The Indeterministic Character of Evolutionary Theory: No "No Hidden Variables Proof" but No Room for Determinism Either
Author(s): Robert N. Brandon and Scott Carson
Source: *Philosophy of Science*, Vol. 63, No. 3 (Sep., 1996), pp. 315-337
Published by: The University of Chicago Press on behalf of the Philosophy of Science Association
Stable URL: http://www.jstor.org/stable/188099
Accessed: 14/10/2011 14:38

Philosophy of Science

September, 1996

THE INDETERMINISTIC CHARACTER OF EVOLUTIONARY THEORY: NO "NO HIDDEN VARIABLES PROOF" BUT NO ROOM FOR DETERMINISM EITHER*

ROBERT N. BRANDON†‡

Departments of Philosophy and Zoology
Duke University

AND

SCOTT CARSON

Department of Philosophy
Duke University

In this paper we first briefly review Bell's (1964, 1966) Theorem to see how it invalidates any deterministic "hidden variable" account of the apparent indeterminacy of quantum mechanics (QM). Then we show that quantum uncertainty, at the level of DNA mutations, can "percolate" up to have major populational effects. Interesting as this point may be it does not show any autonomous indeterminism of the evolutionary process. In the next two sections we investigate drift and natural selection as the locus of autonomous biological indeterminacy. Here we conclude that the population-level indeterminacy of natural selection and drift are ultimately based on the assumption of a fundamental indeterminacy at the level of the lives and deaths of individual organisms. The following section examines this assumption and defends it from the determinists' attack. Then we show that, even if one rejects the assumption, there is still an important reason why one might think evolutionary theory (ET) is autonomously indeterministic. In the concluding section we contrast the arguments we have mounted against a deterministic hidden variable account of ET with the proof of the impossibility of such an account of QM.

*Received July 1995; revised December 1995.

†We wish to thank John Beatty, Alex Rosenberg and Elliott Sober for helpful comments on an earlier draft of this paper.

‡Send reprint requests to the authors, Department of Philosophy, Duke University, Durham, NC 27708.

Philosophy of Science, 63 (September 1996) pp. 315–337. 0031-8248/96/6303-0001$2.00

1. Introduction. With the possible exception of relativity theory, philosophers have devoted more attention to quantum mechanics (QM) and evolutionary theory (ET) than to any other scientific theories. We think this is entirely understandable since they are probably the two most highly confirmed theories in the history of science, and no other scientific theories rival these two in terms of their implications for traditional philosophical topics. Steven Weinberg (1992), who received the Nobel Prize for his work on unified field theory, has speculated on the future shape of physics. He thinks much may change, but he cannot imagine a future physics without QM.[1] The great evolutionary biologist, Theodosius Dobzhansky, has said, "Nothing in biology makes sense except in the light of evolution." Because of its centrality to all of biology we, with much less authority than Weinberg, can hardly imagine a future biology without ET. Philosophers have recognized the potential impact of evolutionary thinking on almost all areas of philosophic concern, ranging from ethics and epistemology to metaphysics. The philosophical implications of QM are less broad, but probably deeper. In particular, the apparently irrefutable indeterminism of QM has affected the way philosophers think about causality, explanation, and the basic aims of science.

ET is also apparently indeterministic; certainly the best and most influential treatments of the probabilistic nature of ET have drawn this conclusion (Beatty 1984, Sober 1984, Richardson and Burian 1992). Moreover, the propensity interpretation of fitness (Brandon 1978, 1990; Brandon and Beatty 1984; Burian 1983; Mills and Beatty 1979; Richardson and Burian 1992), which has been accepted by most philosophers of biology and many working evolutionary biologists, presupposes that natural selection is fundamentally probabilistic. Recently, however, two philosophers, Rosenberg (1988, 1994) and Horan (1994), have questioned this conclusion. They have argued that the statistical character of evolutionary theory is best viewed instrumentally, i.e., that the probabilities involved in evolutionary theory are epistemic—they reflect our ignorance—and that if we were smarter and/or if we had different aims, evolutionary theory could be recast as a purely deterministic theory. In other words they argue that the *process* of evolution is deterministic while, for various reasons, our best *theory* of evolution is indeterministic. This is exactly the sort of position that has been ruled out in QM by Bell's Theorem (Bell 1964, 1966). We want to show that it is also ruled out in ET, though in a less decisive way.

---

[1] Even Einstein, perhaps the foremost critic of QM, said of it, "The formal relations which are given in this theory—i.e., its entire mathematical formalism—will probably have to be contained, in the form of logical inferences, in every useful future theory." (in Schilpp 1949, 667).

2. Bell's Theorem. With the publication in 1964 (also see 1966) of Bell's important proof of the impossibility of hidden-variables accounts of quantum mechanics, the very foundations of quantum theory underwent a paradigm shift: no longer could the mysterious correlations of events at a distance be explained as the results of causes operating in unobservable ways; the implication of Bell's results was that the correlations do not reflect the subjective probabilities borne of the measurement problem but rather are a product of the objectively probabilistic nature of sub-atomic processes. Accounts of quantum uncertainty now routinely refer to these ontological implications of Bell's work, and the task of explaining the mysterious correlations has taken on a whole new character (see, e.g., Cushing and McMullin 1989).

Bell's work marked an important transition in the interpretation of experiments modeled on a thought-problem first proposed by Einstein, Podolsky, and Rosen in 1935.[2] These EPR-type experiments had shown that sub-atomic particles emitted in opposite directions from a common source and measured for some property (such as spin) at distant measurement stations (A and B) exhibit correlations with respect to the measured property. The nature of the correlation, whether positive or negative, depends on the settings at the two measuring stations. Experimenters at each station must adjust the measurement settings prior to the actual measurement event, and can do so at the last moment, in which case it will be in principle impossible (given the assumption of "locality" discussed below) for the experimenter at station A to communicate the setting of her device to the experimenter at station B, and vice versa.

Given this fact about the measurement stations the correlations that are predicted by quantum theory are all the more curious: it seems as though the particle at one measurement station somehow "knows" the property of the particle at the other measurement station. Since this did not seem very likely to the early theorists, various other explanations were offered, including the influence of "hidden variables" operating at the emission source. One of the assumptions underlying the paradoxical character of these correlations is the so-called "locality" assumption, namely the assumption that there can be no transfer of information across some distance d in less than d/c (where c is the speed of light) time, so the possibility that the particles were somehow "communicating" with each other must be

---

[2] Cushing 1989 uses a simplified version of the Einstein-Podolsky-Rosen (EPR) thought experiment due to Bohm 1951; our description of Bell's results are based on the probabilities associated with the Bohm model. Hughes 1989 describes a similar sort of experiment (the so-called Stern-Gerlach experiment) of even earlier provenance (1921, though the history of the problem goes back to 1913 and Rutherford's criticism of Bohr's model of the hydrogen atom). The differences between these two types of experiment are not significant to our discussion, and we use the designation EPR to refer to the sort of experiment that produces the paradoxical results discussed by Bell.

ruled out in order to make QM compatible with special relativity. Different sorts of accounts were given to show why the hidden variables thought to be responsible for the correlations were likely to remain hidden, but Bell was able to show that every sort of hidden-variables account contradicts the predictions of quantum theory, predictions which have been upheld by empirical experiments (Clauser and Shimony 1978).[3] Thus he showed that there was no possible locally deterministic theory capable of dealing with EPR-type setups. It would seem that either locality must go or that the deterministic-hidden-variables-type account of the correlations must go; the parsimonious physicist chooses to save relativity at the expense of what is, after all, a purely speculative explanation of the data, the only virtue of which is that it conforms to some deeply felt a priori metaphysics.

At least part of what is interesting to us about Bell's results is that they render incoherent a particular interpretation of QM, an interpretation motivated by a deeply held belief that the universe *must* be deterministic and which, therefore, sees QM as merely a useful instrument for predicting the phenomena of the microphysical world. We label this interpretation 'local instrumentalism', to contrast it with a global instrumentalism for all of science. Nothing in Bell's results, nor in QM in general, rules out global instrumentalism or any sort of global anti-realism.[4] But in this paper we will not be concerned with these views, since they are not relevant to the hidden variable accounts of the EPR results, which, after all, had realist motivations. They were motivated by the belief/hope that the world was *really* deterministic. Given that belief, and the unquestionable empirical success of QM, local instrumentalism with respect to QM was tempting. Bell's results show that for the global realist—someone concerned with the unity of science—the price for taking this path is simply too high. In what follows, we hope to show that local instrumentalism with respect to ET extracts a similarly high cost.

3. Population Effects of Quantum Uncertainty.

Are biological processes deterministic or indeterministic? For most such processes the honest answer is that we do not know. However, given the *known* indeterminism of micro-physics and the well-supported assumption of the dependency of biology on chemistry and ultimately physics, one might think that biological processes, e.g., ontogenetic development and evolution, are likely to

---

[3] Actually Bell's results apply beyond quantum mechanics proper; see, e.g., Mermin 1989.

[4] As we are using the term 'local instrumentalism', global instrumentalism does not imply it; indeed, as explained in the text above, local instrumentalism is predicated on some sort of global realism. Thus Bell's Theorem, which refutes local instrumentalism, does not touch the global instrumentalist.

be indeterministic as well.[5] At least, one might think, the possibility that some biological processes are indeterministic should be taken seriously. But that is not always the case. As Elliott Sober puts it, "Philosophers and scientists sometimes seem to think that the indeterminism discovered by quantum mechanics is limited to the micro-level. The idea seems to be that organisms and populations are composed of too many fundamental physical particles for chance at the microlevel to "percolate up" to the macro-level that evolutionary biology describes." (1984, 121). Sober goes on to say that he fails to find that reasoning compelling. We agree, and in this section present a simple example showing how quantum uncertainty at the level of a point mutation can have major evolutionary implications.

Consider a population composed of two haploid genotypes, *A* and *a*. Let their relative frequencies be p and q respectively ($p + q = 1$). The fitness of $A = 1 - q$, while the fitness of $a = 1 - p$. (Figure 1) This situation has three equilibrium points, at $p = 1$, at $p = 0$ and at $p = .5$. Points $p = 1$ and $p = 0$ are stable equilibria, small perturbations from

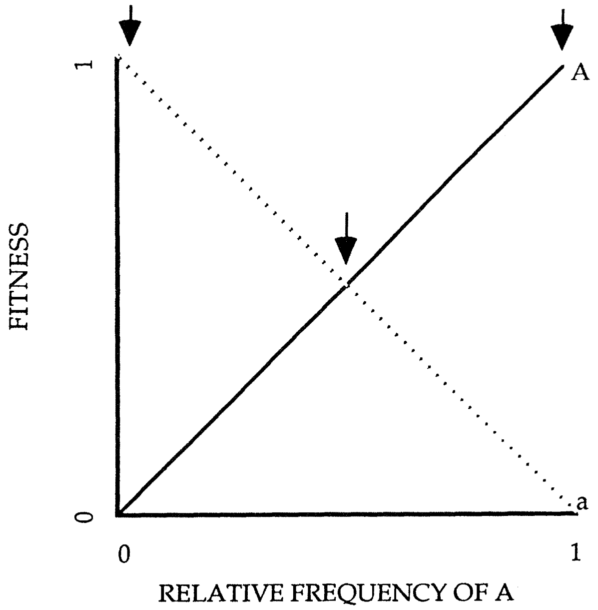

Figure 1. Fitness curves for genotypes *A* (solid line) and *a* (broken line). The three equilibria are marked by arrows.

---

[5] The term "dependency" in this sentence is carefully chosen; it is not meant to imply that biology is reducible to chemistry and physics.

them result in selection pushing the population back to the original point. But $p = .5$ is unstable. Small perturbations from it result in selection driven evolution towards $p = 1$ or $p = 0$, depending on the direction from .5 the perturbation takes. Now suppose a population is at $p = .5$, and that a point mutation turns an A individual into an a. That population will go to $p = 0$. In another population poised at $p = .5$, a mutation turns an a into an A. That population goes to $p = 1$. And so on.

If such mutations are genuine quantum phenomena, then quantum uncertainty would "percolate up" in a powerful way to the level of populations. The evolutionary trajectory of such populations would be genuinely indeterministic. And so evolution, in this way at least, would be indeterministic.

What does this example show? First, even though at present the equations of QM cannot be applied to DNA because of its structural complexity, there is little doubt that at least some mutations are genuine quantum phenomena. Second, the simple example just discussed is not biologically unrealistic in any of its crucial aspects save one. In particular, frequency dependent selection of the sort modeled above is common enough in nature. Thus if the model is not unrealistic and if quantum indeterminacy shows up in some mutations, then it follows that quantum indeterminacy can "infect" the evolution of populations. But the model is unrealistic in one important respect, namely, if such a situation existed in nature we would not expect to find populations at $p = .5$, because drift would surely bump any such population off that unstable point. Of course, among evolutionary biologists the standard way of thinking about drift treats it as an indeterministic phenomenon on a par with quantum mutation. Whether such thinking is justifiable will be considered in the next section.

What we have shown is that quantum indeterminacy can certainly have an amplified effect at the macrolevel. This point is not recognized by Rosenberg[6] or Horan, but they argue that quantum indeterminism cannot be the source of genuine autonomous evolutionary indeterminism. For ET to be autonomously indeterministic it must be indeterministic in a way that does not depend on QM. Thus they think, and we agree, that the crucial question concerns the evolutionary processes of drift and natural selection. It is to these that we now turn our attention.

---

[6] He says, "In general the quantum probabilities involved in biological processes are so small, and the asymptotic approach to determinism of everything physical above the level of the chemical bond so close, that quantum mechanical probability could never explain the probabilistic character, if any, of either evolutionary phenomena or evolutionary theory." (1994, 61). However Rosenberg does recognize that quantum indeterminism can "infect" the macroworld (60).

4. Genetic Drift. Most population genetic textbooks define genetic drift as genetic changes that occur by "sampling error" from the gamete pool (see, e.g., Roughgarden 1979, Ch. 5). All sexually reproducing organisms produce many more gametes than the number of offspring they reproduce. Suppose in a sexually reproducing population of 10 individuals, 5 male and 5 female, the males produce a total of 1,000 sperm cells and the females produce 1,000 ova. Due to some density limiting factors at most 10 offspring can be produced for the next generation. Assuming no relevant physical differences among the sperm (e.g., no differences in viability or in ability to fertilize an ovum) and no such differences among the ova, the sampling of ten sperm from the sperm pool and ten ova form the ova pool will be "random." If there is genetic variation in the gamete pool (the sperm pool + the ova pool) then that variation is unlikely to be transmitted unchanged to the next generation (the zygote pool). For instance, if allele *A* is rare in the gamete pool (say, 1%) then it is unlikely to appear in the zygote pool.[7] Similarly, if 50% of the gamete pool is *A*, it is unlikely that the frequency of *A* in the next generation will be exactly .50. Thus the change in the frequency of *A* from the gamete pool to the zygote pool due to random sampling effects is called genetic drift.

Although this is the standard textbook treatment of drift, it is too restrictive. John Beatty (1984) usefully distinguishes between "indiscriminate parent sampling" and "indiscriminate gamete sampling." Only the latter term fits the standard textbook treatment, but the former is an equally plausible form of "sampling error." We can think of the latter case as one where the gametes go through a restrictive bottleneck and the former as one where the parents themselves go through the bottleneck. For instance, our ten parents may be indiscriminately reduced to two before reproduction. What was said about gamete sampling error holds here as well, for instance an allele rare in the parental generation (say 10%) is unlikely to make it into the next generation. Other forms of drift are probably important in evolution but will not be discussed here.[8] Thus we suggest that genetic drift be characterized as any transgenerational (evolutionary) change in gene or genotype frequencies due to sampling error.

Drawing colored balls from an urn provides a helpful model for thinking about drift.[9] Suppose we have an urn containing a large number of

---

[7] The probability that it will appear equals .02. That is the sum of the probability that it will appear among the sampled sperm (.01) and the probability that it will occur among the sampled ova (.01).

[8] Brandon (1990, Ch. 2) discusses another drift-like phenomenon, the random distribution of different genotypes over heterogeneous selective environments, which certainly occurs and certainly has evolutionary consequences.

[9] Since this model will be crucial to our discussion of drift and selection it is important for us to remove one potential source of misunderstanding. This model does not represent all

colored balls, say 10,000, half of which are red and half black. If we sample that urn by taking ten balls out, the most probable outcome (the outcome we will get most often—the *mode*), is that our sample will contain five red and five black balls. But that most probable outcome occurs only about one quarter of the time, other trials will yield other outcomes (e.g., 3 red, 7 black). The average outcome, the *mean* or the *expected outcome,* corresponds to the ratio of red to black in the sampled urn, independent of the size of our sample. But the "sampling error," the deviation from the ratio in the sampled urn of our sample, does depend on the size of the sample; the larger the sample the smaller the expected error. This is just elementary statistics and it applies directly to genetic drift—the smaller our sample (whether from the gamete pool, parent pool or whatever) the greater the likely deviation from the parental gene/genotype frequencies.

The above suggests that, just as in statistics, the inferences we can make with respect to the occurrence of drift are probabilistic ones: the smaller the population size (and therefore the sample size) the *more likely* drift is to occur.[10] Thus drift is the *probable* outcome of certain biological sampling processes such as reproduction. This is correct as far as it goes, but it is worth noting that there are certain situations, which are easy to produce experimentally and which must occur at least occasionally in nature, that make drift *necessary,* not just highly probable. For instance, suppose there is a population of 1,000 individuals in generation 1 that goes through a severe bottleneck between generations 1 and 2 due to a crash in environmental resources, a sudden explosion of predators, or experimental intervention. The bottleneck results in a population of size 10 in generation 2. Suppose there is an allele *A* that occurs with a frequency of .01 in generation 1. It either will, or more likely, will not make it to generation 2. In either case its frequency will have been changed from .01 to 0 or, less likely, from .01 to .1 or .2, . . . , or 1. It is mathematically impossible for its frequency to remain the same between the two generations and so drift *must* occur. (The mathematical point we make here is obvious; however, to our knowledge, no one has heretofore drawn the connection to drift.)

Roughgarden uses the language of information theory to describe drift. "The gene frequency, *p*, is a 'signal' that must be transmitted from one

---

aspects of *evolution* by drift, or evolution by selection and drift. In particular it provides no representation of reproduction and so none of transgenerational change. It could be modified to do so, but this would be irrelevant to the points we are making. Our points are solely concerned with sampling of parents or gametes for which the urn model is wholly adequate.

[10] Here we are considering drift in the absence of selection and mutation. In more realistic models that include both mutation and selection the likely effect of drift is a function of population size, or more strictly *effective population size,* mutation rates and the strength of selection. For instance, selection dominates drift when $4Ns \ll 1$, where N is the effective population size and s is the selection differential, and drift dominates selection when $4Ns \gg 1$. See Roughgarden 1979, 74–79.

generation's gene pool to the next. The effect of sampling error is to introduce 'noise' into the communication channel. Because of this noise, the signal that is received fluctuates from generation to generation" (1979, 58). Despite the "noisiness" of the process the theory of genetic drift has been developed in a way that allows it to make precise predictions concerning the evolutionary consequences of drift. We have already alluded to one such consequence that is easy to understand: In the absence of other evolutionary forces rare alleles tend to disappear by drift. Indeed, even common alleles disappear under the action of drift. In a computer simulation of five diploid populations of constant size 8 all with initial allele frequency, p, of $A = .5$, all five had drifted to fixation ($p = 1$), or loss ($p = 0$) by generation 41. In a simulation of fifty such populations, almost all had drifted to fixation or loss by generation 20. (Roughgarden 1979, Figs. 5.1 and 5.2). States $p = 1$ and $p = 0$ are called absorbing barriers since in the absence of mutation or immigration once there a population cannot move from that state.

The theory of genetic drift cannot predict which population will evolve to which absorbing state, but it can predict for an ensemble of populations the percent of populations that will evolve to one state and the percent that will evolve to the other and the expected time this evolution will take. (For $p = .5$, 50% of the populations will go to state $p = 1$, for $p = .4$, 40% will go to state $p = 1$, and so on.) Such predictions are of great theoretical importance, for instance in helping to distinguish empirically drift from natural selection and so helping to address the general question of the relative importance of drift versus selection. They are also of great practical importance in areas such as conservation biology, agriculture and medicine. (Does the AIDS virus stand still while we try to design a vaccine?) Thus one should not conclude from the fact that drift is an inherently directionless process that its effects are therefore unpredictable. Just as QM has had tremendous success as a predictive science so too has the theory of genetic drift. Both theories' success comes at the level of ensembles—ensembles of particles in the case of QM, ensembles of populations in the case of drift. What looks like a nonpredictive theory at one level (individual particles, individual organisms, or even individual populations of organisms) is seen to be quite predictive when we focus on the appropriate level. We will return to this point shortly.

The preceding discussion of drift has departed from standard presentations in only two ways. First, we have argued for a more extensive conception of drift than that given in most population genetics texts (but one which we think reflects the view of most evolutionary biologists). Second, we have shown, for the first time as far as we are aware, that certain situations *force* drift to occur as opposed to merely making it more probable. Otherwise, the above represents the received view of genetic drift.

On that view drift is clearly a stochastic or probabilistic or indeterministic phenomenon. The situation with natural selection is a bit more complicated.

5. Natural Selection. Biologists often describe natural selection as a *deterministic* phenomenon. By this they mean its effect is directional (i.e., has a predictable direction) as opposed to drift. They do not mean that it is deterministic in the philosophers' sense (Sober 1984, 110–115). There are deterministic population genetic models of natural selection, but they all involve infinite population sizes and so preclude drift.[11] This suggests the following: *natural selection is indeterministic at the population level because (in real life as opposed to certain formal models) it is inextricably connected with drift.* We will briefly defend this position and then show how it is based on a philosophically controversial supposition that natural selection is indeterministic at the level of the lives and deaths of individual organisms.

Let us return to our urn example. As a model of drift, each ball in the urn has an equal probability of being pulled. Suppose however that different balls had different chances of being pulled, e.g., some balls are big and some small, with the big ones more likely to be chosen; or some balls are sticky and some slippery, with the sticky ones more likely to be pulled. Take the latter case. Suppose a sticky ball is twice as likely to be pulled as a slippery ball. Following the practice of population genetics, let us assign a "relative fitness" to sticky balls of 1, and a "relative fitness" to the slippery balls of .5. Suppose that in our urn of 10,000 balls there is a 50:50 ratio of sticky to slippery. If we pull out ten balls, what is the expected result? Recall that where each ball has an equal probability of being pulled (the drift model) the expected ratio of the types in our sample just equals the overall ratio in the urn. But here (the selection model) different balls have different probabilities of being pulled. In general, with two types the expected frequency of type a in a sample (p_s) equals the fitness of type a (w_a) times the frequency of type a in the sampled population (p) all divided by the mean fitness of the balls. Likewise for type b (where q_s is the expected frequency of type b in the sample and q is the frequency of b in the sampled population—$p_s + q_s = 1$ and $p + q = 1$). Thus:

$$p_s = pw_a/(pw_a + qw_b) \text{ and}$$
$$q_s = qw_b/(pw_a + qw_b)$$

In our numerical example, the expected result is a 2:1 ratio and so with

---

[11] Sober (1984, 111) argues that even in these models with infinite population sizes the inference from starting gene frequencies and selection coefficients to gene frequencies at some later time is still only probabilistic and not deductive. We take no position on this matter. In saying that such models are deterministic we are merely reporting their formal structure, not commenting on whether that formal structure is philosophically justifiable.

sample size of ten we expect 6⅔ sticky balls and 3⅓ slippery balls. Given that we pick whole numbers of balls this expected result will never occur in a sample of size ten.

The urn example with balls that have different probabilities of being chosen is a good model of selection and drift. The differing probabilities associated with sticky and slippery represents selection and the possible deviation from the expected result represents drift. Suppose we pull ten balls from our urn and get 6 sticky balls and 4 slippery ones. This result deviates from the expected result and this deviation is drift. In the section on drift above we presented examples of sampling error based on indiscriminate sampling. What our current example exemplifies is sampling error where the sampling is not indiscriminate (the sampling procedure does discriminate between sticky and slippery). All the conclusions we drew about drift when it acts alone apply here where it acts in concert with selection. In particular, everything else being equal, the smaller the sample size the more likely drift is to occur; and, as our numerical example shows, some situations *force* drift to occur.

Our point in using this example is that adding differing probabilities for being pulled from the urn does not affect the possibility of sampling error when we take a finite sample.[12] Thus the example illustrates how selection and drift are linked in finite populations. We take this to go beyond the epistemic point that it is difficult to distinguish drift from selection.[13] The above example is one where, by hypothesis, we know all there is to know about the situation yet drift can occur, and sometimes must occur, along with selection. But we must now turn to the assumption on which this conclusion is based.

Our example shows that at the population level natural selection is indeterministic because with finite populations selection does not eliminate the drift-effects of sampling error. Put another way, when we are concerned with finite populations, and all real biological populations are obviously finite, it is legitimate to consider drift without selection but not vice versa. Except in cases where the probabilities of being sampled are all either 1 or 0, selection in finite populations always involves the possibility of drift. But our example presupposes that the balls in the urn have real

---

[12] For a given sample size, the greater the selection differentials (the difference between the fitnesses of the types) the less the chance of drift. In the extreme case, if one type of ball has a probability of 1 of being pulled and the other a probability of 0, then drift cannot occur. See footnote 10 above.

[13] Beatty (1984, 196) comes close to making this point, but seems to stop short of it. Concerning an example similar to ours he says, "it seems that we must say of some evolutionary changes that they are to some extent, or in some sense, a matter of natural selection *and* to some extent, or in some sense, a matter of random drift. And the reason (one of the reasons) we must say this is that it is conceptually difficult to distinguish natural selection from random drift . . ." We agree with the first sentence in this quote, but not the second. Conceptually (though often not empirically) drift and natural selection are clearly distinct.

probabilities of being drawn that lie between 0 and 1. As a model of selection, this is just another way of saying that the reproductive output of the lives of individual organisms is a genuinely stochastic affair. Thus that is the basic assumption on which our conclusions concerning drift and selection have been based.

6. Fitnesses—Real Propensities or Useful Instruments? Let us now consider the exceptional case where the probability of being sampled equals either 1 or 0 for every individual. Our previous example presupposed that the balls in the urn have real probabilities of being drawn that lie between 0 and 1. We have seen how that presupposition results in population-level indeterminacy. Suppose, contrary to that presupposition, that there is some *hidden variable,* in this case unequally distributed between slippery and sticky, that takes two values, 0 and 1, and that any ball with the 0-value of that variable has a 0 chance of being pulled while any ball with the 1-value has a probability of being pulled of 1. Even in this case, the "relative fitnesses" we have assigned the two types of balls—sticky and slippery—are by no means useless; they provide the statistical information concerning the distribution of the hidden variable between our two types of balls and they allow us to draw statistical generalizations of the sort discussed above. However, according to the arguments of Rosenberg and Horan, such assignments of relative fitnesses would be merely epistemic, merely useful instruments given our state of ignorance of the hidden variable. The appropriate attitude towards these fitnesses would be instrumentalist; we would ask, "Are they useful in producing predictions?" (as we have seen the answer to this would be "yes"), not "Do they represent the true propensities of our urn-drawing set-up?".

Both Rosenberg and Horan are clear that their arguments are intended to apply specifically to ET. They are not arguing for a global instrumentalism with respect to all science, but rather a local instrumentalism with respect to ET.[14] Indeed their arguments are not terribly interesting from the point of view of the global instrumentalist. Obviously, for such a person, ET, like all of science, should be viewed instrumentally. Thus, Rosenberg and Horan are taking some form of scientific realism as their default position; they are looking for special reasons for instrumentalism with respect to ET. (Thus, in this way their position is exactly analogous to that of the proponents of hidden variable accounts of QM which, as we saw above, are predicated on realist presuppositions.) For the moment, we follow their lead in this, that is, we will assume some form of scientific

---

[14] To be more accurate, Rosenberg's (1994) stated position is realist with respect to ET, but instrumentalist with respect to the probabilities employed by ET. In our view, this makes him an instrumentalist with respect to ET, i.e., we do not see how one can separate the probabilities employed in ET from the theory proper.

realism as a background default position in responding to Rosenberg's and Horan's arguments concerning ET. Indeed, to refute their arguments it is necessary for us to take this stance. We are looking for an ET analogue of the no hidden variables proof in QM; it makes no sense to do that from an instrumentalist standpoint since even that proof cuts no ice against a global instrumentalism.

We think the arguments of Rosenberg and Horan are problematic in at least two ways. First, even if one accepts the supposition of a deterministic hidden variable, it does not follow that fitnesses must be thought of instrumentally. We will deal with this matter in the next section. Second, and this we think is the crux of the matter, the supposition of deterministic hidden variables is itself problematic. This section is devoted to exploring those problems.

Why, other than wishful thinking, should one posit deterministic hidden variables underlying natural selection and drift? We will argue that there is no scientifically justifiable reason to do so, but first we examine the arguments offered by Horan and Rosenberg.

Horan (1994, 84) asserts that the drift effect of sampling error requires indiscriminate (i.e. equiprobable) sampling. (Our slippery/sticky ball example shows this to be simply wrong.) But even the indiscriminate sampling cases are not truly indeterministic since, according to Horan, all biological sampling is deterministic: "Fires, floods, famine and disease sample from a large population by eliminating individuals from the reproducing group. The same fire, the same flood, in the same conditions would create the same sample" (1994, 84). Unless one finds argument-by-sincere-assertion convincing, this fails to convince.[15] What about natural selection? Partly based on her misunderstanding of the requirements for sampling error mentioned above, and partly based on an argument from Rosenberg, Horan says, "If fitness theories describe a supervenient relationship, such that any two organisms with identical phenotypic traits subject to the same selection pressures must have the same level of fitness, then the relationship between phenotypic traits and fitness is deterministic (Rosenberg 1985)" (1994, 85). This assertion seems to beg the question against the proponents of the propensity interpretation of fitness who think: (a) fitness is a supervenient property; and (b) it is a probabilistic, not deterministic, dispositional property (see, e.g. Brandon 1978, 1990; Brandon and Beatty 1984; Mills and Beatty 1979; Burian 1984; Richardson and Burian 1992). Since this assertion is based on arguments by Rosenberg, we now turn to his most recent views on this matter.

---

[15] As John Beatty pointed out to us, it is worth noticing that Horan's assertion deals only with the parent sampling form of drift, not gamete sampling. Is it intuitively obvious that the same meiotic event followed by the same fertilization event would produce the same set of genomes?

Rosenberg (1988 and 1994) does indeed think that natural selection is a deterministic affair, but not simply because of the supervenience of fitness. By definition, if fitness supervenes on organisms in their common selective environment, then two phenotypically identical organisms in the same selective environment will have the same *expected fitness* (or in Brandon's terminology, *adaptedness*). From that it does not follow that there is a deterministic relationship between their expected fitness and their realized fitness.[16] There will not be a deterministic relationship if, as the proponents of the propensity interpretation argue, expected fitness is a probabilistic disposition. Thus it would be a mistake to assert, as Horan does, that this deterministic relationship is implied by supervenience. As far as we can tell, Rosenberg does not make this mistake. Instead, he bases his argument for the deterministic nature of natural selection on the closely related doctrine he calls *mereological determinism.*

Rosenberg offers no definition of mereological determinism, so we will have to come up with one which we think fits his needs. If macro-states $A_1, A_2, \ldots, A_n$ supervene on micro-states $I_1, I_2, \ldots, I_m$, then, by definition, any two systems in the same micro-state I_k must be in the same macro-state A_l. Thus, as stated above, two identical organisms in the same selective environment will have the same expected fitness. But, as shown above, that does not imply that the two organisms will have the same realized fitness. If mereological determinism is to imply the determinism of natural selection, then it must be defined in a stronger way than supervenience. The following definition would serve Rosenberg's needs: *mereological determinism holds of a system S, if and only if the total micro-state description of S at time t determines every macro-property of S at t.* Mereological determinism so defined implies the supervenience of *every* macro-property of S on its micro-states. And so, not only would expected fitness be determined by the properties of the organism in its selective environment, so would its realized fitness. Thus natural selection would be deterministic, or so it seems.

Is the principle of mereological determinism true? One might accept it as an analytic truth, or one might think of it as a candidate empirical truth. On the latter conception, it is difficult to see how to test it empirically, and the reason for that difficulty is precisely the reason why, true or false, it is irrelevant to the question of determinism in ET. ET, like most of science, deals with dynamic processes, but the statement of mereological determinism above is static—it deals solely with time t. What we need to know is whether two systems in identical micro-states at time t will be in

---

[16] On the distinction between expected fitness and realized fitness see Brandon 1990 (Ch. 1) or Richardson and Burian 1992. Basically the probabilistically defined expected fitness represents the chances of various levels of reproductive success of the organism while the realized fitness is simply the organism's actualized reproductive success.

the same macro-state (or micro-state) at time $t + \Delta$. We could define mereological determinism so that it has that implication: *mereological determinism holds of a system S if and only if the total micro-state description of S at t determines every micro-state and hence every macro-property of S for every* $t + \Delta$.[17] So defined we know mereological determinism is in general false, as shown in QM for example by the Stern-Gerlach experiment (see Hughes 1989, 1–8 or Albert 1994, 1–16). Is there any reason to think it is true of the process of natural selection?

Neither Rosenberg nor Horan offer any such reason, so let us turn to a direct examination of the question itself. We need not state the question in terms of mereological determinism. For our purposes a better statement of the question is: Are the processes of natural selection and drift governed by deterministic hidden variables, are they, appearances to the contrary, genuinely deterministic processes? This seems to be an empirical question. If determinism is true in this context, then identical organisms in identical environments should have identical evolutionary fates. Can this idea be experimentally tested?

At first glance, the answer would seem to be "yes." Many organisms (some animals, most plants and probably all protists) are clonable. For instance, grasses grow laterally by producing tillers which are modules composed of a green leafy top and roots. These tillers are capable of becoming physiologically independent plants. Experimentally then it is easy to produce multiple plants that are genetically identical by separating these tillers. Those plants can then be placed in a controlled environmental setting and grown, and their lifetime reproductive success, or some components thereof, can be measured. The determinist prediction is that identical plants in identical environments will have identical reproductive success; indeed will have identical height, weight, flower number, etc.

This experiment is eminently doable; fortunately it has been performed many times so we are saved the trip to the greenhouse. Biologists do not do such experiments in order to test the determinist hypothesis, they do them for various other reasons. To mention just one example, Bever (1994) cloned genotypes of three different grass species in part to see whether these plants performed better in their own soil communities (soils in which those plants had been grown and which therefore contain the various microorganisms associated with them) or the soil communities of other species. These experiments involved growing numerous replicate genotypes in carefully controlled identical soil preparations on the same tabletop in the same greenhouse at the same time. (That would be one treatment. The different carefully controlled soil preparations constituted the

---

[17] As Elliott Sober has pointed out to us, this definition should probably not be termed one of *mereological* determinism, but rather something like *forward-directed* determinism. No matter, our point is to find a definition that suits Rosenberg's needs.

different treatments which were then compared.) It is interesting to note that this experimental procedure assumes an indeterministic response. That is, it assumes that different copies of the same genotype in the same treatment will give different results; otherwise the experiment could be made much smaller with single copies of each genotype for each treatment. Anyone familiar with such experiments will know, without looking at the data, that this assumption was met. Bever measured inflorescence mass, total plant biomass, and root to shoot ratio for each genotype in each treatment and as expected there is considerable variance for each such measure, even though there is a statistically significant effect of most of these treatments.

Bever's study addresses an important question in a novel way. We are simply appropriating his data to address the determinist hypothesis. In that context, there is nothing outstanding about his study, any of the large number of other studies that could have been used to address our question would show the same result. Indeed, as pointed out above, the result is so expectable that it forms the basis of the statistical design of such experiments. Biologists would be shocked into disbelief with any experimental thwarting of this expectation. Do we then have an experimental refutation of the determinist prediction?

Notice that the determinist cannot avail herself of one response to this apparent refutation of their prediction. That response is to attribute the variance in the data to measurement-error. Rosenberg (1994, esp. Ch. 4) seems to place great weight on the importance of measurement-error in ET. For instance he says:

> The application of Newtonian mechanics to astronomy, or for that matter to any system which requires sampling and so may introduce measurement-errors, requires an appeal to statistical considerations and produces results that are probabilistic. Given a probability distribution of the positions and momenta of the bodies in a Newtonian mechanical system, we can predict with a probability as close to one as we like, that probability distributions of position and momentum at any future time will be equally arbitrarily close to those the deterministic equations of Newtonian mechanics leads us to expect when applied to the original probability distribution. Yet these facts have no tendency to show that Newtonian mechanics is probabilistic. (1994, 69)

He goes on to attribute the apparent probabilistic nature of ET to measurement-error that is based on our interests and cognitive limitations, thus Rosenberg's instrumentalism with respect to ET. Measurement-error is, of course, possible in the sort of study discussed above. However, in the first place, it is not all that likely in many cases when the response variable is something like the number of inflorescences. (It would take a pretty sloppy experimenter to count three inflorescences in an *Anthoxanthum odoratum* as four.) In the second place, it is not responsible for the variance in the data. The variance results *not* from multiple measures of the same thing (as in the astronomy case), but rather from independent measures of different entities (the different copies of the same genotype). Put another way, multiple replicates are used in such experiments, not because measurement-error is expected, but rather because different replicates will in reality behave somewhat differently. Again, Bever's study is in no way exceptional in this regard. Indeed, it is safe to say that measurement-error plays no role at all in accounting for the statistical character of experimental evolutionary biology.

The determinist has available one final response that is not so easily defeated. The determinist can blame the unwanted result not on measurement-error, but on another sort of experimental error, *viz.,* experimental error in producing either the supposedly identical organisms or in producing the supposedly identical conditions. For present purposes there is no important difference between these two cases; we will focus on the latter. According to this response there is some hidden variable distributed among our replicated environments that make them deterministically different. If we only had access to this hidden variable we could see how flower pot 1 is different from 2 and thus see why the copy of the genotype in pot 1 performed differently from that in pot 2.

Such a response puts us in the position of QM prior to Bell's Theorem. A hidden variable is postulated that would, if it existed, explain in a deterministic way our apparently indeterministic data. It is hard for us even to imagine a Bell-like result (i.e., a contradiction derived from the supposition of deterministic hidden variables) for ET, and so if we are to counter this determinist gambit it will have to be with less decisive means. Less decisive, but, we hope to show, by no means indecisive.

In science in general, the positing of theoretical entities is taken seriously when (1) the positing of the entity aids the development of theory; and (2) the available empirical evidence supports the posit. It is beyond doubt that the positing of genuinely probabilistic propensities governing the evolutionary fates of individual organisms has been an integral part of the impressive development of evolutionary population genetic theories in this century. And, as we have briefly touched on above, all the available empirical evidence supports this idea. In contrast, the positing of deterministic hidden variables in evolutionary theory serves no theoretical purpose at all, and, insofar as it is allowed to be addressed by data is contradicted by empirical data.

(Indeed, not only does the determinists' posit serve no useful purpose, it raises a serious theoretical puzzle. The posit must, of course, be consistent with the phenomena, which in this case includes phenomena such as no statistically significant relationship between genotype/phenotype and reproductive success in some particular population. That means the posit must be consistent with the macrolevel appearance of chance. But that implies that the micro-causal determinants of reproductive success must arrange themselves so as to be consistent with this apparent stochasticity. Why and how do they do that? This question may be answerable; our point is only that it cannot be ignored by the apologist for determinism. Indeed, it was by dealing with precisely this sort of question—how can the deterministic hidden variables arrange themselves in a way that could account for the phenomena of EPR-type setups?—that led Bell to his famous impossibility result. Although, as stated earlier, we do not envision such a result for ET, we do think that any answer to this question would fail every test of scientific plausibility.)

Were the hidden variables being posited in order to make theory conform to any other preconceived idea this would be a completely compelling case against it. But, such is the power of deterministic ideology that we think further argumentation is in order.

Let us consider a hypothetical case, similar to ones discussed above in the sections on drift and natural selection, where drift is forced to occur because of population contraction. We will suppose that we have experimental control of a haploid population with two genotypes *A* and *a*. Further we suppose that we control a predator of this population that preferentially feeds on *A*, so that it eats two *A*'s for every *a* as long as both are available. The predator eats at a known rate (independent of density and of frequency of the two genotypes) so that we can reduce the population by a set amount by leaving our predator in our population cage for a set amount of time. We start with a population of 1000, with an initial 50:50 ratio of the two types. If we leave the predator in the cage for the amount of time required to reduce the population to 700, what is our expected ratio of *a* to *A*? Since the predator eats 2 *A*'s for every *a* and has eaten 300, we expect our remaining population to consist of 400 *a*'s and 300 *A*'s. For the reasons discussed above, this *expected result* will not always occur, but suppose it does this time. Then the determinist is happy; an assignment of relative fitnesses to the two types (and information about the eating rate and time in cage of our predator) has resulted in a precise prediction for our experimental result.

Now let us run the experiment again from the same starting point, this time leaving the predator in the cage longer, so that the population is ultimately reduced from 1,000 to 400. Let us take a snapshot of the population at $t = 600$, to record the genotype ratio then. What is the expected result? 366⅔ *a*'s and 233⅓ *A*'s. Again, assuming our organisms come in whole numbers, we will not get this result. Suppose instead we get 366 *a*'s and 234 *A*'s. The determinist can, after the fact, slightly readjust the relative fitnesses to predict this result. As justification the determinist can posit a hidden difference in the selective environments of this run and the first experiment. (Note that the increase in time the predator stays in the cage would not normally be thought of as a change in selective environment since the predator's eating behavior is frequency- and density-independent.) Let the experiment continue until $t = 400$. Here the expected result is 300 *a*'s and 100 *A*'s. That result is possible and let us suppose we get it. Now the determinist is forced to readjust the relative fitnesses again, this time back to the originals.

It is unlikely that any set of real life experiments would ever yield results so favorable to the determinist, yet this example is meant to illustrate the scientific absurdity of the determinists' position. Clearly the positing of these hidden differences is purely gratuitous; they are posited for no reason other than to save the deterministic character of the theory. Returning to Bever's experiments, why posit hidden differences among the flower pots? To account for the differences in performance of different copies of the same genotype, the determinist will answer. But that is not the only way of accounting for such differences. The indeterminist accounts for them by positing probabilistic propensities governing the behavior of the plants. Of course, Bever's experiments alone do not show the determinist to be wrong. (Remember that they are not alone; hundreds of experiments show similar results.) Our point is simply that when you remove the bias of deterministic metaphysics and compare the cases to be made for deterministic hidden variables versus probabilistic propensities, there is no contest. The positing of deterministic hidden variables for the sole reason of saving determinism hinders the development of predictive theory.[18]

Thus, the case to be made for genuine indeterminism at the level of the lives and deaths of individual organisms, the case for positing probabilistic propensities for various evolutionary fates, is the standard sort of argument made for the reality of any theoretical entity. Such propensities aid the development of theory and they accord with the available empirical evidence. (Remember, this argument is not designed to convince the global antirealist, only the local instrumentalist.) By this standard, the positing of deterministic hidden variables fails miserably. We conclude that the assumption on which our models of drift and natural selection are based is well-founded.[19]

---

[18] Presumably Rosenberg would agree with this since he thinks probabilistic propensities are a *useful* fiction. Our argument is that they are useful and real.

[19] Space precludes any detailed discussion of another powerful argument, suggested to us by John Beatty, for the reality of the relevant propensities. As discussed by Brandon (1990) and Beatty and Finsen (1989), the work of John Gillespie (1977) has demonstrated that selection can act on the variance in offspring number as well as on mean offspring number. But the variance is, of course, a property of a probability distribution. This strikes us as a compelling argument for the reality of the distribution and hence of the probabilities.

7. Autonomous Statistical Laws. We have argued that the population-level indeterminacy of ET is based on the assumption of indeterminacy at the level of the lives and deaths of individual organisms, and that by the normal standards of science this assumption is well supported. We could stop here, but that would present an incomplete picture of indeterminacy in ET. In particular, there is a reason why one might think that even if the lives and deaths of individual organisms are completely determined, evolution at the population level is still, in an important sense, indeterministic.

Sober (1984) in a section aptly titled, "What Laplace's Demon Would be Missing", argues that the "probability concepts of evolutionary theory have an autonomous explanatory power, whether determinism is true or not" (p. 129). He argues that even if every event had micro-deterministic causes and if there were a being (Laplace's demon) capable of gathering and analyzing the relevant data, certain statistical generalizations important in ET would not fall out of that micro-deterministic picture—Laplace's demon would not discover them.

As an example Sober mentions generalizations from the theory of drift that we have already discussed. Suppose there is a population with a neutral allele, *A*, at a locus that occurs with a frequency of .6. As we have seen, the probability that *A* will go to fixation is .6 while its alternative, *a*, will go to fixation with a probability of .4. If determinism is true, then these probabilities are merely epistemic and there is a deterministic causal story that predicts the real outcome. But, as Sober points out, in an ensemble of such populations we know that *A* will go to fixation in approximately 60% of them, while *a* will go to fixation in the others. This is true irrespective of the micro-causal details of each individual population. Laplace's demon would be lacking this general knowledge.

We will offer another example which we think reinforces Sober's point. The example is Francis Galton's (1889) explanation of the phenomenon of hereditary regression to the mean. Consider a trait like height in human beings. That trait is normally distributed. A normal curve can be characterized by its mean value and its variance (dispersion about the mean). Let one curve represent the parental generation, and a second curve the offspring generation. For Galton, height is *heritable* if taller than average parents have taller than average offspring and shorter than average parents have shorter than average offspring. Put another way, if parental deviation from the (parental) mean is plotted against offspring deviation from the (offspring) mean a positive *regression* between the two indicates heritability. In fact, as Galton showed, height is heritable in humans. Consider outlying parents, i.e., parents on either end of the parental distribution. Their offspring should also show a large deviation from their mean, but a less strong deviation than that of their parents. Why? Consider parents at the tall end of the distribution. If the offspring distribution curve is normal and if these parents produce about as many offspring as other parents, then there is simply more room on the short side of the offspring distribution than to the tall side and so, statistically, most of their children will have to deviate less from the mean than they did. Similarly for extremely short parents and their offspring. And similarly, though with a lesser effect, for parents who are simply tall or short, not extremely tall or short. Thus, regression to the mean is explained.

Galton's explanation of regression to the mean is a excellent illustration of what Hacking (1990) calls an 'autonomous' statistical law. In his historical study of the emergence of such laws he claims that, "*Statistical laws became autonomous when they could be used not only for the prediction of phenomena but also for their explanation*" (p. 182, emphasis in original). We have briefly seen how Galton explained the phenomenon of regression to the mean. But some explanations are better than others, so is it not open to the determinist to argue that a better explanation would come from an understanding of the microphysical details of heredity? We think this is unlikely.

Consider the height of a specific individual, Jones, who is 6′5″. How do we explain Jones's height? One explanation would be in terms of his particular genetic makeup and his particular developmental environment. Another explanation might be in terms of the heights of his parents, who were 6′8″ and 6′10″, the heritability of height and Galtonian regression to the mean. These explanations are not in conflict, both offer complementary explanatory information. But let us now consider the general phenomenon of regression of height to the mean in humans. Galton's explanation works well here. Is there an alternative general microphysical explanation? There could be, but probably not. And if there is one for humans and height will it also apply to all other normally distributed traits in all other species? That is highly unlikely.[20]

We conclude that many of the statistical generalizations in ET are autonomous in Hacking's sense. Thus even if, contrary to our main thesis, the individual-level events of evolution are purely deterministic, the population-level generalizations are probabilistic.

8. Conclusion. It is nice to be able to force one's opponent into a contradiction. That is a neat and simple demonstration of the intellectual deficiencies of your opponent's position. Bell's results in QM do just that:

---

[20] Notice that a proponent of a causal-mechanical model of explanation (sensu Salmon 1989) would argue that the micro-physical explanation of Jones' height is better than the Galtonian, not because it may or may not be deterministic, but because it is more mechanistic than the Galtonian (see Brandon 1996, Chap. 11 for a discussion of Galton's explanation and mechanism). We happen to agree with this, but our point here is that the micro-physical explanation lacks the generality of Galton's—it cannot explain the general phenomena of regression to the mean (on the virtue of generality in explanations see Kitcher 1989).

they show that the posit of locally deterministic hidden variables that would account for the theoretical and empirical phenomena of EPR-type setups involves a contradiction; that is, the price of positing such hidden variables is logical inconsistency. For most, that is simply too high a price to pay. But not all intellectual or conceptual problems in science can be reduced to logical inconsistencies. We have shown that the price to pay for positing underlying deterministic hidden variables to account for the phenomena of natural selection and genetic drift is embracing experimentally meaningless and theoretically obfuscatory hypotheses. That too is a high price.

Have we shown then that one must be a realist with respect to the probabilistic propensities of ET? Of course not. We have not even attempted to address the thoroughgoing scientific antirealist or instrumentalist. What we have shown is that *if* one is a realist in one's attitude towards science—that is, if one thinks that a primary aim of doing science is to develop theories that truly describe the mechanisms producing the phenomena, and if one takes theoretical fruitfulness and experimental confirmation as evidence for the reality of theoretical entities—*then* one should conclude that ET is fundamentally indeterministic. *If,* however, one is a metaphysical determinist with respect to ET—that is, one who has decided for reasons outside of science that the process of evolution is deterministic—*then* one should conclude, along with Rosenberg and Horan, that ET is an instrumental science. We have not tried to adjudicate between these two antecedent starting points; although our prejudices are probably clear.

REFERENCES

Albert, D. Z. (1992), *Quantum Mechanics and Experience.* Cambridge, MA: Harvard University Press.

Beatty, J. (1984), "Chance and Natural Selection", *Philosophy of Science* 51:183–211.

Beatty, J., and Finsen, S. (1989), "Rethinking The Propensity Interpretation", In M. Ruse (ed.), *What the Philosophy of Biology Is: Essays for David Hull.* Dordrecht, Holland: Kluwer, pp. 17–31.

Bell, J. S. (1964), "On the Einstein-Podolsky-Rosen Paradox", *Physics* 1:195–200.

———. (1966), "On the Problem of Hidden Variables in Quantum Mechanics", *Reviews of Modern Physics* 38:447–452.

Bever, J. D. (1994), "Feedback Between Plants and Their Soil Communities in an Old Field Community", *Ecology* 75:1965–1977.

Bohm, D. (1951), *Quantum Theory.* Englewood Cliffs, NJ: Prentice Hall.

Brandon, R. N. (1978), "Adaptation and Evolutionary Theory", *Studies in the History and Philosophy of Science* 9:181–206.

———. (1990), *Adaptation and Environment.* Princeton: Princeton University Press.

———. (1996), *Concepts and Methods in Evolutionary Biology.* Cambridge: Cambridge University Press.

Brandon, R. N., and Beatty, J. (1984), "The Propensity Interpretation of 'Fitness': No Interpretation is no Substitute", *Philosophy of Science* 51:342–347.

Burian, R. (1983), "Adaptation", in M. Grene (ed.), *Dimensions of Darwinism.* Cambridge and New York: Cambridge University Press, pp. 287–314.

Clauser, J. F., and Shimony, A. (1978), "Bell's Theorem: Experimental Tests and Implications", *Reports on Progress in Physics* 41:1881–1927.

Cushing, J. T. (1989), "A Background Essay", In Cushing and McMullin (1989), pp. 1–24.

Cushing, J. T. and McMullin, E. (1989), *Philosophical Consequences of Quantum Theory: Reflections on Bell's Theorem: Studies in Science and the Humanities from the Reilly Center for Science, Technology, and Values,* Volume 2. Notre Dame, IN: University of Notre Dame Press.

Galton, F. (1889), *Natural Inheritance.* London and New York: Macmillan.

Gillespie, J. H. (1977), "Natural Selection for Variances in Offspring Number: A New Evolutionary Principle", *American Naturalist* 111:1010–1014.

Hacking, I. (1990), *The Taming of Chance.* Cambridge: Cambridge University Press.

Horan, B. (1994), "The Statistical Character of Evolutionary Theory", *Philosophy of Science* 61:76–95.

Hughes, R. I. G. (1989) *The Structure and Interpretation of Quantum Mechanics.* Cambridge, MA: Harvard University Press.

Kauffman, S. A. (1993), *The Origins of Order: Self-Organization and Selection in Evolution.* New York: Oxford University Press.

Kitcher, P. (1989), "Explanatory Unification and the Causal Structure of the World", In Kitcher and Salmon (1989), pp. 410–505.

Kitcher, P. and Salmon, W. C. (eds.) (1989), *Scientific Explanation.* Minneapolis: University of Minnesota Press.

Mermin, N. D. (1989), "Can You Help Your Team Tonight by Watching on TV? More Experimental Metaphysics from Einstein, Podolsky, and Rosen", In Cushing and McMullin (1989), pp. 38–59.

Mills, S., and Beatty, J. (1979), "The Propensity Interpretation of Fitness", *Philosophy of Science* 46:263–286.

Richardson, R. C., and Burian, R. M. (1992), "A Defense of the Propensity Interpretation of Fitness", in D. Hull, M. Forbes and K. Okruhlik (eds.), *PSA 1992,* Vol. 2. East Lansing: Philosophy of Science Association.

Rosenberg, A. (1985), *The Structure of Biological Science.* Cambridge, Cambridge University Press.

———. (1988), "Is the Theory of Natural Selection a Statistical Theory?", *Canadian Journal of Philosophy (Suppl.)* 14:187–207.

———. (1994), *Instrumental Biology or the Disunity of Science.* Chicago: University of Chicago Press.

Roughgarden, J. (1979), *Theory of Population Genetics and Evolutionary Theory: An Introduction.* New York: Macmillan.

Salmon, W. (1989), "Four Decades of Scientific Explanation", In Kitcher and Salmon (1989), pp. 3–219.

Schilpp, P. A. (ed.) (1949), *Albert Einstein: Philosopher-Scientist.* La Salle, IL: Open Court.

Sober, E. (1984), *The Nature of Selection: Evolutionary Theory in Philosophical Focus.* Chicago: University of Chicago Press.

Weinberg, S. (1992), *Dreams of a Final Theory: The Search for the Fundamental Laws of Nature.* New York: Pantheon.