, not in the world. Synchronicity is a psychological principle grafted onto a naturalistic conception of the world which leaves scientific naturalism intact. Jung is therefore correct to admit that synchronicity explains nothing—the notion is scientifically sterile.

9. Explaining Coincidence

Instead of explaining coincidence, the scientific naturalist explains it away. To stay true to his scientific naturalism, he has two courses. In the one case he refuses to admit that σ-coincidences even occur. In the other he admits that they occur, but then refuses to say anything more. The first course is consistent with LSP, but leads to difficulties in the face of well-attested σ-coincidences. The second course rejects LSP, but offers no rationale why σ-coincidences should occur. Is there really a dilemma? If the evidence for σ-coincidence is cogent, why not simply discard LSP? And if it isn't, why not simply reject all supposed σ-coincidences as spurious and keep LSP intact? Neither a facile rejection of LSP, nor a blanket rejection of σ-coincidence can, however, be sustained. LSP is sound gambling strategy—the same gambling strategy regularly employed by scientists in formulating their theories.[81] On the other hand, the evidence for σ-coincidence is not easy to dismiss.

There really is a problem for the scientific naturalist. Unfortunately, up to now I've concentrated on toy examples based on random sampling from a suitably stylized chance mechanism (e.g., coin tossing). If all well-attested σ-coincidences are without exception trivial, it is legitimate to ask whether my efforts might not have been expended more usefully elsewhere. If focusing on toy examples weren't enough, all previous examples of σ-coincidence have been stated counterfactually. Hence there was never a question of verifying the occurrence of this or that σ-coincidence; instead our analysis centered on the conditions for the possibility of a σ-coincidence. Add to this an air of science-fiction surrounding many of the previous examples, and it's not clear whether I'm addressing a real problem at all. Perhaps the hard-boiled scientific naturalist is correct. Perhaps all reputed σ-coincidences are so inconsequential or ill-supported that they need not be admitted into rational discussion. Perhaps Jung's "psychic" patients are not to be trusted—Jung's evidence is after all anecdotal.

If explaining σ-coincidence is a more reputable enterprise than explaining unicorns and goblins, we need a stock of well-attested σ-coincidences. I therefore offer the following catalogue of σ-coincidences which I personally find compelling. Each of the examples is worth an entire book—some an entire library. It follows that I can offer only the merest sketch of why I accept these examples as genuine σ-coincidences. I should point out, however, that the first step is the biggest. Once the conviction dawns that even one σ-coincidence is legitimate, it becomes easier to accept others. For instance, because I already accept the miraculous healings attributed to Jesus, I have no difficulty accepting Alexis Carrel's report of miraculous healings. In the same vein, Jung's examples of synchronicity become plausible once the first big step is taken.

The Resurrection. Consider the specification (Luke 9:22, Luke 24:6–7). Luke 9:22 records the following description:

> And [Jesus] said, "The Son of Man must suffer many things and be rejected by the elders, chief priests and teachers of the law, and he must be killed and on the third day be raised to life."

Luke 24:6–7 records the corresponding event:

> He [Jesus] is not here; he has risen! Remember how he told you, while he was still with you in Galilee: "The Son of Man must be delivered into the hands of sinful men, be crucified and on the third day be raised again."

Is (Luke 9:22, Luke 24:6–7) a σ-coincidence? If these events happened in the order and manner outlined in the Gospel of Luke, then the following observations hold: (1) Luke 9:22 predicts the event recorded in Luke 24:6–7 and is therefore asserted with zero-information about that event; (2) physics accounts for the event recorded in Luke 24:6–7 as a thermodynamic accident of incredibly low probability (much less than the surreality index 10^{-150}). Hence if we accept Luke's Gospel as an historically accurate record of the events surrounding Jesus' life, then (Luke 9:22, Luke 24:6–7) is a σ-coincidence.[82] «««»»

Miracles. Miracles typically involve a specification. At one point when Jesus predicts his death and resurrection, he states, "I have told you now before it happens, so that when it does happen you will believe." (John 14:29) The Resurrection is epistemically significant precisely because it was predicted.[83] The faith Jesus expects from his disciples derives not merely from the event of the Resurrection, but from that event in conjunction with its specification. Using our technical apparatus, we might say it is important for belief that miracles comprise σ-coincidences. For a miracle simply to occur apart from any specification disconnects it from human understanding. It would be instructive to catalogue the miracles of Scripture which are in some way specified. Here is a sampler: Moses told Pharaoh precisely which plague would strike Egypt next if Pharaoh did not obey God by letting Israel depart; only after Elijah called down fire from heaven did it actually fall; in many acts of supernatural healing recorded in the New Testament, a prayer precedes the actual healing—this prayer can be taken as a specification.

Miraculous healings preceded by prayer are widely reported even in this age. Consider the following remarks by Dr. Alexis Carrel. Alexis Carrel received the 1912 Nobel Prize in Physiology and Medicine for his work in transfusion, suturing blood vessels, and transplantation of organs. During Word War I he and Henry Dakin developed a method for treating wounds.[84] Unless therefore Carrel is willfully deceiving his reader, it is hard to attribute the following account to misunderstanding or lack of expertise:

> In all countries, at all times, people have believed in the existence of miracles, in the more or less rapid healing of the sick at places of pilgrimage, at certain sanctuaries. But after the great impetus of science during the nineteenth century, such belief ... disappeared. It was generally admitted, not only that miracles did not exist, but that they could not exist. As the laws of thermodynamics make perpetual motion impossible, physiological laws oppose miracles. Such is still the attitude of most physiologists and physicians. However, in view of the facts observed during the last fifty years this attitude cannot be sustained. The most important cases of miraculous healing have been recorded by the Medical Bureau of Lourdes. Our present conception of the influence of prayer upon pathological lesions is based upon the observation of patients who have been cured almost instantaneously of various affections, such as peritoneal tuberculosis, cold abscesses, osteitis, suppurating wounds, lupus, cancer, etc. The process of healing changes little from one individual to another. Often, an acute pain. Then a sudden sensation of being cured. In a few seconds, a few minutes, at the most a few hours, wounds are cicatrized, pathological symptoms disappear, appetite returns.... The miracle is chiefly characterized by an extreme acceleration of the processes of organic repair.[85]

Carrel concludes on a theological note:

> The only condition indispensable to the occurrence of the phenomenon is prayer. But there is no need for the patient himself to pray, or even to have any religious faith. It is sufficient that some one around him be in a state of prayer.[86]

Unfortunately, like Simon Magus of Acts 8, charlatans claiming supernatural powers are always ready to entice and bilk a gullible public. A skeptic like James Randi (professional magician, debunker of the paranormal, and author of *The Faith Healers*) does the public genuine service when he exposes the tricks of television faith healers. However, because his skepticism derives from his scientific naturalism, even if Randi were presented with evidence of a miraculous healing that satisfied his stringent standards, he would dismiss it as spontaneous remission/regression of the disease. His book *The Faith Healers* does a wonderful job exposing fraud on the part of sleazy evangelists.[87] But he never takes the phenomenon described by Carrel seriously. Despite his extensive, though vain, efforts to obtain proof positive of miracles (from the sleazy faith healers as well as from Lourdes), Randi's mind is made up. He agrees with Ellen Bernstein who writes,

> Miracles ... are conditional; they depend on time, place, what is known, and what is not known. As medical sophistication increases, miracles necessarily decrease, which may mean that the days of "miraculous cures" at Lourdes are numbered.[88]

This is nothing less than faith that the scientific picture of the world will eventually accommodate everything, including miracles. Whether this requires more faith than religion, I don't know. But at this point, the scientific picture of the world still contains substantial gaps, as even Randi must admit when he quotes from *Acta Orthopaedica Scandinavica*:

> A histologically confirmed malignant, primary bone tumour in the [left] pelvis, presumably an osteosarcoma, underwent spontaneous regression. The large tumour was inoperable and gave rise to severe pain as well as difficulty in walking. After 2 years of progression, with increasing destruction of the pelvic bones, the clinical and radiological condition improved spontaneously, and at present the patient is alive, almost symptom-free, after 6 years follow-up.[89]

Is this regeneration of the pelvic bones a miracle? The preceding report makes no mention of prayer. But who can doubt that the sufferer did not recall the halcyon days of robust health and desire a return to good health? This very desire constitutes a specification. I would therefore include this case among the miracles referred to by Carrel.

Although I believe that miracles, and particularly miracles of healing, refer to genuinely occurring phenomena, I also believe a conservative attitude toward them is warranted. Miracles are not to be found in every nook and cranny (unlike medieval hagiography). Carrel has some sobering words:

> Miraculous cures seldom occur. Despite their small number, they prove the existence of organic and mental processes that we do not know.... They are stubborn, irreducible facts, which must be taken into account. The author knows that miracles are as far from scientific orthodoxy as [mysticism].... *But science has to explore the entire field of reality*. [The author] has attempted to learn the characteristics of this mode of healing, as well as of the ordinary modes. He began this study in 1902, at a time when the documents were scarce, when it was difficult for a young doctor, and dangerous for his future career, to become interested in such a subject.... There is a slowly growing literature about miraculous healing. Physicians are becoming more interested in these extraordinary facts. Several cases have been reported at the Medical Society of Bordeaux by professors of the medical school of the university and other eminent physicians.[90] «««»»

Prophecy. Fulfilled prophecy always includes two elements: the prediction of an event and the event itself. By equating prediction with description, we see therefore that prophecies are always specifications. Moreover, since predictions by definition always precede the events foretold, the second defining condition for coincidence—(ZInf)—is always satisfied. If therefore the predicted event can properly be considered the output of a chance mechanism, the first defining condition for coincidence—(Ch)—will also be satisfied and the prophecy will constitute a coincidence. If in addition the event is sufficiently improbable, the prophecy will constitute a σ-coincidence.

Prophecies can be distinguished from miracles by the type of event specified. The type of event predicted in a prophecy accords with the scientific picture of the world. The type of event foretold in a miracle does not fit into the scientific picture of the world except as a bizarre accident—an incredibly improbable thermodynamic or quantum mechanical accident. If I predict how a coin will land in a thousand tosses, that's prophecy. If I predict that the coin will spontaneously toss itself a thousand times, that's a miracle. The probabilities associated with miracles are many orders of magnitude less than the surreality index 10^{-150}. In this vein Richard Dawkins offers an amusing, but instructive example:

> If a marble statue of the Virgin Mary suddenly waved its hand at us we should treat it as a miracle.... In the case of the marble statue, molecules in solid marble are continuously jostling against one another in random directions. The jostlings of the different molecules cancel one another out, so the whole hand of the statute stays still. But if, by sheer coincidence, all the molecules just happened to move in the same direction at the same moment, the hand would move. If they then all reversed direction at the same moment the hand would move back. In this way it is *possible* for a marble statue to wave at us. It could happen.

> The odds against such a coincidence are unimaginably great but they are not incalculably great. A physicist colleague has kindly calculated them for me. The number is so large that the entire age of the universe so far is too short a time to write out all the noughts! [The probability in question is 1/M where M is a number whose common logarithm is way beyond a trillion—$\log_{10} M \geq 10^{12}$. Now that's big!] ... We can *calculate* our way into regions of miraculous improbability far greater than we can *imagine* as plausible.[91]

Unlike miracles, prophecies involve events which by themselves don't surprise us—the surprise, if any, lies in the coincidence between prediction and event. The question therefore remains whether any known instances of fulfilled prophecy count as σ-coincidences. The problem with many prophecies is their vagueness. What I want in a prophecy is unambiguous prediction—a clear way to decide where, when, and how the prophecy was fulfilled. Even many of the Scriptural prophecies do not satisfy this requirement (controversies in the history of theology over the millennium make this abundantly clear).

Another problem with prophecy is the lack of clear probability assignments to the foretold events. Peter Stoner's assignment of likelihoods to the Old Testament messianic prophecies is a case in point: his assignments resulted from having a class of undergraduates vote on what they thought was the probability of a given messianic prophecy accurately predicting an event in the life of an arbitrary individual. By combining all these probabilities via the probability calculus, Stoner computed a surreal probability of 10^{-157} that any individual could satisfy 48 of these Old Testament prophecies. Stoner's conclusion was that these prophecies referred to Jesus and confirmed him in his role as messiah.[92]

Despite the obvious weaknesses in Stoner's argument generally, I believe there is an inner core to his argument that can be salvaged. Just as in a court of law where the consistent testimony of independent witnesses grows stronger with the number of witnesses, so the consistent testimony of Old Testament prophecy about the messiah grows stronger with the number of prophets (as long as the prophets are saying basically the same thing). Alone Isaiah 53—Isaiah's prophecy about the suffering servant—contains so many precise details that a quantitative analysis in terms of probabilities does not seem far-fetched. I believe the messianic prophecies of the Old Testament refer to Jesus and taken jointly constitute a σ-coincidence. A precise argument to support this assertion has yet to be formulated. «««»»

Parapsychology. Parapsychology is the scientific study of a certain type of coincidence. Although parapsychological coincidences are often described tendentiously as extrasensory perception (ESP) or psychokinesis (PK), descriptive terms like these suggest a supernatural causal explanation which is neither philosophically nor scientifically warranted. Unlike most sciences which begin with a universally acknowledged object of study, parapsychology is largely concerned with determining whether there is an object of study at all. The literature is broad and there are strong feelings on both sides, with some holding that parapsychology constitutes a genuine science, with others holding that it constitutes a pseudoscience. The academic community as a whole looks askance at parapsychology. Parapsychology tends not to be represented in the psychology departments of North American universities. The *Journal of Parapsychology*, the unequaled source for research in parapsychology, is not considered a prestigious journal within academic circles. For instance, when Princeton University cut back its library budget in 1988, it canceled its subscription to this journal.

Parapsychology experiments can be schematized as follows: The experimenter randomly samples from a chance mechanism $\mathbf{C} = (\Omega,\Sigma,\mathbf{P})$ and obtains an outcome $\varphi \in \Omega$. Meanwhile, the subject, having been given a descriptive language $\mathbf{D} = (\Lambda,\tau)$ for $\mathbf{C}$, selects a description R^* from Λ. Probabilistically, the important thing about such experiments is the reported p value. The p value is the probability of the event $\tau(R^*)$ described by the subject, i.e., $p = \mathbf{P}(\tau(R^*))$. From the vantage of parapsychology the experiment is successful if the following conditions are met:

(1) $(R^*,\{\varphi\})$ is a specification, i.e., $\varphi \in \tau(R^*)$.
(2) $(R^*,\{\varphi\})$ satisfies (Ch) and (ZInf), i.e., $(R^*,\{\varphi\})$ is also a coincidence.
(3) $p = \mathbf{P}(\tau(R^*))$ is small—the smaller the better.

(1) and (3) are easy to assess. (2) is more difficult and centers on methodology and experimental design. Ruth Reinsel summarizes the crucial elements of this methodology as follows:

> No matter what the form of psi [the factor or faculty supposedly responsible for parapsychological coincidences] being investigated, double-blind methods are essential. That is, the assistant who prepares the targets [in our terminology, the events φ] has no contact with the subjects. The experimenter who interacts with the subjects does not know what the targets are on any given trial. Scoring is double-checked by an assistant who is blind to the hypothesis of the experiment, had no contact with the subjects, and does not know to which experimental group or condition the subjects belonged.[93]

Having described the proper way to conduct a parapsychological experiment and thereby satisfy the three preceding conditions, Reinsel throws a sop to the scientific community. To keep peace with the scientific community she adds:

> Finally, it is important in parapsychological research, as in other fields of scientific endeavor, to await independent replication of a finding before drawing any definite conclusions.[94]

In making this comment, Reinsel vitiates the whole parapsychological enterprise. The desire for replication in parapsychology is simply misplaced. Suppose the p value of a well designed and rigorously controlled parapsychology experiment is less than the surreality index σ ($p < \sigma$). Then this one experiment would provide conclusive grounds for rejecting the Law of Small Probability (i.e., $(R^*,\{\varphi\})$ would constitute a well-confirmed σ-coincidence). A single parapsychological experiment would therefore suffice to overturn scientific naturalism's confidence in LSP.

Given a well-confirmed σ-coincidence, to demand experimental confirmation through "independent replication" is self-defeating—like requiring a lottery winner to win additional lotteries for his first win to be credible. One such win is adequate, especially if the payoff is big enough. When the p value is a surreal probability, the payoff is sufficient for all of space and time. Replication is otiose for parapsychology experiments whose p value is a surreal probability (assume the usual proviso about the experiment being well designed and rigorously controlled). In the language of lotteries and specificational resources, σ is so small that despite the combined effort of all lottery players in the universe, the probability of the lottery having a winner remains miniscule. With parapsychology experiments one success at $p < \sigma$ suffices to establish parapsychology as a valid scientific enterprise. Further successes are gravy.

Now it is common in the parapsychology literature to see very small p values reported for ESP trials—$p < 10^{-60}$ is common. This is bigger than my σ index 10^{-150}, but smaller than Borel's 10^{-50}. Unfortunately, my knowledge of the parapsychology literature is sparse and unsystematic. I do not know whether a well designed and rigorously controlled parapsychological experiment has produced a σ-coincidence for $\sigma = 10^{-150}$—σ around 10^{-80} is the best I've seen. One way around my ignorance would be to augment the surreality index. Thus I might argue that $\sigma = 10^{-150}$ is ultra-conservative (which I believe it is) and that a much larger σ will work equally well. For instance, it is unthinkable that before the heat death of the universe humans should perform more than 10^{30} scientific experiments, much less 10^{30} parapsychological experiments. Hence a surreality index on the order of 10^{-30} seems not unreasonable.

In fine, I believe parapsychology constitutes a valid scientific enterprise. Its object of inquiry is a certain type of σ-coincidence, often labeled psi, ESP, or PK. Although these labels suggest a causal explanation of the coincidence which transcends the natural order, the actual coincidence can be analyzed without presupposing any particular causal interpretation. What I should like to see is a comprehensive catalogue of parapsychological experiments indexed by their p values together with a critical review of each corresponding experimental setup. «««»»

Anthropic Coincidence. Stephen Hawking defines the anthropic principle as follows: "We see the universe the way it is because if it were different, we would not be here to observe it."[95] Richard Swinburne characterizes the anthropic principle this way: "Unless the universe were an orderly place, men would not be around to comment on the fact."[96] The anthropic principle is supposed to block all metaphysical explanations of design and coincidence. To see that the principle fails this mission, consider Swinburne's delightful example of a mad kidnapper:

> Suppose that a madman kidnaps a victim and shuts him in a room with a cardshuffling machine. The machine shuffles ten packs of cards simultaneously and then draws a card from each pack and exhibits simultaneously the ten cards. The kidnapper tells the victim that he will shortly set the machine to work and it will exhibit its first draw, but that unless the draw consists of an ace of hearts from each pack, the machine will simultaneously set off an explosion which will kill the victim, in consequence of which he will not see which cards the machine drew. The machine is then set to work, and to the amazement and relief of the victim the machine exhibits an ace of hearts drawn from each pack. The victim thinks that this extraordinary facts needs an explanation in terms of the machine having been rigged in some way. But the kidnapper, who now reappears, casts doubt on this suggestion. "It is hardly surprising," he says, "that the machine [drew] only aces of hearts. You could not possibly see anything else. For you would not be here to see anything at all, if any other cards had been drawn." But of course the victim is right and the kidnapper is wrong. There is indeed something extraordinary in need of explanation in ten aces of hearts being drawn. The fact that this peculiar order is a necessary condition of the draw being perceived at all makes what is perceived no less extraordinary and in need to explanation.[97]

If the reader is uncomfortable with science fiction examples, he need merely reflect on a well-known game of chance which fortunately is barred from the casinos—Russian roulette. The relevance of both examples to our discussion is this: Just because a chance event is the condition for the possibility specifying that event does not obviate the need to explain the coincidence between specification (i.e., description) and event. If the specification is merely read off the chance event, there is no coincidence and nothing needs to be explained. But if the specification is formulated with zero information about the chance event—even if that event is responsible for my continued existence—the call for explanation remains in force. The (ZInf) condition for coincidence is more subtle than strict causal independence between specification and event. Causality is a philosophically intractable notion. Indeed, if God is the creator of everything, it is not clear whether any two things are causally independent—all causal chains lead back to the same God.

Consider now the problem of cosmogony—how did the universe come to be the place it is? Scientifically, the most reputable model for the origin of the universe currently is the big bang scenario. At the moment of the big bang the proportions among atomic elements as well as the fundamental constants of physics were decided. According to the big bang model these proportions and fundamental constants could have been different. Simply to assert that this is a contingent universe, however, misses the point. Not just life as we know it, but life generally—any sort of organized complexity—appears impossible if we jiggle even slightly these proportions and fundamental constants. True, if the proportions and constants were different, we would not be here to comment on that fact. But we are here and we know what proportions and constants must obtain for us to be here. Swinburne's case of the mad kidnapper applies. In the big bang science understands a chance event. In ascertaining the proportions and fundamental constants that obtain in the actual world, science formulates a specification. The chance event is responsible for the existence of the universe, and is therefore a precondition for this specification. But the specification—the proportions and constants—can be discovered without presupposing any cosmogonal model. The zero-information condition obtains. There is a coincidence that needs to be explained. In fact I believe it is a σ-coincidence. My knowledge of astrophysics and cosmology is too scanty, however, to assign definite probabilities. Anyone with the proper expertise should be able to assign probabilities and perform the necessary calculations.[98]

«««»»

The Origin of Life. With the origin of life we come full circle. Life has always been the key instance of design. Indeed, a design argument that is barred from using life is destined for extinction: When the living systems that propound design arguments no longer consider themselves objects of intelligent design, they cease to find intelligent design anywhere else in the universe. Notice that we are back to our original metaphysical question, Is the fundamental principle of reality intelligent or unintelligent? Zoologist Richard Dawkins believes life is designed in the sense of being specified. Moreover, he believes that specification requires "a special effort of explanation."[99] But he also believes that the specification involved in life's origin derives from a blind, not an intelligent watchmaker, i.e., from nature. Dawkins' special effort of explanation is neo-Darwinism.

Before analyzing the origin of life in reference to design, let us contrast the origin of life with the origin of species. The origin of life is concerned with the transition from inorganic matter to living systems. The origin of species is concerned with the transitions between living systems. Darwin's theory of natural selection is a theory of competition among already living entities: the best survive and have the most offspring. Hence his title, *The Origin of Species*. Natural selection needs living things on which to operate. Most scientists agree that selection pressures are real and induce long term changes in living systems. Just how extensive and radical are such changes, however, is an open question.[100] Darwin and his successors, the neo-Darwinists, believe that natural selection explains the full range of living systems once a single organism is postulated. Life according to this view is totally plastic. Varying degrees of plasticity are of course possible, down to the absolute fixity of the species, a view common before Darwin. The point to recognize is that the origin of life is a fundamentally different problem from the origin of species. The latter requires that life already be present, the former that no life be present, at least not initially. The origin of life is generally regarded as the more difficult of the two problems.

Let us formulate the problem of life's origin within the framework of chance mechanisms and descriptive languages. The relevant chance mechanism is the universe—call it **C**. The critical event output by **C** is CARBOLIFE. CARBOLIFE is ordinary carbon-based life utilizing DNA/RNA as the replication mechanism. The relevant descriptive language is one that describes all possible configurations of the universe—call it **D**. The critical description from **D** is LIFE*. LIFE* describes not only CARBOLIFE, but all other life forms capable of attaining consciousness. (This requirement is necessary to exclude trivial life forms whose complexity barely exceeds a game of tic-tac-toe.[101] LIFE* describes life forms whose intellectual potential at least matches that of humans. Since humans actually exist, this requirement has empirical support.) LIFE* is therefore a description which leaves the possibility of other life forms open. It should be pointed out, however, that CARBOLIFE is the only life with which we are actually acquainted—SILICOLIFE (i.e., silicon-based life) for all we know might be the null event. The specification which therefore concerns us is (LIFE*, CARBOLIFE). The big question is whether (LIFE*, CARBOLIFE) is a σ-coincidence.

First observe that before life emerged, the universe comprised a vast array of inorganic chemistry experiments. Each of these experiments (much like Urey and Miller's experiment in the 1950's shooting electric charges through ammonia solutions) can be viewed as sampling a chance mechanism. The composite of all these chance mechanisms is therefore the grand chance mechanism we call the universe, or equivalently **C**. CARBOLIFE is the random output of one of those chemistry experiments/chance mechanisms all of which jointly constitute the universe. CARBOLIFE therefore satisfies defining condition (Ch) of coincidence. Next observe that although LIFE* depends for its formulation on CARBOLIFE (in typical anthropic fashion, we have to exist in order to say we exist), LIFE* requires no information about CARBOLIFE. The meanest savage can affirm LIFE* in total ignorance of the event CARBOLIFE. In fact the savage's ignorance about CARBOLIFE is matched by the scientific community's. To be sure, stories about life emerging from a tranquil ponds abound.[102] But the scientific community can neither reconstruct the event that gave birth to the first life (the evidence is all gone), nor determine the exact organism that got the ball rolling. This is zero-information. Condition (ZInf) is therefore satisfied as well. (LIFE*, CARBOLIFE) is a coincidence.

Richard Dawkins agrees. The problem for him is not whether (LIFE*, CARBOLIFE) is a coincidence, but whether it is a coincidence of sufficiently small likelihood. Dawkins addresses this problem in *The Blind Watchmaker* in a chapter entitled "Origins and Miracles." It is enough for him to show that the probabilities favor life originating at least once (once is enough—as soon as life is on the scene, natural selection can take over). At the root of his probabilistic argument is what he calls the spontaneous generation probability, or SGP. The SGP is "the probability ... that life will originate on any randomly designated planet of some particular type."[103] He goes on to say,

> It is the SGP that we shall arrive at if we sit down with our chemistry textbooks, or strike sparks through plausible mixtures of atmospheric gases in our laboratory, and calculate the odds of replicating molecules springing spontaneously into existence in a typical planetary atmosphere.[104]

Logically, the next step in the argument should be a computation—actually compute SGP. But here Dawkins pleads ignorance: "Chemists don't know the answer to this question."[105] Without an answer, without even an upper bound on SGP, Dawkins abruptly ends the argument:

> There are probably more than a billion billion [10^{18}] available planets in the universe. If each of them lasts as long as Earth, that gives us about a billion billion billion [10^{27}] planet-years to play with. That will do nicely! A miracle is translated into practical politics by a multiplication sum.[106]

Apparently a billion billion billion is so big as to validate any argument. The gap in Dawkins' logic is transparent.[107] A billion billion billion years to perform chemistry experiments on planets as hospitable as the earth need not render life probable. Dawkins' numbers aren't even all that startling given the numbers we've considered earlier. A billion billion billion years is 10^{27} years or approximately 10^{34} seconds. The total surface area of the earth is under 10^{21} square millimeters. Hence if each square millimeter of the earth's surface could be utilized every second to perform an experiment that might issue in life, there would be no more than 10^{55} ($= 10^{27} \times 10^{21}$) such experiments. This number isn't even close to the universal specification bound $\Pi = 10^{150}$.

What is the probability that any of these 10^{55} experiments will issue in life? What if there were 10^{150} such experiments—more experiments than could possibly be packed into the universe? Would this many additional experiments render life probable? Even this number of experiments cannot overcome the extreme improbabilities that arise when concrete numbers are assigned to probabilities like Dawkins' SGP. The improbabilities are truly staggering. Fred Hoyle, for instance, computes that a single cell might on the basis of chance be expected every 10^{40000} years if the entire universe were filled with a prebiotic liquid (an assumption he means to be generous).[108] Bernd-Olaf Küppers, commenting merely on a certain subunit of a virus, writes:

> The RNA sequence that codes for the virus-specific subunit of the replicase complex consists of approximately a thousand nucleotides, ... so that it already possess $\lambda^n = 4^{1000} \sim 10^{600}$ alternative sequences.... The spontaneous synthesis [of this system] ... is therefore extremely improbable.[109]

He concludes that probability theory, even when supplemented by Karl Popper's propensity theory, "does not bring us a single step further as regards the statistical aspect of the origin of life."[110] Lecomte du Noüy found similarly wild improbabilities back in the 1940's.[111]

I'm neither a chemist nor a molecular biologist. Hence I don't know if such wild improbabilities are the final word. If they are only approximately correct, then (LIFE*, CARBOLIFE) is certainly a σ-coincidence. If on the other hand molecular biologists should discover a ratchet principle whereby inorganic matter can successively build itself up into a living system without unduly straining the probabilities, then (LIFE*, CARBOLIFE) would remain a coincidence, but lose its status as a surreal coincidence. What is the likelihood of (LIFE*, CARBOLIFE)? Completely convincing figures are not available. What rough estimates there are, however, strongly support that (LIFE*, CARBOLIFE) is a σ-coincidence. «««»»

What has become of our fundamental intuition? Early in the game Borel told us that "events whose probability is sufficiently small never occur."[112] This he called The Single Law of Chance. What's more, he gave us a concrete number, $\sigma = 10^{-50}$, to characterize his sufficiently small probabilities. In the course of our analysis we found much that needed tidying up in Borel's original formulation. The question of probabilistic resources needed to be addressed: with 10^{50} opportunities to observe an event, 10^{-50} is hardly small enough. The question of specification needed to be addressed: events of probability 10^{-50}, 10^{-500}, and even 10^{-5000}, happen all the time (just start flipping a coin); specified events, however, don't. The related question of specificational resources needed to be addressed: how many specifications can the community with an interest in the specified event formulate? Questions about gambling strategies, the role of community vis-a-vis logic, and the size of the universe needed to be addressed as well. Our analysis culminated in a universal specification bound $\Pi = 10^{150}$ which in turn facilitated a precise reformulation of Borel's Single Law of Chance. We called it the Law of Small Probability (LSP): For $\sigma < 1/(2\Pi)$ no σ-specification is a coincidence.

G. K. Chesterton wrote, "The most incredible thing about miracles is that they happen."[113] Miracles happen. σ-coincidences happen. If this is true, what becomes of the Law of Small Probability? Logic seems to demand that LSP be false. What's more, the preceding catalogue provides solid evidence that LSP is in fact false. To say that LSP is false and leave it at that, however, is to abort our analysis before it comes to fruition. The important thing to understand is not that LSP is false, but why LSP is false and how it can be salvaged. LSP is false because σ-coincidences really do happen. Fine. But why do σ-coincidences happen? The scientific naturalist when pressed to accept the reality of σ-coincidences wants either to leave off explanation or to invoke a modified science that can accommodate σ-coincidence. The first is simply a bald denial that any explanation is required. The second introduces bogeymen like action at a distance and reverse causality to preserve the scientific picture of the world, though not the science that is universally recognized.

The scientific naturalist runs into problems with LSP because he is a gambler at heart. Scientific induction, inference to the best explanation, and the propounding of hypotheses are all gambles. Evidence is never univocal, only probable. The scientific naturalist is a betting man and feels obliged to bet where the probabilities are greatest. Hence he encounters difficulties whenever required to bet in ways that are not optimal, at least from the vantage of probability. LSP is a bet he desperately wants to make. This is the reason for all the skeptical and rationalist and humanist societies intent on debunking the paranormal. This is the reason why accepting one's first σ-coincidence is often a conversion experience. If, however, the conversion is merely an acceptance of σ-coincidence that leaves naturalism intact, then the problem remains what to do about LSP? LSP is not merely sound betting strategy. From the vantage of probability, LSP is more likely to be right than its negation, in fact much more likely as we shrink σ. Why then did LSP fail?

Is it important to answer this question? If all the σ-coincidences in the world were as significant as tossing a sequence of 1000 heads, then no answer would be required. The problem, however, is that coincidences are rarely value neutral. Consider the following comment by the theologian William Hordern of what was then Garrett Biblical Institute:

> I couldn't care less whether Caesar crossed the Rubicon or not. It doesn't make any difference to me. I'm not going to lead my life any differently tomorrow either way; nothing stands or falls with it. Perhaps if I made my living out of history, and was battling with some other colleague, we might have ourselves a real battle among historians over precisely such questions. There is hardly anything that has happened in past history that doesn't get debated by historians at some time or other. Most of us couldn't care less, however; we have no real involvement with this. But [in the resurrection of Jesus] we have a story that comes to us from two thousand years ago, and if it is true, then my destiny not only here but hereafter depends upon this story—and you ask me to believe it on the basis only of the generally unreliable historical data?[114]

Hordern clearly isn't convinced that the Resurrection is a bona fide σ-coincidence. He does, however, make a startling admission: if the Resurrection were a genuine σ-coincidence, his and the destiny of the human race would depend on it.

If LSP fails, the reason it fails must be sought in the coincidences that make it fail. The scientific naturalist is uncomfortable with LSP failing. Someone like myself doesn't experience the same discomfort. Unlike the scientific naturalist, I'm not a gambler at heart. As a Christian I have a different set of commitments. For me LSP is not final. Instead another law is final, the Law of the Sovereign God (LSG):

> [God's] dominion is an everlasting dominion, and his kingdom is from generation to generation. And all the inhabitants of the earth are reputed as nothing; and he doeth according to his will in the army of heaven, and among the inhabitants of the earth; and none can stay his hand, or say unto him, What doest thou?[115]

The Resurrection leads me to believe that God was incarnate in Jesus, reconciling the world to himself. The miracles and prophecies reported in Scripture lead me to believe that God is an active participant in history, that nothing happens apart from him, and that coincidences are not autonomous accidents. Modern miracles of healing lead me to believe that God still cares for and loves his creation. Parapsychology leads me to believe that physicalist reductions of man are invalid. The anthropic coincidences lead me to believe that a creator of tremendous wisdom and power is responsible for this universe. Finally the origin of life leads me to believe that God is an incredibly capable designer.

My aim here isn't to proselytize, but to indicate why LSP doesn't distress me. LSP is a problem for scientific naturalism. LSP is not a problem for Christian theology. There are a great many other conceptual frameworks, however, which can as well accommodate the breakdown of LSP. Any brand of monotheism will do (e.g., Judaism and Islam). Any dualistic religion that postulates competing gods of good and evil will do (e.g., Manichaeism and Zoroastrianism). A dualistic philosophy of the neo-Platonic stripe will work. An anthropology which postulates "hidden powers of the mind" can at least accommodate parapsychology. A pantheistic view that permits action at a distance and reverse causality (i.e., effects preceding their causes) will work. The demarcation is between naturalism on the one hand, and non-naturalism (I won't say supernaturalism) on the other.

An analogy with mathematics might clarify the problem of reconciling LSP, coincidence, and scientific naturalism. Before Gödel proved his famous incompleteness theorem, Hilbert had affirmed that "every definite mathematical problem must necessarily be susceptible of an exact settlement"[116]; alternatively, every mathematical proposition is decidable. Hilbert thereby affirmed the completeness of mathematics, that in principle every mathematical question humans might pose could be settled by humans (given sufficient computational resources). Now there is a straightforward way to undermine this claim to completeness: simply find a proposition which is undecidable and can be proven to be undecidable. Inspired by the liar paradox, Gödel constructed what is now called a Gödel sentence, i.e., a sentence for which it can be shown that neither it nor its negation has a mathematically acceptable proof.

If we now shift gears from math to physics, we find that LSP makes a similar completeness claim about physics. Namely, LSP asserts that we won't witness any bizarre thermodynamic or quantum mechanical accidents. Thus we won't see Dawkins' statue of the Virgin Mary wave at us, we won't see people spontaneously combust as in the movie *This is Spinal Tap*, we won't see guns materializing from nowhere and shooting our favorite enemies, and we won't see dead men come back to life. LSP makes it possible to assert that physics is a complete, coherent account of reality. Now just as the Gödel sentence undermined the completeness of mathematics, so a single σ-coincidence undermines the completeness of physics: any bona fide σ-coincidence falsifies LSP once and for all.

The incompleteness of physics is a problem for scientific naturalism, but not for other world views (Christian theism in fact demands the incompleteness of physics). What then is the effect of a design argument which takes the Law of Small Probability and the fact of σ-coincidences, and plays the two against each other—a principle of reason against a matter of fact? Kant was right: such an argument does not establish the truth of the Christian God. It does, however, demonstrate the incompleteness of physics and thereby undercut the pretensions of scientific naturalism. It does not settle the apologetic question about which religion, if any, is correct. Its effect is negative rather than positive. It shows that any system of thought that on first principles eschews religion is going to have problems living up to those principles. I think this is all we can expect from a design argument; to demand more is misleading. If you want more, you'll have to look elsewhere—in the case of Christianity, to history.[117]

Appendix: Salvaging the Law of Small Probability

What then becomes of the Law of Small Probability? The Law of Small probability states categorically that no specified σ-event happens by chance, i.e., there are no σ-coincidences. However, given a σ-specification (R*,S), we said in Section 6 that it is an empirical question whether (R*,S) is also a coincidence. For (R*,S) to be a coincidence it must satisfy the two defining conditions for coincidence: (Ch) and (ZInf), respectively the chance condition (i.e., S resulted through random sampling) and the zero-information condition (i.e., R* was formulated without any information about S). Now the empirical evidence relevant to these conditions falls roughly into three categories:

(1) Positive evidence in favor of both (Ch) and (ZInf) being satisfied.
(2) Positive evidence against at least one of (Ch) or (ZInf) being satisfied.
(3) Ambivalent or insufficient evidence about (Ch) and (ZInf) being satisfied.

In the first instance we are obliged to believe that (R*,S) is a σ-coincidence, contrary to LSP. In the second, we have proof positive that (R*,S) isn't a coincidence. But in the third, do we side with LSP or do we leave open the possibility that (R*,S) is a coincidence?

An example might help us decide the last question. Suppose a Las Vegas gambling house is required to keep records of all its results. The gaming commissioner at year's end examines the casino's blackjack winnings and discovers that the casino won more than its expected number of wins. By itself this is unexceptional. Chance being what it is, he expects that the casino will win sometimes more, sometimes less than the probabilistically expected value. The problem, however, is that the casino won unduly many times at blackjack. The commissioner discovers that when he performs the calculation, the probability that the casino's wins exceeded by so great a margin its expected number of wins is 10^{-200}, a surreal probability. His conclusion: either the casino was cheating or the records are in error. This uncharitable conclusion is entirely in keeping with the Law of Small Probability. The commissioner does not think the casino just happened to be lucky.

There are legal ramifications to the Law of Small Probability. In discrimination cases it is usually enough to prove that the proportions are all wrong, even if affirmative action is not actively employed. The burden of proof is with the employer. The bigoted employer cannot pretend that the total lack of minority representation at his firm is the result of chance. Of course, there is a very small, positive probability that no minority applicant should ever contact his firm's employment office. But appealing to this small probability will not form the basis of his legal defence if the ACLU takes him to court.

The point then is this: *Assume LSP unless there is very strong evidence to the contrary.* This is plain common sense. Solomon encouraged this attitude: "The simple believeth every word; but the prudent man looketh well to his going."[118] Scientists rightly adopt this skeptical stance all the time. Consider, for instance, R. A. Fisher's analysis of Gregor Mendel's data on peas: Fisher thought Gregor Mendel's data were falsified. Why? It is thought that "Mendel's data were massaged," as one statistics text puts it, because the observations he made matched his theory too closely. Interestingly, the coincidence that elicited this accusation of data falsification was a specified event whose probability was no more extreme than 10^{-5}, a probability which is big by our standards (cf. Borel's

10^{-50} and my 10^{-150}). Fisher concluded his analysis of Mendel's experiment by charging Mendel's gardening assistant with deception.[119]

The Law of Small Probability is unavoidable whenever there is a question of data falsification. I recall hearing about a psychologist who was about to be dismissed from his post for lifting data from one of his articles and transporting it into another. The give-away was a 2 x 2 table which made identical appearances in both articles. Each of the table's four blocks contained a three digit number. If the researcher was honest, the odds would therefore have been better than one in a trillion that this same table should appear twice in his research. He resigned in shame rather than defend a 10^{-12} improbability.

The Law of Small Probability is inverse to the statistical problem of hypothesis testing. Writing about hypothesis testing, Ian Hacking observes,

> An adequate theory of testing must consider not only the statistical hypothesis under test, but also rivals to it. This may be common-sense: "don't reject something unless you've something better." It involves a conception now becoming general in the philosophy of science, and which is currently striving to oust the former idea that an hypothesis could be rejected independently of what other theories are available.[120]

LSP falls under what Hacking calls "the former idea": LSP rejects the chance occurrence of a specified σ-event without offering an alternative explanation, statistical or otherwise.

In statistics one is generally given a sample space Ω, a family of distributions $\Psi = \{\mathbf{P}_\theta \mid \theta \in \Theta\}$ on Ω (the cardinality of Ψ is strictly greater than 1), and a random sample $X_1 = x_1, X_2 = x_2, \ldots, X_n = x_n$, and from these one has to decide which $\mathbf{P}_\theta$ represents the population distribution (cf. the empirical distribution function).[121] On the other hand, to say that specified events don't happen when the probability is too small is to start out with a sample space Ω, a single distribution $\mathbf{P}$, and a specified event $E \subset \Omega$; and then, if E happens, to reject $\mathbf{P}$ without supplying an alternative hypothesis. Statistically speaking this is what LSP is doing—it is strictly in the business of rejecting $\mathbf{P}$, not accepting anything else.

To appreciate the difference between hypothesis testing and LSP, imagine that a die is to be thrown six million times. Hypothesis testing is interested in two hypotheses: H_o, the null hypothesis, which asserts that the die is fair (i.e., each face has probability 1/6); and H_1, the alternate hypothesis, which asserts that the die is in some way loaded. Suppose now that the die is thrown six million times and that each face appears precisely one million times. A chi-square goodness of fit test has no possibility of rejecting the null hypothesis.[122] Hypothesis testing therefore accepts, or if you will fails to reject, the null hypothesis H_o. The probability of so close a fit with expected values, however, is a surreal probability. LSP therefore leads to the opposite conclusion: the null hypothesis H_o is false because the probability that six million throws of a fair die match expected values so closely is virtually infinitesimal.

NOTES

[1]See Bowler [1989: 52ff.] for a nice overview of the design argument. It is significant that such overviews are common in texts on evolution. The reason is simple. Prior to Darwin the design argument was a scientifically respectable way to infer God from nature. Historical accounts of evolution therefore treat the design argument as a precursor to modern evolutionary accounts of origins.

[2]See Quine [1961].

[3]Psalm 14:1 and Psalm 53:1.

[4]Romans 1:19–20.

[5]Schaff and Wace [1989: 290].

[6]*Summa Theologiae* I, 2 (3) (see Aquinas [1270: 13]).

[7]*Summa Theologiae* I, 2 (3) (see Aquinas [1270: 14]). This quote—that nature works for a determinate end—begins Thomas's reply to the following objection: "It seems that everything we see in the world can be accounted for by other principles, supposing God did not exist. For all natural things can be reduced to one principle, which is nature, and all voluntary things can be reduced to one principle, which is human reason, or will. Therefore there is no need to suppose God's existence." [Aquinas 1270: 12] This objection is thoroughly modern, and though easy to dispense with in an Aristotelian framework, is difficult to handle when teleology is not presupposed. In fact this objection formulates the central problem facing Christian apologetics in modern times. By relying on a defunct peripatetic science, Aquinas offers the modern apologist little help.

[8]See "The Specific Nature of the Philosophic Order" in Gilson [1990].

[9]Monod [1972: 21].

[10]Paley [1802: 352-53].

[11]Swinburne [1979: 135].

[12]Swinburne's argument from design falls within this pattern. What I'm calling laws of nature he calls regularities of succession.

[13]I am of course oversimplifying, but the general point is correct. Note that I have said nothing about verifying the underlying fact of the matter. Such verification will be anything but straightforward.

[14]Swinburne [1979: 148].

[15]Alvin Plantinga [1967: chapter 4; 1974: 81–84] seems not to appreciate this point in his criticism of the teleological argument. Indeed, Plantinga appears to accept Hume's criticism as decisive: "Hume's criticism seems correct. The conclusion to be drawn, I think, is that the teleological argument ... is unsuccessful"—Plantinga [1974: 84].

[16]Dawkins [1987: 5–6].

[17]Hume [1779: 67].

[18]In fairness to Hume I should mention that he does not appear to endorse this argument. Hume begins the argument with the following qualification: "What if I should revive the old Epicurean hypothesis? This is commonly, and I believe justly, esteemed the most absurd system that has yet been proposed. . . ." Hume [1779: 67].

[19]There are good reasons—scientific, philosophical, and theological—to suppose that these finiteness conditions are true. I am deliberately sidestepping steady state, oscillating, and inflationary universes: the scientific facts seem to go against the first two, whereas last makes most of the universe causally irrelevant to the piece of the universe we inhabit. See Jaki [1989] and Ross [1989] for more details.

[20]See Kant [1929: 518–524], the passage entitled "The Impossibility of the Physico-theological Proof."

[21]Swinburne [1979: 141].

[22]See Reid [1986: 72]. Hilbert was responding to Emil duBois-Reymond's pessimism about the progress of knowledge epitomized in the catchword *ignoramus et ignorabimus*—we are ignorant and shall remain ignorant. Hilbert thoroughly repudiated this pessimism: "Every definite mathematical problem must necessarily be susceptible of an exact settlement, either in the form of an actual answer to the question asked, or by the proof of the impossibility of its solution and therefore the necessary failure of all attempts." [Reid 1986: 13]

[23]Swinburne [1979] makes much of subjective/epistemic probabilities.

[24]Even Dawkins [1987: 158] has flights of whimsy in which he admits that the phenomena of nature seem to require a miracle: "Does it sound to you as though it would need a miracle to make randomly jostling atoms join together into a self-replicating molecule? Well, at times it does to me too." For a recent treatment of the standard design argument which exploits our naive wonder at God's handiwork, see the chapter entitled "Why Can't an

Accident Design a Machine" in Gange [1986]. Although Gange's argument won't satisfy the philosopher, I doubt it will ever lose its appeal to the man on the street.

[25]Paley [1802: 352].

[26]Although I'm using the language of computer science, my aim is simply to motivate my subject. That a computational device should input a design as program and output a designatum is not, as we shall see, the fundamental idea underlying design.

[27]Philosophers understand this fine point as a type-token distinction: σ is the type, σ_p and σ_d are the tokens.

[28]This was in fact Michelangelo's "philosophy of sculpting": God had placed statues in the rocks; the sculptors role was find them.

[29]For a general introduction to cryptography, see Patterson [1987].

[30]See Mood, Graybill, and Boes [1974: chapter 9] for a general discussion of hypothesis testing.

[31]A sane jury would surely think this defence attorney needs as much restraint as the defendant. I owe the "thermodynamic gun" to Charlie Huenemann.

[32]The quote is from Wilson and Weldon [1978: 321]. They are in turn referring to Borel [1962: chapter 3]. The "basic law of probability" referred to is Borel's Single Law of Chance which he defines as "phenomena with very small probabilities do not occur" [Borel 1962: 1] or equivalently "events whose probability is sufficiently small never occur" [Borel 1962: 25].

[33]In fairness to Wilson and Weldon [1978: 321] it must be noted that they don't say an event with probability 1 in 10^{50} "cannot happen," but rather that "we can state with certainty that [it] will *never* happen." Thus they might be taken to say that future predicted events of such small probability cannot happen. Although still not adequate, this interpretation is consistent with our subsequent account of specification.

[34]Dawkins [1987: 8]. Emphasis added.

[35]Dawkins [1987: 9]. Emphasis added.

[36]Dawkins [1987: 15].

[37]See Bauer [1981: chapter 6] for a theoretical justification. We tacitly assume that repetitions are stochastically independent. Of course an event with zero probability never occurs, no matter how many repetitions.

[38]Wilson and Weldon [1978: 321].

[39]Hume [1779: 155]. Emphasis added.

[40]See Grimmett and Stirzaker [1982: 38] for their discussion of the geometric distribution. $P(E(n)) = 1-(1-p)^n$ is the maximum probability of observing E given our probabilistic resources. P(E(n)) is p when n = 1 and increases to 1 as n goes to infinity.

[41]The frequentist view of probability is championed in von Mises [1981], the subjectivist view in Keynes [1973]. I am adopting the view of Borel, which is often called an a priori view. I prefer to think of it as an intrinsic or inherentist view of probability: the probabilities are inherent in the chance set-up.

[42]Borel [1963: 28].

[43]Hughes [1989] is a convenient introduction to quantum mechanics for non-physicists, especially for mathematicians and philosophers.

[44]This is not to reject frequentist and subjective approaches to probability out of hand. They simply irrelevant to our investigation.

[45]See the chapter on measurement in Hughes [1989].

[46]Since τ is a one-to-one correspondence, it may seem that the descriptive language Λ is redundant—every statement about Λ is equivalent to a statement about Σ. We retain the distinction between description and event for two reasons: first, it expedites our subsequent analysis; second, when Ω is infinite, the relation between description and event becomes an interesting problem for recursion theory.

[47]See the introduction to Martinich [1985].

[48]Although the reader will not find me advocating a metaphysical doctrine of chance in this essay, my position is probably closest to that of Thomas Aquinas. For my theory of randomness, see Dembski [1991].

[49]Here and in the next example my reference to causality is philosophically innocuous. I'm using causality in a strictly common sense way, not presupposing any metaphysical doctrine of causality.

[50]See Borel [1962: chapter 3].

[51]Peirce [1878: 1313].

[52]Peirce [1878: 1313–1314].

[53]Peirce [1878: 1316].

[54]Peirce [1878: 1315].

[55]Peirce [1878: 1316].

[56]For the lottery paradox and its relation to possible worlds semantics see Stalnaker [1984: 90-91].

[57]Peirce [1878: 1315].

[58]For a friendly introduction to Shannon's ideas see Hamming [1986].

[59]The question of a universal time bound for computation in the actual world is at issue here. As the smallest time unit physicists are able to obtain by juggling their fundamental constants, the Planck time must currently be regarded as the ne plus ultra for computation. For electronic computers Wegener [1987: 2] has a larger bound: 5.6×10^{-33} s. I personally am comfortable with 10^{-24} s as a universal time bound. Known as the *chronon*, this is the time it takes for light to traverse the nucleus of an atom. It is regarded as the smallest *measurable* time—see Jaki [1966: 265].

[60]Observe that the argument for determining Π depends only on the actual world having (1) a bounded number of bits available for storage at any moment—call it β; (2) a bounded switching frequency for the bits—call it ν; (3) a bounded lifetime—call it λ. Under these conditions Π equals the product $\beta\nu\lambda$. The concrete Π we computed used $\beta = 10^{80}$, $\nu = 10^{45}\ s^{-1}$, and $\lambda = 10^{25}$ s. I don't foresee these estimates needing adjustment in the near future, if ever.

[61]A σ level of 10^{-150} isn't bad at all—one has to flip a coin only 500 times to attain such an improbability.

[62]Pletcher [1990: 205–206].

[63]Quoted in Pletcher [1990: 210] from Gardner [1972].

[64]Quoted in Pletcher [1990: 210] from Skinner [1977].

[65]Cf. Jung's notion of synchronicity or "meaningful coincidence." See Jung [1973].

[66]This quote appears as the closing sentence in Hubble's work on galaxies, *The Realm of the Nebulae*. See Hubble [1936: 202].

[67]Cf. Arthur Fine's Natural Ontological Attitude (NOA) in Fine [1986: 112–135].

[68]See Hume [1748].

[69]*The Skeptical Inquirer* is published by the Committee for the Scientific Investigation of Claims of the Paranormal. The address is Box 229, Central Park Station, Buffalo, New York 14215.

[70]Cf. Jaki [1989: 134–139].

[71]If this sounds more like science fiction than science, you are not alone.

[72]Hume [1989: 67].

[73]Jaki [1989: 28].

[74]See Jastrow [1980].

[75]Cf. Paul's admonition to Timothy: "Foolish and unlearned questions avoid, knowing that they do gender strifes." [2 Timothy 2:23]

[76]Dawkins [1987: 5].

[77]See Jung [1973] and Koestler [1972: chapter 3].

[78]Jung [1973: 4].

[79]Jung's book entitled *Synchronicity* is in fact subtitled *An Acausal Connecting Principle*.

[80]Jung [1973: 115].

[81]Any model of scientific inference (e.g., induction, abduction, and the hypothetico-deductive method) is to some degree based on the probabilistic ideas that gave rise to LSP.

[82]The historicity of the Gospels, and the Bible more generally, is of course another question, one I've answered to my own satisfaction, but one which nonetheless needs to be thought through individually. There has been plenty of work in this area. The apologetics of John Warwick Montgomery is a good place to start.

[83]Cf. Dembski [1990: 23–25].

[84]*New Columbia Encyclopedia*, 4th ed., s.v. "Alexis Carrel."

[85]Carrel [1935: 148–149].

[86]Carrel [1935: 149].

[87]The exposure of Peter Popoff is a marvelous piece of detective work. See chapter 9 of Randi [1987].

[88]Randi [1987: 29].

[89]Randi [1987: 28–29].

[90]Carrel [1935: 148 note]. I've placed in italics Carrel's assertion "But science has to explore the entire field of reality." My article "Inconvenient Facts: Miracles and the Skeptical Inquirer" takes up Carrel's claim—see Dembski [1990].

[91]Dawkins [1987: 159–160].

[92]See Stoner's chapter on prophecy and probability in *Science Speaks* [Stoner 1952]. Stoner's argument is widely known in evangelical circles through the apologetic work of Josh McDowell [1986: 166–167; 1990: chapter 19].

[93]Reinsel [1990: 194].

[94]Reinsel [1990: 194].

[95]Hawking [1988: 183].

[96]Swinburne [1979: 137].

[97]Swinburne [1979: 138].

[98]See the articles by Robert Newman in Part 2 of Montgomery [1991] as well as Barrow and Tipler [1986].

[99]Dawkins [1987: 15].

[100]See chapter 4 entitled "A Partial Truth" in Denton [1986].

[101]The artificial life of Christopher Langton [1989] is a case in point.

[102]This was Darwin's speculation about the origin of life.

[103]Dawkins [1987: 144].

[104]Dawkins [1987: 144]. Since no planets have been discovered outside our solar system, talk of "typical planets" is empirically unjustified. Dawkins should probably limit himself to "planets like our own."

[105]Dawkins [1987: 144].

[106]Dawkins [1987: 145].

[107]If I were to play amateur psychologist, I would attribute the fault in Dawkins' reasoning to an incompatible set of commitments: He agrees that (LIFE*, CARBOLIFE) is a coincidence. He demands an explanation for this coincidence—chance will not do. Finally, he refuses to consider supernatural design a valid explanation: "To explain the origin of the DNA/protein machine by invoking a supernatural Designer is to explain precisely nothing, for it leaves unexplained the origin of the Designer. You have to say something like 'God was always there', and if you allow yourself that kind of lazy way out, you might as well just say 'DNA was always there', or 'Life was always there', and be done with it." [Dawkins 1987: 141]

[108]See Hoyle and Wickramasinghe [1981: 1–33, 130-141], Hoyle [1982: 1–65], and the appendix by Herman Eckelmann in Montgomery [1991].

[109]Küppers [1990: 68]. Küppers is a pupil of Manfred Eigen.

[110]Küppers [1990: 68].

[111]See chapter 3 of du Noüy [1947].

[112][Borel 1962: 25].

[113]Chesterton [1950: 11].

[114]Quoted from Montgomery [1965: 106-107].

[115]Daniel 4:34–35.

[116]Reid [1986: 13].

[117]Although this is the logical place to end our investigations, it is also the logical place to begin a historical apologetic on behalf of Christianity. The work of John Warwick Montgomery is a good place to start.

[118]Proverbs 14:15.

[119]See Freedman, Pisani, and Purves [1978: 426-427] and Fisher [1965: 53].

[120]Hacking [1965: 89].

[121]See Mood, Graybill, and Boes [1974: chapter 9].

[122]See Freedman, Pisani, and Purves [1978: chapter 28].

REFERENCES

Adler, Mortimer. J. and Robert M. Hutchins, editors
1952 *Great Books of the Western World*. Chicago: Encyclopaedia Britannica.

Aquinas, Thomas
c.1270 *Summa Theologica*, translated by Fathers of the English Dominican Province, revised by D. J. Sullivan. In Adler and Hutchins [1952: vol. 19].

Barrow, John D. and Frank J. Tipler
1986 *The Anthropic Cosmological Principle*. Oxford: Oxford University Press.

Bauer, Heinz
1981 *Probability Theory and Elements of Measure Theory*. London: Academic Press.

Borel, Emile
1962 *Probabilities and Life*, translated by M. Baudin. New York: Dover.
1963 *Probability and Certainty*, translated by D. Scott. New York: Walker.

Bowler, Peter J.
1989 *Evolution: The History of an Idea*, revised edition. Berkeley, Calif.: University of California Press.

Bush, L. Russ, editor
1983 *Classical Readings in Christian Apologetics, A.D. 100–1800*. Grand Rapids, Mich.: Zondervan.

Carrel, Alexis
1936 *Man the Unknown*. New York: Harper & Brothers.

Chesterton, Gilbert K.
1950 *The Innocence of Father Brown*. New York: Penguin Books.

Dawkins, Richard
1987 *The Blind Watchmaker*. New York: W. W. Norton.

Dembski, William A.
1990 "Inconvenient Facts: Miracles and the Skeptical Inquirer." *Bulletin of the Evangelical Philosophical Society*, **13**: 18–45.
1991 "Randomness by Design." *Nous*, **25**(1): 75–106.

Denton, Michael
1986 *Evolution: A Theory in Crisis*. Bethesda, Maryland: Adler & Adler.

du Noüy, Lecomte
1947 *Human Destiny*. New York: Longmans, Green and Company.

Fine, Arthur
1986 *The Shaky Game: Einstein Realism and the Quantum Theory*. Chicago: University of Chicago Press.

Fisher, Ronald A.
1965 *Experiments in Plant Hybridisation*. Edinburgh: Oliver and Boyd.

Freedman, David, Robert Pisani, and Roger Purves
1978 *Statistics*. New York: Norton.

Gange, Robert
1986 *Origins and Destiny*. Waco, Texas: Word Books.

Gardner, Martin
1972 "Arthur Koestler: Neoplatonism Rides Again" (review of *The Roots of Coincidence*). *World* (1Aug1972): 87–89.

Gilson, Etienne
1990 *Methodical Realism*, translated by P. Trower. Front Royal, Virginia: Christendom Press.

Grim, Patrick, editor
1990 *Philosophy of Science and the Occult*, 2nd edition. Albany, N.Y.: State University of New York Press.

Grimmett, Geoffrey R. and David R. Stirzaker
1982 *Probability and Random Processes*. Oxford: Oxford University Press.

Hacking, Ian
1965 *Logic of Statistical Inference*. Cambridge: Cambridge University Press.

Hamming, Richard W.
1986 *Coding and Information Theory*, 2nd ed. Englewood Cliffs, N. J.: Prentice-Hall.

Hawking, Stephen W.
1988 *A Brief History of Time: From the Big Bang to Black Holes*. New York: Bantam.

Hoyle, Fred
1982 *Cosmology and Astrophysics*. Ithaca, N.Y.: Cornell University Press.
Hoyle, Fred and Chandra Wickramasinghe
1981 *Evolution from Space*. New York: Simon & Schuster.
Hubble, Edwin P
1936 *The Realm of the Nebulae*. New Haven, Conn.: Yale University Press.
Hughes, R. I. G.
1989 *The Structure and Interpretation of Quantum Mechanics*. Cambridge, Mass.: Harvard University Press.
Hume, David
1748 *Of Miracles*. La Salle, Ill.: Open Court, 1985.
1779 *Dialogues Concerning Natural Religion*. Buffalo, N.Y.: Prometheus Books, 1989.
Jaki, Stanley L.
1966 *The Relevance of Physics*. Chicago: University of Chicago Press.
1989 *God and the Cosmologists*. Washington, D.C.: Regnery Gateway.
Jastrow, Robert
1980 *God and the Astronomers*. New York: Warner Books.
Jung, Carl G.
1973 *Synchronicity*, translated by R. F. C. Hull. Princeton, N.J.: Princeton University Press.
Kant, Immanuel
1929 *Critique of Pure Reason*, translated by N. K. Smith. New York: St. Martin's Press.
Keynes, John Maynard
1973 *A Treatise on Probability*. Cambridge: Cambridge University Press.
Koestler, Arthur
1972 *The Roots of Coincidence: An Excursion into Parapsychology*. New York: Random House.
Küppers, Bernd-Olaf
1990 *Information and the Origin of Life*. Cambridge, Mass.: MIT Press.
Langton, Christopher G.
1989 *Artificial Life* (The proceedings of an interdisciplinary workshop on the synthesis and simulation of living systems held September, 1987 in Los Alamos, New Mexico). In *The Santa Fe Institute Studies in the Sciences of Complexity*, volume 6. Redwood City, Calif.: Addison-Wesley, Advanced Book Program.
Martinich, A. P., editor
1985 *The Philosophy of Language*. New York: Oxford University Press.
McDowell, Josh
1986 *Evidence that Demands a Verdict*, volume 1. San Bernardino, Calif.: Here's Life Publishers.
1990 *A Ready Defense: The Best of Josh McDowell*. San Bernardino, Calif.: Here's Life Publishers.
Monod, Jacques
1972 *Chance and Necessity: An Essay on the Natural Philosophy of Modern Biology*. New York: Random House.
Montgomery, John W.
1979 *History and Christianity*. Downers Grove, Illinois: InterVarsity Press.
Montgomery, John W., editor
1991 *Evidence for Faith: Deciding the God Question*. Dallas, Texas: Probe.
Mood, Alexander M., Franklin A. Graybill, and Duane C. Boes
1974 *Introduction to the Theory of Statistics*, 3rd edition. New York: McGraw-Hill.
Newman, James R.
1988 *The World of Mathematics*, in 4 volumes. Redmond, Wash.: Tempus.
Paley, William
1802 *Natural Theology or Evidences of the Existence and Attributes of the Deity Collected from the Appearances of Nature*. In Bush [1983].
Patterson, Wayne
1987 *Mathematical Cryptology for Computer Scientists and Mathematicians*. Totowa, N.J.: Rowman & Littlefield.
Peirce, Charles S.
1878 "The Red and the Black." In Newman [1988: 1311–1317]. Originally titled "The Doctrine of Chances." See *Philosophical Writings of Peirce*, edited by J. Buchler, New York: Dover, 1955.

Plantinga, Alvin
1967 *God and Other Minds*. Ithaca, N.Y.: Cornell University Press.
1974 *God, Freedom, and Evil*. New York: Harper & Row.
Pletcher, Galen K.
1990 "Coincidence and Explanation." In Grim [1990].
Quine, Willard
1961 "On What There Is." In *From a Logical Point of View*, 2nd edition, pp. 1–19. New York: Harper & Row.
Randi, James
1987 *The Faith Healers*. Buffalo, N.Y.: Prometheus.
Reid, Constance
1986 *Hilbert–Courant*. New York: Springer.
Reinsel, Ruth
1990 "Parapsychology: An Empirical Science." In Grim [1990].
Ross, Hugh
1989 *The Fingerprint of God*. Orange, Calif.: Promise Publishing Co.
Schaff, Philip and Henry Wace
1989 *Nicene and Post-Nicene Fathers*, series 2, volume 7. Grand Rapids, Mich.: Eerdmans.
Skinner, Burrhus F.
1977 "The Force of Coincidence." *The Humanist*, **37**(3): 10–11.
Stalnaker, Robert C.
1984 *Inquiry*. Cambridge, Mass.: MIT Press.
Stoner, Peter W.
1952 *Science Speaks*. Wheaton, Ill.: Van Kampen Press.
Swinburne, Richard
1979 *The Existence of God*. Oxford: Oxford University Press.
von Mises, Richard
1981 *Probability, Statistics and Truth*, 2nd revised English edition. New York: Dover.
Wegener, Ingo
1987 *The Complexity of Boolean Functions*. Stuttgart: Wiley-Teubner.
Wilson, Clifford and John Weldon
1978 *Close Encounters: A Better Explanation*. San Diego: Master Books.