Proceedings of the Pittsburgh Workshop in History and Philosophy of Biology, Center for Philosophy of Science, University of Pittsburgh, March 23-24 2001

Session 4: Evolutionary Indeterminism

- **Robert Brandon (Duke University) The Difference Between Drift and Selection**
- **Alan Love (Pittsburgh) Commentary**
- **Paul Griffiths (Pittsburgh) Commentary**
- **Frederic Bouchard (Duke) Summary of Discussion**
- **Workshop Participants**

The difference between drift and selection: A Reply to Millstein

Robert N. Brandon

Roberta Millstein (2002) correctly identifies a serious problem with the view that natural selection and random drift are not conceptually distinct (a view seemingly defended by Beatty, 1984). The problem is that one of the major disputes in 20th, and now 21st, Century evolutionary biology—the selectionist/neutralist debate—is incoherent if that view is correct.[1] Biologists do make mistakes and do, sometimes, pursue blind alleys, but Millstein is loath to say that this huge area of work is all based on a conceptual mistake. I completely agree. But then it is incumbent on one to clearly draw this distinction. That is the primary goal of her article. She thinks that once we distinguish between process and outcome, that the *processes* of natural selection and drift can be sharply distinguished, even though the *outcomes* may not be clearly distinguishable. Unfortunately, her solution to this very real problem does not work. She identifies another "solution" to this problem; one that she thinks fares less well than hers and that she attributes to Brandon and Carson (1996). I agree with her that this "solution", which treats the theory of evolution by natural selection as deterministic, is unattractive for a number of reasons. Fortunately for me, I do not now hold, nor have I ever held, a position remotely close to that attributed to me (unless 180° away counts as close). More fortunately still, there is a real solution to the problem.

1. Exegetical Matters.

The primary point of Brandon and Carson (1996) is that the process of evolution by natural selection is autonomously indeterministic, i.e., indeterministic in a way that follows directly from the theory of evolution by natural selection[2]. That is a strong conclusion and the reader of this discussion is not asked to accept it in total.[3] The argument for it is complicated, but the form of the argument is familiar to philosophers of science. It is the standard argument for the reality of theoretical entities. When the positing of such entities plays a crucial role in the development of a body of a predictive and explanatory theory, and when all available experimental evidence supports the posit, then (if, of course, you are a realist) you should accept the reality of the posited entity. The better confirmed the theory, the stronger this argument.

Our conclusion, one that we state very clearly, is that the probabilistic propensities attributed by the theory of natural selection are real. In short, fitness values are real propensities, not useful instruments. So our conclusion is not about the *theory* of evolution by natural selection, but rather about the *process.* However the major premise of this argument is that the *theory,* which is explanatorily and predictively successful, contains, in an essential way, these probabilistic propensities. So if the theory of natural selection has at its foundation these probabilistic propensities in what sense could it be deterministic?

In a paragraph from which Millstein quotes we answer that question. "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." (Brandon and Carson 1996, p. 324). That last sentence is followed by a reference to Sober (1984, 110-115), which I will discuss shortly, but first let me elaborate on the last sentence. As philosophers use the

---

[1] Also a considerable amount of theoretical work showing that the evolution of certain traits requires the interaction of selection and drift would be misguided. See, e.g., Rausher and Englander 1987 on the evolution of evolutionarily stable strategies for habitat selection under soft selection, or Eshel and Feldman 1982 on the sex-ratio evolution.

[2] This sort of autonomy does not commit us the truth, or even the sensibility, of the following counterfactual: If quantum mechanics were deterministic then the evolutionary process would still be indeterministic. I for one have no idea how to evaluate the truth of that statement.

[3] The negative points made here do not depend on that conclusion at all. The positive ones do depend on a part of that conclusion, i.e., that the probabilistic propensities assigned to individuals are real. The autonomy of those probability values is irrelevant to the points made here.

term, a *theory* is deterministic if from a complete *state description* of a system at time *t* one can derive a complete state description of that system at some later time *t'*. (One could complicate this characterization enormously, but this suffices for present purposes.) A *process* is deterministic then if there is such a theory that truly describes it. No biologist I know thinks of natural selection as deterministic in this sense. When they do use the term 'deterministic' to describe natural selection they do so to contrast natural selection with random drift. Both processes can result in changes in gene and genotype frequencies across generational time, but our best theories of these processes differ in that theories of selection can predict the direction, as well as rate, of cross-generational change, while our theories of drift can predict only the rate, but not the direction, of such change. That, and only that, is what is meant by "deterministic" natural selection in contrast to drift.

I really do not think we could have been much clearer in the aforementioned paragraph, but evidently our reference to Sober mislead Millstein. We referred to Sober simply to give the reader a more extended discussion of the point we were making quickly. Upon rereading Sober, I'm not sure how he could have been clearer either. Millstein quotes the following sentence from Sober (a sentence we do not quote—we simply refer to pp. 110-115): "*When it acts alone, the future frequencies of traits in a population are logically implied by their starting frequencies and the fitness values of the various genotypes.*" (Sober 1994, 110, italics in the original; quoted by Millstein, 48). However, reading only two paragraphs further one finds that Sober thinks this to be false. He says, referring to the above claim, "We can now correct the characterization offered above of why natural selection is a deterministic evolutionary force. Fitness values plus starting frequencies do *not* permit the deduction of changes in gene frequencies, even on the supposition of infinite population size. They do, however, permit a probability inference of almost unbeatable strength . . ." (italics in the original, Sober 1994, 111). Of course, when we move to more realistic models with finite population size, this probabilistic inference loses strength (how much depends on the population size and the selection differentials). Not only are such models more realistic, they are much more interesting when one is concerned with the relationship between drift and selection, as Carson and I certainly were.

The theory of natural selection is not to be identified with any one of its many models. But given that each such model is probabilistic, and given that I have for a long time defended the view that the very foundation of the theory of evolution by natural selection is the explicitly probabilistic Principle of Natural Selection,[4] it seems highly dubious that I would think that the theory is deterministic. That doubt is well placed. I don't.

2. Process and Outcome.

The key to Millstein's approach to differentiating selection from drift is the distinction between *process* and *outcome*. In general such a distinction is useful. In this case, can we, as Millstein claims, distinguish between the processes of drift and selection without regard to the outcomes? As I will show, we cannot, at least not in a way that allows these two concepts to play their normal role in evolutionary theory.

Drift, as Millstein recognizes, can take many forms. She focuses primarily on parental sampling; that is, on the process by which organisms from one generation are "sampled" to be the parents of the next generation. For ease of exposition I will follow her in this. Focusing on this particular form of drift will result in no loss of generality; and in fact the major points I wish to make can be illustrated initially without delving into biology at all.

---

[4] The Principle of Natural Selection, as I have explicated it, is an instance of the Principle of Direct Inference from probability theory. I state it as follows: If *a* is better adapted than *b* in environment *E*, then (probably) *a* will have more offspring than *b* in *E*. The notion of relative adaptedness contained in this principle is defined in terms of objective probabilities of different levels of reproductive success. I argue that these probabilities should be given a propensity interpretation. The final definition of adaptedness (or expected fitness) is then the expected (in the mathematical sense) reproductive success discounted by a function of the variance, the exact form of the function depending on the nature of the variance. See Brandon 1990, pp. 9-24.

According to Millstein, drift as a process can be characterized as an *indiscriminate* sampling process. The contrast, of course, is with a *discriminate* sampling process. And the difference between discriminate and indiscriminate sampling processes is itself easy to understand. A sampling process is indiscriminate if and only if each entity in the pool to be sampled has an equal probability of being chosen. For example, if we are pulling balls from an urn, that process is an indiscriminate sampling process if and only if each ball in the urn has an equal probability of being pulled. In the biological case, imagine a population of 1,000 individuals, only 100 of which will get to mate and reproduce. If each of the 1,000 individuals has an equal probability of becoming a parent then that process is indiscriminate sampling. Less abstractly, if there are no physical differences that make the difference between those 100 who are successful vs. the remaining 900 who are not, then this parental sampling is indiscriminate.

In contrast, suppose that there are physical differences that make a difference with respect to the probability of becoming a parent. Then the sampling is discriminate and, according to Millstein, is a case of natural selection. But notice that all that is required to differentiate these processes, discriminate vs. indiscriminate sampling, is the difference between equiprobable results and results that are not equiprobable. Unfortunately that distinction does not map well onto the ways biologists differentiate drift from selection.

Again consider sampling balls from an urn. Suppose that there are 100 balls in the urn, 50 red and 50 black. Further suppose that each ball has an equal probability of being chosen on any one draw (P = .01). (The equal number of red and black together with the equal probability of each ball being chosen results in the P(red) = P(black) = .5.) A sampling here consists in pulling a single ball from the urn, recording its color, and then replacing the ball. Let's now imagine a trial of four such samplings. The most probable outcome, the one that will occur most often in a long sequence of such trials, i.e., the *mode,* is 2 red and 2 black. That outcome is also the *expected* or *mean* outcome, i.e., the one that corresponds to the overall frequency of the two types in the urn.[5] But there are four other possible outcomes: all red; all black; 3 red, 1 black; and 1 red, 3 black. Here are the probabilities of each of these possible outcomes:

Example 1

P (all red) = .0625
P (3 red, 1 black) = .25
P (2 red, 2 black) = .375
P (1 red, 3 black) = .25
P (all black) = .0625

Notice that while the expected outcome is indeed the modal outcome, it occurs in slightly fewer than 4 out of 10 trials, or conversely, outcomes that deviate from the expectation occur in slightly more than 6 out of 10 trials.

Now let us change the setup so as to change the probabilities of the results (red or black). There are two ways to do this. First we could change the relative frequency of the two colors in the urn, e.g., if we replace five of the black balls with five red, then we would get: P(red) = .55 and P(black) = .45. But in this case the sampling would still be indiscriminate, i.e., every ball would still have the same probability of being drawn. To make the sampling discriminate we need to change the probabilities of individual balls being drawn so that they are no longer equal. So let's repaint the balls using slightly sticky red paint and slightly less sticky black paint, but retaining the equal numbers of red and black. Let us suppose that this change in the physical characteristics of the two types of balls results in the following probabilities: each red ball now has P = .011 of being drawn and each black ball has P = .009 (where before both types had P =

---

[5] The expected outcome equals the overall frequency of the two types in the urn only under the assumption of equiprobability. When we discard that assumption, as we will, the expected outcome equals the sum of the products of the probabilities of the two types times their frequencies. Although we are here talking about the simple case of drawing balls from an urn, that expected value is just two complications away from my definition of expected fitness. See footnote 2 above. (The two complications are: 1. In biology reproduction is not all or nothing, there are different levels of reproductive success; and 2. The expectation needs to be discounted by the variance.)

.01). Now the process of sampling, i.e., of drawing balls from the urn, is discriminate. The probability of red is now 50 ? .011 = .55 and the probability of black is 50 ? .009 = .45. Let us again sample (with replacement) four balls. Again there are five possible outcomes and we can again calculate the probability of each:

Example 2.
P (all red) = .091506
P (3 red, 1 black) = .299475
P (2 red, 2 black) = .367538
P (1 red, 3 black) = .200475
P (all black) = .041006

Notice that the outcome, 2 red, 2 black, is still the most likely outcome, but its probability has been lowered as has the difference between it and the next most likely outcome. Also notice the *expected* outcome, 55% red and 45% black, is of course not possible with a sample size of four.[6]

One more example: here we make the red paint stickier still so that the resulting probability for each red ball getting picked is .015 and the probability for each black ball is .005. Again there are 50 red and 50 black balls. Now the probability of red is .75 and of black .25. Sampling four balls we get the following distribution of probabilities for the 5 possible results:

Example 3.
P (all red) = .316406
P (3 red, 1 black) = .421875
P (2 red, 2 black) = .210938
P (1 red, 3 black) = .046875
P (all black) = .003906

In this case, as in the first example, the expected result (75% red, 25% black) is also the modal result.

What we have done in these three examples is start with the one distribution of probabilities that corresponds to indiscriminate sampling, namely the equiprobable distribution. The next two examples move away from that distribution towards increasing differences between the probabilities for red and for black. Probability Theory allows for an infinite number of distributions of the probabilities for each of the 100 balls being drawn. The only constraint is that these 100 probability values sum to 1. But if we continue moving in the direction of our examples, there is a distribution of maximal probability difference. We will discuss it shortly, but first let ask the question: has the change in our experimental setup, a change from indiscriminate sampling to discriminate sampling, resulted in a qualitative change with respect to drift?

We cannot satisfactorily answer that question until we jettison Millstein's approach to defining drift as a process. I will argue explicitly for this move presently, but for now let us, tentatively, adopt an outcome-oriented conception of drift, namely, that *drift is any deviation from the expected result due to sampling error.*[7] If that is what drift is, then the difference between example 1 and examples 2 and 3 is quantitative, not qualitative. Consider examples 1 and 3. In example 1 the probability of the expected result is .375. That means that in a long sequence of such trials, approximately 37.5 % would yield exactly the expected result while 62.5 % would yield a result that deviated, more or less, from the expectation. In example 3 the expected result will occur more frequently in a long series of trials, approximately 42% of the time. Both setups will regularly lead to results that deviate from the expectation; the difference between them is quantitative, not qualitative.

---

[6] Brandon and Carson (1996) discuss examples such as this where drift is forced to occur. I will not stress that point here.

[7] This is the characterization of drift that Brandon and Carson (1996) adopt. It is standard in the biological literature, see e.g., Roughgarden, (1979, chap. 5).

Let us be more explicit about the analogy between drawing balls from an urn and biology. Getting drawn from the urn corresponds to parental sampling, i.e., becoming one of the organisms that reproduces. The equiprobable distribution represents the case where all organisms in a population have equal fitness. Differences in probabilities represent selection differentials. So we have just illustrated the well known result form biological studies of drift: *everything else being equal, the greater the selection differentials the smaller the expected effect of drift.*

With respect to drift the equiprobable distribution does not stand out among the infinite number of possible distributions. However another sort of distribution, what I will call distributions of *maximal probability difference* (MPD for short), does. This set of distributions is qualitatively distinct with respect to drift. In our simple urn case, where there are 100 balls in the urn, the MPD occurs when one of the balls has a probability of 1 of being drawn, and all other balls have a probability of 0. Thus here there are 100 such distributions. Given the nature of our setup, only one ball could have a probability = 1 of being drawn. One can imagine different setups. Suppose we were to simultaneously grab four balls out of the urn. Then a distribution where four particular balls have a probability of 1 of being drawn while the other 96 have a probability of 0 is an MPD distribution. In general we get the maximal probability difference when all the probabilities equal either 0 or 1, and at least some equal 0 and some 1.[8] In the biological case a maximal fitness value, which is an expected number of offspring, say 15.75, is normalized to 1. So the biological analogue of MPD is where all the fitnesses in a population equal either 1 or 0, again with at least some of each value. Call this a population with *maximal fitness difference.*

Distributions of MPD are qualitatively distinct with respect to drift. Drift *cannot* occur with such a distribution. Balls are sampled from an urn, or organisms are sampled from a population, but with all probabilities equaling either 0 or 1 there can be no sampling error, no deviation from the expected result. The expected result occurs with probability 1. Notice that it follows from what was said above about drift, that with the MPD distribution, the expected effect of drift is minimized. That is true, but under-informative in that it masks the qualitative difference between such distributions and the infinite number of other possible probability distributions. With an MPD distribution, drift is not just highly unlikely, it is impossible.[9]

What about natural selection? If it is to be distinguished from drift, then we cannot identify it with the process of discriminate sampling, since, as we have seen, that process does not differ qualitatively from indiscriminate sampling with respect to drift (except in the extreme case of MPD). To get anywhere in our exploration of drift we were forced to an outcome based characterization of drift. A similar move will be needed for natural selection.

I have characterized natural selection as follows: natural *selection is differential reproduction that is due to differential adaptedness (or fitness) to a common selective environment* (see, e.g., Brandon 1990, chaps. 1 & 2). This identifies an outcome—differential reproduction—but it has an explicitly causal component to it as well. Natural selection is not just any case of differential reproduction, but is those cases that are *due* to differential adaptedness (or fitness). The whole point of the propensity interpretation of fitness, or adaptedness (Brandon 1978, Mills and Beatty 1979), at least as I have developed it, is to provide the conceptual machinery adequate for an explanatory theory of natural selection. Among other things this requires being able to distinguish those cases of differential reproduction due to differential adaptedness from cases of differential reproduction that are drift. This, in essence, is the so-called "tautology problem".

With that conception of natural selection in mind let us return to the continuum of probability distributions discussed above. As we saw, the equiprobable distribution allows for drift (it does not make drift

---

[8] For any setup of the sort we are considering there will be exactly one equiprobable probability distribution. In our simplest urn drawing model there are exactly 100 distributions of maximal probability difference, different setups will have different combinatorial possibilities, but as long as the number of objects is finite there will only be a finite number of MPD distributions. The import of this will be discussed later.

[9] This is in contrast to the case of infinite population size, where drift is *highly* unlikely, but not impossible. See discussion in Sec. 1 above.

necessary). But selection *cannot* occur under such a regime, because selection requires differential fitness or adaptedness—in our urn model that translates into probability differences. So although the equiprobable distribution is not qualitatively distinct from the infinity of other possible distributions with respect to drift; it is qualitatively distinct with respect to selection. Selection requires fitness differences, so this is the single probability distribution that precludes selection. Put in other words, with indiscriminate sampling, selection cannot occur.

On the other hand, the set of MPD distributions forces selection to occur, while, as we saw above, precluding drift. Selection must occur with an MPD distribution because, to use our simple model, balls with a probability of 1 of being drawn will be drawn, and those with a probability of 0 will not be, otherwise those probability values are not truly 1 and 0. Just to make the analogy explicit: organisms with a fitness of 0 will not reproduce. Organisms with a fitness of 1 will. This is quite clear if we make the simplifying assumption that reproduction is all or nothing, i.e., an organism either reproduces exactly *n* offspring or 0 offspring. Then the urn model fits exactly. But when we allow for the possibility of producing different numbers of offspring, as real organism are want to do, then two organisms could achieve the same expected number of offspring in different ways. For example, type A could always have 2 offspring, while type B could sometimes have 1 and sometimes 3, but with a mean of 2. Type C has 0 fitness. So, if you have read Gillespie (1977), you will know that selection can occur here between A and B (with A being favored) even though this seems to be the analogue of an MPD distribution.[10] But, in fact it is not. When appropriately discounted, the fitness of B is lower than A, and so this is not an MPD-type distribution.

Notice the asymmetry between drift and natural selection vis-à-vis the two poles of our probability distribution continuum. At the equiprobable end *drift is possible and selection is impossible.* At the MPD end *selection is necessary (not just possible) and drift is impossible.*

Putting that slight asymmetry aside, could we use these two poles to make a process-oriented distinction between selection and drift? That is, could we equate the process of drift with indiscriminate sampling and selection with *maximally* discriminate sampling? No, for two reasons. First, from a mathematical point of view, the equiprobable distribution and the finite set of MPD distributions represent an infinitesimally small fraction of the possible probability distributions. Were we to categorize sampling with an equiprobable distribution as drift and sampling with an MPD distribution as selection, we would be leaving all but an infinitesimally small fraction of cases uncategorized. Biologically things are probably worse still for this sort of strategy. Has any real biological population in the history of life on Earth ever realized one of these two extremes? I will not pretend to know the answer to that question, but I would not be surprised if it were no. Certainly it is safe to say that no population ever studied has met these conditions. Which means that the biological action is in the distributions that fall between the equiprobable and the MPD. It is here that we need to be able to distinguish drift from selection. And we can.

3. Moths, Reference Classes and Selective Environments.

Every day moths get eaten by birds. Now I'm as fond of moths as the next person, but generally this fact about moths does not upset me. However, in an otherwise insightful article published almost 20 years ago John Beatty (1984) unleashed into the philosophical literature a population of moths that confused him and that seem to have caused considerable confusion since. It is time to clarify this confusion and to let these poor moths rest in peace.

Though fictional, Beatty's moths are modeled on Kettlewell's well-known studies (Kettlewell 1955, 1956). They come in two forms--light and dark--and inhabit a forest that contains light and dark trees in a ratio of 40:60. It is clear that Beatty assumes, though he never explicitly state this, that neither form of the moth behaviorally discriminates between light and dark trees. In other words, the moths land at random on trees and so have a 40% chance of landing on a light tree and a 60% chance of landing on a dark tree. Birds then prey on the moths based on their conspicuousness relative to their background. Beatty then assumes that the fitness distributions of the two types overlap in that environment (195).

---

[10] For further discussion of this see Brandon 1990, pp. 18-22.

Now Beatty asks us to suppose that, *by chance*, more of the dark moths happen to land on light trees than would be expected, but not vice versa, i.e., the light moths behaved more or less as expected. As a result the frequency of the light moths increases in the next generation, contrary to expectation. What are we to make of this? Beatty says:

> Is the change in frequency of genes and genotypes in question a matter of natural selection, or a matter of random drift? That is, is the change in question the result of sampling discriminately or indiscriminately with regard to fitness differences? It is not easy to maintain that the sampling was entirely indiscriminate with regard to differences in survival and reproductive ability. At least it is difficult to maintain that the death by predation of *conspicuously* dark moths in this environment is indiscriminate sampling, whereas the death of *conspicuously* light moths in the same environment is selection. On the other hand, it is also difficult to maintain that selection alone is the basis of the change. At least, it is difficult to maintain that the fittest were selected. (195-96)

Beatty concludes with the following:

> In other words, 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, especially where *improbable results of natural selection* are concerned (196).

This statement is so qualified, so wishy-washy; that one might think it would be hard to disagree with it. But as Millstein correctly notes, were we to accept it we would be forced to say that the large amount of work that has gone into the selectionist-neutralist debate is based on a conceptual mistake. Millstein's solution, as we have seen, is to say that outcomes don't matter and that so long as the sampling was discriminate then this is selection. And so she unambiguously classifies Beatty's case as selection. We have also seen that her approach is a non-starter.

Notice that Beatty's problem would not have arisen had the woods been 100% dark or 100% light. The problem arose because dark moths by chance landed on light trees more often than they should have, and this problem could not have arisen had the environment been homogeneous with respect to background color. This matters because fitness is obviously relative to an environment and Beatty is unsure as to which environment he should relativize these unlucky moth's fitness. Should it be the environment that is characterized by the statistical distribution of the colors of the trees in the woods (60% dark, 40% light), or the environment that the moth happened on just prior to its demise (light)? This is a genuine biological problem and one of the utmost importance to the theory of evolution by natural selection. But before sketching its solution, let me point to its analogue in probability theory.

The dominant objectivist interpretation of probability during the 20th century was the limit-relative-frequency interpretation. According to it, the probability of attribute *A* just is the limit of the relative frequency of *A* in an infinite series of trials. But which infinite series of trials? The problem becomes particularly acute when we try to assign a probability to a single case. For instance, suppose we are about to flip a particular US quarter and want to know the probability of its landing heads on that particular flip. Unfortunately that particular flip belongs to an indefinite number of potentially infinite sequences of events, e.g., flips of any sort of coin, flips of any sort of US coin, flips of US quarters, flips that occur on Tuesdays, flips that start with head side up, and so on. To which sequence do we assign it? That is the *reference class problem.* If we could unambiguously assign it to some particular sequence, then we could just assign the limit of the relative frequency of heads in that sequence as the probability in that particular case. But if there is no objective way of assigning the event to one particular sequence, then it seems that we cannot give an objective probability to the single case.

One might think that the propensity interpretation of probability fares better with respect to the problem of assigning a probability to a single case since according to it, a probability just is some set of physical features of an object (e.g., coin) in a particular *chance setup* (e.g., tossing devise). But Weslely Salmon has argued (e.g., Salmon 1970, p. 38-40) that the reference class problem arises just as acutely when we try to specify the nature of the chance setup.[11] Does any toss of any coin count as part of our chance set up? Or just tosses of US quarters? Or, etc.? A Pyrhric victory could be achieved by saying that it is just this single toss that counts as our chance setup. But then we could never compile reliable statistics to validate any claim about the probability.

Hans Reichenbach, perhaps the leading exponent of the frequency interpretation, thought the reference class problem lacked a fully adequate solution. The best he could do was to recommend that we adopt as the reference class "the narrowest class for which reliable statistics can be compiled" (Reichenbach 1949, p. 374). But this is a pragmatic recommendation that seeks to maximize two variables--narrowness and reliability--that are at odds with each other. The narrower our class the less reliable our statistics, and the more reliable we seek to make our statistics the less narrow our reference class. Thus in the end Reichenbach refused to use the word 'probability' to apply to single events, instead in such cases he used the word 'weight', and argued that 'probability' applied literally only to sequences (Salmon 1970, p. 41).

Salmon, a student of Reichenbach, succeeded where Reichenbach failed. To solve the reference class problem he developed the concept of *homogeneity*. A class is homogeneous with respect to some attribute (or outcome) *A* if and only if there is no *place selection* of the class that is statistically relevant to *A*. The concept of *place selection* comes from von Mises, and is any partition of the reference class *not* made in terms of the attribute (or outcome) in question (Salmon 1970, pp. 42-43). In other words, we can partition a sequence of tosses of a coin any way we like (e.g., tosses of this quarter vs. some other coin; tosses on Tuesday vs. on other days; tosses that start with the head up vs. those that start tail up, etc.) just so long as we don't do so in terms of the outcome itself. Thus tosses that yield heads vs. tosses that yield tails is not a place selection and is not allowable. (Obviously that partition is statistically relevant, but it is cheating.) Salmon then adopts the following *reference class rule:* "choose the broadest homogeneous reference class to which the single event belongs" (Salmon 1970, p. 43).

Before moving on let me note one important point. The notion of homogeneity as defined above is an ontological concept not an epistemic one. One can, and Salmon does, define *epistemic homogeneity*--a class is epistemically homogeneous with respect to attribute *A* if and only if there is no *known* place selection of it that is statistically relevant to *A* (Salmon 1970, p. 44). But we are here interested in objective homogeneity.

I have treated the biological problem that Beatty stumbled on elsewhere (see Brandon 1986 and 1990, chap. 2, and Antonovics, Ellstrand and Brandon 1988), and so will be brief here. When dealing with natural selection the relevant notion of the environment is what I've termed the *selective environment.* Intuitively, the selective environment is the arena within which selection occurs. It is measured in terms of the relative actualized reproductive success of two or more competing types. What it is defined as, i.e., what is being measured, is the relative fitnesses (expected reproductive success) of two or more competing types across space, or time, or some other suitable dimension. For the sake of concreteness, consider a spatial scale. A region of space is selectively homogeneous with respect to types $A_1, A_2, \ldots A_n$ if and only if the relative fitnesses of types $A_1, A_2, \ldots A_n$ are constant within that region. In other words, a region of space is selectively homogeneous if and only if it cannot be partitioned by means of a place selection in a way that is statistically relevant to the relative fitness of the competing types.

A region of space is *heterogeneous* if it is not homogeneous. In other words, a spatial gradient, say a transect up a mountainside, is selectively heterogeneous if and only if the relative fitness of the competing types changes along the transect. Empirically this would be indicated by a genotype-environment [G x E] interaction. We have lots of evidence from ecological genetics that selective environments are indeed heterogeneous. But, of course, the scale of such heterogeneity depends on all sorts of ecological factors; it may be a matter of millimeters (e.g., in the annual plant *Erigeron annuus*[12]) or of hundreds of kilometers.

---

[11] I don't completely agree with Salmon. I think that the problem is less serious for the propensity interpretation. However, Salmon is surely right that the propensity interpretation still faces the reference class problem.

[12] See Straddon and Bennington 1998.

What about Beatty's moths? The case is fictional and Beatty does not supply enough information to answer that question. But we can fill in details in ways that would allow us to address the question. First, we can describe the case where the whole wooded area in question is selectively homogeneous. That may sound counterintuitive. Surely, one might think, a light colored tree is not part of the same selective environment as a dark tree. The "relative fitness" of light and dark moths differs dramatically on these two trees. But this objection mistakenly treats fitness as an instantaneous property of an organism rather than as a property of the organism's life history.[13] Here is where behavior becomes crucial. Plants stand still and wait to be counted; moths do not. So it is quite conceivable that in some particular populations of *Erigeron* one spot in a field is in one selective environment, while 5 millimeters away is a different selective environment. That is not conceivable for moths. Whatever the spatial scale of selective heterogeneity for moths it is going to be considerably greater than that of a plant like *Erigeron*. Moths fly around and land on many different trees. Their probability of being devoured by a bird depends on the match, or lack thereof, of their color and the statistical average color of the background that they create by their behavior. Thus if the two tree types are distributed randomly about the woods and both types of moths show no behavioral preference for one type of tree over the other, then the woods in question are selectively homogeneous. This is fully consistent with Beatty's story. Then this is a case where both selection and drift occurred. (We cannot quantify the effect of each since the story is not quantitative.)

On the other hand we could certainly fill out the story in other ways so that the woods are selectively heterogeneous. Suppose the size of the woods is considerably greater than the average size of the territory a moth typically inhabits. Suppose further that the two tree types are not distributed randomly about the woods but rather are patchy, with the ratio of light to dark differing significantly among patches. Finally suppose that the size of these patches is on the order of, or greater than, the size of typical moth habitats. Then the woods are selectively heterogeneous, and to fully understand the evolutionary dynamics of this population one would need a model of what I have termed *compound selection*. Compound selection involves selection *within* selectively homogeneous environments and distribution *among* such environments (see Brandon 1990, pp. 71-77). But in any realistic scenario there will again be both selection and drift.

The details of Beatty's story are under-specified, and the appropriate conceptual tools are not utilized, so it is impossible to say anything definitive about it. Let those moths rest in peace.

4. Process and Outcome Redux.

In the urn case there is a process of sampling balls from the urn. In the biological case there is a process of sampling from among the members of one generation to get the parents of the next. In both cases the sampling may be governed by the single equiprobable distribution, in which case it is indiscriminate, or, much more likely, it may be discriminate. Ignoring MPD distributions, in either case the result of the sampling may, or may not, be the expected one. That is, there is a single process, sampling. In any given case, whether or not that process generates the expected result can be known, but can be known only after the fact. (Of course, prior to the sampling process we can make probabilistic prediction about this.) Thus we can distinguish drift from selection, using the conceptions of drift and selection tentatively suggested above, but only after the fact. (While again we can make good probabilistic predictions prior to the fact.) This means that the outcome, deviation or lack thereof from expectation, is a necessary component of our conceptions of drift and selection. It also means that, as we saw above, characterizations of sampling processes alone will be incapable of making the selection/drift distinction.[14]

---

[13] Biologists regularly measure the relative fitnesses of organisms during one part of their life cycle, and as we make that period shorter and shorter we would converge on instantaneous fitness. Such a measure is sufficient for understanding the ecological process of selection that occurs during that part of the life cycle. But that measure corresponds to evolutionary fitness only under the assumption that selection occurs *only* during that part of the life cycle.

[14] This implies that Sober's (1984) description of evolutionary theory as a theory of forces, with selection and drift being separable forces, must be wrong. Selection and drift are products of the exact same process

I've just claimed that we can know, after the fact, whether or not drift was involved in the transmission of gene and genotype frequencies from one generation to the next. Moreover, I claim, we can quantify the extent to which it is involved. Granted this will be practically difficult in real biological cases, but is it possible in principle?

First consider the common analogy between the evolutionary process and information theory. In information theory we have a source (S), a communication channel and a receiver (R). A message (information) is sent from S to R via the communication channel. Noise in the channel is a measure of the independence of the received message from the sent message, in other words, a measure of the difference between the received message from that which was sent. Can we know, quantitatively, this difference? Yes. After the transmission we can compare the original message to that which was received, and we can quantify the extent to which they differ. Prior to the transmission of some particular message we can quantify the "noisiness" of the channel, either based on past experience, or on some more direct knowledge of the processes involved in transmission. And given the level of noisiness we can predict, probabilistically, the extent to which our message will be transformed. But whether or not noise has an effect in a particular case can be known only after the fact. There is here a single process, transmission over a noisy channel, which sometimes does, and sometimes does not, result in a transformation of the message.

How is this analogous to the process of evolution by natural selection? The signal to be sent is the product of genotypic frequencies and their associated fitnesses (where sampling is discriminate), or in the limiting case of indiscriminate sampling, is just the genotypic frequencies themselves. "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." (Roughgarden 1979, p. 58, emphasis added). Just as in the information theory case, here too we can, after the fact, know qualitatively whether or not drift has occurred, i.e., know whether or not the next generation's genotypic frequencies match those predicted by the past generation's frequencies and fitnesses. And if they do not match we can quantify this lack of fit, i.e., we can quantify the effect of drift. As should be clear by now, this *post hoc* quantification of drift requires knowledge of the fitnesses of the individuals of the last generation.

There are two obvious ways to object to the above account of the difference between selection and drift. First, one could deny the existence of individual fitnesses. And one could deny the existence of such things for at least two quite different reasons. One could pursue an instrumentalist stance towards fitness and say that they are not real, but are merely useful instruments. But then one would probably still accept my approach to differentiating selection and drift--because it is useful--and just not give it the realist interpretation I give.[15] Or one could deny that the probabilities that make up an individual's fitness attach to individuals at all. Instead, on this view, the probabilities that play a role in evolutionary theory attach only to ensembles of individuals.[16] The major point of Brandon and Carson (1996) was that a realist stance towards individual fitness was justified. It is my hope that this article has added further evidence for that claim.

Second, one might admit the possible existence of individual fitness, but deny the possibility of our knowledge of fitness. My comment on this stratagem applies equally to the denial of the reality of fitness as well. It is that there is a whole field of evolutionary biology, ecological genetics, which has as its basic goal the measurement of individual fitnesses. Given that biologists have successfully measured individual fitness in the field and in the lab since, at least, 1898 (see Weldon 1898) I'm not much moved by the claim that we cannot measure it. In my view, the actual is always possible.

---

and so cannot be separate forces (keeping in mind the distinctions we can make between MPD distributions and the equiprobable distribution). Space limitations prevent me from pursuing that further here.

[15] Footnote to Rosenberg etc.

[16] Footnote to Ariew et al.

5. Conclusions.

I have offered definitions of selection and drift that are outcome-oriented and so can be definitively applied only after the fact. *Drift is any deviation from the expected levels of reproduction due to sampling error. Selection is differential reproduction that is due to (and in accord with) expected differences in reproductive success.* As we have seen, a process-oriented approach cannot make the appropriate distinction between selection and drift because both result from the same process—sampling. Thus Millstein gets things exactly wrong.

On my approach drift and selection are outcomes that can be distinguished given the appropriate probabilities. Assigning probabilities to individuals' lives, deaths and levels of reproductive success depends, at least, on solving the reference class problem with respect to biology. That has been solved with the concept of selective environmental homogeneity. In the last section I briefly mentioned, but did not pretend to fully address, other potential philosophical objections to my approach. All I can say here is that one better hope that my approach is the correct one, otherwise the explanatory power of evolutionary theory is greatly diminished.[17]

---

[17] See Brandon (forthcoming).

Comments on Robert Brandon's paper

Alan C. Love

Brandon's response to Millstein on the distinction between natural selection [NS] and random drift [RD] is nested in a rapidly growing literature that involves a host of related topics, including evolutionary indeterminism and fitness.[1] My aim is not to defend or criticize Brandon's specific response to Millstein but rather to put pressure on related themes less prevalent in their debate (and the wider literature), such as the individuation of causal processes, contrastive explanation, the desire for theoretical generality, and the use of fictionalized examples.

Millstein's aim is to keep NS and RD *conceptually* distinct, taking this distinction as a prerequisite for ascertaining whether or not either has been more prevalent *or* more important in the history of life. Toward this aim she demarcates the *processes* of NS and RD from the *outcomes* of these processes. (Millstein 2002, 34-38) NS is a *discriminate* sampling process, where the discrimination depends on fitness differences playing a causal role. RD is an *indiscriminate* sampling process, where the lack of discrimination is because fitness differences do not play a causal role. (Cf. Beatty 1984) NS and RD are conceptually distinguishable as processes because NS involves sampling with respect to the causal role of fitness differences while RD does not, regardless of whether *outcomes* are in accord with expectation values. Ignoring outcomes is desirable because the same process can lead to different outcomes and different processes can generate the same outcome.[2] That actual outcomes of the causal processes of NS and RD in natural populations are difficult to distinguish is an epistemological worry that does not infect our ability to conceptually distinguish them as processes.

Following an earlier paper (Brandon and Carson 1996), Brandon's response is straightforward. In short, NS and RD cannot be distinguished as processes because they are both instances of the same process, namely sampling.[3] The modifiers "indiscriminate" and "discriminate" misconstrue the biologist's aims since probability theory marks this distinction only via the presence or absence of equiprobable results. Quantitative deviations from equiprobability are not enough to distinguish NS and RD. As he remarks: "If [NS] is to be distinguished from drift, then we cannot identify it with the process of discriminate sampling, since, … that process does not differ qualitatively from indiscriminate sampling with respect to drift …we [are] forced to an outcome based characterization of drift." (Brandon 2002, 13) Millstein's dismissal of outcome also misses the riches of evolutionary theory in predictively capturing the rate *and* direction of NS in contrast to only being able to discern the rate of RD.

---

[1] For example, some have argued that RD is best construed as claim of ignorance about causal processes, a place-holder for the future discovery of what is actually going on or an indicator of measurement error. (Graves, Horan and Rosenberg 1999; Horan 1994; Rosenberg 1988, 1994) This approach stems from a commitment to evolutionary processes being asymptotically deterministic (?), with the probabilities reflecting our epistemic shortcomings rather than being objective features of the world, thereby implying that evolutionary phenomena are for all practical purposes deterministic. (The claim is *not* that evolutionary processes *are* deterministic.) Others have attempted to locate the source of stochasticity *in nature itself* (thereby implying that evolutionary processes *are* indeterministic), due to quantum mechanical percolation (Brandon and Carson 1996; Stamos 2001), thermodynamic processes (Rosenberg 2001), random foraging behavior (Glymour 2001), or the finite size of natural populations (Brandon 1990; Brandon and Carson 1996; Sober, 1993 [1984])

[2] Millstein's account is similar to John Endler's: "the definition of the process of natural selection takes no account of the details of its outcome." (Endler 1992, 222) In Endler's terminology, the following condition is fulfilled for NS but not for RD: "a consistent relationship between [a variable] trait and mating ability, fertilizing ability, fertility, fecundity, and/or survivorship (fitness variation)." (Endler 1992, 220; cf. Endler 1986, ch. 1)

[3] Shanahan accepts this and develops a view that NS and RD are idealized concepts found as endpoints on a continuum of biological sampling processes, where pragmatic factors determine our distinction between them. (Shanahan 1992)

NS is central, if not at the center of contemporary evolutionary theory [ET], and the notion of RD is critically related to it.[4] What exactly are they? Almost everyone agrees that NS is a *causal process*, although there has been some recent dissent. (Walsh 2000; Walsh, Lewens and Ariew 2002)[5] RD should also be considered a causal process (or *not*) for similar reasons. (Beatty 1992; Millstein 1996). Sober's influential treatment cashes out causal process as *force*, articulating ET in terms of source and consequence laws analogous to classical mechanics. (Sober 1993 [1984]) Others want to retain the causal process talk but eschew "force" terminology. (Brandon 1990; Brandon and Carson 1996; Endler 1986; Hodge 1987, 1992; Millstein 2002) Historically, "causal process" is interpretable when connected with the *vera causa* ideal of the 19th century (Hodge 1987, 1992), but in the present context some further explication is required, either in terms of what distinguishes a causal process from a *non*-causal process (Salmon 1984, 1998) or in terms of how to individuate one causal process from another. (Glymour 1998) It is this latter issue that is in view, contrary to any worries about the nature of causation. That is, how does one distinguish the *causal* process of NS from the *causal* process of RD?

Both Millstein and Brandon (and others holding to a causal interpretation) agree that NS and RD refer to a *variety* of distinct causal processes. The concept of NS can refer to sexual selection or nonsexual selection, the latter encompassing a broad range of phenomena including mortality, fertilizing ability, fertility, and fecundity. (Endler 1986, 1992)[6] NS also is understood to have different effects on populations (e.g., habitat diversity and choice or trait covariance) as well as on trait frequency distributions (directional, disruptive, and stabilizing) that may be continuously varying, discontinuously varying, or have quantitative thresholds. (Endler 1986, ch. 1) The concept of RD is at least decomposable into indiscriminate parent sampling, indiscriminate gamete sampling, the founder effect, the random distribution of different genotypes over heterogeneous selective environments, and rate fluctuation in processes such as migration or mutation. (Beatty 1992; Brandon 1990, ch. 2) Neither author attempts to explicitly consider all these processes when making the distinction between NS and RD, and yet their results claim as much. The issue at stake is whether or not the results of distinguishing particular causal processes labeled as NS and RD can be *generalized* to a distinction between NS and RD that leaves these causal processes unspecified. Sewall Wright labeled selection a "wastebasket category" when generally construed as any determinate direction and amount of change in gene frequency. (Wright 1955, 20) Brandon and others have worried about the issue of generalization (Brandon 1990, 140-2; Glymour 1999; Lennox and Wilson 1994; Shanahan 1992), and I will return to it in my concluding remarks.[7]

Another issue related to the distinction between NS and RD is whether or not explanations in ET are *contrastive*. I think it is safe to say that in ET the aim to conceptually distinguish NS and RD results from a desire to say things like, "NS accounts for the evolutionary changes in this population *rather than* RD," in our explanations of actual outcomes, even if both processes were in operation.[8] In practice, population biologists have specific causal processes in view. The gap between a general NS/RD distinction and a specific contrast between mortality selection and indiscriminate parental sampling may be greater than previously suspected. But is the contrast class a double? Why not a triple or some *n*-tuple? The typical factors named for causing a departure from Hardy-Weinberg expectations are NS, RD, migration, and mutation (among others). Why have we not worried about distinguishing *fecundity* selection from *im*migration? Notice that this is not a worry about whether or not we can represent these

---

[4] The attempts to distinguish NS and RD *in* "ET" are usually carried out in relation to population biology. Advocates of RD as a reflection of ignorance argue from the commendable desire to seek out the actual causal structure of evolutionary processes. (Horan 1994; Rosenberg 1994) (This move also has historical precedent. See Beatty 1992; Gigerenzer *et al*. 1989, ch. 4.) Notions of chance in other explicitly evolutionary biological disciplines, such as paleobiology, are not usually considered (but see Millstein 2000). The inability of population genetics to capture accurately or robustly the historical causal structure of evolutionary processes has recently been reviewed by Lewontin. (Lewontin 2000; cf. Horan 1994) We should be cautious of what exactly ET is composed of, and qualify that what we are really aiming at is a distinction between NS and RD in population biology with respect to changes in gene frequencies. Once made, its significance for ET is a further question.

[5] The attempt to understand natural selection as a "law" is no longer fashionable. (Cf. Byerly 1983; Reed 1981)

[6] Here I am focusing on "subprocesses" rather than "sequential" processes, such as phenotypic selection and evolutionary response. (Cf. Endler 1986, 1992)

[7] My point is *not* about a NS/RD distinction in *theory* versus a NS/RD distinction in *practice*. The issue is at what *level* of detail in ET a NS/RD distinction is drawn (hence the issue of *generality*), not about the pragmatic ability to recover distinctions among causal processes as one finds in fitness component analysis.

[8] In fact, I take it be an *explanatory presupposition* in ET, *sensu* Hitchcock 1999.

variables in the math of theoretical population genetics (which we can). If we find this worry tractable (e.g., immigration does not typically result in evolutionary change), should we be including other distinct causal processes as part of the contrast class? Is evolutionary change in a population due to mortality selection *or* indiscriminate gamete sampling *or* disruptive selection on fecundity *or* a founder effect *or* sexual selection on mating ability *or…*? Why are we not more worried about distinguishing one causal process from another among a heterogeneous class of processes, especially if our aim is to isolate relevant causes of differential reproduction? (Cf. Horan 1994, 82) The outcomes of these processes are also variable, providing further potential for distinguishing NS from RD in actual populations. The delineation of a contrast class is also directly related to modeling causal processes, including the choice of appropriate null models. (Wimsatt 1987) It has already been shown that this is not straightforward for NS and RD. (Beatty 1987)

Related to the issue of specifying appropriate contrast classes is whether or not the statistical explanations given in ET can be contrastive at all.[9] Until recently (Glymour 1998; Hitchcock 1999), statistical explanations were purchased at the cost of constrastiveness. (Salmon 1984) Bruce Glymour's account of contrastive explanation adequacy provides for individuating the etiologies of causal *traces* via their sequence of causal relations. (Glymour 1998, 455-460) Causal trace *kinds* are causal *processes*. Causal *mechanisms* are composed of all relevant causal processes in the production of some outcome and effect frequencies can be determined by the component causal process frequencies. Glymour's framework also allows us to make contrastive explanations that adduce relative importance, not just presence or absence, determining which, if any, particular causal process predominates to produce an outcome. No one else has offered a competing account that provides this level of resolution for individuating causal processes.

Connected with my worry about ignoring the distinctness of causal processes subsumed under the labels NS and RD is a concern about the use of fictitious examples. This charge is easier to level here because Brandon has consistently kept actual examples in view, such as the case of heavy metal tolerance in different species of grass. (Brandon 1990) In the literature there are gray and brown squirrels with color discriminating or colorblind predators (Millstein 2002), long and short necked giraffes being poached (Rosenberg 1988, 1994), idealized peppered moth scenarios (Beatty 1984),[10] pink and yellow snails with the latter having greater sunlight tolerance (Millstein 1996), and genetically identical twins, one of whom is zapped by lightning. (Beatty 1984; Shanahan 1992) Although there is a role for simplified examples to reveal important aspects of a problem (cf. Hodge 1987, 251-2), there is insufficient causal detail in all of these examples. Philosophers generate conflicting intuitions by implicitly or explicitly specifying these details, which then better elucidates what distinct kinds of causal processes might have been occurring. Brandon's strategy in Section 3 of his reply to Millstein is to deploy further distinctions about the causal role of the environment (cf. Brandon 1990, ch. 2) to resolve ambiguities in Beatty's peppered moth scenario. (Beatty 1984) (Millstein's reconstruction of Beatty's idealized peppered moths (Millstein 2002) and her challenge to Rosenberg's account of drift in the giraffe scenario (Millstein 1996) are similar.) Shanahan puts it well: "I suspect that the appearance of a conceptual difficulty stems in part from the very real practical difficulty of correctly analyzing the moth example in the absence of the necessary causal information, not from any essential conceptual indistinguishability between drift and selection." (Shanahan 1989, 488) But we should also be worried about selecting balls out of urns or a series of coin tosses; these examples are meant to suppress the possibility that further details of causal structure are at issue, such as the micro-causal description of the table that a coin is tossed upon. In conclusion, rather than adjudicating between Brandon and Millstein, I want to stress that *both* are operating with a very abstract notion of process (and outcome) in attempting to draw a fully general distinction between NS and RD. It is a further issue to what degree my comments provide ammunition to push their debate to the next level. I have been advocating one side of a set of tradeoffs with respect to the explanatory ideals of unification and mechanism, succinctly stated by Lennox and Wilson.

"An overly abstract theory may obscure differences between types of causal processes underlying evolutionary phenomena, differences that are relevant to a satisfactory explanation of those phenomena. On the other hand, if one distinguishes the types of causal processes underlying the phenomena and exploits these differences in giving

---

[9] As an aside related to *explanation*, upon returning to Sober's classic discussion of evolutionary theory (Sober 1984) it is clear that a number of aspects deserve revisiting but are beyond the scope of this commentary. (See, e.g., chapters 4 and 5)

[10] Further reason to be chary of peppered moth scenarios derives from the reevaluation of Kettlewell's experiments. (Coyne 1998; Hagen 1999; Majerus 1998; Millar and Lambert 1999; Rudge 1999)

explanations, it may no longer be clear whether the explanations fall within the domain of a single, general theory." (Lennox and Wilson 1994, 79-80; see also Shanahan 1992)

If we recall the difference between ET and the world ET describes (recently stressed in Rosenberg 2001; cf. Graves, Horan and Rosenberg 1999; Rosenberg 1988, 1994), generalized accounts of NS and RD that abstract away from the details of actual causal processes upon which the NS/RD distinction is made increase the distance between epistemology and ontology. In order to keep ET close to actual evolutionary processes and outcomes, generalization may have to be sacrificed in order to capture the necessary detail to provide accurate contrastive explanations of causal processes generating evolutionary change. This may prove to be promising ground for making *conceptual* distinctions between NS and RD, as well as assessing the direction, magnitude, and relative significance of these *processes* in the history of life (i.e. the *outcomes*).

References

Beatty, J. (1984), "Chance and Natural Selection", *Philosophy of Science* 51: 183-211.
---. (1987), "Natural Selection and the Null Hypothesis", in J. Dupré (ed.) *The Latest on the Best: Essays on Evolution and Optimality*. Cambridge, MA: MIT Press, 53-75.
---. (1992), "Random Drift", in E.F. Keller and E.A. Lloyd (eds.) *Keywords in Evolutionary Biology*. Cambridge, Mass.: Harvard University Press, 273-81.
Brandon, R. (1990), *Adaptation and Environment*. Princeton: Princeton University Press.
---.2002. "The Difference Between Drift and Selection: A Reply to Millstein", Paper presented at the Biology and Philosophy Workshop, University of Pittsburgh, March 23-24, 2002.
Brandon, R. and S. Carson (1996), "The indeterministic character of evolutionary theory: no "hidden variables proof" but no room for determinism either", *Philosophy of Science* 63: 315-337.
Byerly, H.C. (1983), "Natural Selection as a Law: Principles and Process", *American Naturalist* 121(5): 739-745.
Coyne, J.A. (1998), "Not black and white", *Nature* 396: 35-36.
Endler, J.A. (1986), *Natural Selection in the Wild*. Princeton: Princeton University Press.
---. (1992), "Natural Selection: Current Usages", in E.F. Keller and E.A. Lloyd (eds.) *Keywords in Evolutionary Biology*. Cambridge, Mass.: Harvard University Press, 220-24.
Gigerenzer, G., Z. Swijtink, T.M. Porter, L. Daston, J. Beatty and L. Krüger (1989), *The Empire of Chance: How probability changed science and everyday life*. Cambridge: Cambridge University Press.
Glymour, B. (1998), "Contrastive, Non-Probabilistic Statistical Explanation", *Philosophy of Science* 65: 448-471.
---. (1999), "Is Pure *R*-Selection Really Selection?" *Philosophy of Science* 66(Proceedings): S185-S195.
---. (2001), "Selection, Indeterminism, and Evolutionary Theory", *Philosophy of Science* 68: 518-535.
Graves, L., B.L. Horan and A. Rosenberg (1999), "Is Indeterminism the Source of the Statistical Character of Evolutionary Theory?" *Philosophy of Science* 66: 140-157.
Hagen, J.B. (1999), "Retelling Experiments: H.B.D. Kettlewell's Studies of Industrial Melanism in Peppered Moths", *Biology and Philosophy* 14: 39-54.
Hitchcock, C. (1999), "Contrastive Explanation and the Demons of Determinism", *British Journal for Philosophy of Science* 50: 585-612.
Hodge, M.J.S. (1987), "Natural Selection as a Causal, Empirical, and Probabilistic Theory", in L. Krüger, G. Gigerenzer and M.S. Morgan (eds.) *The Probabilistic Revolution. Volume 2: Ideas in the Sciences*. Cambridge, MA: The MIT Press, A Bradford Book, 233-270.
---. (1992), "Natural Selection: Historical Perspectives", in E.F. Keller and E.A. Lloyd (eds.) *Keywords in Evolutionary Biology*. Cambridge, Mass.: Harvard University Press, 212-19.
Horan, B.L. (1994), "The Statistical Character of Evolutionary Theory", *Philosophy of Science* 61: 76-95.
Lennox, J.G. and B.E. Wilson (1994), "Natural Selection and the Struggle for Existence", *Studies in the History and Philosophy of Science* 25: 65-80.
Lewontin, R.C. (2000), "What Do Population Geneticists Know and How Do They Know It?" in R. Creath and J. Maienschein (eds.) *Biology and Epistemology*. Cambridge: Cambridge University Press, 191-214.
Majerus, M.E.N. (1998), *Melanism: Evolution in Action*.
Millar, C. and D. Lambert (1999), "Industrial Melanism - a classic example of another kind?" *BioScience* 49(12): 1021-1023.
Millstein, R.L. (1996), "Random Drift and the Omniscient Viewpoint", *Philosophy of Science* 63(Proceedings): S10-S18.
---. (2000), "Chance and Macroevolution", *Philosophy of Science* 67: 603-624.
---. (2002), "Are Random Drift and Natural Selection Conceptually Distinct?" *Biology and Philosophy* 17: 33-53.

Reed, E.S. (1981), "The Lawfulness of Natural Selection", *American Naturalist* 118(1): 61-71.
Rosenberg, A. (1988), "Is the Theory of Natural Selection a Statistical Theory?" in M. Matthen and B. Linsky (eds.) *Philosophy & Biology (Supplementary Volume 14 of the Canadian Journal of Philosophy)*. Calgary, Alberta: University of Calgary Press, 187-207.
---. (1994), *Instrumental Biology or The Disunity of Science*. Chicago: University of Chicago Press.
---. (2001), "Discussion Note: Indeterminism, Probability, and Randomness in Evolutionary Theory", *Philosophy of Science* 68: 536-544.
Rudge, D.W. (1999), "Taking the Peppered Moth with a Grain of Salt", *Biology and Philosophy* 14: 9-37.
Salmon, W.C. (1984), *Scientific Explanation and the Causal Structure of the World*. Princeton: Princeton University Press.
---. (1998), "Causality: Production and Propagation" in *Causality and Explanation*. New York: Oxford University Press, 285-301.
Shanahan, T. (1989), "Discussion: Beatty on Chance and Natural Selection", *Philosophy of Science* 56: 484-489.
---. (1992), "Selection, drift and the aims of evolutionary theory", in P.E. Griffiths (ed.) *Trees of Life: Essays in Philosophy of Biology*. Dordrecht: Kluwer Academic Publishers, 133-161.
Sober, E. (1993 [1984]), *The Nature of Selection: Evolutionary Theory in Philosophical Focus*. Chicago and London: University of Chicago Press.
Stamos, D.N. (2001), "Quantum Indeterminism and Evolutionary Biology", *Philosophy of Science* 68: 164-184.
Walsh, D.M. (2000), "Chasing Shadows: Natural Selection and Adaptation", *Studies in the History and Philosophy of the Biological and Biomedical Sciences* 31C(1): 135-154.
Walsh, D.M., T. Lewens and A. Ariew (2002), "The Trials of Life: Natural Selection and Random Drift", *Philosophy of Science*(forthcoming).
Wimsatt, W.C. (1987), "False Models as Means to Truer Theories", in M.H. Nitecki and A. Hoffman (eds.) *Neutral Models in Biology*. New York: Oxford University Press, 23-55.
Wright, S. (1955), "Classification of the factors of evolution", *Cold Spring Harbor Symposia in Quantitative Biology* 20: 16-24.

Comments on Robert Brandon’s paper

Paul Griffiths

I’m going to concentrate here on Brandon’s response to the treatment of evolutionary indeterminacy by Denis Walsh, Tim Lewins and Andre Ariew. Their paper seemed to me to rest heavily on a misunderstanding of Elliot Sober’s treatment of this issues. They identify his view with what they call the ‘dynamical conception’ of evolutionary theory:

> “Evolutionary theory is commonly described as a theory of forces… natural selection, random drift, mutation, recombination and migration. … When a population changes it’s trait frequencies it is supposed that some combination of evolutionary forces may be acting concurrently. Hence one can ask how much of some observed change owes itself to each force. … We shall call this the *dynamical* conception of evolutionary theory.[1]” (Walsh, Lewins and Ariew, In Press, 2-3), Footnote 1 cites (Sober, 1984)

Walsh et al argue that the dynamical conception is incorrect because the interactions between the various causal processes that determine inter-generational changes in gene frequencies cannot be adequately represented as a set of vectors that produce a resultant force. But although Sober he is very attached to the language of forces, he is quite clear the analogy with mechanics as just that – an expository analogy. He freely admits that it cannot be turned into a fully adequate representation of the evolutionary process. In particular, he agrees with Walsh at al that drift – sampling error – cannot be treated as a force acting on a population:

> “Which, we might ask, has been the more powerful force in influencing the frequency of a character…? The question…sounds like the question we posed in the previous section about selection and mutation. But there is a difference. Not to put too fine a point on it, the difference is that the question about selection and drift makes no sense. …
> One cannot compare two populations, one of which ends up further from the expected value than the other, and say that the former was more influenced by drift. The two populations, after all, were influenced by the very same pair of causal factors.” (Sober, 1987, 111-112)

The real issue in discussions of evolutionary indeterminacy is not a clash between fundamentally different ‘dynamical’ and ‘ statistical’ interpretations of evolutionary theory, but rather how to interpret the statistical constructs of population biology. The specific issue that separates Jon Hodge and Bob Brandon’s papers at this meeting, it seems to me, is whether to give ontological import to some of these constructs, namely to fitnesses construed as the propensities to survive and reproduce of individual organisms in a population. I think some light can be thrown on that dispute my some of my earlier work on levels of explanation in biology.

I have argues in several places that we have different ways of classifying biological traits, giving rise to different levels of explanation in biology, because this allows us to bring traits under different, complementary sets of generalizations (Griffiths, 1994; Griffiths, 1996a; Griffiths, 1996b). For example, we group traits by homology so as to collect all those with a particular history, and thus all those that have a particular adaptive-historical explanation. We group homologies by analogy, so as to draw attention to ecological generalizations that apply to traits with different specific histories. Grouping traits by fitness can be seen as the limiting case of this process of abstraction. Just as analogous traits fall under generalizations more abstract that those that apply to homologous traits, traits which resemble one another merely in having the same relative fitness (or, with frequency dependent selection, the same fitness function) obey a still more abstract set of generalizations – the principles of population genetics and evolutionary game theory. Hodge’s intuition that explanations at this level are not really causal explanations arises because, unlike natural historical and ecological generalizations, the principles of population genetics do not have obvious empirical content. They look very much like simple applications of algebra, either to the empirical facts of genetics, or, in evolutionary game theory, to the minimal empirical fact that parents resemble their offspring! Hodge therefore resists the idea that the algebraic values that figure in explanations at the population-genetic level denote real, causally active properties of organisms – propensities.

In reply to Hodge, it seems to me that Brandon has three arguments for the ontological import of propensity talk. First, although each individual organism lives or dies and reproduces or fails to reproduce because of some individual ensemble of causal influences, this process gives rise to predictable patterns at the population level. Organisms of a given type prosper or fail to prosper in a predictable manner. So there must be some *robust* explanation of reproductive success, and propensities of survive and reproduce provide that explanation. Second, the idea that statistical explanation is a summary of diverse, specific causal facts and not, itself, a causal fact is supported by the fact that statistical explanations of the fate of individuals proceed by assigning those individuals to a narrowest references class. Typically, the choice of reference class reflects the extent of our knowledge of the causal processes involved, making the particular explanation we offer a function of our ignorance. But Brandon has shown that the biological version of the narrowest reference class - the set of organisms that share a common 'selective environment' *sensu* (Brandon, 1990) - can be discovered through experiment. This suggests that an organism's membership of this class is an empirical property of that organism that we try to estimate through experiment. Finally, as Brandon stressed in today's presentation, propensities can be measured independently of their population level effects, by a number of empirical 'direct measurement' methods. Hence, one can argue, those propensities cannot be mere artifacts of the statistical methods we use to sum up the individual fates of organisms.

References

Brandon, R.: 1990, *Adaptation and Environment*, Princeton University Press, Princeton.
Griffiths, P. E.: 1994, 'Cladistic classification & functional explanation,' *Philosophy of Science* 61 (2), 206-227.
Griffiths, P. E.: 1996a, 'Darwinism, Process Structuralism and Natural Kinds,' *Philosophy of Science* 63, S1-S9.
Griffiths, P. E.: 1996b, 'The Historical Turn in the Study of Adaptation,' *British Journal for the Philosophy of Science* 47, 511-532.
Sober, E.: 1984, *The Nature of Selection*, M.I.T. Press.
Sober, E.: 1987. 'What is adaptationism?', in, *The Latest on the Best*, Dupre, J. (Ed), MIT Press, pp. 105-118.
Walsh, D. M., Lewins, T. and Ariew, A.: In Press, 'The trials of life: Natural selection and random drift,' *Philosophy of Science*.

Summary of Discussion of Robert Brandon's paper

Frederic Bouchard

Reply by Brandon to commentators:

Brandon addressed some of Love's comments concerning the Matthen and Ariew paper. The main point was that Matthen and Ariew have a distorted view of the relationship between individual fitness and population fitness. Brandon pointed out that if fitness component analysis is true (and it appears to be as it is successfully used in biology today), Matthen and Ariew are wrong about their description of evolutionary biology. the main problem is that they do not recognize the important significance of ecological genetics. Antonovics for example uses fitness component analysis to individuate fitnesses. We then build up overall fitness from these various individual fitnesses. Ecological genetics alone won't tell us much about evolution (it's scope being too restricted), and it is agreed that evolutionary theory aims at generality, but that generality shouldn't be obtained at the price of rejecting important aspects of evolutionary biology.

Brandon went on to address on some of Griffiths' comments. Brandon agrees completely with Griffiths although some points need to be clarified. Salmon assumes nothing about determinism or indeterminism; he assumes only objectivity of probabilities. Another important point is that Salmon doesn't prejudge the size of the possible homogeneous reference classes. Maybe nature won't lend itself easily to science (i.e. very small homogeneous reference classes) maybe it will. In the same way, Brandon doesn't assume anything about the broadness or narrowness of reference classes, although he does point out that adaptation requires certain broad selective environments. Without these 'neighborhoods' adaptation couldn't operate, for organisms would simply hop from one patch to the other without having to respond to the selective pressures of their original patch.

Discussion Period

Question by Hodge: Hodge pointed out that you have to go beyond the biology to get to the philosophy. Contrary to what Brandon said, you should have a causal difference to explain the non-equiprobable, to justify a difference in the math. The match is a means not an end directing the inquiry. In a small population, the math for drift and the math for a selection could 'look' the same. By pure chance a drift pattern could look like a selection pattern. The math alone cannot tell you what is happening. Only via ontology can a difference between the two be explained. Hodge went on to say that he didn't believe Brandon really addressed Millstein's concerns (see paper in January 2002 issue of Biology and Philosophy). The Math can only describe outcomes. The question of processes hasn't been addressed.

Reply by Brandon. Brandon disagreed with this description. With his urn examples (and the sampling of balls), he wanted to used a thoroughly *causal* example. That is why physical properties like 'stickiness' were used to show the causal chain. The other point Brandon disagreed about was about the Hodge's example itself. Brandon believes that we should want values that don't just fit a short run. A lucky break in a small population cannot give you a good prediction model. In the long run, it is highly unlikely that the pattern of drift and selection could be interchangeable. Brandon agrees with the tenor of Hodge's comment. The math alone cannot give you causal explanations, but Brandon thinks it can track it more accurately than Hodge does.

As for Millstein's concerns, Brandon does think he addressed them. If you take seriously the idea of probabilistic causes, the same process can result in different outcomes. Selection and drift are the same process but different outcomes. You can't look at the process alone to explain fitness differences.

Question by Millstein: She points out that Brandon and her are focusing on different questions. Brandon is focused on the pratice of population genetics where any deviation from the selection is drift. Hodge and she wish to step back from these models to try to understand what fitness *really* is. Brandon deals with biological practice while Millstein want to focus biological concepts.

Reply by Brandon: He first points out that he is interested in ecological genetics and population genetics. He's interested in real measures of real organisms. But since we are dealing with probabilistic capacities, there will be deviations which will need to be explained. He went on to describe at Krietman example where you have a known philogeny and you look at change in spatially indiviuated population. Even in this case you have selection in genome selection. Selection is weak but still present. Although there is the possibility for absolute randomness, the overwhelming majority of actual cases will be between the two ends of the continuum.

Reply by Millstein: She asked whether that means the selection-neutralist debate has been resolved in favor of selection because it seems always present.

Reply by Brandon: He doesn't think that follows. He doesn't exclude the possibility for neutrality of traits, but he thinks our best data show tells us that there are no selection neutral traits. Krietman's resutlt are far from perfect but they should be taken seriously until we could use a more refined method.

Question by Jim Bogen: In the spirit of Griffiths' remarks Bogen wants to come back to the question of drift and causation. He points out that drift cannot give you causation (to which Brandon agrees). If that is the case what kind of explanation of drift provides.

Reply by Brandon: Suppose two carbon 15 atoms, and you observe them for one year and neither one decays. An improbable event occurs. You would expect at least one to decay. If it doesn't happen you use drift. If the world is indeterministic, probable things happen but the flip side is that improbable things happen as well. Drift is needed to account for these improbable occurrences.

SOMEONE points out that explaining is subsuming under the relevant generalization. Sometimes you meet the expectation sometimes you do not. That is the pure Salmon

Question by Lennox: Lennox points out that Hodge and Brandon seem to agree about the ontological status of fitness. The confusion is concerning general claims vs. abstract claims. Does math abstract properties shared by populations?

Reply by Brandon: You can give an abstract characterization of X but is that X what is ? no. What X is is this set of property and relations in a selective environment. Hodge should be happy (Hodge agrees with this statement). There is a process in nature with the components of fitness that give us the mathematical results. That process is in nature. Scientists need to figure out how these capacities are put together in nature.

Lennox goes on to point out that, that was what Love wanted instead of math model.

Reply by Brandon: Looking at a standard Fitness analysis of altruism,

$W_s = 1 + bx$
$W_a = 1 + bx - c$

This is a Metapopulation model. We can mathematically write this down but there is no reason to believe that there is a general mathematical theory about whether we should be adding, or multiplying (or some other operation) components. It s not that clean.

Matthen and Ariew's suggestion is even more 'dirty'. They mention a possible analogy with statistical mechanics, but even if we could operate that analogy (which we probably couldn't because we are interested in individual fitness values and not just ensemble properties) it's far from clear that it could be helpful since we don't know what's really happening in the case of statistical mechanics. We don't know what is the relationship between individual properties and ensemble properties.

Question by Lennox: In Love's comments, you can get philosophical mileage in trying to understand causal complexity. This complexity seems absent from your account.

Love: Love points out that he wasn't going for a theory-application distinction but to a conceptual clarification question. This probably doesn't illuminate the debate between Brandon vs. Millstein, but it seems Brandon isn't going at the conceptual level at all.

Reply by Brandon : Brandon clarifies one of his examples. Using the information theoretic example we see how we can know how much noise was used. We might not know how much REAL noise happened because of flip-flop (noise could cancel itself between two monitors). We could be wrong but there are ways of looking closer. Counterbalancing or flip-flop can be examined upon closer examination. We CAN look and compare different fitness components contra what Matthen and Ariew said.

Question by Mitchell: In response to Griffiths. With fitness you can talk abstractly. You can choose your level of abstraction. The problem with dispositions is that you lose empirical content. So you need multiple ways to pick out these properties. The way Glymour breaks down processes to what happens to individuals doesn't make the property at higher level any less real. The level you chose is relative to what you want to explain.

Brandon agrees. If you aren't interested in evolutionary consequences, the whole life history is ok but if you are interested in evolutionary consequences you need to look at the components contra what Matthen and Ariew claim.

Question by Schwartz: You can't get away of embedded assumption (e.g. what is a species). I don t know if there can be a general theory of the appearance of novelty. It's not clear you can get a general theory that will explain all organisms. How universal do you think this problem is?

Reply by Brandon: Brandon points out that this is exclusively about micro-evolutionary biology and not about general evolutionary change. Differences in flowering times which affect reproduction, then speciation. Micro evolution is what interests Brandon in this talk

Question by Pfeiffer: She point out that the homogeneity requirement is too strong. Does it have to be homogeneous with regard to any outcome. A weaker standard would be more appropriate.

Brandon thinks this is not the way to characterize his position. Homogeneity is relative to relative fitness.

Pfeiffer: It's not necessary for adaptation to have homogeneity. It just has to be homogeneous with respect to the specific adaptation.

Brandon. He points out that he didn't require strict homogeneity he just required the absence of ordinal changes. It depends on the thing you are looking at. Look at Butterflies. Pupa has a small relevant environment whereas the butterfly will have a much larger homogeneous environment. It doesn't have to be constant. Strict homogeneity is not required. Selective neighborhood are necessary. You can have relative differences but not ordinal differences. Adaptive evolution is evidence that life is not necessarily chaotic and that you have broader selective environments. We all agree that the theory is statistical, but is it emergentist? No, because it is derived from probabilities given to individuals. This is a pragmatic realist argument. If it works, is supported by your best science etc. You believe it.

Question by someone. Looking at the bar graph in the "selection is possible" partition. Selection differences are in effect. If you use the full force of probabilistic causation (e.g. Eels or Cartwright) you can't say selection was operating when you have the fitness differences. Selection occurred here and drift occurred here is that right?

Reply by Brandon: The signal example meant to show that Brandon wants us to quantify the relative importance of selection and drift. The processes are the same. Sampling processes with intermediate probabilities. Processes are the same, the results are the same (but you need to believe in probalitistic causation). Brandon accepts that he is assuming quite a lot: propensities exist. But he thinks this gives you a good consistent picture.

Question by Linquist: In response to Bogen. Looking at Conservation geneticists. A population is undergoing a meltdown when a bottleneck exists for so long that bad recessive alleles accumulate. These are periods when drift was overwhelmingly strong. You need to draw the distinction between the 2.

Reply by Brandon: Small samples will have large deviation form the expectation. Part of the process is the probability values the other part is sample size. The bottleneck is the sampling process.

Comment by Hodge: Brandon has specified a sense in which selection and drift are the same process. They both are erroneous sampling. In the next generation there is a sample from the previous sampling. The only case where there is no error is if the generations are the same size.

Brandon disagrees

Hodge: Selection is erroneous sampling due to non-fortuitous causes, drift is erroneous sampling with fortuous sampling. We have a match with what Lloyd talked about the other day. Why assume that selection is what we should expect and not drift. We could have pure drift models. Looking at a pure drift model, drift cannot be deviation since we aren't expecting selection in the first place. We want egalitarianism. Brandon is sexist in a sense. In Brandon's story selection is superior and drift is deviant.

Reply by Brandon: He describes a population model.
$P1+Q1=1$
$P2 = P1\ Wa$
$Q2 = Q1\ Wb$

Where
Pi = frequency of type a in generation i
Qi = frequency of type b in gen I
Wa is fitness of type a.
Wb is fitness of type b
We have expectations whether or not Wa = Wb. This doesn't force drift or not. Hodge is assuming we have no expectations when the generations are the same but that's false. We are all deviant…

Workshop Participants

Arnold, Karen (Pittsburgh)
Bickle, John (Cincinnati)
Bogen, James (Pittsburgh)
Bouchard, Frederic (Duke)
Brandon, Robert (Duke)
Brigandt, Ingo (Pittsburgh)
Delahanty, Megan (Pittburgh)
Fabrega, Horacio (Pittsburgh)
Fagan, Melinda (Texas)
Feest, Uljana (Pittsburgh)
Griffiths, Paul (Pittsburgh)
Harris, Dehlia (Melbourne)
Guildenhuis, Peter (Northwestern)
Hodge, Jon (Leeds)
Hourdequin, Marion (Duke)
Lennox, James (Pittburgh)
Linquist, Stephan (Duke)
Lloyd, Elisabeth (Indiana)
Love, Alan (Pittsburgh)
Machamer, Peter (Pittsburgh)
Matsumoto, Shunkichi (Tokai)
McClelland, Jay (CMU)
McGuire, James (Pittsburgh)
Millstein, Roberta (CSU Hayward)
Mirus, Chris (Notre Dame)
Mitchell, Sandra (Pittsburgh)
Olby, Robert (Pittsburgh)
Parker, Wendy (Pittsburgh)
Paul, Diane (UM Boston)
Pfeiffer, Jessica (UM Baltimore)
Piccinni, Gualtiero (Pittsburgh)
Quartz, Steven (CalTech)
Richardson, Robert (Cincinnati)
Ruetsche, Laura (Pittsburgh)
Scarantino, Andrea (Pittsburgh)
Schwartz, Jeffrey (Pittsburgh)
Skipper, Robert (Cincinnati)
Steel, Daniel (Pittsburgh)
Sterrett, Susan (Duke)
Sullivan, Jackie (Pittsburgh)
Tabery, James (Pittsburgh)