

The Design Argument

Elliott Sober

The design argument is one of three main arguments for the existence of God; the others are the ontological argument and the cosmological argument. Unlike the ontological argument, the design argument and the cosmological argument are a posteriori. And whereas the cosmological argument can focus on any present event to get the ball rolling (arguing that it must trace back to a first cause, namely God), design theorists are usually more selective.

Design arguments have typically been of two types – *organismic* and *cosmic*. Organismic design arguments start with the observation that organisms have features that adapt them to the environments in which they live and that exhibit a kind of *delicacy*. Consider, for example, the vertebrate eye. This organ helps organisms survive by permitting them to perceive objects in their environment. And were the parts of the eye even slightly different in their shape and assembly, the resulting organ would not allow us to see. Cosmic design arguments begin with an observation concerning features of the entire cosmos – the universe obeys simple laws, it has a kind of stability, its physical features permit life and intelligent life to exist. However, not all design arguments fit into these two neat compartments. Kepler, for example, thought that the face we see when we look at the moon requires explanation in terms of intelligent design. Still, the common thread is that design theorists describe some empirical feature of the world and argue that this feature points towards an explanation in terms of God’s intentional planning and away from an explanation in terms of mindless natural processes.

The design argument raises epistemological questions that go beyond its traditional theological context. As William Paley (1802) observed, when we find a watch while walking across a heath, we unhesitatingly infer that it was produced by an intelligent designer. No such inference forces itself upon us when we observe a stone. Why is explanation in terms of intelligent design so compelling in the one case, but not in the other? Similarly, when we observe the behavior of our fellow human beings, we find it irresistible to think that they have minds that are filled with beliefs and desires. And when we observe non-human organisms,

the impulse to invoke mentalistic explanations is often very strong, especially when they look a lot like us. When does the behavior of an organism – human or not – warrant this mentalistic interpretation? The same question can be posed about machines. Few of us feel tempted to attribute beliefs and desires to hand calculators. We use calculators to help us add, but they don't literally figure out sums; in this respect, calculators are like the pieces of paper on which we scribble calculations. There is an important difference between a device that *we* use to help us think and a device that *itself* thinks. However, when a computer plays a decent game of chess, we may find it useful to explain and predict its behavior by thinking of it as having goals and deploying strategies (Dennett, 1987b). Is this merely a useful fiction, or does the machine really have a mind? And if we think that present-day chess-playing computers are, strictly speaking, mindless, what would it take for a machine to pass the test? Surely, as Turing (1950) observed, it needn't look like us. In all these contexts, we face *the problem of other minds* (Sober, 2000a). If we understood the ground rules in this general epistemological problem, that would help us think about the design argument for the existence of God. And conversely – if we could get clear on the theological design argument, that might throw light on epistemological problems that are not theological in character.

What is the Design Argument?

The design argument, like the ontological argument, raises subtle questions concerning what the logical structure of the argument really is. My main concern here will not be to describe how various thinkers have presented the design argument, but to find the soundest formulation that the argument can be given.

The best version of the design argument, in my opinion, uses an inferential idea that probabilists call *the likelihood principle*. This can be illustrated by way of Paley's (1802) example of the watch on the heath. Paley describes an observation that he claims discriminates between two hypotheses:

- (W) O1: the watch has features G1 . . . Gn.
- W1: the watch was created by an intelligent designer.
- W2: the watch was produced by a mindless chance process.

Paley's idea is that O1 would be unsurprising if W1 were true, but would be very surprising if W2 were true. This is supposed to show that O1 *favours* W1 over W2; O1 supports W1 more than it supports W2. Surprise is a matter of degree; it can be captured by the concept of conditional probability. The probability of O given H – $\Pr(O | H)$ – represents how unsurprising O would be if H were true. The likelihood principle says that comparing such conditional probabilities is the way to decide what the direction is in which the evidence points:

- (LP) Observation O supports hypothesis H1 more than it supports hypothesis H2 if and only if $\Pr(O | H1) > \Pr(O | H2)$.

There is a lot to say on the question of why the likelihood principle should be accepted (Hacking, 1965; Edwards, 1972; Royall, 1997; Forster and Sober, 2003; Sober, 2002); for the purposes of this essay, I will take it as a given.

We now can describe the likelihood version of the design argument for the existence of God, again taking our lead from one of Paley's favorite examples of a delicate adaptation. The basic format is to compare two hypotheses as possible explanations of a single observation:

- (E) O2: the vertebrate eye has features F1 . . . Fn.
E1: the vertebrate eye was created by an intelligent designer.
E2: the vertebrate eye was produced by a mindless chance process.

We do not hesitate to conclude that the observations strongly favor design over chance in the case of argument (W); Paley claims that precisely the same conclusion should be drawn in the case of the propositions assembled in (E).¹

Clarifications

Several points of clarification are needed here concerning likelihood in general and the likelihood version of the design argument in particular. First, I use the term "likelihood" in a technical sense. Likelihood is not the same as probability. To say that H has a high likelihood, given observation O, is to comment on the value of $\Pr(O | H)$, not on the value of $\Pr(H | O)$; the latter is H's *posterior probability*. It is perfectly possible for a hypothesis to have a high likelihood and a low posterior probability. When you hear noises in your attic, this confers a high likelihood on the hypothesis that there are gremlins up there bowling, but few of us would conclude that this hypothesis is probably true.

Although the likelihood of H (given O) and the probability of H (given O) are different quantities, they are related. The relationship is given by Bayes's theorem:

$$\Pr(H | O) = \Pr(O | H)\Pr(H)/\Pr(O).$$

$\Pr(H)$ is the *prior probability* of the hypothesis – the probability that H has before we take the observation O into account. From Bayes's theorem we can deduce the following:

$$\Pr(H1 | O) > \Pr(H2 | O) \text{ if and only if } \Pr(O | H1)\Pr(H1) > \Pr(O | H2)\Pr(H2).$$

Which hypothesis has the higher posterior probability depends on how their likelihoods are related, but also on how their prior probabilities are related. This

explains why the likelihood version of the design argument does not show that design is more probable than chance. To draw this further conclusion, we'd have to say something about the prior probabilities of the two hypotheses. It is here that I wish to demur (and this is what separates me from card-carrying Bayesians). Each of us perhaps has some subjective degree of belief, before we consider the design argument, in each of the two hypotheses E1 and E2. However, I see no way to understand the idea that the two hypotheses have *objective* prior probabilities. Since I would like to restrict the design argument as much as possible to matters that are objective, I will not represent it as an argument concerning which hypothesis is more probable.² However, those who have prior degrees of belief in E1 and E2 should use the likelihood argument to update their subjective probabilities. The likelihood version of the design argument says that the observation O2 should lead you to increase your degree of belief in E1 and reduce your degree of belief in E2.

My restriction of the design argument to an assessment of likelihoods, not probabilities, reflects a more general point of view. Scientific theories often have implications about which observations are probable (and which are improbable), but it rarely makes sense to describe them as having objective probabilities. Newton's law of gravitation (along with suitable background assumptions) says that the return of Halley's comet was to be expected, but what is the probability that Newton's law is true? Hypotheses have objective probabilities when they describe possible outcomes of a chance process. But as far as anyone knows, the laws that govern our universe were not the result of a chance process. Bayesians think that *all* hypotheses have probabilities; the position I am advocating sees this as a special feature of *some* hypotheses.³

Just as likelihood considerations leave open what probabilities one should assign to the competing hypotheses, they also don't tell you which hypothesis you should *believe*. I take it that belief is a dichotomous concept – you either believe a proposition or you do not. Consistent with this is the idea that there are three attitudes one might take to a statement – you can believe it true, believe it false, or withhold judgment. However, there is no simple connection of the matter-of-degree concept of probability to the dichotomous (or trichotomous) concept of belief. This is the lesson I extract from the lottery paradox (Kyburg, 1961). Suppose 100,000 tickets are sold in a fair lottery; one ticket will win and each has the same chance of winning. It follows that each ticket has a very high probability of not winning. If you adopt the policy of believing a proposition when it has a high probability, you will believe of each ticket that it will not win. However, this conclusion contradicts the assumption that the lottery is fair. What this shows is that high probability does not suffice for belief (and low probability does not suffice for disbelief). It is for this reason that many Bayesians prefer to say that individuals have *degrees* of belief. The rules for the dichotomous concept are unclear; the matter-of-degree concept at least has the advantage of being anchored to the probability calculus.

In summary, likelihood arguments have rather modest pretensions. They don't tell you which hypotheses to believe; in fact, they don't even tell you which hypotheses are probably true. Rather, they evaluate how the observations at hand discriminate among the hypotheses under consideration.

I now turn to some details concerning the likelihood version of the design argument. The first concerns the meaning of the intelligent design hypothesis. This hypothesis occurs in W1 in connection with the watch and in E1 in connection with the vertebrate eye. In the case of the watch, Paley did not dream that he was offering an argument for the existence of *God*. However, in the case of the eye, Paley thought that the intelligent designer under discussion was God himself. Why are these cases different? The bare bones of the likelihood arguments (W) and (E) do not say. What Paley had in mind is that building the vertebrate eye and the other adaptive features that organisms exhibit requires an intelligence far greater than anything that human beings could muster. This is a point that we will revisit at the end of this essay.

It is also important to understand the nature of the hypothesis with which the intelligent design hypothesis competes. I have used the term "chance" to express this alternative hypothesis. In large measure, this is because design theorists often think of chance as the alternative to design. Paley is again exemplary. *Natural Theology* is filled with examples like that of the vertebrate eye. Paley was not content to describe a few cases of delicate adaptations; he wanted to make sure that even if he got a few details wrong, the weight of evidence would still be overwhelming. For example, in Chapter 15 he considers the fact that our eyes point in the same direction as our feet; this has the convenient consequence that we can see where we are going. The obvious explanation, Paley (1802, p. 179) says, is intelligent design. This is because the alternative is that the direction of our eyes and the direction of our gait were determined by chance, which would mean that there was only a $\frac{1}{4}$ probability that our eyes would be able to scan the quadrant into which we are about to step.

I construe the idea of chance in a particular way. To say that an outcome is the result of a *uniform chance process* means that it was one of a number of *equiprobable* outcomes. Examples in the real world that come close to being uniform chance processes may be found in gambling devices – spinning a roulette wheel, drawing from a deck of cards, tossing a coin. The term "random" becomes more and more appropriate as real-world systems approximate uniform chance processes. However, as R. A. Fisher once pointed out, it is not a "matter of chance" that casinos turn a profit each year, nor should this be regarded as a "random" event. The financial bottom line at a casino is the result of a large number of chance events, but the rules of the game make it enormously probable (though not certain) that casinos end each year in the black. All uniform chance processes are probabilistic, but not all probabilistic outcomes are "due to chance."

It follows that the two hypotheses considered in my likelihood rendition of the design argument are not exhaustive. Mindless uniform chance is one alternative

to intelligent design, but it is not the only one. This point has an important bearing on the dramatic change in fortunes that the design argument experienced with the advent of Darwin's (1859) theory of evolution. The process of evolution by natural selection is *not* a uniform chance process. The process has two parts. Novel traits arise in individual organisms "by chance;" however, whether they then disappear from the population or increase in frequency and eventually reach 100 percent representation is anything but a "matter of chance." The central idea of natural selection is that traits that help organisms survive and reproduce have a better chance of becoming common than traits that hurt. The essence of natural selection is that evolutionary outcomes have *unequal* probabilities. Paley and other design theorists writing before Darwin did not and could not cover all possible mindless natural processes. Paley addressed the alternative of uniform chance, not the alternative of natural selection.⁴

Just to nail down this point, I want to describe a version of the design argument formulated by John Arbuthnot. Arbuthnot (1710) carefully tabulated birth records in London over 82 years and noticed that, in each year, slightly more sons than daughters were born. Realizing that boys die in greater numbers than girls, he saw that this slight bias in the sex ratio at birth gradually subsides until there are equal numbers of males and females at the age of marriage. Arbuthnot took this to be evidence of intelligent design; God, in his benevolence, wanted each man to have a wife and each woman to have a husband. To draw this conclusion, Arbuthnot considered what he took to be the relevant competing hypothesis – that the sex ratio at birth is determined by a uniform chance process. He was able to show that if the probability is $\frac{1}{2}$ that a baby will be a boy and $\frac{1}{2}$ that it will be a girl, then it is enormously improbable that the sex ratio should be skewed in favor of males in each and every of the years he surveyed (Stigler, 1986, pp. 225–6).

Arbuthnot could not have known that R. A. Fisher (1930) would bring sex ratio within the purview of the theory of natural selection. Fisher's insight was to see that a mother's mix of sons and daughters affects the number of *grand*offspring she will have. Fisher demonstrated that when there is random mating in a large population, the sex ratio strategy that evolves is one in which a mother invests equally in sons and daughters (Sober, 1993, p. 17). A mother will put half her reproductive resources into producing sons and half into producing daughters. This equal division means that she should have more sons than daughters, if sons tend to die sooner. Fisher's model therefore predicts the slightly uneven sex ratio at birth that Arbuthnot observed.⁵

My point in describing Fisher's idea is not to fault Arbuthnot for living in the eighteenth century. Rather, the thing to notice is that what Arbuthnot meant by "chance" was very different from what Fisher was talking about when he described how a selection process might shape the sex ratio found in a population. Arbuthnot was right that the probability of there being more males than females at birth in each of 82 years is extremely low, if each birth has the same chance of producing a male as it does of producing a female. However, a male-biased sex ratio in the population is extremely probable, if Fisher's hypothesized process is

doing the work. Showing that design is more likely than chance leaves it open that some third, mindless, process might still have a higher likelihood than design. This is not a defect in the design argument, so long as the conclusion of that argument is not overstated. Here the modesty of the likelihood version of the design argument is a point in its favor. To draw a stronger conclusion – that the design hypothesis is more likely than *any* hypothesis involving mindless natural processes – one would have to attend to more alternatives than just design and (uniform) chance.⁶

I now want to draw the reader's attention to some features of the likelihood version of the design argument (E) concerning how the observation and the competing hypotheses are formulated. First, notice that I have kept the observation O2 conceptually separate from the two hypotheses E1 and E2. If the observation were simply that "the vertebrate eye exists," then since E1 and E2 both entail this proposition, each would have a likelihood of unity. According to the likelihood principle, this observation does not favor design over chance. Better to formulate the question in terms of explaining the properties of the vertebrate eye, not explaining why the eye exists. Notice also that I have not formulated the design hypothesis as the claim that God exists; this existence claim says nothing about the putative designer's involvement in the creation of the vertebrate eye. Finally, I should point out that it would do no harm to have the design hypothesis say that God created the vertebrate eye; this possible reformulation is something I'll return to later.

Other Formulations of the Design Argument, and Their Defects

Given the various provisos that govern probability arguments, it would be nice if the design argument could be formulated deductively. For example, if the hypothesis of mindless chance processes entailed that it is *impossible* that organisms exhibit delicate adaptations, then a quick application of *modus tollens* would sweep that hypothesis from the field. How ever much design theorists might yearn for an argument of this kind, there apparently are none to be had. As the story about monkeys and typewriters illustrates, it is *not* impossible that mindless chance processes should produce delicate adaptations; it is merely very *improbable* that they should do so.

If *modus tollens* cannot be pressed into service, perhaps there is a probabilistic version of *modus tollens* that can achieve the same result. Is there a law of improbability that begins with the premise that $\text{Pr}(O \mid H)$ is very low and concludes that H should be rejected? There is no such principle (Royall, 1997, chapter 3). The fact that you won the lottery does not, by itself, show that there is something wrong with the conjunctive hypothesis that the lottery was fair and a million tickets were sold and you bought just one ticket. And if we randomly drop a very

sharp pin onto a line that is 1,000 miles long, the probability of its landing where it does is negligible; however, that outcome does not falsify the hypothesis that the pin was dropped at random.

The fact that there is no probabilistic *modus tollens* has great significance for understanding the design argument. The logic of this problem is essentially comparative. To evaluate the design hypothesis, we must know what it predicts and compare this with the predictions made by other hypotheses. The design hypothesis cannot win by default. The fact that an observation would be very improbable if it arose by chance is not enough to refute the chance hypothesis. One must show that the design hypothesis confers on the observation a higher probability, and even then the conclusion will merely be that the observation *favours* the design hypothesis, not that that hypothesis *must be true*.⁷

In the continuing conflict (in the United States) between evolutionary biology and creationism, creationists attack evolutionary theory, but never take even the first step in developing a positive theory of their own. The three-word slogan “God did it” seems to satisfy whatever craving for explanation they may have. Is the sterility of this intellectual tradition a mere accident? Could intelligent design theory be turned into a scientific research program? I am doubtful, but the present point concerns the logic of the design argument, not its future prospects. Creationists sometimes assert that evolutionary theory “cannot explain” this or that finding (e.g., Behe, 1996). What they mean is that certain outcomes are *very improbable* according to the evolutionary hypothesis. Even this more modest claim needs to be scrutinized. However, if it were true, what would follow about the plausibility of creationism? In a word – *nothing*.

It isn't just defenders of the design hypothesis who have fallen into the trap of supposing that there is a probabilistic version of *modus tollens*. For example, the biologist Richard Dawkins (1986, pp. 144–6) takes up the question of how one should evaluate hypotheses that attempt to explain the origin of life by appeal to strictly mindless natural processes. He says that an acceptable theory of this sort can say that the origin of life on earth was somewhat improbable, but it must not go too far. If there are n planets in the universe that are “suitable” locales for life to originate, then an acceptable theory of the origin of life on earth must say that that event had a probability of at least $\frac{1}{n}$. Theories that say that terrestrial life was less probable than this should be rejected. How does Dawkins obtain this lower bound? Why is the number of planets relevant? Perhaps he is thinking that if α is the actual frequency of life-bearing planets among “suitable” planets (i.e., planets on which it is possible for life to evolve), then the true probability of life's evolving on earth must also be α . There is a mistake here, which we can uncover by examining how actual frequency and probability are related. With small sample-size, it is perfectly possible for these quantities to have very different values (consider a fair coin that is tossed three times and then destroyed). However, Dawkins is obviously thinking that the sample size is very large, and here he is right that the actual frequency provides a good estimate of the true probability. It is interesting that Dawkins tells us to reject a theory if the probability it assigns is

too *low*, but why doesn't he also say that it should be rejected if the probability it assigns is too *high*? The reason, presumably, is that we cannot rule out the possibility that the earth was not just *suitable* but was *highly conducive* to the evolution of life. However, this point cuts both ways. Although α is the *average* probability of a suitable planet's having life evolve, it is still possible that different suitable planets might have different probabilities – some may have values greater than α while others may have values that are lower. Dawkins's lower bound assumes that the earth was above average; this is a mistake that might be termed the "Lake Wobegon fallacy."

Some of Hume's ([1779] 1980) criticisms of the design argument in his *Dialogues Concerning Natural Religion* depend on formulating the argument as something other than a likelihood inference. For example, Hume at one point has Philo say that the design argument is an argument from analogy, and that the conclusion of the argument is supported only very weakly by its premises. His point can be formulated by thinking of the design argument as follows.

Watches are produced by intelligent design.
 Organisms are similar to watches to degree p.
 p[=====]
 Organisms were produced by intelligent design.

Notice that the letter "p" appears twice in this argument. It represents the degree of similarity of organisms and watches, and it represents the probability that the premises confer on the conclusion. Think of similarity as the proportion of shared characteristics. Things that are 0 percent similar have no traits in common; things that are 100 percent similar have all traits in common. The analogy argument says that the more similar watches and organisms are, the more probable it is that organisms were produced by intelligent design.

Let us grant the Humean point that watches and organisms have relatively few characteristics in common (it is doubtful that there is a well-defined totality consisting of all the traits of each, but let that pass). After all, watches are made of metal and glass and go "tick tock"; organisms metabolize and reproduce and go "oink" and "bow wow." This is all true, but entirely irrelevant, if the design argument is a likelihood inference. It doesn't matter how overall-similar watches and organisms are. With respect to argument (W), what matters is how one should explain the fact that watches are well adapted for the task of telling time; with respect to (E), what matters is how one should explain the fact that organisms are well adapted to their environments. Paley's analogy between watches and organisms is merely heuristic. The likelihood argument about organisms stands on its own (Sober, 1993).

Hume also has Philo construe the design argument as an inductive argument, and then complain that the inductive evidence is weak. Philo suggests that for us to have good reason to think that our world was produced by an intelligent designer, we'd have to visit other worlds and observe that all or most of them

were produced by intelligent design. But how many other worlds have we visited? The answer is – not even one. Apparently, the design argument is an inductive argument that could not be weaker; its sample size is zero. This objection dissolves once we move from the model of inductive sampling to that of likelihood. You don't have to observe the processes of intelligent design and chance at work in different worlds to maintain that the two hypotheses confer different probabilities on your observations.

Three Possible Objections to the Likelihood Argument

There is another objection that Hume makes to the design argument, one that apparently pertains to the likelihood version of the argument that I have formulated and that many philosophers think is devastating. Hume points out that the design argument does not establish the attributes of the designer. The argument does not show that the designer who made the universe, or who made organisms, is morally perfect, or all-knowing, or all-powerful, or that there is just one such being. Perhaps this undercuts some versions of the design argument, but it does not touch the likelihood argument we are considering. Paley, perhaps responding to this Humean point, makes it clear that his design argument aims to establish the *existence* of the designer, and that the question of the designer's *characteristics* must be addressed separately.⁸ My own rendition of the argument follows Paley in this regard. Does this limitation of the argument render it trivial? Not at all – it is *not* trivial to claim that the adaptive contrivances of organisms are due to intelligent design, even when details about this designer are not supplied. This supposed “triviality” would be *big* news to evolutionary biologists.

The likelihood version of the design argument consists of two premisses – $\text{Pr}(O \mid \text{Chance})$ is very low and $\text{Pr}(O \mid \text{Design})$ is higher. Here O describes some observation of the features of organisms or some feature of the entire cosmos. The first of these claims is sometimes rejected by appeal to a theory that Hume describes under the heading of the Epicurean hypothesis. This is the monkeys-and-typewriters idea that if there are a finite number of particles that have a finite number of possible states, then, if they swarm about at random, they eventually will visit all possible configurations, including configurations of great order.⁹ Thus, the order we see in our universe, and the delicate adaptations we observe in organisms, in fact had a high probability of eventually coming into being, according to the hypothesis of chance. Van Inwagen (1993, p. 144) gives voice to this objection and explains it by way of an analogy: Suppose you toss a coin 20 times and it lands heads every time. You should not be surprised at this outcome if you are one among millions of people who toss a fair coin 20 times. After all, with so many people tossing, it is all but inevitable that some people should get 20 heads. The outcome you obtained, therefore, was not improbable, according to the chance hypothesis.

There is a fallacy in this criticism of the design argument, which Hacking (1987) calls “the inverse gambler’s fallacy.” He illustrates his idea by describing a gambler who walks into a casino and immediately observes two dice being rolled that land double-six. The gambler considers whether this result favors the hypothesis that the dice had been rolled many times before the roll he just observed or the hypothesis that this was the first roll of the evening. The gambler reasons that the outcome of double-six would be more probable under the first hypothesis:

$$\begin{aligned} &\Pr(\text{double-six on this roll} \mid \text{there were many rolls}) > \\ &\Pr(\text{double-six on this roll} \mid \text{there was just one roll}). \end{aligned}$$

In fact, the gambler’s assessment of the likelihoods is erroneous. Rolls of dice have the *Markov property*; the probability of double-six on this roll is the same ($\frac{1}{36}$), regardless of what may have happened in the past. What is true is that the probability that a double-six will occur *at some time or other* increases as the number of trials is increased:

$$\begin{aligned} &\Pr(\text{a double-six occurs sometime} \mid \text{there were many rolls}) > \\ &\Pr(\text{a double-six occurs sometime} \mid \text{there was just one roll}). \end{aligned}$$

However, the *principle of total evidence* says that we should assess hypotheses by considering *all* the evidence we have. This means that the relevant observation is that *this* roll landed double-six; we should not focus on the logically weaker proposition that a double-six occurred *sometime*. Relative to the stronger description of the observations, the hypotheses have identical likelihoods.

Applying this point to the criticism of the design argument that we are presently considering, we must conclude that the criticism is mistaken. It *is* highly probable (let us suppose), according to the chance hypothesis, that the universe will contain order and adaptation somewhere and at some time. However, the relevant observation is more specific – *our* corner of the universe is orderly and the organisms now on earth are well adapted. These events *do* have very low probability, according to the chance hypothesis, and the fact that a weaker description of the observations has high probability on the chance hypothesis is not relevant (see also White, 2000).¹⁰

If the first premise in the likelihood formulation of the design argument – that $\Pr(O \mid \text{Chance})$ is very low – is correct, then the only question that remains is whether $\Pr(O \mid \text{Design})$ is higher. This, I believe, is the Achilles heel of the design argument. The problem is to say how probable it is, for example, that the vertebrate eye would have features $F_1 \dots F_n$, if the eye were produced by an intelligent designer. What is required is not the specification of a single probability value, or even a precisely delimited range of values. All that is needed is an argument that shows that this probability is indeed higher than the probability that chance confers on the observation.

The problem is that the design hypothesis confers a probability on the observation only when it is supplemented with further assumptions about what the designer's goals and abilities would be if he existed. Perhaps the designer would never build the vertebrate eye with features $F_1 \dots F_n$, either because he would lack the goals or because he would lack the ability. If so, the likelihood of the design hypothesis is zero. On the other hand, perhaps the designer would want above all to build the eye with features $F_1 \dots F_n$ and would be entirely competent to bring this plan to fruition. If so, the likelihood of the design hypothesis is unity. There are as many likelihoods as there are suppositions concerning the goals and abilities of the putative designer. Which of these, or which class of these, should we take seriously?

It is no good answering this question by assuming that the eye was built by an intelligent designer and then inferring that he must have wanted to give the eye features $F_1 \dots F_n$ and that he must have had the ability to do so since, after all, these are the features we observe. For one thing, this pattern of argument is question begging. One needs *independent* evidence as to what the designer's plans and abilities would be if he existed; one can't obtain this evidence by *assuming* that the design hypothesis is true (Sober, 1999). Furthermore, even if we assume that the eye was built by an intelligent designer, we can't tell from this what the probability is that the eye would have the features we observe. Designers sometimes bring about outcomes that are not very probable given the plans they have in mind.

This objection to the design argument is an old one; it was presented by Keynes (1921) and before him by Venn (1866). In fact, the basic idea was formulated by Hume. When we behold the watch on the heath, we know that the watch's features are not particularly improbable on the hypothesis that the watch was produced by a designer who has the sorts of *human* goals and abilities with which we are familiar. This is the deep disanalogy between the watchmaker and the putative maker of organisms and universes. We are invited, in the latter case, to imagine a designer who is radically different from the human craftsmen with whom we are familiar. But if this designer is so different, why are we so sure that he would build the vertebrate eye in the form in which we find it?

This challenge is not turned back by pointing out that we often infer the existence of intelligent designers when we have no clue as to what they were trying to achieve. The biologist John Maynard Smith tells the story of a job he had during World War II inspecting a warehouse filled with German war material. He and his co-workers often came across machines whose functions were entirely opaque to them. Yet, they had no trouble seeing that these objects were built by intelligent designers. Similar stories can be told about archaeologists who work in museums; they often have objects in their collections that they know are artifacts, although they have no idea what the makers of these artifacts had in mind.

My claim is not that design theorists must have independent evidence that singles out a specification of the exact goals and abilities of the putative intelligent

designer. They may be uncertain as to which of the goal-ability pairs GA-1, GA-2, . . . , GA-n is correct. However, since

$$\Pr(\text{the eye has } F_1 \dots F_n \mid \text{Design}) = \sum_i \Pr(\text{the eye has } F_1 \dots F_n \mid \text{Design} \ \& \ \text{GA-i})\Pr(\text{GA-i} \mid \text{Design}),$$

they do have to show that

$$\sum_i \Pr(\text{the eye has } F_1 \dots F_n \mid \text{Design} \ \& \ \text{GA-i})\Pr(\text{GA-i} \mid \text{Design}) > \Pr(\text{the eye has } F_1 \dots F_n \mid \text{Chance}).$$

I think that Maynard Smith in his warehouse and archaeologists in their museums are able to do this. They aren't sure exactly what the intelligent designer was trying to achieve (e.g., they aren't certain that GA-1 is true and that all the other GA pairs are false), but they are able to see that it is not terribly improbable that the object should have the features one observes if it were made by a human intelligent designer. After all, the items in Maynard Smith's warehouse were symmetrical and smooth metal containers that had what appeared to be switches, dials, and gauges on them. And the "artifacts of unknown function" in anthropology museums likewise bear signs of human handiwork.

It is interesting in this connection to consider the epistemological problem of how one would go about detecting intelligent life elsewhere in the universe (if it exists). The Search for Extraterrestrial Intelligence (SETI) project, funded until 1993 by the US National Aeronautics and Space Administration and now supported privately, dealt with this problem in two ways (Dick, 1996). First, the scientists wanted to send a message into deep space that would allow any intelligent extraterrestrials who received it to figure out that it was produced by intelligent designers (namely, us). Second, they scan the night sky hoping to detect signs of intelligent life elsewhere.

The message, transmitted in 1974 from the Arecibo Observatory, was a simple picture of our solar system, a representation of oxygen and carbon, a picture of a double helix representing DNA, a stick figure of a human being, and a picture of the Arecibo telescope. How sure are we that if intelligent aliens find these clues, they will realize that they were produced by intelligent designers? The hope is that this message will strike the aliens who receive it as evidence favoring the hypothesis of intelligent design over the hypothesis that some mindless physical process (not necessarily one involving uniform chance) was responsible. It is hard to see how the SETI engineers could have done any better, but still one cannot dismiss the possibility that they will fail. If extraterrestrial minds are very different from our own – either because they have different beliefs and desires or because they process information in different ways – it may turn out that their interpretation of the evidence will differ profoundly from the interpretation that human beings would arrive at, were they on the receiving end. To say anything more precise about this, we'd have to be able to provide specifics about the aliens'

mental characteristics. If we are uncertain as to how the mind of an extraterrestrial will interpret this evidence, how can we be so sure that God, if he were to build the vertebrate eye, would endow it with the features we find it to have?

When SETI engineers search for signs of intelligent life elsewhere in the universe, what are they looking for? The answer is surprisingly simple. They look for narrow-band radio emissions. This is because human beings have built machines that produce these signals and, as far as we know, such emissions are not produced by mindless natural processes. The SETI engineers search for this signal, not because it is “complex” or fulfills some a priori criterion that would make it a “sign of intelligence,” but simply because they think they know what sorts of mechanisms are needed to produce it.¹¹ This strategy may not work, but it is hard to see how the scientists could do any better. Our judgments about what counts as a sign of intelligent design must be based on empirical information about what designers often do and what they rarely do. As of now, these judgments are based on our knowledge of *human* intelligence. The more our hypotheses about intelligent designers depart from the human case, the more in the dark we are as to what the ground rules are for inferring intelligent design. It is imaginable that these limitations will subside as human beings learn more about the cosmos. But for now, we are rather limited.

I have been emphasizing the fallibility of two assumptions – that we know what counts as a sign of extraterrestrial intelligence and that we know how extraterrestrials will interpret the signals we send. My point has been to shake a complacent assumption that figures in the design argument. However, I suspect that SETI engineers are on much firmer ground than theologians. If extraterrestrials evolved by the same type of evolutionary process that produced human intelligence, that may provide useful constraints on conjectures about the minds they have. No theologian, to my knowledge, thinks that God is the result of biological processes. Indeed God is usually thought of as a *supernatural* being who is radically different from the things we observe *in* nature. The problem of extraterrestrial intelligence is therefore an intermediate case; it lies between the watch found on the heath and the God who purportedly built the universe and shaped the vertebrate eye, but is much closer to the first. The upshot of this point for Paley’s design argument is this: *design arguments for the existence of human (and human-like) watchmakers are often unproblematic; it is design arguments for the existence of God that leave us at sea.*

I began by formulating the design hypothesis in argument (E) as the claim that an intelligent designer made the vertebrate eye. Yet, I have sometimes discussed the hypothesis as if it asserted that *God* is the designer in question. I don’t think this distinction makes a difference with respect to the objection I have described. To say that some designer or other made the eye is to state a disjunctive hypothesis. To figure out the likelihood of this disjunction, one needs to address the question of what each putative designer’s goals and intentions would be.¹² The theological formulation shifts the problem from the evaluation of a disjunction to the evaluation of a disjunct, but the problem remains the same. Even supposing

that God is omniscient, omnipotent, and perfectly benevolent, what is the probability that the eye would have features $F_1 \dots F_n$, if God set his hand to making it? He *could* have produced those results if he had wanted. But why think that this is what he *would* have wanted to do? The assumption that God can do anything is part of the problem, not the solution. An engineer who is more limited would be more predictable.

There is another reply to my criticism of the design argument that should be considered. I have complained that we have no way to evaluate the likelihood of the design hypothesis, since we don't know which auxiliary assumptions about goal-ability pairs we should use. But why not change the subject? Instead of evaluating the likelihood of design, why not evaluate the likelihood of various conjunctions – (Design & GA-1), (Design & GA-2), etc.? Some of these will have high likelihoods, others will have low, but it will no longer be a mystery what likelihoods these hypotheses possess. There are two problems with this tactic. First, it is a game that two can play. Consider the hypothesis that the vertebrate eye was created by the mindless process of electricity. If I simply get to *invent* auxiliary hypotheses without having to *justify* them independently, I can just stipulate the following assumption: if electricity created the vertebrate eye, the eye must have features $F_1 \dots F_n$. The electricity hypothesis now is a conjunct in a conjunction that has maximum likelihood, just like the design hypothesis. This is a dead end. My second objection is that it is an important part of scientific practice that conjunctions be broken apart (when possible), and their conjuncts scrutinized (Sober, 1999, 2000). If your doctor runs a test to see whether you have tuberculosis, you will not be satisfied if she reports that the likelihood of the conjunction “you have tuberculosis & auxiliary assumption 1” is high while the likelihood of the conjunction “you have tuberculosis & auxiliary assumption 2” is low. You want your doctor to address the first *conjunct*, not just the various *conjunctions*. And you want her to do this by using a test procedure that is *independently* known to have small error probabilities. Demand no less of your theologian.

My formulation of the design argument as a likelihood inference, and my criticism of it, have implications concerning the problem of evil. It is a mistake to try to *deduce* the non-existence of God from the fact that so much evil exists. Even supposing that God is all-powerful, all-knowing, and entirely benevolent, there is no contradiction in the hypothesis that God allows various evils to exist because they are necessary correlates of greater goods, where we don't understand in any detail what these correlations are or why they must obtain (Plantinga, 1974). The status of the problem changes, however, when we think of it as *non-deductive* in character (Madden and Hare, 1968; Rowe, 1979; Plantinga, 1979). Within the framework of likelihood inference, there are two quantities we must evaluate: What is the probability that there would be as much evil as there is, if the universe were produced by an all-powerful, all-knowing, and entirely benevolent God? And what is the probability of that much evil's existing, if the universe were produced by mindless natural processes? The logical observation that saves

theism from the attempt to deduce the non-existence of God comes back to haunt the theistic hypothesis in this new context. If the ways of God are so mysterious, we have no way to evaluate the first of these likelihoods. The theistic hypothesis is saved from disconfirmation by the fact that it is untestable.

The Relationship of the Organismic Design Argument to Darwinism

Philosophers who criticize the organismic design argument often believe that the argument was dealt its death blow by Hume. True, Paley wrote after Hume, and the many Bridgewater treatises elaborating the design argument appeared after Hume's *Dialogues* were published posthumously. Nonetheless, for these philosophers, the design argument after Hume was merely a corpse that could be propped up and paraded. Hume had taken the life out of it.

Biologists often take a different view. Dawkins (1986, p. 4) puts the point provocatively by saying that it was not until Darwin that it was possible to be an intellectually fulfilled atheist. The thought here is that Hume's skeptical attack was not the decisive moment; rather, it was Darwin's development and confirmation of a substantive scientific explanation of the adaptive features of organisms that really undermined the design argument (at least in its organismic formulation). Philosophers who believe that theories can't be rejected until a better theory is developed to take its place often sympathize with this point of view.

My own interpretation coincides with neither of these. As indicated above, I think that Hume's criticisms largely derive from an empiricist epistemology that is too narrow. However, seeing the design argument's fatal flaw does not depend on seeing the merits of Darwinian theory. The likelihood principle, it is true, says that theories must be evaluated comparatively, not on their own. But for this to be possible, each theory must make predictions. It is at this fundamental level that I think the design argument is defective.

Biologists often present two criticisms of creationism. First, they argue that the design hypothesis is untestable. Second, they contend that there is plenty of evidence that the hypothesis is false. Obviously, these two lines of argument are in conflict.¹³ I have already endorsed the first criticism, but I want to say a little about the second. A useful example is Stephen Jay Gould's (1980) widely read article about the panda's thumb. Pandas are vegetarian bears who have a spur of bone (a "thumb") protruding from their wrists. They use this device to strip bamboo, which is the main thing they eat. Gould says that the hypothesis of intelligent design predicts that pandas should *not* have this inefficient device. A benevolent, powerful, and intelligent engineer could and would have done a lot better. Evolutionary theory, on the other hand, says that the panda's thumb is what we should expect. The thumb is a modification of the wrist bones found in the common ancestor that pandas share with carnivorous bears. Evolution by

natural selection is a tinkerer; it does not design adaptations from scratch, but modifies pre-existing features, with the result that adaptations are often imperfect.

Gould's argument, I hope it is clear, is a likelihood argument. I agree with what he says about evolutionary theory, but I think his discussion of the design hypothesis leads him into the same trap that ensnared Paley. Gould thinks he knows what God would do if he built pandas, just as Paley thought he knew what God would do if he built the vertebrate eye. But neither of them knows this. Both help themselves to *assumptions* about God's goals and abilities. However, it is not enough to make assumptions about these matters; one needs independent evidence that these auxiliary assumptions are true. Paley's problem is also Gould's.

Anthropic Reasoning and Cosmic Design Arguments

Evolutionary theory seeks to explain the adaptive features of organisms; it has nothing to say about the origin of the universe as a whole. For this reason, evolutionary theory conflicts with the organismic design hypothesis, but not with the cosmic design hypothesis. Still, the main criticism I presented of the first type of design argument also applies to the second. I now want to examine a further problem that cosmic design arguments sometimes encounter.¹⁴

Suppose I catch 50 fish from a lake, and you want to use my observations O to test two hypotheses:

- O: All the fish I caught were more than 10 inches long.
- F1: All the fish in the lake are more than 10 inches long.
- F2: Only half the fish in the lake are more than 10 inches long.

You might think that the likelihood principle says that F1 is better supported, since

$$(1) \Pr(O | F1) > \Pr(O | F2).$$

However, you then discover how I caught my fish:

- (A1) I caught the fish by using a net that (because of the size of its holes) can't catch fish smaller than 10 inches, and I left the net in the lake until there were 50 fish in it.

This leads you to replace the analysis provided by (1) with the following:

$$(2) \Pr(O | F1 \& A1) = \Pr(O | F2 \& A1) = 1.0.$$

Furthermore, you now realize that your first assessment, (1), was based on the erroneous assumption that

(A0) The fish I caught were a random sample from the fish in the lake.

Instead of (1), you should have written

$$\Pr(O \mid F1 \ \& \ A0) > \Pr(O \mid F2 \ \& \ A0).$$

This inequality is true; the problem, however, is that (A0) is false.

This example, from Eddington (1939), illustrates the idea of an *observational selection effect* (an OSE). When a hypothesis is said to render a set of observations probable (or improbable), ask yourself what assumptions allow the hypothesis to have this implication. The point illustrated here is that the procedure you use to obtain your observations can be relevant to assessing likelihoods.¹⁵

One version of the cosmic design argument begins with the observation that our universe is “fine-tuned.” That is, the values of various physical constants are such as to permit life to exist, but if they had been even slightly different, life would have been impossible. McMullin (1993, p. 378) summarizes some of the relevant facts as follows:

If the strong nuclear force were to have been as little as 2 percent stronger (relative to the other forces), all hydrogen would have been converted into helium. If it were 5 percent weaker, no helium at all would have formed and there would be nothing but hydrogen. If the weak nuclear force were a little stronger, supernovas could not occur, and heavy elements could not have formed. If it were slightly weaker, only helium might have formed. If the electromagnetic force were stronger, all stars would be red dwarfs, and there would be no planets. If it were a little weaker, all stars would be very hot and short-lived. If the electron charge were ever so slightly different, there would be no chemistry as we know it. Carbon (¹²C) only just managed to form in the primal nucleosynthesis. And so on.

I’ll abbreviate the fact that the values of these physical constants fall within the narrow limits specified by saying that “the constants are right.” A design argument can now be constructed, one that claims that the constants’ being right should be explained by postulating the existence of an intelligent designer, one who wanted life to exist and who arranged the universe so that this could occur (Swinburne, 1990a). As with Paley’s organismic design argument, we can represent the reasoning in this cosmic design argument as the assertion of a likelihood inequality:

$$(3) \Pr(\text{constants are right} \mid \text{Design}) > \Pr(\text{constants are right} \mid \text{Chance}).$$

However, there is a problem with (3) that resembles the problem with (1). Consider the fact that

(A3) We exist, and if we exist the constants must be right.

We need to take (A3) into account; instead of (3), we should have said:

$$(4) \quad \Pr(\text{constants are right} \mid \text{Design \& A3}) = \Pr(\text{constants are right} \mid \text{Chance \& A3}) \\ = 1.0.$$

That is, given (A3), the constants must be right, regardless of whether the universe was produced by intelligent design or by chance.

Proposition (4) reflects the fact that our observation that the constants are right is subject to an OSE. Recognizing this OSE is in accordance with a *weak anthropic principle* – “what we can expect to observe must be restricted by the conditions necessary for our presence as observers” (Carter, 1974). The argument involves no commitment to *strong anthropic principles*. For example, there is no assertion that the correct cosmology must entail that the existence of observers such as ourselves was inevitable, nor is it claimed that our existence *explains* why the physical constants are right (Barrow, 1988; Earman, 1987; McMullin, 1993).¹⁶

Although this point about OSEs undermines the version of the design argument that cites the fact that the physical constants are right, it does not touch other versions. For example, when Paley concludes that the vertebrate eye was produced by an intelligent designer, his argument cannot be refuted by claiming that:

$$(A4) \quad \text{We exist, and if we exist vertebrates must have eyes with features } F1 \dots Fn.$$

If (A4) were true, the likelihood inequality that Paley asserted would have to be replaced with an equality, just as (1) had to be replaced by (2) and (3) had to be replaced by (4). But fortunately for Paley, (A4) is false. However, matters change if we think of Paley as seeking to explain the modest fact that organisms have at least one adaptive contrivance. If this were false, we would not be able to make observations; indeed, we would not exist. Paley was right to focus on the details; the more minimal description of what we observe does not sustain the argument he wanted to endorse.

The issue of OSEs can be raised in connection with other cosmic versions of the design argument. Swinburne (1990b, p. 191) writes that “the hypothesis of theism is that the universe exists because there is a God who keeps it in being and that laws of nature operate because there is a God who brings it about that they do.” Let us separate the *explananda*. The fact that the universe exists does *not* favor design over chance; after all, if the universe did not exist, we would not exist and so would not be able to observe that it does.¹⁷ The same point holds with respect to the fact that the universe is law-governed. Even supposing that lawlessness is possible, could we exist and make observations if there were no laws? If not, then the lawful character of the universe does not discriminate between design and chance. Finally, we may consider the fact that our universe is governed by one set of laws, rather than another. Swinburne (1968) argues that the fact that our universe obeys *simple* laws is better explained by the hypothesis

of design than by the hypothesis of chance. Whether this observation also is subject to an OSE depends on whether we could exist in a universe obeying alternative laws.

Before taking up an objection to this analysis of the argument from fine-tuning, I want to summarize what it has in common with the fishing example. In the fishing example, the source of the OSE is obvious – it is located in a device outside of ourselves. The net with big holes insures that the observer will make a certain observation, regardless of which of two hypotheses is true. But where is the device that induces an OSE in the fine-tuning example? There is none; rather, it is the observer's own existence that does the work. Nonetheless, the effect is the same. Owing to the fact that we exist, we are bound to observe that the constants are right, regardless of whether our universe was produced by chance or by design.¹⁸

This structural similarity between fishing and fine-tuning may seem to be undermined by a disanalogy. In the latter case, we know that proposition (3) is correct – the probability that the constants will be right if the universe is created by a powerful deity bent on having life exist is greater than it would be if the values of the constants were set by a uniform chance process. This inequality seems to hold, regardless of how or whether we make our observations. The fishing example looks different; here we know that proposition (1) is false. There is no saying whether a likelihood inequality obtains until we specify the procedure used to obtain the observations; once we do this, there *is* no likelihood inequality. Thus, in fine-tuning, we have an inequality that is true because it reflects the metaphysical facts; in fishing, we have an inequality that is false for epistemic reasons. My response is that I agree that this point of difference exists, but that it does nothing to save the argument from fine-tuning. Although proposition (3) is true, we are bound to observe that the constants are right, regardless of whether our universe arose by chance or by design. My objection to proposition (3) is not that it is false, but that it should not be used to interpret the observations; (4) is the relevant proposition to which we should attend.

To visualize this point, imagine that a deity creates a million universes and that a chance process does the same for another million. Let's assume that the proportion of universes in which the constants are right is greater in the former case. Doesn't it follow that if we observe that the constants are right in our universe, this observation favors the hypothesis that our universe arose by design? In fact, this does not follow. It *would* follow if we had the same probability of observing any of the first million universes if the design hypothesis were true, and had the same probability of observing any of the second million universes if the chance hypothesis were true. But this is not the case – our probability of observing a universe in which the constants are right is unity in each case.

What this means is that a full understanding of the workings of OSEs must acknowledge that there are two stages at which a bias can be introduced. There is first the process by which the system described by the hypotheses under test generates some state of the world that we are able to observe. Second, there is the

process by which we come to observe that state of the world. This two-step process occurs in fishing and fine-tuning as follows:

Composition of the lake → Contents of the net → We observe the contents of the net.

Origin of the universe → Constants are right → We observe that the constants are right.

The OSE in the fishing example arises in the first step; the OSE in fine-tuning crops up in the second step.

Leslie (1989, pp. 13–14, 107–8), Swinburne (1990a, p. 171), and van Inwagen (1993, pp. 135, 144) all defend the fine-tuning argument against the criticism I have just described. Each mounts his defense by describing an analogy with a mundane example. Here is Swinburne’s rendition of an example that Leslie presents:

On a certain occasion the firing squad aim their rifles at the prisoner to be executed. There are twelve expert marksmen in the firing squad, and they fire twelve rounds each. However, on this occasion all 144 shots miss. The prisoner laughs and comments that the event is not something requiring any explanation because if the marksmen had not missed, he would not be here to observe them having done so. But of course, the prisoner’s comment is absurd; the marksmen all having missed is indeed something requiring explanation; and so too is what goes with it – the prisoner’s being alive to observe it. And the explanation will be either that it was an accident (a most unusual chance event) or that it was planned (e.g., all the marksmen had been bribed to miss). Any interpretation of the anthropic principle which suggests that the evolution of observers is something which requires no explanation in terms of boundary conditions and laws being a certain way (either inexplicably or through choice) is false.

First a preliminary clarification – the issue isn’t whether the prisoner’s survival “requires explanation” but whether this observation provides evidence as to whether the marksmen intended to spare the prisoner or shot at random.¹⁹

My response takes the form of a dilemma. I’ll argue, first, that if the firing squad example is analyzed in terms of the likelihood principle, the prisoner is right and Swinburne is wrong – the prisoner’s survival does not allow him to conclude that design is more likely than chance. However, there is a different analysis of the prisoner’s situation, in terms of the *probabilities* of hypotheses, not their *likelihoods*. This second analysis says that the prisoner *is* mistaken; however, it has the consequence that the prisoner’s inference differs fundamentally from the design argument that appeals to fine-tuning. Each horn of this dilemma supports the conclusion that the firing squad example does nothing to save this version of the design argument.

So let us begin. If we understand Swinburne’s claim in terms of the likelihood principle, we should read him as saying that

- (L1) $\Pr(\text{the prisoner survived} \mid \text{the marksmen intended to miss}) >$
 $\Pr(\text{the prisoner survived} \mid \text{the marksmen fired at random}).$

He thinks that the anthropic principle requires us to replace this claim with the following irrelevancy:

- (L2) $\Pr(\text{the prisoner survived} \mid \text{the marksmen intended to miss} \ \& \ \text{the prisoner survived}) = \Pr(\text{the prisoner survived} \mid \text{the marksmen fired at random} \ \& \ \text{the prisoner survived}) = 1.0.$

This equality would lead us to conclude (Swinburne thinks mistakenly) that the prisoner's survival does not discriminate between the hypotheses of design and chance.

To assess the claim that the prisoner has made a mistake, it is useful to compare the prisoner's reasoning with that of a bystander who witnesses the prisoner survive the firing squad. The prisoner reasons as follows: "given that I now am able to make observations, I must be alive, whether my survival was due to intelligent design or chance." The bystander says the following: "given that I now am able to make observations, the fact that the prisoner is now alive is made more probable by the design hypothesis than it is by the chance hypothesis." The prisoner is claiming that he is subject to an OSE, while the bystander says that he, the bystander, is not. Both, I submit, are correct.²⁰

I suggest that part of the intuitive attractiveness of the claim that the prisoner has made a mistake derives from a shift between the prisoner's point of view and the bystander's. The bystander is right to use (L1) to interpret his observations; however, the prisoner has no business using (L1) to interpret his observations since he, the prisoner, is subject to an OSE. The prisoner needs to replace (L1) with (L2). My hunch is that Swinburne thinks the prisoner errs in his assessment of likelihoods because we bystanders would be making a mistake if we reasoned as he does.²¹

The basic idea of an OSE is that we must take account of the procedures used to obtain the observations when we assess the likelihoods of hypotheses. This much was clear from the fishing example. What may seem strange about my reading of the firing squad story is my claim that the prisoner and the bystander are in different epistemic situations, even though their observation reports differ by a mere pronoun. After the marksmen fire, the prisoner thinks "I exist" while the bystander thinks "he exists;" the bystander, but not the prisoner, is able to use his observation to say that design is more likely than chance, or so I say. If this seems odd, it may be useful to reflect on Sorensen's (1988) concept of *blindspots*. A proposition p is a blindspot for an individual S just in case, if p were true, S would not be able to know that p is true. Although some propositions (e.g., "nothing exists," "the constants are wrong") are blindspots for everyone, other propositions are blindspots for some people but not for others. Blindspots give rise to OSEs; if p is a blindspot for S , then if S makes an observation to

determine the truth value of p , the outcome must be that not- p is observed. The prisoner, but not the bystander, has “the prisoner does not exist” as a blindspot. This is why “the prisoner exists” has an evidential significance for the bystander that it cannot have for the prisoner.²²

To bolster my claim that the prisoner is right to think that likelihood does not distinguish between chance and design, I want to describe a slightly different problem. Suppose that a firing squad always subjects its victims to the same probabilistic process, which has the result that the prisoner either survives or is killed. 1,000 prisoners who have one by one each survived the firing squad are assembled and are asked to pool their knowledge and estimate the value of an unknown probability. What is the probability that a prisoner will survive if the firing squad fires? The standard methodology here is *maximum likelihood estimation*; one finds the value of the parameter of interest that maximizes the probability of the observations. This is why, if a coin lands heads 512 out of 1,000 tosses, the “best” estimate of the probability that the coin will land heads when it is tossed is 0.512. Those who believe that the single prisoner has evidence about his firing squad’s intentions are obliged to conclude that the best estimate in this new problem is that the probability is unity. However, those persuaded that the single prisoner is subject to an OSE will want to maintain that the 1,000 prisoners are in the same boat. These skeptics will deny that the observations provide a basis for estimation. Isn’t it *obvious* that testimony limited to survivors provides no evidence on which to base an estimate of the probability that someone will survive the firing squad’s shooting? And if this is true of a 1,000 survivors, how can a *single* survivor be said to know that design is more likely than chance?

I now turn to a different analysis of the prisoner’s situation. The prisoner, like the rest of us, knows how firing squads work. They always or almost always follow the order they receive, which is almost always to execute someone. Occasionally, they produce fake executions. They almost never fire at random. What is more, firing squads have firm control over outcomes; if they want to kill (or spare) someone, they always or almost always succeed. This and related items of background knowledge support the following *probability* claim:

$$(Pf) \quad \Pr(\text{the marksmen intended to spare the prisoner} \mid \text{the prisoner survived}) > \Pr(\text{the marksmen intended to spare the prisoner}).$$

Firing squads rarely intend to spare their victims, but the survival of the prisoner makes it very probable that his firing squad had precisely that intention. The likelihood analysis led to the conclusion that the prisoner and the bystander are in different epistemic situations; the bystander should evaluate the hypotheses by using (L1), but the prisoner is obliged to use (L2). However, from the point of view of probabilities, the prisoner and the bystander can say the same thing; both can cite (Pf).²³

What does this tell us about the fine-tuning version of the design argument? I construed that argument as a claim about likelihoods. As such, it is subject to an

OSE; given that we exist, the constants must be right, regardless of whether our universe was produced by chance or by design. However, we now need to consider whether the fine-tuning argument can be formulated as a claim about probabilities. Can we assert that

- (Pu) $\Pr(\text{the universe was created by an intelligent designer} \mid \text{the constants are right})$
 $> \Pr(\text{the universe was created by an intelligent designer})?$

I don't think so. In the case of firing squads, we have frequency data and our general knowledge of human behavior on which to ground the probability statement (Pf). But we have neither data nor theory to ground (Pu). And we cannot defend (Pu) by saying that an intelligent designer would ensure that the constants are right, because this takes us back to the likelihood considerations we have already discussed. The prisoner's conclusion that he can say nothing about chance and design *is* mistaken if he is making a claim about probabilities. But the argument from fine-tuning can't be defended as a claim about probabilities.

The rabbit/duck quality of this problem merits review. I have discussed three examples – fishing, fine-tuning, and the firing squad. If we compare fine-tuning with fishing, they seem similar. This makes it intuitive to conclude that the design argument based on fine-tuning is wrong. However, if we compare fine-tuning with the firing squad, *they* seem similar. Since the prisoner apparently has evidence that favors design over chance, we are led to the conclusion that the fine-tuning argument must be right. This shifting gestalt can be stabilized by imposing a formalism. The first point is that OSEs are to be understood by comparing the *likelihoods* of hypotheses, not their *probabilities*. The second is that it is perfectly true that the prisoner can assert the *probability* claim (Pf). The question, then, is whether the design argument from fine-tuning is a likelihood argument or a probability argument. If the former, it is flawed because it fails to take account of the fact that there is an OSE. If the latter, it is flawed, but for a different reason – it makes claims about probabilities that we have no reason to accept; indeed, we cannot even *understand* them as objective claims.²⁴

A Prediction

It was obvious to Paley and to other purveyors of the organismic design argument that if an intelligent designer built organisms, that designer would have to be far more intelligent than any human being could ever be. This is why the organismic design argument was for them an argument for the existence of *God*. I predict that it will eventually become clear that the organismic design argument should never have been understood in this way. This is because I expect that human beings will eventually build organisms from non-living materials. This achievement will not close down the question of whether the organisms we

observe were created by intelligent design or by mindless natural processes; on the contrary, it will give that question a practical meaning, since the organisms we will see around us will be of both kinds.²⁵ However, it will be abundantly clear that the fact of organismic adaptation has nothing to do with whether God exists. When the Spanish conquistadors arrived in the New World, several indigenous peoples thought these intruders were gods, so powerful was the technology that the intruders possessed. Alas, the locals were mistaken; they did not realize that these beings with guns and horses were merely *human* beings. The organismic design argument for the existence of God embodies the same mistake. Human beings in the future will be the conquistadors, and Paley will be our Montezuma.

Notes

I am grateful to Martin Barrett, Nick Bostrom, David Christensen, Ellery Eells, Branden Fitelson, Malcolm Forster, Alan Hajek, Daniel Hausman, Stephen Leeds, William Mann, Lydia McGrew, Derk Pereboom, Roy Sorensen, and Richard Swinburne for useful comments. I have used portions of this chapter in seminars and lectures that I have given at many colleges and universities, too numerous to list here. My thanks to participants for their stimulating and productive discussion.

- 1 Does this construal of the design argument conflict with the idea that the argument is an *inference to the best explanation*? Not if one's theory of inference to the best explanation says that observations influence the assessment of explanations in this instance via the vehicle of likelihoods.
- 2 Another reason to restrict the design argument to likelihood considerations is that it is supposed to be an *empirical* argument. To invoke prior probabilities is to bring in considerations *besides* the observations at hand.
- 3 In light of the fact that it is possible for a hypothesis to have an objective likelihood without also having an objective probability, one should understand Bayes's theorem as specifying how the quantities it mentions are related to each other, *if all are well defined*. And just as hypotheses can have likelihoods without having (objective) probabilities, it also is possible for the reverse situation to obtain. Suppose I draw a card from a deck of unknown composition. I observe (O) that the card is the four of diamonds. I now consider the hypothesis (H) that the card is a four. The value of $\Pr(H \mid O)$ is well defined, but the value of $\Pr(O \mid H)$ is not.
- 4 Actually, Paley (1802) *does* consider a "selective retention" process, but only very briefly. In Chapter 5 (pp. 49–51) he explores the hypothesis that a random process once generated a huge range of variation, and that this variation was then culled, with only stable configurations surviving. Paley argues against this hypothesis by saying that we should see unicorns and mermaids if it were true. He also says that it mistakenly predicts that organisms should fail to form a taxonomic hierarchy. It is ironic that Darwin claimed that his own theory *predicts* hierarchy. In fact, Paley and Darwin are both right. Darwin's theory includes the idea that all living things have common ancestors, while the selection hypothesis that Paley considers does not.
- 5 More precisely, Fisher said that a mother should have a son with probability p and a daughter with probability $(1 - p)$, where the effect of this is that the expected

expenditures on the two sexes are the same; the argument is not undermined by the fact that some mothers have all sons while others have all daughters.

- 6 Dawkins (1986) makes the point that evolution by natural selection is not a uniform chance process by way of an analogy with a combination lock. This is discussed in Sober (1993, pp. 36–9).
- 7 Dembski (1998) construes design inference as allowing one to argue in favor of the design hypothesis, and “sweep from the field” all alternatives, without the design hypothesis ever having to make a prediction. For criticisms of Dembski’s framework, see Fitelson, Stephens, and Sober (1999).
- 8 Paley (1802) argues in Chapter 16 that the benevolence of the deity is demonstrated by the fact that organisms experience more pleasure than they need to (p. 295). He also argues that pain is useful (p. 320) and that few diseases are fatal; he defends the latter conclusion by citing statistics on the cure rate at a London hospital (p. 321).
- 9 For it to be certain that all configurations will be visited, there must be infinite time. The shorter the time-frame, the lower the probability that a given configuration will occur. This means that the estimated age of the universe may entail that it is very *improbable* that a given configuration will occur. I set this objection aside in what follows.
- 10 It is a standard feature of likelihood comparisons that O_w sometimes fails to discriminate between a pair of hypotheses, even though O_s is able to do so, when O_s entails O_w . You are the cook in a restaurant. The waiter brings an order into the kitchen – someone ordered bacon and eggs. You wonder whether this information discriminates between the hypothesis that your friend Smith ordered the meal and the hypothesis that your friend Jones did. You know the eating habits of each. Here’s the probability of the order’s being for \pm bacon and \pm eggs, conditional on the order’s coming from Smith and conditional on the order’s coming from Jones:


Pr(– Smith)			Pr(– Jones)		
	Eggs			Eggs	
	+	–		+	–
+	0.4	0.1	+	0.1	0.4
Bacon			Bacon		
–	0.2	0.3	–	0.5	0

The fact that the customer ordered bacon does not discriminate between the two hypotheses (since $0.5 = 0.5$). And the fact that the customer ordered eggs doesn’t help either (since $0.6 = 0.6$). However, the fact that the customer ordered bacon *and* eggs favors Smith over Jones (since $0.4 > 0.1$).

- 11 The example of the SETI project throws light on Paley’s question as to why we think that watches must be the result of intelligent design, but don’t think this when we observe a stone. It is tempting to answer this question by saying that watches are “complicated” while stones are not. However, there are many complicated natural processes (like the turbulent flow of water coming from a faucet) that don’t cry out for explanation in terms of intelligent design. Similarly, narrow-band radio emissions

may be physically “simple” but that doesn’t mean that the SETI engineers were wrong to search for them.

- 12 Assessing the likelihood of a disjunction involves an additional problem. Even if the values of $\Pr(O \mid D1)$ and $\Pr(O \mid D2)$ are known, what is the value of $\Pr(O \mid D1 \text{ or } D2)$? The answer is that it must be somewhere in between. But exactly where depends on further considerations, since $\Pr(O \mid D1 \text{ or } D2) = \Pr(O \mid D1)\Pr(D1 \mid D1 \text{ or } D2) + \Pr(O \mid D2)\Pr(D2 \mid D1 \text{ or } D2)$. If either God or a superintelligent extraterrestrial built the vertebrate eye, what is the probability that it was God who did so?
- 13 The statement “ p is both false and untestable” is logically consistent (assuming that the verificationist theory of meaning is mistaken). However, the *assertion* of this conjunction is paradoxical, akin to Moore’s paradoxical statement “ p is true but I don’t believe it.” Both conjunctions embody pragmatic, not semantic, paradoxes.
- 14 To isolate this new problem from the one already identified, I’ll assume in what follows that the design hypothesis and the chance hypothesis with which it competes have built into them auxiliary assumptions that suffice for their likelihoods to be well defined.
- 15 This general point surfaces in simple inference problems like the ravens paradox (Hempel, 1965). Does the fact that the object before you is a black raven confirm the generalization that all ravens are black? That depends on how you gathered your data. Perhaps you sampled at random from the set of *ravens*; alternatively, you may have sampled at random from the set of *black ravens*. In the first case, your observation confirms the generalization, but in the second it does not. In the second case, notice that you were bound to observe that the object before you is a black raven, regardless of whether all ravens are black.
- 16 Although weak and strong anthropic principles differ, they have something in common. For example, the causal structure implicitly assumed in the weak anthropic principle is that of two effects of a common cause:



In contrast, one of the strong anthropic principles assumes the following causal arrangement:

$$\text{(SAP) We exist now} \rightarrow \text{Origin of the universe} \rightarrow \text{Constants now are right.}$$

Even though (WAP) is true and (SAP) is false, both entail a *correlation* between our existence and the constants now having the values they do. To deal with the resulting OSEs, we must decide how to take these correlations into account in assessing likelihoods.

- 17 Similarly, the fact that there is something rather than nothing does not discriminate between chance and design.
- 18 The fishing and fine-tuning examples involve *extreme* OSEs. More modest OSEs are possible. If C describes the circumstances in which we make our observational determination as to whether proposition O is true, and we use the outcome of this determination to decide whether $H1$ or $H2$ is more likely, then a *quantitative* OSE is present precisely when

$$\begin{aligned} &\Pr(O \mid H1 \ \& \ C) \neq \Pr(O \mid H1) \text{ or} \\ &\Pr(O \mid H2 \ \& \ C) \neq \Pr(O \mid H2). \end{aligned}$$

A *qualitative* OSE occurs when taking account of C alters the likelihood ordering:

$$\begin{aligned} &\Pr(O \mid H1 \ \& \ C) > \Pr(O \mid H2 \ \& \ C) \text{ and } \Pr(O \mid H1) \not> \Pr(O \mid H2) \text{ or} \\ &\Pr(O \mid H1 \ \& \ C) = \Pr(O \mid H2 \ \& \ C) \text{ and } \Pr(O \mid H1) \neq \Pr(O \mid H2). \end{aligned}$$

Understood in this way, an OSE is just an example of *sampling bias*.

- 19 There is a third possibility – that the marksmen intended to kill the prisoner. But for the sake of simplicity (and also to make the firing squad argument more parallel with the argument from fine-tuning), I’ll ignore this possibility.
- 20 The issue, thus, is not whether (L1) or (L2) are true (both are), but which an agent should use in interpreting the bearing of observations on the likelihoods of hypotheses. In this respect the injunction of the weak anthropic principle is like the principle of total evidence – it is a pragmatic principle, concerning which statements should be used for which purposes.
- 21 In order to replicate in the fine-tuning argument the difference between the prisoner’s and the bystander’s points of view, imagine that we observe through a telescope another universe in which the constants are right. We bystanders can use this observation in a way that the inhabitants of that universe cannot.
- 22 Notice that “I exist” when thought by the prisoner, is a priori, whereas “the prisoner exists,” when thought by the bystander, is a posteriori. Is it so surprising that an a priori statement should have a different evidential significance than an a posteriori statement?

I also should note that my claim is that the proposition “I am alive” does not permit the prisoner to conclude that design is more likely than chance. I do not say that there is no proposition he can cite after the marksmen fire that discriminates between the two hypotheses. Consider, for example, the observation that “no bullets hit me.” This favors design over chance, even after the prisoner conditionalizes on the fact that he is alive. Notice also that if the prisoner were alive but riddled with bullets, it is not so clear that design would be more likely than chance.

- 23 I have argued that the prisoner should assign the same likelihoods to chance and design, but that he is entitled to think that his survival lowers the probability of chance and raises the probability of design. On its face, this contradicts the following consequence of Bayes’s theorem:

$$\frac{\Pr(\text{Chance} \mid \text{I survive}) \Pr(\text{I survive} \mid \text{Chance}) \Pr(\text{Chance})}{\Pr(\text{Design} \mid \text{I survive}) \Pr(\text{I survive} \mid \text{Design}) \Pr(\text{Design})} .$$

If the ratio of posterior probabilities is greater than the ratio of priors, this must be because the two likelihoods have different values.

The reason my argument implies no such contradiction is that I have argued, first, that the relevant likelihoods are *not* the ones displayed above, but are ones that take account of the presence of an OSE. I further imagined that the prisoner possesses knowledge (inferred from frequencies) that the two posterior probabilities displayed

above are, respectively, low and high. This inference might be called “direct” since it proceeds without the prisoner’s having to assign values to likelihoods. Bayes’s theorem describes how various quantities are related when each is well defined; it does not entail that all of them are well defined in every situation (Sober, 2002). It is a familiar point made by critics of Bayesianism that likelihoods can be well defined even when prior and posterior probabilities are not. This severing of the connection between likelihoods and probabilities, or something like it, arises in the firing squad problem. The prisoner can know that chance is improbable and that design is highly probable, given his observation after the firing squad fires that he exists, even though his evaluation of likelihoods should focus on likelihoods that are identical in value.

- 24 The hypothesis that our universe is one among many has been introduced as a possible explanation of the fact that the constants (in our universe) are right. A universe is here understood to be a region of spacetime that is causally closed. See Leslie (1989) for discussion. If the point of the multiverse hypothesis is to challenge the design hypothesis, on the assumption that the design hypothesis has already vanquished the hypothesis of chance, then the multiverse hypothesis is not needed. Furthermore, in comparing the multiverse hypothesis and the design hypothesis, one needs to attend to the inverse gambler’s fallacy discussed earlier. This is not to deny that there may be other evidence for the multiverse hypothesis; however, the mere fact that the constants are right in our universe does not favor that hypothesis.
- 25 As Dennett (1987a, pp. 284–5) observes, human beings have been modifying the characteristics of animals and plants by *artificial selection* for thousands of years. However, the organisms thus modified were not *created* by human beings. If the design argument endorses a hypothesis about how organisms were brought into being, then the work of plant and animal breeders, per se, does not show that the design argument should be stripped of its theological trappings.

References

- Arbuthnot, John (1710) An Argument for Divine Providence, Taken from the Constant Regularity Observ’d in the Births of Both Sexes. *Philosophical Transactions of the Royal Society of London*, 27, 186–90.
- Barrow, John (1988) *The World Within the World*. Oxford: Clarendon Press.
- Behe, Michael J. (1996) *Darwin’s Black Box*. New York: Free Press.
- Carter, Brandon (1974) Large Number Coincidences and the Anthropic Principle in Cosmology. In M. S. Longair (ed.), *Confrontation of Cosmological Theories with Observational Data* (pp. 291–8). Dordrecht: Reidel.
- Darwin, Charles (1859) *On the Origin of Species*. London: John Murray. Cambridge, Mass.: Harvard University Press, 1964.
- Dawkins, Richard (1986) *The Blind Watchmaker*. New York: Norton.
- Dembski, William A. (1998) *The Design Inference*. Cambridge: Cambridge University Press.
- Dennett, Daniel C. (1987a) Intentional Systems in Cognitive Ethology – the “Panglossian Paradigm” defended. In *The Intentional Stance* (pp. 237–86). Cambridge, Mass.: MIT Press.
- Dennett, Daniel C. (1987b) True Believers. In *The Intentional Stance* (pp. 13–42). Cambridge, Mass.: MIT Press.

- Dick, Stephen J. (1996) *The Biological Universe – the Twentieth-Century Extraterrestrial Life Debate and the Limits of Science*. Cambridge: Cambridge University Press.
- Earman, John (1987) The Sap Also Rises – a Critical Examination of the Anthropic Principle. *American Philosophical Quarterly*, 24, 307–17.
- Eddington, Arthur (1939) *The Philosophy of Physical Science*. Cambridge: Cambridge University Press.
- Edwards, A. W. F. (1972) *Likelihood*. Cambridge: Cambridge University Press.
- Fisher, R. A. (1930) *The Genetical Theory of Natural Selection*. New York: Dover. Second edn. 1957.
- Fitelson, Branden, Stephens, Christopher, and Sober, Elliott (1999) How Not to Detect Design – a Review of William Dembski’s *The Design Inference*. *Philosophy of Science*, 66, 472–88. Also available at the following URL: <http://philosophy.wisc.edu/sober>.
- Forster, Malcolm R. and Sober, Elliott (2003) Why likelihood? In M. Taper and S. Lee (eds.), *The Nature of Scientific Evidence*. Chicago: University of Chicago Press. Also available at the following URL: <http://philosophy.wisc.edu/forster>.
- Gould, Stephen Jay (1980) *The Panda’s Thumb*. New York: Norton.
- Hacking, Ian (1965) *The Logic of Statistical Inference*. Cambridge: Cambridge University Press.
- Hacking, Ian (1987) The Inverse Gambler’s Fallacy: The Argument from Design. The Anthropic Principle Applied to Wheeler Universes. *Mind*, 96, 331–40.
- Hempel, Carl G. (1965) Studies in the Logic of Confirmation. In *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. New York: Free Press.
- Hume, David ([1779] 1980) *Dialogues Concerning Natural Religion*. Indianapolis: Hackett Publishing.
- Keynes, John Maynard (1921) *A Treatise on Probability*. London: Macmillan.
- Kyburg, Henry (1961) *Probability and the Logic of Rational Belief*. Middletown, Conn.: Wesleyan University Press.
- Leslie, John (1989) *Universes*. London: Routledge.
- Madden, Edward and Hare, Peter (1968) *Evil and the Concept of God*. Springfield: Charles Thomas.
- McMullin, Ernan (1993) Indifference principle and anthropic principle in cosmology. *Studies in the History and Philosophy of Science*, 24, 359–89.
- Paley, William (1802) *Natural Theology, or, Evidences of the Existence and Attributes of the Deity, Collected from the Appearances of Nature*. London: Rivington.
- Plantinga, Alvin (1974) *The Nature of Necessity*. New York: Oxford University Press.
- Plantinga, Alvin (1979) The Probabilistic Argument from Evil. *Philosophical Studies*, 35, 1–53.
- Rowe, William L. (1979) The Problem of Evil and Some Varieties of Atheism. *American Philosophical Quarterly*, 16, 335–41.
- Royall, Richard M. (1997) *Statistical Evidence – a Likelihood Paradigm*. London: Chapman and Hall.
- Sober, Elliott (1993) *Philosophy of Biology*. Boulder, Colo.: Westview Press.
- Sober, Elliott (1999) Testability. *Proceedings and Addresses of the American Philosophical Association*, 73, 47–76. Also available at the following URL: <http://philosophy.wisc.edu/sober>.
- Sober, Elliott (2000a) Evolution and the Problem of Other Minds. *Journal of Philosophy*, 97, 365–86.

- Sober, Elliott (2000b) Quine's Two Dogmas. *Proceedings of the Aristotlean Society*, 74, 237–80.
- Sober, Elliott (2002) Bayesianism – Its Scope and Limits. In R. Swinburne (ed.), *Bayesianism*. London: British Academy. Also available at the following URL: <http://philosophy.wisc.edu/sober>.
- Sorensen, Roy (1988) *Blindspots*. Oxford: Oxford University Press.
- Stigler, Stephen M. (1986) *The History of Statistics*. Cambridge, Mass.: Harvard University Press.
- Swinburne, Richard (1968) The Argument from Design. *Philosophy*, 43, 199–212.
- Swinburne, Richard (1990a) Argument from the Fine-Tuning of the Universe. In J. Leslie (ed.), *Physical Cosmology and Philosophy* (pp. 160–179). New York: Macmillan.
- Swinburne, Richard (1990b) The Limits of Explanation. In D. Knowles (ed.), *Explanation and its Limits* (pp. 177–93). Cambridge: Cambridge University Press.
- Turing, Alan (1950) Computing Machinery and Intelligence. *Mind*, 59, 433–60.
- van Inwagen, Peter (1993) *Metaphysics*. Boulder, Colo.: Westview Press.
- Venn, John (1866) *The Logic of Chance*. New York: Chelsea.
- White, Roger (2000) Fine-Tuning and Multiple Universes. *Noûs*, 34, 260–76.

Suggested Further Reading

- Pennock, Robert T. (ed.) (2001) *Intelligent Design Creationism and its Critics*. Cambridge, Mass.: MIT Press.
- Pennock, Robert T. (2002) *Tower of Babel: The Evidence Against the New Creationism*. Cambridge, Mass.: MIT Press.
- Ruse, Michael (2003) *Darwin and Design: Does Evolution have a Purpose?* Cambridge, Mass.: Harvard University Press.