CHANCE IN EVOLUTIONARY THEORY: FITNESS, SELECTION, AND GENETIC DRIFT IN PHILOSOPHICAL AND HISTORICAL PERSPECTIVE

A Dissertation

Submitted to the Graduate School

of the University of Notre Dame

in Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

by

Charles H. Pence

_____

Grant Ramsey, Director

Graduate Program in History and Philosophy of Science

Notre Dame, Indiana

April 2014

CHANCE IN EVOLUTIONARY THEORY: FITNESS, SELECTION, AND GENETIC DRIFT IN PHILOSOPHICAL AND HISTORICAL PERSPECTIVE

Abstract

by

Charles H. Pence

Discussions of the foundations of evolutionary theory – especially natural selection, fitness, and genetic drift – are saturated with terms referring to various kinds of chance, stochasticity, randomness, unpredictability, and so forth. This dissertation examines these uses of chance in philosophical and historical perspective.

I begin by arguing that, both in the contemporary and historical arenas, the current state of the literature on chance is deeply troubling. Work in the philosophy of biology (i) often conflates various clearly distinct notions of chance, and (ii) often approaches the analysis of chance from the perspective of a debate (on the causal potency of natural selection and genetic drift) that does not in fact profitably engage evolutionary theory. Historically, as well, the most common way of analyzing the development of the use of chance in evolutionary theory does not engage the actual research of historical actors, a point I make by exploring the work of Karl Pearson and W.F.R. Weldon at the turn of the twentieth century.

I thus propose a new guiding question for research into the role of chance in evo-

lutionary theory: what is the relationship between our statistical biological theories and the processes in the world those theories aim to describe? I then offer a novel framework for determining the answer to this question, derived from a deeply biologically-informed understanding of fitness, selection, and drift. This view combines core insights from work in philosophy on the propensity interpretation of fitness with cutting-edge biological treatments of population modeling. Chance enters this model at only a single point – the distribution over the various possible lives that an organism might live – and this single source can explain the influence of chance throughout fitness, natural selection, and genetic drift. This framework, I claim, constitutes a fruitful way to understand both the foundations of evolutionary theory and the role of chance in those foundations.

Sometimes they call this "necessity," because nothing can be other than it is determined in the "fated" (if I may) and immutable sequence of its eternal order. But occasionally they call it "chance" because it brings about many things we find unforeseen and unexpected owing to the obscurity of their causes (or our ignorance of them).

— Marcus Tullius Cicero, *Academica,* I.vii., 45 BC

I have hitherto sometimes spoken as if the variations – so common and multiform in organic beings under domestication, and in a lesser degree in those in a state of nature – had been due to chance. This, of course, is a wholly incorrect expression, but it serves to acknowledge plainly our ignorance of the cause of each particular variation.

— Charles Darwin, *On the Origin of Species,* p. 131, 1859 AD

for julia

without whom i am far too likely

to sow my isn't and reap my same

CONTENTS

LIST OF FIGURES

PREFACE

The dissertation that follows is, as philosophers are wont to do, arranged in something that I hope resembles logical or conceptual order. But I am both a philosopher of biology and a fan of the works of Nietzsche, so I think this preface provides a wonderful place to pursue a *genealogical* explanation of my project.

One of my first introductions to the philosophy of biology came via a term paper early in my graduate career. I'd written my senior thesis on the idea of laws of nature, and wanted to take a swing at looking at whether or not there might be a reasonable application of the concept of law in evolutionary theory. The paper wasn't very good, but it got me rapidly acquainted with general "principles of natural selection," like those proposed by Richard Lewontin (1970), Robert Brandon (1978), and the duo of Frédéric Bouchard and Alex Rosenberg (Bouchard and Rosenberg, 2004). I was immediately drawn to the thought that there might be "foundations" work in the philosophy of biology, analogous to the bread-and-butter work of philosophers of physics on the conceptual bases of quantum mechanics and general relativity.

Since philosophers of biology give (probably too much) weight to the process of natural selection in their discussions of evolutionary theory, any attempt to build a foundational program in the philosophy of biology has to offer a high-level, general accounting of just what natural selection is. And whenever anyone does this, they – of necessity – invoke at least

one, and probably more than one, concept of *chance.* Whether in the principle of natural selection itself, the concept of fitness, population sampling, random genetic drift, or any of dozens of other sources, this approach to biological theory is laden with references to chance, statistics, probability, uncertainty, unpredictability, indeterminism, and a vast collection of other "chancy" phenomena.

Having spent a fair bit of time as an undergraduate contemplating quantum mechanics, this, too, was particularly interesting. And so I went looking for a synthetic review of just what all this "chance" talk in biology is supposed to be doing for us. What do we really mean by invoking chance in the evolutionary process? The entanglements, pun intended, with other general issues in the philosophy of science are numerous and exciting – inter-level relationships, reductionism, explanation, causation, and the interpretation of probability are only a few examples. Surely someone, I thought, would have taken the time to spell out at least some of these connections.

I was wrong – or, at least, mostly wrong. There's been a fair amount of journal literature on the subject, and a few great doctoral dissertations with relevant material. But no overarching, broad, synthetic study like the one I had hoped to find yet exists. I had a project.

Unfortunately, however, you are not presently holding an overarching, broad, synthetic study of chance in evolution. In what I take to be a not-uncommon occurrence, you are holding instead a discussion of a whole host of basic philosophical and historical issues that we have to get clear on *before* we can actually write the kind of study I wanted to write, covering the vast and unruly landscape of chance as it's actually invoked in biological theory.

And that, then, is the descent-with-modification of my dissertation project.

There are many ways to divide and recombine the various threads that I develop here in

order to describe a "narrative arc" for the dissertation, but one is perhaps the most perspicuous. My goal in the first four chapters is to consider the ways in which the debate over the role of chance in evolution has been framed within the philosophical and historical communities. I argue in chapters one and four that, in the philosophical community, debate over chance has largely stagnated within the last decade, as a result of two factors: an extensive conflation of various clearly distinct concepts related to chance (chapter one) and a contemporary debate over the causal potency of natural selection and genetic drift which has deflected research away from an important set of unresolved issues in our interpretation of evolutionary theory (chapter four).

We might, in the spirit of integrating the history and philosophy of science, hope that an analysis of the ways in which chance was introduced into the vocabulary of evolutionary theory could help to clarify its contemporary usage. But in history, as well, I argue that standard ways of understanding the development of chance within evolutionary theory do not respond to a whole host of interesting problems and questions raised by historical actors (chapters 2 and 3).

In the process of this analysis, however, I introduce a new question – a new framework for our inquiry into the study of chance in evolution. We ought, I claim, to focus on the *relationship between statistical biological theories and the processes in the world those theories aim to describe.* This question, presented in chapter three, grounds the project of the last two chapters. In chapter five, I introduce a novel interpretation of the fitness of individual organisms, developed with Grant Ramsey. This view, as I show in chapter six, constitutes an incredibly fruitful way in which to approach the role of chance in fitness, natural selection, and genetic drift. While I lack the space to pursue all the various avenues for future research

arising from this program, I am confident that the combination of this new picture of chance in evolution with future work examining the biological sciences themselves will result in something not too dissimilar from the kind of understanding of chance I had hoped to find some five years ago.

ACKNOWLEDGMENTS

I should begin by acknowledging my great debt to my advisor, Grant Ramsey. The fact that I've wound up working in the philosophy of biology is almost certainly a causally overdetermined event, with far more explanations available than required. But perhaps the most compelling is that work within the philosophy of biology these days is, quite simply, a whole lot of fun – and it's Grant's perspective, not just as an advisor but also as a co-author, that is more than anything responsible for my coming to realize that. It's been a real pleasure to have so many opportunities for great collaboration and off-the-wall idea development, on a whole host of projects of all sorts, from traditional philosophy articles to mathematical biology and the digital humanities – an experience that not many graduate students are fortunate enough to have.

I also owe a great deal of thanks to my committee: Anjan Chakravartty, Phil Sloan, and Fritz Warfield. Each of them has provided me with extensive assistance over the course of the preparation of my dissertation, and they have my utmost gratitude for their help. Worthy of particular mention is Phil's influence – it is undeniably his enthusiasm, extensive help, and immense body of knowledge that turned me from a reluctant student of history into a forceful advocate for the integration of the historical and philosophical approaches to understanding the sciences. Chapters two and three (and, with them, a whole facet of my career) would certainly not exist without him.

Next, most of the chapters here come with their own sets of acknowledgments – several have been published independently, and all have benefited from extensive discussion and debate with a whole host of philosophers both within and beyond Notre Dame.

For chapter 1, thanks to Robert Audi, Nevin Climenhaga, Christopher French, Matthew Lee, Francesca Merlin, Ben Rossi, and Jesse Schupack. Thanks to the participants in the "Chance in Evolution" seminar at the University of Notre Dame for sparking my thoughts on this subject, and the students in the "Philosophy of Biology" graduate seminar for further discussion. Thanks also to an audience at ISHPSSB 2011, especially Marshall Abrams and Jon Hodge, an audience at the Indiana Philosophical Association's Spring 2012 Meeting, especially for detailed and insightful comments by Ioan Muntean, and an audience at Louisiana State University, especially Jon Cogburn, Jeff Roland, and Husain Sarkar.

Chapter 2 is reprinted and adapted from Pence, C.H., forthcoming, "The Early History of Chance in Evolution," *Studies in History and Philosophy of Science,* with the kind permission of Elsevier B.V. Thanks to audiences at the Notre Dame HPS Graduate Workshop, at HOPOS 2012, especially Erik Peterson, and at Integrated History and Philosophy of Science 4 (&HPS4, 2012), especially Don Howard, Theodore Arabatzis, and John Norton. Figure 2.1 is reprinted from Pearson (1930, p. 352) by the kind permission of Cambridge University Press. Figure 2.2 is taken from figure 7 of Galton (1889, p. 63) (public domain). Figure 2.3 is the frontispiece of Pearson (1906) (public domain).

Chapter 3 is reprinted and adapted from Pence, C.H., 2011, "'Describing Our Whole Experience': The Statistical Philosophies of W.F.R. Weldon and Karl Pearson," *Studies in History and Philosophy of Biological and Biomedical Sciences* 42:475–485, with the kind permission of Elsevier B.V. Thanks to Tom Stapleford for invaluable help with broader connections

to the history of statistics, and to Theodore Porter for comments on Pearson. I've also had several very useful exchanges with Gregory Radick on the topic of W.F.R. Weldon. Thanks as well to several anonymous referees for comments on earlier versions. Finally, this chapter also benefited from audiences at ISHPSSB 2009 and the 2009 MBL/ASU Summer History of Biology Seminar at Woods Hole.

For chapter 4, thanks to Mohan Matthen for a helpful discussion about his views on the definition of causal processes. Finally, thanks to an audience at PhiloSTEM-3, especially Mark Jordan, Timothy Fuller, and Bernd Buldt. Figure 4.1 is after figure 5 from Shapiro and Sober (2007).

Chapter 5 is reprinted and adapted from Pence, C.H. and Ramsey, G., 2013, "A New Foundation for the Propensity Interpretation of Fitness," *British Journal for the Philosophy of Science* 64(4):851–881, with the kind permission of my coauthor Grant Ramsey and Oxford University Press. We would like to thank Robert Brandon, John Endler, Agustín Fuentes, Douglas J. Futuyma, James Justus, Elliott Sober, and several anonymous referees for comments throughout the drafting of this paper. They should be held even less responsible for our errors than usual, as the paper changed significantly over the course of its development.

Finally, I would of course be remiss not to mention my family, who have managed collectively to keep me sane during the trying process of getting a doctorate. My parents have unfailingly supported me throughout graduate school, never questioning whether or not a Ph.D. in philosophy was really a good idea. (For this, I'm not sure whether to thank them or have their heads examined.) A pair of two-week summer writing sabbaticals in Texas are responsible for the existence of at least two of these chapters – parental food, drink, and general merriment was enough to sustain quite a few epic eight- or ten-hour writing days.

Most impressively, my mother, a professional editor, has managed (in spite of my best efforts) to dramatically improve the quality of the next several hundred pages of prose, and even picked up enough history and philosophy of biology to make some great suggestions of her own.

On the home front, I would certainly have drifted off the deep end were it not for Julia. (And the cats.) If you don't already know the depth of my appreciation for all she's done for me, my words won't be able to adequately convey it here.

CHAPTER 1

THE CONFLATION OF "CHANCE" IN EVOLUTION

Discussions of evolutionary theory are saturated with references to 'chance'. These references appear in areas as diverse as macroevolution (Conway Morris, 2009; Desjardins, 2011), mutation (Stamos, 2001; Dietrich, 2006; Merlin, 2010), foraging theory (Glymour, 2001), and environmental variation (Lenormand et al., 2009). As philosophers of science, how can we come to terms with this diversity of usage of what is widely recognized to be an extremely difficult and thorny concept, rife with philosophical issues? We should begin by surveying the landscape of contemporary work on chance within the philosophy of biology. It will be the project of this chapter, then, to establish a fairly destructive claim: that current work on chance largely fails to distinguish various notions in the vicinity of 'chance' which, I argue, should be kept clearly distinct.

Perhaps the most hotly debated instances center on the appropriate interpretation of natural selection, fitness, and genetic drift, where we find a particularly large number of concepts related to 'chance.' Fitness, for example, is considered probabilistically because individual organisms in the same environment will have, as Darwin himself noted, a differing "chance of leaving offspring" (Darwin, 1861, p. 88). Genetic drift, as well, is described as "random," often on the basis of our inability to predict in advance, for any given population,

the precise impact of genetic drift on its evolution – Brandon calls drift a "prediction without a direction" (2006, p. 325).

In recent years, two main positions on the nature of selection and drift have crystallized. One approach, deriving from the propensity interpretation of fitness (Brandon, 1978; Mills and Beatty, 1979), considers selection and drift to be probabilistic causal processes (Hodge, 1987; Stephens, 2004; Abrams, 2009a; Otsuka et al., 2011; Abrams, 2013; Millstein, 2013; Ramsey, 2013b,c). The other approach, which has come to be known as the "statisticalist" interpretation of evolutionary theory (Matthen and Ariew, 2002; Walsh et al., 2002; Ariew and Lewontin, 2004; Krimbas, 2004; Walsh, 2007; Ariew and Ernst, 2009; Walsh, 2010, 2013), claims that, on the contrary, selection and drift are merely statistical trends, abstracted from genuinely causal events in the lives and deaths of individual organisms. We will return to this debate in detail in chapter four, but for now it suffices to note that issues about chance are central here as well – the choice between the statisticalist and causalist approaches is often framed as one between objective and subjective notions of 'chance.'[1]

Usage of 'chance' in the philosophy of biology, therefore, is highly multivocal. I will use the phrase "'chance'-like concepts" to refer to the set of quite different notions deployed in discussions of chance – including 'chance', 'randomness', 'stochasticity', 'probability', '[un]predictability', and so forth.[2] All of these various concepts find frequent use in the biological sciences, and have thus been imported into our philosophical discussions of evolutionary

1. For example, in Pigliucci and Kaplan (2006, p. 4), where the contrast is presented as one between "statistical shadows" and the underlying mechanisms which generate them.

2. Nothing should be read into the choice of 'chance'-like as opposed to 'stochasticity'-like, 'randomness'-like, or any of the other possible options – I choose chance merely for the breadth of its application in natural language.

theory. I do not mean to imply that they refer somehow to "varieties of" chance – indeed they do not – but these terms all clearly are related to 'chance' and, though I lack the space to argue the point here, all tend to appear interchangeably in biological work on chance in evolution.

Given the frequent appearance of these 'chance'-like concepts, one would expect that there would be an extensive literature discussing the relationship between them and their impact on the interpretation of evolutionary theory. Surprisingly enough, however, such a literature is almost entirely absent.[3] I hope to take a first step toward providing an analysis of these concepts here.

I will begin by quickly establishing three distinctions that we may use to separate a working set of four of these 'chance'-like concepts. I will then turn to two detailed case studies. The first is an exchange from the causal-interpretation camp: a paper published by Brandon and Carson (1996), and a reply by Graves et al. (1999). We will see that both the original article and the response to it are dramatically undercut by a thoroughgoing conflation of these four concepts. We can then put our distinctions to good use: a plausible reconstruction of Brandon and Carson's argument that makes sense of many of their claims.

The second case study comes from the statisticalist camp, a paper by Matthen and Ariew (2002). Once again, its main argument is entirely undermined by a conflation of our four 'chance'-like notions: its apparent plausibility derives from a notion of 'chance' that fails entirely to engage with the argument's causalist target. While the distinctions we will draw, then, are well-known, it seems they are not well-understood, and not utilized with sufficient

3. Notable exceptions include Millstein (2000, 2011) and, in the biological literature, Lenormand et al. (2009).

caution.

1.1. Four 'Chance'-Like Concepts

For the purposes of examining our case studies, we will develop four distinct concepts related to 'chance'. These four concepts can be separated by pointing out three distinctions, the project of this section. While this section will move quickly, I will elucidate the various 'chance'-like concepts more fully as required during the discussions of the case studies below. The mere recognition of these four notions and the conceptual differences between them will suffice for the arguments in the remainder of the chapter.

1.1.1. Process Properties versus Outcome Properties

The first distinction separates properties of *sequences of outcomes* from properties of the underlying mechanism or *process* which generates those outcomes.[4] We will only require one outcome-level concept, randomness, well-characterized in the mathematical literature:[5]

> *randomness:* the property (of a sequence of outcomes) of either (a) being unable to be proven by any test to be *non*-random (Martin-Löf, 1966), or (b) being maximally 'complex' according to some measure of informational complexity (Martin-Löf, 1969; Kolmogorov and Uspenskii, 1987).

4. In separating 'process' and 'outcome', I mean only that a process is whatever it is that generates the sequence of outcomes – no more complicated notion of causal process is required.

5. Traditionally, this distinction (between process and outcome properties) is described as one between 'chance' and 'randomness' (Eagle, 2013). It should be clear from the following that describing all process-based 'chance'-like notions as 'chance' is problematic, but this distinction is well-recognized nonetheless.

Randomness, here, is a property of sequences of outcomes. On definition (a), a random sequence is one that has no discernible pattern within its outcomes. Consider a binary random number generator. "All ones" is clearly a pattern, as is "alternating between one and zero." Even complex patterns, such as "the digits of pi," are still patterns.[6] A sequence is only random if no pattern can be detected within it. On definition (b), a random sequence is one that is maximally complex or entropic, meaning that it cannot be shortened or compressed – the shortest and simplest possible way to describe the sequence is just to list its elements. A long sequence of alternating ones and zeros, for example, could be described very quickly as "ten thousand alternating ones and zeros." A truly random sequence, though, could only be listed as "first a one, then six zeros, then four ones, …."

For our purposes, we must note only that randomness, as an outcome property, is distinct from any of the process properties that we will go on to consider. Examples can be found both of non-chancy processes that generate random sequences (such as a computer program functioning as a pseudo-random number generator[7]), and of non-random sequences generated by chancy processes (such as a quantum-mechanical random number generator that, despite the improbability of such an occurrence, produces the binary digits of pi). Randomness is therefore quite clearly conceptually distinct from any of our 'chance'-like concepts that might apply to the process generating the random sequence.[8]

---

6. The statistical detection of patterns within putatively random sequences is an active area of research; see, for example, the battery of computerized tests developed by Marsaglia (1995).

7. A pseudo-random number generator can only produce finitely many finite random sequences, but the point still stands for the sequences that it does produce.

8. For the problems inherent in defining "chance" at the outcome level or "randomness" at the process level, see Eagle (2013).

1.1.2. Subjective versus Objective

Our second distinction will separate *subjective* properties of processes from *objective* ones.[9] Again, for our purposes, I will point out only one subjective property of processes, unpredictability:

> *unpredictability:* the inability of a particular agent, with a given epistemic state (at a particular time), to be able to accurately predict the precise outcome of a given process.

Unpredictability, however, is a slippery concept, which may be clarified along any of several different axes. First, we can construct notions of unpredictability that are either subjective or objective. The notion that I will engage with in the remainder of this discussion is a subjective one, as defined above – a subjective claim about a given epistemic agent and her present relationship to her evidence base.

Objective notions of unpredictability, however, do certainly exist, and are more nuanced. Unpredictability *in principle,* for example, might be taken to refer to facts about all possible observers of a given type or within a given community, including cognitive makeup, available scientific equipment, and so forth. In this sense, for example, the outcome of a deterministically chaotic system might be unpredictable for a human, but predictable for a computer, or, the arrival of a storm might have been unpredictable for a sixteenth-century navigator, but predictable for a twenty-first-century meteorologist.

---

9. One might hold an interpretation of probabilities on which there exists *no* variety of objective chance whatsoever. For example, all "objective" probabilities for Jeffrey's imaginary arch-Humean "are simply projections of robust features of judgmental [subjective] probabilities from our minds out into the world, whence we hear them clamoring to get back in" (Jeffrey, 2004, p. 19). We assume for the sake of argument, then, that the reader countenances some variety of objective probability.

Alternatively, unpredictability in principle might refer to the impossibility of *any* given observer or apparatus (or Laplacean demon) to predict the outcome of a given system. This sense is familiar from discussions of quantum mechanics, where the claim is often made that, regardless of what apparatus we might construct, some outcomes will remain unpredictable.[10]

Third, we might consider unpredictability in principle as a veiled reference to objective chances – that is, if a system is indeed objectively chancy, it will thus be impossible to make certain kinds of predictions about its outcomes.

Finally, to add one more wrinkle concerning the notion of unpredictability, we might be speaking either of the ability to provide *deterministic* predictions of a given system's outcome (i.e., predictions of its precise results), or of the ability to provide *probabilistic* predictions of that outcome (i.e., the odds that a given outcome will obtain). In evolutionary biology, where probabilistic evidence is common, we usually mean that the particular, deterministic outcomes are unpredictable – but even this is not always made clear.

In the following, then, we will restrict our focus to the notion of unpredictability defined above: a subjective concept that describes the present relationship between an individual agent, her evidence base, and the precise outcome of the (biological) process at issue. The discussion here stands, in any case, as yet another example of the difficulties that a lack of clarity can engender in work on chance in evolution.

---

10. Yet another notion of unpredictability is possible, on which, while an outcome is theoretically knowable, it is uncomputable, in the sense that a maximally efficient computer constructed from all available matter in the universe would still be unable to arrive at the required prediction. This concept, too, is familiar from quantum mechanics, where the dynamics of many (e.g., large) quantum systems are unpredictable in this sense (see, e.g., Laughlin et al., 2000).

1.1.3. Collapsible versus Non-Collapsible

The third and final distinction we require separates two classes of objective, process-based 'chance'-like concepts: what I will here call *collapsible* probabilities and *non-collapsible* probabilities. Let's begin with a simple presentation of the idea at work, and then turn to its application to process-based notions of chance.

Consider a conditional probability statement for the probability of some event E conditionalized upon some other evidence C, $\Pr(E \mid C) = x$. Now, assume that the event E is a macro-level event (i.e., is not described in fundamental physical terms), and the value x is non-extremal (not equal to 1 or 0). Two things might be true of this conditional probability value. For some events (the outcomes of quantum-mechanical experiments, most commonly), it will be the case that *whatever facts* about the history and present state of the universe one adds to the condition C, the value x of the conditional probability will *always* remain non-extremal.[11] Such conditional probabilities are often described as tracing back to causal indeterminism – but however we understand their source, there is *no* fact of the matter that determines the outcome of a system like this with certainty.

Alternatively, it may be the case that for some amount of facts added to the condition C, the probability value *becomes extremal.* For example, consider an unbiased coin, flipped in a deterministic universe. All we mean by claiming that the coin is unbiased is something like $\Pr(\text{heads} \mid \text{coin flipped}) = 0.5$.[12] But since the universe in this case is, by hypothesis,

11. Conditioning on *future* facts, of course, would cause many such probabilities to become extremal, and so they must be excluded. Thanks to Nevin Climenhaga for pointing this out.

12. Assuming that by "coin flipped" we mean a relatively vague state of affairs that includes a reasonable

deterministic, it is also the case that Pr(heads | coin flipped ∧ precise flip details) = 1 or 0. Thus, adding more information causes the value to "collapse" to either one or zero. Call conditional probabilities for which this occurs *collapsible,* and conditional probabilities for which it does not *non-collapsible.*

This, then, serves as our last distinction concerning varieties of objective chance:

> Objective chances are said to be *non-collapsible* if no further amount of conditionalization on past or present facts would cause them to adopt extremal values.[13] Otherwise, such chances are *collapsible.*

I will assume that the account of non-collapsible objective chances is relatively clear – these are the probabilities, for example, that are often posited to appear in quantum mechanics. The question, then, is whether or not there is a coherent concept of *objective, collapsible probabilities* – objective "chances" that could exist even in a deterministic universe.

Why would we believe that there *are* any objective, collapsible probabilities? While I certainly don't intend to resolve any issues in the interpretation of probability within the scope of this chapter, we can consider two possible defenses of this position prominent in current literature. First, Glynn (2010) has argued in detail that chances are, in an important sense, *level-relative.* A macro-level chance, then, may be objective without objective chances being present at the micro-level, creating exactly the sort of objective, collapsible probabilities we discuss here.

Second, a number of philosophers have performed substantial technical work with the intent of grounding a notion of probability that could play this role. Sober (2010) has argued that such macro-level probabilities play the same role in many of our scientific theories as other

range of linear and angular velocity. See the model of coin-flipping in Diaconis (1998), brought to my attention by Sober (2010).

13. For a process to exhibit non-collapsible chance, it must, of course, assign non-collapsible probabilities to at least two of its outcomes.

objective, repeatably and reliably measurable properties of objects, and should be considered objective on this basis. In a different vein, several other authors (including Ismael (2009, 2011), Abrams (2012), Rosenthal (2010), and Strevens (2011, 2013)) have attempted to define an objective notion of probability grounded in various properties of causal systems which may obtain at the macro-level regardless of the underlying determinism or indeterminism of the micro-level. Nor is the idea restricted to technical work in the philosophy of science – Harman argued in the early 1970s that a roulette wheel, for example,

> instantiates a nondeterministic automaton. Perhaps, under some complex description, it also instantiates a deterministic automaton. Our explanations of various outcomes presuppose the nondeterministic way of looking at the wheel without committing us either way with respect to the deterministic viewpoint. (Harman, 1973, pp. 51–52)

We do not even require, however, that any of these defenses of the existence of objective, collapsible probability succeeds.[14] Rather, we require only that the *conceptual* distinction between collapsible and non-collapsible objective probabilities is a genuine one. If this is true, then for a given system featuring objective probabilities, those probabilities are either collapsible or not.

1.1.4. Four 'Chance'-Like Concepts

We are, therefore, left with four quite distinct concepts related to 'chance' (see figure 1.1, where the "branching points" are our three distinctions, and our four resulting 'chance'-like

14. If all these defenses fail, it will merely be the case that when a causal theorist of natural selection speaks of an objective, collapsible probability, they are not referring to a feature of the actual world.

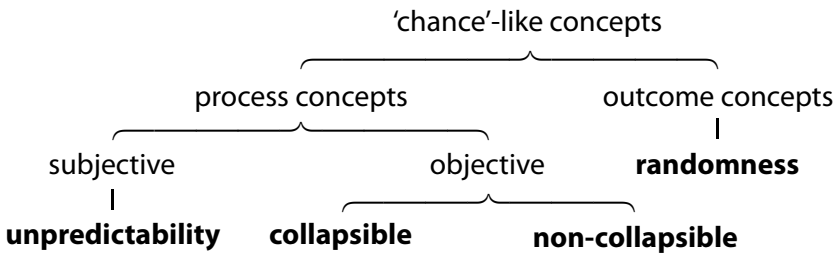

Figure 1.1: Four 'chance'-like concepts

concepts are bold-face).[15] We have randomness (as applied to sequences), unpredictability (for our purposes, a variety of subjective chance), collapsible objective chances, and non-collapsible objective chances.

It is worth emphasizing that I do not intend this to be (anywhere close to) a complete partitioning of all possible 'chance'-like concepts, or even all the ones relevant to the philosophy of biology. Most importantly absent is the concept of 'historicity', so central to the study of the history of life (Beatty, 2006b; Desjardins, 2011). What I hope to have laid out so far is a set of four obviously distinct concepts, all of which are at play in work on chance and randomness in evolution, and all of which we may readily agree should be kept separate in our philosophical work on "chancy" evolutionary phenomena.

Sufficient awareness of the distinctness of these four notions, then, is precisely what I hope to show has been absent from much of the literature. I will provide two examples – one argument and response from the causal theorists of selection and drift, and another argument from the statisticalists. And these examples are not merely made unclear by such confusion. Rather, they are entirely undermined by the persistent conflation of these four

15. The four notions are *distinct,* but not *independent* – for example, many of them, when identified in a given system, may provide us with some *prima facie* reason to believe that the others are present. But the important point stands that inference from one to the others is *not* licensed without further evidence.

distinct concepts.

1.2. In What Sense is Natural Selection "Chancy?"

Other than the often-repeated claim that fitness must be interpreted probabilistically in order to prevent the "survival of the fittest" from coming out tautologous, authors working within the causal interpretation of natural selection rarely discuss what exactly it is that is supposed to make fitness and selection "chancy." One attempt, however, to answer this question is provided by Brandon and Carson (1996, henceforth BC).[16]

BC's paper is incredibly rich, and it offers several arguments about quantum causal indeterminism, the nature of genetic drift, the character of natural selection, and the outcomes of biological experiments. For our purposes here, however, I will draw out one particular argument in their work – an argument for the claim that natural selection is "chancy:"

P1. Drift is unpredictable.
P2. Drift is an autonomous statistical law.
C1. Drift is chancy (*a fortiori* from both P1 and P2).
P3. Natural selection and drift are "inextricably connected" (BC, p. 324).
C2. Natural selection is (objectively) chancy (from C1 and P3).

As we will see below, this argument is problematic for several reasons. First, BC's premise P2 fails to demonstrate anything "chancy" about genetic drift, due to a mistaken construal of Hacking's concept of autonomy. C1 is thus based entirely on the unpredictability of drift. Their argument, therefore, can do nothing more than establish the *subjective* "chanciness" of

16. I can find two other attempts to solve this problem within the causalist camp: Ramsey (2006) and Pence and Ramsey (2013) offer another solution, in terms of (presumably, objective modal facts about) the possible lives of individual organisms, and Pfeifer (2005) ascribes the chanciness, at least in part, to our (presumably subjective) choice to ignore certain environmental features. (On the statisticalist side, Matthen (2009) appears to agree with Pfeifer.)

genetic drift. They hope to show, however, that natural selection is *objectively* chancy, a claim that, in this argument, can be supported only by a conflation of several of our 'chance'-like concepts.

1.2.1. Genetic Drift as Unpredictable

BC are defenders of what might be called the "sampling error" school of genetic drift: "we suggest," they write, "that genetic drift be characterized as any transgenerational (evolutionary) change in gene or genotype frequencies due to sampling error" (BC, p. 321). There exists, on this view, an expected outcome – the outcome we would predict were a given scenario governed merely by natural selection, with offspring populations produced in proportion to fitness values. Genetic drift is then defined as any deviation from this expected outcome due to what they call "random sampling effects" (sampling of gametes to form offspring, of survivors of a population bottleneck, etc.).[17]

Their discussion of genetic drift, however, leaves it unclear exactly what sort of 'chance' is to be found there. They claim that "drift clearly is a stochastic or probabilistic or indeterministic phenomenon" (BC, p. 324). As we saw in the last section, this is an obvious conflation: "stochastic" generally refers to randomness, "probabilistic" might refer to any process-based 'chance'-like concept, and "indeterministic" usually applies only to non-collapsible objective chances. What sense of chance, then, do they consider drift to

17. Brandon's position has since become more complex, though that complexity is absent from BC. We will return to his view later. Within the causalist camp, this picture of drift is to be contrasted with those who identify drift with a certain set of indiscriminate sampling processes (Hodge, 1987).

exemplify?

Deciphering a response to this question in their arguments is quite difficult. The most plausible reconstruction points toward *unpredictability:* many of their assertions about drift are phrased in terms of "the inferences we can make" (BC, p. 322), or what genetic drift "can predict" or "cannot predict" (BC, p. 323). I take their premise here, therefore, to be P1 as stated above: genetic drift is unpredictable.

1.2.2. The "Autonomy" of Genetic Drift

Next, turn to BC's premise P2, that genetic drift is an autonomous statistical law. Near the end of their paper, they make the claim that even if events at the level of individual organisms are in fact "purely deterministic, the population-level generalizations are probabilistic" (BC, p. 335) because of genetic drift's status as an autonomous statistical law, *sensu* Hacking (1990). They seem to point toward a notion of 'chance' similar to the collapsible, objective probabilities discussed in the last section: even if population-level generalizations were grounded in deterministic events in the lives of individual organisms, there may remain a sense in which those higher-level generalizations are objectively probabilistic.

The heavy lifting, however, is done here by Hacking's criterion of "autonomy," which should certainly not be used in this manner. Hacking defines autonomous statistical laws as those which "could be used not only for the prediction of phenomena but also for their explanation" (Hacking, 1990, p. 182). But what grounds the ability of these laws to make such explanations in addition to predictions is not described in Hacking's work. The criterion as it

stands is under-specified, if not circular: autonomous statistical laws are those that can explain, and the laws that can explain are, perhaps, the autonomous ones.[18] Hacking's criterion of autonomy cannot, therefore, be used as an independent method of sorting statistical laws into two kinds: autonomous and non-autonomous. There is nothing *about genetic drift* that makes it autonomous. At best, there might be facts *about the way drift is used* that ground such a claim. Premise P2 does nothing to establish whether or not genetic drift itself is "chancy."

1.2.3. From Drift to Selection

Given that P2 fails to go through, and P1 refers merely to unpredictability, the only possible interpretation of BC's first conclusion, C1, is that genetic drift is chancy insofar as it is unpredictable. That is, their first two premises make reference (or make reference *successfully*) only to unpredictability (in the sense defined in the second section).

Turning to selection and fitness, then, BC argue that "natural selection is indeterministic at the population level because (in real life as opposed to certain formal models) it is inextricably connected with drift" (BC, p. 324), what I call premise P3. An analysis of this premise would take us too far afield, so I will accept it for the sake of argument. BC provide extensive support for this claim, including their discussion of the "necessity" of genetic drift in certain circumstances (BC, pp. 321–322).[19]

---

18. Notably, this causes no problems for Hacking's own analysis: his project is a descriptive historical one, locating the first time a scientist or philosopher *used* a statistical law *as though it were autonomous* (Francis Galton is the culprit).

19. An account of drift other than drift as sampling error would, however, make this premise highly problematic. See Millstein (2005) for a skeptical discussion of Brandon's characterization of drift.

The chance present in drift, therefore, will be used to ground a very robust sense of "indeterminacy" throughout the evolutionary process. They also quite clearly intend this use of "indeterminacy" to refer to some sort of objective chance – they argue that "if one is a realist …then [on the basis of their argument] one should conclude that [evolutionary theory] is fundamentally indeterministic" (BC, p. 336). This, then, is the reason that my reconstructed conclusion C2 includes a parenthetical "objectively."

This presents an obvious problem. If C1 has established only the subjective unpredictability of genetic drift, and C2 is intended to claim that natural selection as a whole is chancy in some sort of objective manner, then BC's argument has failed outright. And it has done so precisely because of a conflation of several of our 'chance'-like notions. While the "autonomy" claim seems to defend something like the existence of collapsible objective chances, it fails – the only premise that goes through is one based on unpredictability, yet the conclusion is intended to evince some unspecified sort of objective chance in natural selection. BC have thus conflated collapsible objective chances, unpredictability, and possibly (depending on the reading of their C2) non-collapsible objective chances – three of our four 'chance'-like concepts. In order for their argument to succeed, at least their premise P1 and their two conclusions, C1 and C2, would have to use 'chance' in the same sense, and they do not.

Before turning away from the initial discussion of BC, it is worth pausing to consider a potential objection. It is, of course, conceivable that the looseness in terminology here is entirely intentional – that is, that rather than an outright conflation or equivocation between these various notions of chance, BC aim to invoke a broader concept (recall the use of "stochastic or probabilistic or indeterministic" above) that does not distinguish precisely

between these chance-like concepts. In that case, then my charge of equivocation clearly fails to apply – but my reconstruction of their argument, to which we will turn in the next section, still offers, I claim, the most promising way in which to reconceive their argument using more precise and finer concepts of chance.

1.2.4. Unpredictability and Objective Chance

Let's consider, then, a response to BC published by Graves, Horan, and Rosenberg (1999, henceforth GHR). To begin, I should note that just as BC's paper is quite expansive, GHR's critique touches on many of BC's arguments that we do not address here.[20] GHR do, however, directly engage with the very argument we have outlined in the preceding. This particular criticism of BC, as we will see, appears quite plausible. But yet more conflation of our 'chance'-like notions results in this reply failing to engage with BC's argument. This conflation thus not only renders BC's work suspect, but GHR's reply as well.

They begin by agreeing with (or, at least, accepting for the sake of argument) BC's characterization of drift as sampling error. For GHR, however, "the following question remains: Are the probabilities employed in the theory epistemic [i.e., subjective] or not?" (GHR, p. 146).[21] They believe the important question to be whether the chance at work in evolutionary theory is subjective or objective. Arguing against BC, they claim that a being as

20. In particular, GHR advance several provocative arguments about the role of quantum indeterminism in evolutionary systems, and provide a critique of BC's closing example involving botany experiments.

21. While the two are conflated by GHR (as well as BC and Matthen and Ariew later), it is important to note that 'epistemic' is not equivalent to 'subjective' in this instance. Epistemic probabilities (on general epistemological usage) reflect something like (possibly objective or ideally rational) degrees of confirmation, while the notion that GHR are after is clearly subjective in the sense laid out earlier.

well-informed as Laplace's demon would not need chances at all in evolution, and that genetic drift is therefore (in this sense) predictable: "if all this information [about some biological system] were available to us, and we had the computational abilities needed to process it, the theory would not rely on probabilities" (GHR, p. 147). They note that BC's argument on this point fails, as we just saw, and they conclude for this reason that BC's argument "*presupposes* objective chanciness" (GHR, p. 150, original emphasis), rather than arguing for it. Begging the question is, GHR claim, the only manner in which BC's conclusion could follow from their premises.

Interestingly, however, GHR have shifted the notion of 'unpredictability' at work in their response. While I have argued that BC employ 'unpredictability' in the sense of the inability of an individual agent to predict the precise outcome of a given system, GHR shift to an unpredictability claim based on a *maximally informed being,* one that could obtain complete knowledge of "genetic drift, mutation, migration, and the forces governing chromosomal segregation, about which we have inadequate information" (GHR, p. 147). Insofar as this doesn't seem to be a *human* observer, this is a much different sort of unpredictability. It seems to be not a reference to epistemic characteristics, as GHR claim it to be, but rather (as discussed earlier in the introduction of unpredictability) to *objective* facts about the system itself.

GHR's response, therefore, seems to be pointing to some variety of *objective* chance. More precisely, it seems quite likely that they intend to refer to non-collapsible objective chances in particular. They return repeatedly to the example of quantum mechanics, discussing why it is the case that "pure probabilistic propensities are viewed as an uncomfortable but unavoidable conclusion in quantum mechanics" and calling these propensities "disposi-

tions ungrounded by further manifest or occurrent properties" (GHR, p. 154), apparently a paradigm example of non-collapsible probabilities in our sense.

This makes trouble for GHR's response, however. Because, while we accused BC of conflating at least two 'chance'-like notions in their argument, it is not clear that they even invoke non-collapsible probabilities at all. (Recall that non-collapsible probabilities *may* feature in their second conclusion, C2.) Thus, GHR's response quite possibly fails to engage with BC's argument.

The conflation of the four 'chance'-like notions I detail here has thus already claimed two casualties. BC's argument conflates two or three of the four notions, and fails to be sound as a result. Further, GHR's response to BC misunderstands BC's use of 'unpredictability', and thus fails to engage with the argument they intend to critique.

1.2.5. Reconstructing Brandon and Carson: Collapsible Objective Chance

I want to close this section, as mentioned in the introduction, by attempting to reconstruct BC's argument – putting our four 'chance'-like notions to use. In particular, if we take all BC's conflated references to 'chance' to refer to collapsible objective chances, then their article – and several of the least clear points within it – becomes far more intelligible.

We may support this reading of BC's argument in several ways. First, and most importantly, Brandon's other writings on drift are consistent with a reinterpretation of BC's premise P1 in terms of collapsible objective chances, rather than unpredictability. In Brandon (2005), he lays out what we might call a hybrid view of drift – one which includes the "sampling

error" component mentioned above, but adds to it mention of the process generating the sampling error. In Brandon's words, this view of drift is "outcome-oriented," but "notice that it does refer to a process, viz. sampling, and so here drift is defined in terms of both process and outcome" (2005, pp. 158–159). While we must indeed arrive at certain outcomes for drift to take place (i.e., there must be actual, realized sampling error), the process generating genetic drift must also have certain characteristics – and I argue that the most consistent way to understand these characteristics of the process is by way of collapsible objective chances.

More recently, Brandon has described drift as an element of the "Zero Force Evolutionary Law," which states that in all evolutionary systems (with variation and heredity), "there is a tendency for diversity and complexity to increase" (Brandon, 2010, p. 708). His toy model of this law consists of a particle diffusing away from an initial position with given transition probabilities (Brandon, 2010, p. 703). After a few iterations, we are able to specify the particle's state only probabilistically, by specifying the odds that the particle will be found in each of the possible end-states.

Importantly, however, there is for Brandon *a causal explanation* underlying this diffusion process. He argues that "the drift producing potential of the sampling processes that are constitutive of the evolutionary process is a fundamental part of the causal structure of our world," featuring in a special class of causal explanations he calls "default-causal" explanations (Brandon, 2006, p. 329).

In a section titled "Drift as a Causal Concept" in his book co-authored with Daniel McShea, they write that an important requirement in coming to understand scientific explanation is "to develop an adequate account of probabilistic causation that would ground such probabilistic explanations" (McShea and Brandon, 2010, p. 106). Indeed, they seem to

directly endorse a macro-level, probabilistic-causation understanding of genetic drift when they write that

> Some might think of drift as the absence of cause. But, as we see it, the relevant causal understanding is the full set of objective probabilities that govern the entities to be sampled. …causal understanding is achieved when we assemble the relevant probabilities governing the events in question. (McShea and Brandon, 2010, p. 107)

While it may be problematic to identify "the relevant causal understanding" with simply the "set of objective probabilities," this certainly seems to be an invocation of what we have called collapsible objective chances.[22]

Returning to BC, we can thus reconstruct their argument in the following manner:

P1*. Drift ascribes collapsible objective chances to its outcomes.
C1*. Drift is objectively chancy (*a fortiori* from P1).
P3. Natural selection and drift are "inextricably connected" (BC, p. 324).
C2. Natural selection is (objectively) chancy (from C1 and P3).

This argument, unlike the one present in their paper, is valid – all of its references to 'chance' are references to the existence of collapsible objective chances.[23]

Further, it makes sense of another odd feature of BC's work: the peculiar reference later in the paper to arguments from "hidden variables." BC claim that the appropriate way to overturn their claim of "indeterminacy" (their C2) would be to demonstrate that there were unequally distributed hidden variables governing the evolutionary process, and then that "assignments of relative fitnesses would be merely epistemic [i.e., subjective], merely useful instruments given our state of ignorance of the hidden variable[s]" (BC, p. 326). Initially, this

22. Unless all the objective probabilities that McShea and Brandon identify within genetic drift trace to non-collapsible probabilities (arising, say, from quantum mechanics), a claim that is both implausible and explicitly disavowed by Brandon and Carson.

23. As mentioned above, the possibility of trouble with P3 entails that the reconstructed argument may or may not be sound. P1* is also contentious, and would clearly be rejected by an advocate of the statistical interpretation of genetic drift.

argument seems quite perplexing. Assuming, for the sake of argument, that the (biological) world may be approximated reasonably well by Newtonian mechanics,[24] there exists a quite obvious set of hidden variables that would do the job: namely, the positions and momenta of all the particles in the universe. These would clearly serve as "hidden variables" grounding a non-chancy causal description of future biological outcomes, and our assignments of relative fitnesses (just like our assignments of all quantities other than particle positions and momenta) would be the result of our ignorance of these hidden variables.

What does our reinterpretation do for this argument, the *prima facie* interpretation of which is so confusing? If BC are referring to collapsible probabilities in our sense, it becomes readily explicable. What would be required to refute their C2 in this case would not be deterministic hidden variables within some other theory upon which evolution depends (molecular biology, chemistry, physics, etc.).[25] Rather, defeating BC here would require the existence of deterministic hidden variables *within the very causal processes* of natural selection and genetic drift (i.e., at the macro-level), driving these evolutionary outcomes (hidden variables of which they are right to claim that we have no evidence). To put this another way, BC claim that refuting their C2 would require a demonstration that the conditional probabilities at issue in evolutionary theory are collapsible *not* by conditionalizing on *micro-level* detail (this we already know, by the definition of collapsible probabilities), but by conditionalizing on more *macro-level, biological* evidence – showing that the biological processes at issue are not

24. Or, just as well, one could consider what the evolutionary process would look like were it instantiated in a classical universe.

25. I follow Brandon and Carson's lead in choosing the relatively innocuous "dependency" to describe the relationship between biology and "lower-level" scientific theories like chemistry and physics. In particular, it "is not meant to imply that biology is reducible to chemistry and physics" (BC, p. 319).

in fact probabilistic at all. The apparent strangeness of the argument, on this reinterpretation, comes from a conflation of collapsible and non-collapsible probabilities. If we resolve this conflation, BC's position on hidden variables is quite reasonable. Reconstructed in this way, BC's hidden-variables argument is at least plausible, and their overall argument, with a revised premise P1*, doesn't obviously beg the question. Sufficient understanding of the distinctions we have drawn here can actively work to *resolve* the confusion present within these arguments.

1.3. The Statisticalist Position and "Chance"

In their initial articulation of what has come to be known as the "statisticalist" interpretation of evolutionary theory, Matthen and Ariew (2002) claim that there exists no way, in principle (as opposed to in specific empirical examples), to distinguish which evolutionary outcomes are the result of natural selection and which are the result of genetic drift, if both are considered probabilistically. One portion of this argument is phrased in terms of the role of 'chance' in various evolutionary explanations. Let's unpack it in detail.

In general, Matthen and Ariew (henceforth MA) aim to demonstrate that believing, as proponents of the "causal process" notion of genetic drift do, that natural selection and genetic drift may be separated as distinct causal influences on the evolutionary process "violates sound probabilistic thinking" (MA, p. 62), and that it does not make sense "to say that drift is a *force* or, more generally, a cause of change that acts independently of selection" (MA, p. 60). They begin with the following example:

> Consider this analogy. You toss a coin four times. What would explain the outcome *two heads?* Answer: the physical setup of the coin-tossing trials. What would explain the outcome *four heads?* Answer: the same thing. Although the

> second result is less probable, the same setup explains both outcomes. (MA, pp. 60–61)

They go on to offer an almost identical instance of this argument as it applies to a biological population. Consider two heterozygous populations (Aa), one of which evolves to homozygosity at each allele (one each AA and aa). One of these alleles, however, is more fit than the other in the given environment. "What explains this difference" in outcome, they ask? "The answer, just as in the case of the coin, might well be: *exactly the same thing,* the same positioning of the traits in the adaptive landscape" (MA, p. 61).

The argument, therefore, seems to go something like this. We have a particular random sequence (be it the results of a coin-tossing experiment, or a set of population-level evolutionary outcomes). Defenders of the causal interpretation of genetic drift argue that there are two, separate influences that together determine outcomes in situations like these: genetic drift and natural selection. This amounts to the claim, according to MA, that one could take these outcomes and "partition the homogeneous reference class to which [they] belong by *improbability* or *chance*" (MA, p. 61). A bit later, they provide a more detailed version of this claim:

> Although it is six times more probable that two heads will turn up in a run of four tosses of a coin than that four will, chance does not play any more of a role in a particular run of four heads than in a particular run of two and two. Thus, one cannot in general differentiate between individual events on the basis of how much they are attributable to chance. (MA, p. 64)

We can formalize this argument as:

> P1. The causal interpretation requires us to separate "more chancy" and "less chancy" events.
> P2. This is impossible.
> C. The causal interpretation fails.

Given our discussion thus far, one particular feature of this argument should stand out: what is it that MA mean by "chance" in these premises? Clearly they believe that this notion is

amenable to comparisons – not only does this ground premise P1, but they claim explicitly that "it makes sense" in certain cases "to quantify the role of chance" (MA, p. 64). Which of our four 'chance'-like notions might fill this role?

First, we may convincingly rule out two of them: unpredictability and non-collapsible probabilities. MA state very explicitly that it is their intent to demonstrate that "there is, in general, no *objective, as opposed to epistemic* [i.e., *subjective*], apportioning of causal responsibility to selection as against drift in a concrete evolutionary history" (MA, p. 68, emphasis added). This statement confirms that MA are not addressing the notion of subjective unpredictability.[26] They also state explicitly that while one situation "in which the notion of probabilistic causation is invoked" is "in quantum mechanics, where it is claimed that the interaction of certain variables is irreducibly indeterministic [i.e., an instance of non-collapsible objective chance] …one would not be justified in claiming that the individual events above [in the examples already mentioned] contained elements of indeterminacy" (MA, pp. 62–63, fn.). They thus also disavow any interest in the issue of non-collapsible probabilities.

There are thus two remaining obvious referents for MA's use of "chance," and both indicate trouble for their argument. First, we might read MA as concerned with the issue of collapsible probabilities in biological systems. On this interpretation, P1 is not coherent, for there is no sense in which the processes of selection and drift can be "less chancy" and "more chancy" in this sense. Most importantly, the existence of collapsible probabilities does not come in degrees. On our definition, either a process produces probabilities that are objective

26. Pfeifer (2005) has argued that a plausible critique of their argument may be marshalled using subjective notions of unpredictability, a possibility that I will not pursue here.

and collapsible, or it does not. If it does not, then it fails to meet our definition outright. There is no sense in which one process can be "more probabilistic" than another.

Further, even if there were a sense in which two causal processes could be compared on this axis, it is not obvious that the causal theorist of drift would be committed to the claim that drift was the "more chancy" of the two processes. Certainly, natural selection is a *biased* process – biased, of course, by fitness values. But while drift lacks this precise sort of bias, it is biased by a host of other variables, including population size as well as many other environmental effects.

One might resolve this difficulty by arguing that drift is unbiased *if* we conditionalize on these effects of population size, environmental effects, and so forth.[27] But if we allow ourselves the ability to conditionalize on any biologically relevant evidence, then we can make selection look unbiased as well – simply conditionalize on fitness values (as well as environment, genetics, and other ingredients of selection), and selection also produces unbiased, collapsible probabilities. This reinforces the point that the 'randomness' or 'bias' of these two processes is a problematic notion – it can change fairly dramatically depending on how we phrase our theories. Given this, it doesn't seem to be the sort of thing that can support MA's distinction between "more chancy" and "less chancy" with respect to collapsible/non-collapsible objective chances. Interpreting MA's use of "chance" as a reference to collapsible probabilities, despite several indications that this is how they intend their argument to be interpreted, makes their premise P1 impossible to satisfy.

This leaves us, I claim, with only one notion of "chance" left: the randomness of the

27. Thanks to Christopher French for bringing this possibility to my attention.

outcomes of the processes of drift and selection. If we interpret MA's use of chance in this way, then premise P2 is obviously correct. While it is possible to say of two given *sequences* that one is more random than the other (for example, that one sequence is more complex than another), it is impossible to do what they claim the causalist is committed to – to take an *individual outcome* from a sequence and claim that it, as opposed to some other outcome, is "more" or "less" random. And premise P1 is no longer *necessarily* false: the causal theory *could* require that we separate outcomes into those that are more and less random – though this would, as MA's argument notes, lead to a serious problem.

But does the causal theory in fact commit itself to such a separation? That is, is premise P1 true of the causal theory? Unfortunately for MA, it is not. Many causal theorists, perhaps most conspicuously Millstein (see, for example, her (2002) response to Beatty (1984)), have noted that looking at a particular set of evolutionary outcomes tells us nothing about the processes that generate those outcomes. On the most common definitions (within the causal-interpretation camp) of selection and drift, she notes that "although random drift and natural selection can be distinguished when they are conceived as processes, they cannot be distinguished fully when conceived as outcomes" (Millstein, 2002, p. 46).[28] While the causal theory could require that we sort outcomes in the way indicated by MA's premise P1, it in fact does not. MA's argument thus fails on this interpretation as well.

This points to the same philosophical trouble in which the BC/GHR debate found itself. MA look to attack the coherence of the entire causal view of selection and fitness,

28. I should note that I do not intend to claim that there is no way in which to separate the overall empirical effects of natural selection and genetic drift – there is a large literature that would lead one to conclude otherwise (Richardson, 2006; Millstein, 2008; Brandon and Ramsey, 2007; Millstein et al., 2009). The claim here is merely that for some arbitrary sequence of evolutionary outcomes, it is impossible in general to claim that a given outcome is "due to selection" or "due to drift."

the token example of which they take to be Sober's force metaphor (MA, pp. 58*ff.*). Their argument, however, trades on a conflation of randomness and the existence of collapsible probabilities. The causal theory itself conceives of 'chance' in an objective manner. Setting aside the 'chance'-like concepts that MA explicitly disclaim leaves us with only one objective notion: collapsible objective chance. Interpreted in this way, MA's references to 'chance' make little sense – the features of 'chance' that they appeal to in their premise P1 simply do not hold of collapsible probabilities. They rely for their argument's plausibility, then, on the fact that their premises *do* seem reasonable if we interpret 'chance' as randomness. On this interpretation, however, the argument fails to engage with the causal theory. Once again, the conflation of our 'chance'-like concepts has undermined a central argument in the debate over the nature of drift and selection.

1.4. Conclusion

I began by quickly describing four distinct 'chance'-like notions with relatively well-understood meanings: (1) randomness of outcomes, (2) subjective unpredictability, (3) collapsible objective chances, and (4) non-collapsible objective chances. We then saw two examples in which the conflation of these four notions directly undermines arguments in the debate over the causal role of natural selection and genetic drift. In order for Brandon and Carson's argument to work, we must resolve a systematic conflation of three 'chance'-like concepts, and Graves, Horan, and Rosenberg respond to Brandon and Carson using a notion of 'chance' to which it is not clear that BC are committed at all. The argument of Matthen and

Ariew, then, uses this same conflation to render plausible an argument which, in fact, does not engage the causal theory whatsoever. These widely cited arguments are thus significantly undermined. Further, we saw that they, at least in some cases, can be salvaged by resolving this conflation.

It is certainly the case that the broader interpretation of chance and randomness is a difficult philosophical problem with a long and storied history. The same is true for the closely allied problem of the interpretation of probability. But detecting the problems with these arguments has not required that we resolve all of the many philosophical issues that come with these other, difficult topics. The four separate 'chance'-like concepts derived in the first section are, in the vast majority of cases, part of the agreed-upon presuppositions of these debates – they are in and of themselves relatively uncontroversial. Almost all of us, then, recognize these four concepts to be distinct. This recognition, though, fails to be reflected in our arguments. We ignore this fact, as I have shown, at our own peril.

If we are searching for a clearer understanding of the role of chance in evolutionary theory, we cannot therefore begin with a literature embroiled in as much difficulty as this. Perhaps, then, we should take our inspiration from the history of science, and consider the manner in which various notions of chance and statistics were introduced into evolutionary theory during its development in the late nineteenth century. It is this project to which I now turn.

CHAPTER 2

THE EARLY HISTORY OF CHANCE IN EVOLUTION

If the philosophical literature does not offer enough help to us in developing a synthetic understanding of the roles of chance in evolutionary theory, then we owe it to ourselves to answer a vitally important historical question before we continue. How were the various chance-like concepts described in the last chapter introduced into the study of evolutionary theory? Coming to terms with this question will be the project of the next two chapters. In this chapter, we will begin by laying out one influential way of understanding the development of chance in evolutionary theory. On this view, advanced most prominently by Depew and Weber (1995) but found throughout the history and philosophy of biology, the introduction of chance is encapsulated by two crucial historical events. First, when and how did biological theories become statistical? Second, when and how did biological theory come to be taken to describe "genuinely chancy" processes in the world?[1]

Elucidating this view is the project of the first section of this chapter. Francis Galton, it is generally recognized, is responsible for the first shift – it was Galton's work on the statistically derived law of ancestral heredity that first introduced statistics into biology. The

1. The referent for "genuinely chancy" is as confused as it was throughout the last chapter; I take it for the remainder of this chapter to refer to some sort of unspecified variety of objective chance. The point will not matter substantially for us, as I will not consider how the second question should be answered.

second shift originates in Sewall Wright's shifting balance theory, which required a much more significant role for a chancy process of genetic drift than the theories which had come before it.

After introducing Depew and Weber's view, we will then explore it in more detail. Section 2.2 will return to Darwin's own works, to establish the now-standard interpretation that Darwin believed evolution to be a *non*-statistical theory of *non*-chancy processes in the world. We then turn to Francis Galton in section 2.3, where I describe his role in the development of the first statistical elements of evolutionary theory. Rather than moving on to Wright, however, we will look in section 2.4 at two of Galton's students at the end of the nineteenth century, Karl Pearson and W.F.R. Weldon. On Depew and Weber's view, we will see that these two would be minor characters.

Why, then, consider Pearson and Weldon at all? It is their work that will serve as our point of departure from considering the introduction of chance in terms of these two focal historical moments. For far from being minor players, I will argue in the next chapter that a vitally important development can be detected in Weldon and Pearson's writings on the philosophical justification for the use of chance in evolution. If we focus only on the two events of the Depew and Weber view, we will entirely fail to recognize an important difference between Pearson and Weldon in conceptions of chance. We must look, then, for a new context for the introduction of chance in evolution – a new driving question on which we are able to understand this philosophical rift that develops between the two men. I will argue that this distinction can be best exposed by considering the *relationship between our statistical biological theories and the processes which those mathematical frameworks are intended to describe.*

As regards this new question, then, a more mathematical, more positivist school of thought, with Pearson at its head, takes these statistics to be a tool for glossing over the (complex, indifferently deterministic or indeterministic) causal details of biological systems. On the other side, a more empiricist, experimentally inclined school, with Weldon at its head, takes statistics to be an essential way of grasping the full causal detail of biological systems. We can see here, I claim, a dramatic difference in the understanding of the connection between evolutionary theory and the evolutionary process, positions that are better comprehended not by way of the "reification" or "objectification" of chance, but by considering the relationship between biological theory and the biological world.

2.1. Two Focal Events

We will begin, then, by discussing the view of the historical development of chance laid out in Depew and Weber's *Darwinism Evolving* (1995) and echoed throughout the literature in the history and philosophy of biology. The second part of their book is devoted to describing the relationship between the advance of a new variety of Darwinism grounded in the developing science of genetics and what they call the "probability revolution" – the same broad historical process that Hacking called the "taming of chance" (Depew and Weber, 1995, p. 202). While they sometimes refer to this revolution as a singular event, they often helpfully break it up into two parts. The first is a "statistical revolution," the introduction of statistics as a tool "for collecting and analyzing quantifiable data," initially in the social and then in the scientific realm (Depew and Weber, 1995, p. 203). Later, with the addition of a

robust probability theory, "the idea arose that probabilities [derived from these statistics] are based on objective propensities of real things" (Depew and Weber, 1995, p. 206). These two ingredients combined to make the probability revolution complete.

When we look at how this revolution may have impacted the biological realm, we are led to investigate the two historical events mentioned at the beginning of the chapter: what was the first time that the statistical revolution was reflected in biological theory, and what was the first time that probability in the genuine, objective sense was utilized?

Depew and Weber go on to describe what have come to be the standard explanations of these two events. For the first, they point to the work of Francis Galton. "Galton," they note, "contributed less to the continuity of the Darwinian tradition by his substantive views … than his conceptual and methodological ones" (Depew and Weber, 1995, p. 201). They make extensive use of the analysis of Hacking, who, as we saw in the last chapter, persuasively argued that Galton was the first not just to use a statistical law for the description of phenomena, but also as "autonomous," as a law "serviceable for explanation" of those phenomena by itself, without having to invoke a large array of supposed underlying, small causes (Hacking, 1990, p. 186). Depew and Weber note that this, as well, is the first time that statistics is used in a positive manner for the support of Darwinian theory, rather than as a way to attack natural selection.

On this point, there is broad agreement throughout the literature in the history of biology. Provine, for example, also locates the introduction of statistical methods with Galton, claiming that Galton "opened the door to a statistical analysis of correlations of characters, an analysis which was to have immense influence upon evolutionary thought" (Provine, 1971, pp. 22–23). Gayon, as well, notes that "the English biometrical school looked to Galton as

the inventor of the statistical methods it used to justify the Darwinian hypothesis of natural selection" (Gayon, 1998, p. 105). Porter (1986, pp. 135, 284–285) and Radick (2011, p. 133) both discuss Galton's modeling of inheritance using drawings from an urn, which directly links Galton's biological work with the mathematical derivation of the law of error.

In the case of the second event – the introduction of an objective, reified, or "genuine" notion of chance in evolutionary theory – Depew and Weber argue that "Sewall Wright opened up this Pandora's box" (1995, p. 287). Wright's turn toward chance, they write, was a way of enhancing the ability of the evolutionary process to create novelty:

> Wright saw in the probability revolution a way to reduce the gap between phenomenal science and metaphysical reality. In brushing Laplacean determinism aside, a probabilistic universe offered even more openings for creative initiations to Wright than to Maxwell or Fisher. […] Wright thought [systems could be driven to novel trajectories] without relying so exclusively on the deterministic intervention of natural selection. Statistical processes themselves did some of the work. (Depew and Weber, 1995, p. 285)

Wright, therefore, completes the probability revolution in the biological sciences. While Fisher, they argue, saw chance as merely a source of mathematical noise, a difficulty in theorizing which needed to be overcome and factored out, it was Wright who first argued that evolution could be phrased as a theory of genuinely chancy processes – including random drift, the chanciness of which occasionally pushed organisms down an adaptive peak and enabled them to reach a higher neighboring optimum. On this view, we have a shift toward 'chance' precisely because chance is, for the first time, an active force which can be implicated in certain sorts of population change (namely, change which runs contrary to fitness gradients).

Discussion of Wright in the context of chance is, perhaps unsurprisingly, rarer. The interpretation of Wright's work is notoriously difficult, and the interpretation of the philosophical commitments of that work all the more so. Hodge provides perhaps the most penetrating

analysis, arguing that Wright was motivated by "the conflict, especially as Bergson had identified it, between the Laplacean determinism presupposed by much scientific thought, and the phenomenon of consciousness itself" (Hodge, 1992b, p. 263). But he goes on to note that

> So many probabilistic themes have proven relevant to the case of Fisher and Wright that the only conclusion to be confidently drawn is that there is no single conclusion. […] [T]he scientific implications of statistical thinking have never been straightforward in their interpretation. …for [Wright and Fisher both,] statistical techniques were a way to further the traditional aim of science in finding the causes, often hidden causes, behind phenomena. Indeterministic or probabilistic causation was never reducible to statistical correlation. (Hodge, 1992b, pp. 287–288)

Matters are, clearly, quite complicated. For our purposes here, however, I will leave the issue underdeveloped. As we will see, whether or not Wright was indeed the correct answer to the second question is immaterial to my project.

It is notable that these are also the two primary foci for the study of chance in evolution in the philosophy of biology. The statistical nature of biological theory has often been seen as a fact demanding a deeper explanation (for example, in the causalist/statisticalist debate, mentioned in the last chapter), and we have seen that much of the broader philosophical debate over chance in biology concerns whether or not there indeed exists some sense of objective chance to be found within the evolutionary process. These two historical events thus parallel two primary strands of philosophical research on chance in evolutionary theory.

Before continuing, I should note that by offering a new, third focus for our historical work on chance in evolution here, I do not at all intend to disparage either this pair of questions or the explanations offered for them. Indeed, both mark significant and important developments in the history of biology, ones which we are right to single out for extra scrutiny. I will argue for the remainder of this chapter, however, that if we restrict ourselves to only looking at the development of chance through these lenses, we run the risk of missing

significant and important developments in the way that chance was understood by practicing biologists.[2] It is this worry – and the example of the philosophical dispute between Pearson and Weldon, which clearly fails to fit within these categories – that drives me toward producing a novel approach to understanding the development of chance in evolution.

2.2. Darwin's View

Now, let us rewind and consider Darwin's position with respect to the two primary historical events laid out above: is Darwin's theory statistical, and does it purport to describe objectively chancy processes?

2.2.1. Darwin on Statistics

Is Darwin positing a statistical theory? One difficulty makes the question itself slightly problematic – the tools of statistics were in such an undeveloped state at the time of the development of evolutionary theory that Darwin wasn't really given the choice to use or reject them. This issue aside, however, Darwin's relationship to statistics is fairly clear. While Darwin did have a copy of Quetelet's *Sur l'homme et le développement de ses facultés* in his library (Rutherford, 1908, p. 69), he did not directly utilize statistical methods in his own work. As Manier notes, Darwin seemed to be unable to apply even a slightly statistical conclusion,

2. I also do not claim that Depew and Weber themselves argued that our focus should be exclusive in this way, or that they failed to notice the problems that would result. They even come close to foreshadowing the account I will develop in section 2.5 when they claim that "what was at stake in the conflict between Fisher and Wright was how many of the conceptual resources of statistical models are relevant to causal explanations of biological processes" (1995, p. 286).

such as his reference to the distribution of general adaptations in birds arriving in a new environment (Darwin, 1837, B 55e), "without deprecating it as a facade which concealed our ignorance" (Manier, 1978, pp. 122–123). Porter rightly notes that Darwin's work "can only in retrospect be construed as statistical" (1986, p. 134). He goes on to describe a series of letters between Karl Pearson and Francis Galton (with input from several of Darwin's descendants). Pearson had hoped to show that Darwin's own work ought truly be considered to be statistical (in line with Pearson's own predilections), but Galton, after consulting with the Darwins, wrote that "I fear you must take it as a fact that Darwin had no liking for statistics" (Porter, 1986, pp. 134–135*nn.*).

Thus we have, throughout the *Origin,* the pervasive feeling that natural selection is intended to be a theory that utilizes only traditional, non-statistical, even largely deterministic sorts of explanations – explanations that are intended to be justifiable on Herschel's Newtonian-derived *vera causa* standard. Several authors, particularly Hodge, have argued that Darwin's theory was explicitly modeled on the ideal for scientific theorizing depicted in Herschel's *Preliminary Discourse.*[3] It is for this reason that Darwin was particularly stung by Herschel's dismissal of the *Origin.* "I have heard," Darwin wrote in a letter, "by a round-about channel, that Herschel says my book 'is the law of higgeldy-piggeldy.' What exactly this means I do not know, but it is evidently very contemptuous. If true this is a great blow and discouragement" (Hull, 1973, p. 7). Darwin was no radical on this score – he had hoped that his theory would be fully legitimate by Herschel's largely Newtonian and deterministic lights.[4]

---

3. Hodge's contribution is a remarkable series of papers: (1977; 1987; 1989; 1992a; 2000; 2009). For others, see also Lennox (2005); Lewens (2009); Waters (2009); Hull (2009).

4. I lack the space to explore it here, but Depew and Weber offer an interesting sociological explanation for Darwin's mistrust of statistics. Statistics, they claim, "had been an uninvited, or at least potentially obnoxious,

2.2.2. Darwin on Chance

What about Darwin's relationship to some sort of objective chance? Within the evolutionary process, Darwin identifies two loci for the operation of chance. The first is the role of chance in the generation of the variation upon which natural selection is supposed to act. Frequently, Darwin argues for the existence of this variation by extrapolation from our experience with domesticated plants and animals. "Can it, then, be thought improbable," he asks, "seeing that variations useful to man have undoubtedly occurred, that other variations useful to some being in the great and complex battle of life, should *sometimes* occur in the course of thousands of generations?" (Darwin, 1861, p. 80, emphasis added). Elsewhere he notes that horticulture, throughout the ages, "has consisted in always cultivating the best known variety, sowing its seeds, and, when a slightly better variety has *chanced to appear,* selecting it, and so onwards" (Darwin, 1861, p. 37, emphasis added).

Further, he often draws analogies with how such variation would work in natural populations. He wonders in the notebooks "whether species may not be made by a little more vigour being given to the chance offspring who have any slight peculiarity of structure, hence seals take victorious seals, hence deer victorious deer, hence males armed & pugnacious [in] all order[s]" (Darwin, 1838b, c 61). Further, one of the goals of the analogy with artifical selection developed in the first chapters of the *Origin* is to bolster our faith that the variations mentioned above in the artificial-selection context also occur in nature. Darwin argues, for example, that changes in environmental conditions in wild populations – as is already known

guest in the respectable, largely Whiggish halls of the BAAS" (1995, p. 151), a harbinger of Marxism, fatalism, or a lack of personal responsibility. The conservative British scientific establishment reacted accordingly, and eventually pushed statistical study out to the Statistical Society of London (Hilts, 1978).

in the case of domestic populations – can increase the chances that beneficial variations will be produced (Darwin, 1861, p. 82).

He seems, however, to be uncomfortable with the prominent role of chance here – a suspicion confirmed all the more by Darwin's incessant work on pangenesis, which would have offered a deterministic, atomistic explanation for the generation of evolutionary variation (Hodge, 1985). At one point in the notebooks, discussing strength in blacksmiths, he writes that in addition to the inheritance of acquired characters, "the other principle of those children, which *chance?* produced with strong arms, outliving the weaker ones, may be applicable to the formation of instincts, independently of habits" (Darwin, 1838c, N 42). The emphasis here is Darwin's own – he seems to be a bit incredulous that chance can be the proper explanation for the appearance of variation, though he at the time has no better story to offer. Throughout the development of evolutionary theory it is "[m]ere chance, as we may call it, [that] might cause one variety to differ in some character from its parents" (Darwin, 1861, p. 111).

The second role Darwin sees for chance in the process of evolution derives from the fact that natural selection is not a perfect discriminator – it is merely the case that a profitable variation "will *tend* to the preservation of that individual" which bears it, and this will lead that individual's offspring to "thus have a *better chance* of surviving" (Darwin, 1861, p. 61, emphasis added). It must surely be the case, he argues, that "individuals having any advantage, however slight, over others, would have the *best chance* of surviving and of procreating their kind" (Darwin, 1861, p. 81, emphasis added).[5] Nothing, however, guarantees a particular

5. References to organisms' "chance of surviving" or "chance of leaving offspring" are one of Darwin's most frequent refrains, and are *incredibly* common throughout Darwin's work. For only a small (!) cross-section of examples, see Darwin (1838d, E 137), Darwin (1861, pp. 5, 88, 90–92, 104, 109, 114, 127, 136, 176, 235, 388),

individual's success – the best the evolutionary process has to offer is the promise of higher fitness. In a passage which nicely exhibits both of Darwin's senses of chance, he writes that natural selection is the process by which "every slight modification, which in the course of ages *chanced to arise,* and which in any way favoured the individuals of any of the species, by better adapting them to their altered conditions, would *tend to be preserved*" (Darwin, 1861, p. 82, emphasis added).

As an aside, three other invocations of chance by Darwin are interesting. First, he does refer – though only twice, and only in the notebooks – to 'chance'-like influences from the environment. He notes that "chance & unfavourable conditions to parent may become favourable to offspring" (Darwin, 1838d, E 26e), and that motives are affected by "effect of heretidary constitution,— education under the influence of others— varied capability of receiving impressions— *accidental* (so called like chance) circumstances" (Darwin, 1838e, O 25). We might read these references as something like an early precursor to the notion of genetic drift. Second, in the *Descent,* Darwin frequently makes a distinction between "mere" chance and *choice* in sexual selection – for example, "if, then, the females do not prefer one male to another, the pairing must be left to mere chance, and this does not appear to me a probable event" (Darwin, 1871, pp. 1:400).[6] Finally, and perhaps most peculiarly, Darwin makes a similar analogy with free will. He muses in Notebook M that "thinking over these things, one doubts existence of free will every action determined by hereditary constitution … we may easily fancy there is [free will], as we fancy there is such a thing as chance.— chance governs the descent of a farthing, free will determines our throwing it up.— equall true the

Darwin (1871, pp. 161, 265, 319–320, 406, 414).

6. See also Darwin (1871, pp. 1:421, 2:117, 2:124, 2:273, 2:358).

two statements. —" (Darwin, 1838a, M 27, original punctuation and spelling). He emphasized the following passage later in the same notebook: "I verily believe free-will & chance are synonymous.— Shake ten thousand grains of sand together & one will be uppermost: — so in thoughts, one will rise according to law" (Darwin, 1838a, M 31).

Returning to the two uses of chance that matter most for us here – chance in the appearance of variation and in the survival of organisms – what does Darwin actually *mean* by the term 'chance'? He only rarely speaks of 'chance' at an abstract level, and when he does, he almost always has something like the law of large numbers in mind. For example, when discussing the issue of sex ratio in the *Descent,* he considers a population with an excess of males. Some organisms in such a population would happen to produce more female offspring. "On the doctrine of chances," he writes, "a greater number of the offspring of the more productive pairs would survive; and these would inherit a tendency to procreate fewer males and more females" (Darwin, 1871, p. 316). In the notebooks, in a discussion of the adaptation of birds to as many various circumstances as possible, Darwin wrote vertically up the side of the page that "law of chance would cause this to have happened in all" (1837, B 55e). The "doctrine" or "law" of chance here is just an expression that in the long run, a slight mathematical advantage in survival or reproduction will win out.[7] We have here no apparent invocation of objective chance.

And in particular examples rather than in the abstract, Darwin only rarely discusses what he takes the correct interpretation of chance to be. One of his only sustained considerations of the issue, at the beginning of the fourth chapter of the *Origin,* is commonly

7. Applications of this law of large numbers are also found throughout Darwin's work. For example, "the chance of [favorable variations'] appearance will be much increased by a large number of individuals being kept" (Darwin, 1861, p. 41) under artificial selection. See also Darwin (1861, pp. 102, 110, 177).

cited:

> I have hitherto sometimes spoken as if the variations – so common and multiform in organic beings under domestication, and in a lesser degree in those in a state of nature – had been due to chance. This, of course, is a wholly incorrect expression, but it serves to acknowledge plainly our ignorance of the cause of each particular variation. (Darwin, 1861, p. 131)

This is as direct an expression of a subjective, unpredictability, or ignorance interpretation of chance as we might hope to find. Darwin explains that whenever he makes reference to "chance," it is merely an indication that we lack knowledge or predictive power with respect to the particular causes of the phenomenon at issue. He goes on to note that one might ascribe the source of variation to the reproductive system, the conditions of life of the parents, climate, food, and so forth. All of these are, that is, possible *true* causes of variation – we simply lack the precision to determine which is genuinely responsible for variation in a given case (or even in the majority of cases).

In the revised 6th edition of the *Origin,* Darwin is even more explicit in his insistence that all apparent "chance" variations are in fact deterministically caused (1876, pp. 6–8).[8] "Each of the endless variations which we see in the plumage of our fowls must have had some efficient cause," he claims (Darwin, 1876, p. 6). Some of this variability is "indefinite" – that is, is probably unpredictable in principle.[9] This is, for example, the sort of variation which even can occur between closely related organisms living in identical environmental conditions. "Definite" variability, on the other hand, is variability induced as the result of varying the environment. This kind of variability, for example, is the cause of the failure of many captive

8. He also considers the issue of the causes of variation at some length in Darwin (1875, pp. 260–282).

9. "Unpredictable in principle," that is, in the sense that predicting the variations are likely forever beyond the cognitive capacities of human naturalists. See chapter 1 for more about the varying senses of unpredictability in principle.

animals to produce offspring (Darwin, 1876, p. 7). Again, however, whether unpredictable in principle or only in practice, both of these sorts of variation are chancy only insofar as we lack complete knowledge of biological systems.

Finally, we have Darwin's famous discussion of chance from the *Variation of Animals and Plants Under Domestication.* He considers the objection, by that point quite familiar, that "selection explains nothing, because we know not the cause of each individual difference in the structure of each being" (Darwin, 1875, p. 427). To reply to this objection, Darwin asks us to consider an analogy. When rock falls from the face of a cliff, he argues, we might call the shape of the fragments that result accidental,

> but this is not strictly correct; for the shape of each depends on a long sequence of events, all obeying natural laws; on the nature of the rock, on the lines of deposition or cleavage, on the form of the mountain, which depends on its upheaval and subsequent denudation, and lastly on the storm or earthquake which throws down the fragments. (Darwin, 1875, p. 427)

We then imagine assembling a structure from these stone fragments. Of course, Darwin argues, an omniscient creator could foresee all these events. But ought we really infer that all the natural laws that caused the stone to take its current shape exist *for the sake of* the structure that the builder eventually builds from them? Clearly not, he implies. It is in this sense that the shape of the stones is *accidental.* And natural selection works in the same way. Many of the variations in organisms are not useful or pleasing to either man or to the animal itself (and many of the artifically selected variations which are pleasing to man are deleterious to the organisms). They are the result of lawlike causal processes, but there is no sense – divine or otherwise – in which the laws are they way they are *for the sake of* the development of some particular character in some particular organism. There is no overarching pattern to find, and for this reason, and only in this sense, can we view the evolutionary process as

"chancy."

All these examples are traditionally cited when discussing Darwin's understanding of chance, and the view that Darwin held an ignorance interpretation of chance can fairly be called the standard interpretation of Darwin on chance.[10] But this agreement masks the interesting depth of Darwin's thought on the matter. We see throughout these quotes an interplay of three distinct ways of understanding chance. First is simple subjective unpredictability as described in the last chapter. Second, and much more important for Darwin, is the concept of 'accident' – the lack of any sort of overarching cause or design, any "for the sake of which" or final cause. Finally we have objective chance, which for Darwin (and his interpreters, as we will see below) is consistently read as some sort of *lack of causation.*

This last sense – objective chance – is categorically rejected by Darwin. It is clear that Manier is correct when he states that Darwin "attributed no causal force to chance itself" (1978, p. 121). All causes, in Darwin's view, are clearly still perfectly Newtonian; both the stone building example in the *Variation* and the variation example in the sixth edition of the *Origin* are very clear about this position. Variation, in general, is more about unpredictability for Darwin – it is the bulk material, viewed throughout the *Origin* as a black box, a fact that provides a necessary and empirically well-confirmed (if inexplicable) *input* to the evolutionary process.

When Darwin discusses the possibility of chance in the process of natural selection, on the other hand, his worry is with *design,* and hence he is primarily concerned in this arena

10. Commentators to argue for such a view include Hull (1973, pp. 62, 426–427), Hodge (1987, p. 243), Depew and Weber (1995, p. 113), and Beatty (2006a, p. 630). Manier slightly complicates the picture, by arguing that chance as used in the notebooks refers to an Aristotelian notion of chance as coincidence (1978, pp. 117–123). We still, however, have chance events both as ignorance and as caused, with no independent, causal notion of chance.

with chance in the sense of *accident.* While it may be a consequence of selection that many features of organisms are accidental, Darwin did not want them to be "chancy" in any more robust sense than this. Hodge compellingly argues that Darwin was only willing to consider variation as generated by an unpredictable, black-box process once he had introduced into his theory the highly *non-chancy* pseudo-agency of natural selection, "analogous to the skilled practice of the breeder's quasi-designing art" (1987, p. 246). We find evidence for this claim in many parts of Darwin's work: in his insistence on the use of the analogy between artificial and natural selection (Sterrett, 2002), in the metaphor of selection's "daily and hourly scrutinising, throughout the world, every variation" (Darwin, 1861, p. 84), and, more broadly, in the connections that have been drawn between Darwin's work and the German Romanticist tradition (Richards, 1992, 2009). We need not read some sort of active agency into Darwin's depiction of selection to see that he held selection's ability to produce the appearance of design in very high esteem.

To return to our original question, however – whether or not Darwin thought the evolutionary process was objectively chancy – we have ample evidence to answer this question firmly in the negative. Evolution is a theory of non-chancy processes for Darwin. But Darwin did not make the distinction between these various senses of chance particularly clear – as we will now see, the various authors to respond to Darwin were often quite confused as regards the interplay of these various meanings of chance.

2.2.3. The Response to Darwin

Darwin's efforts to craft a strictly *non*-causal, ignorance interpretation of chance and then sequester this notion of chance as far as possible from his non-statistical theory of evolution by natural selection certainly go far to explain why Herschel's rejection of Darwin's theory as the "law of higgeldy-piggeldy" stung Darwin so badly. As Depew and Weber note, we must not think that Herschel believed that it was *merely* chance that explained adaptation and speciation. Rather, he was claiming that "natural selection," precisely as Darwin described it, "cannot count as a law of nature or as a true cause" (1995, p. 149). Despite Darwin's efforts to the contrary, they argue, Herschel reads the *Origin* and comes away certain that "what Darwin had done was push Quetelet's social arithmetic down into the biological world and then claim that he had found in natural selection a law of nature" (Depew and Weber, 1995, p. 152). Darwin, as we have seen, had intended to do no such thing.

Herschel, however, was not alone in finding fault even with Darwin's limited use of chance. F.W. Hutton, writing in *The Geologist,* claims that "the greatest objection" to Darwin's theory "is its reliance on *natural* causes and *chance* in effecting the changes" (Hull, 1973, p. 299, original emphasis). Mivart claims that "the theory of Natural Selection may (though it need not) be taken in such a way as to lead men to regard the present organic world as formed, so to speak, *accidentally,* beautiful and wonderful as is confessedly the hap-hazard result" (Hull, 1973, p. 388, original emphasis). Von Baer retells the story of Gulliver's visit to Lagado, where they fill books with words randomly generated by machine (Hull, 1973, p. 419). Even Darwin's defenders got in on the act – Chauncey Wright chided him, in a response to Mivart,

as Darwin

> has not in his works repeated with sufficient frequency his faith in the universality of the law of causation.… [I]n referring any effect to 'accident,' he only means that its causes are like particular phases of the weather, or like innumerable phenomena in the concrete course of nature generally, which are quite beyond the power of finite minds to anticipate or account for in detail, though none the less really determinate or due to regular causes. (Hull, 1973, p. 388)

It is worth noting, however, precisely what is at stake when these authors argue over the concept of "chance" in Darwin's theory. In fact, Wright's riposte to Mivart is off the mark – however correct it may sound on a contemporary conception of objective chance. For the intent of those like Mivart and von Baer is not to claim that Darwin has abandoned causation. It is a shared premise among all these authors, Darwin and critics included, that the truly salient feature of the "chance" which Darwin's theory introduces is a lack of design – *accident,* without overarching cause, as we described it above. Thus did Darwin become (in)famous: not as the opponent of causation, but as the opponent of design.

Such, then, is the state of affairs as of 1859. Darwin has proposed the theory of evolution by natural selection, a non-statistical theory of non-chancy processes in nature. Complex processes, to be sure – processes the details of which may forever escape our knowledge. But the theory itself is intended to conform to Herschel's *vera causa* ideal: the ideal which, according to Herschel, grounds the explanatory power and prowess of Newtonian mechanics. While Darwin may have been far more willing to appeal to (again, his sense of) chance than many of his contemporaries, and while he may have placed much more of the living world under the guidance of an accidental process free of final causes than those who had come before him, we don't see a drastic shift in the role of either statistical theorizing or objective chance in Darwin's work. As of yet, we have seen neither of the historical events for which we are searching. Let us then move forward, to Francis Galton.

2.3. Statistical Theories: Francis Galton

What was the main driving force behind the shift from a non-statistical to a statistical theory of evolution by natural selection? As it turns out, it was an old problem. As early as the "Sketch" of 1842, Darwin was worried about the problem of *blending inheritance.* He writes that "if in any country or district all animals of one species be allowed freely to cross, any small tendency in them to vary will be constantly counteracted" (Darwin, 1909, p. 3), destroying the power of natural selection to alter the species. The point was made far more serious in the review of the *Origin* by the engineer Fleeming Jenkin (1867).

Gayon notes that the thrust of this paper is often misunderstood (1998, pp. 96–97). Jenkin is not merely concerned with the apparent reliance of Darwin's theory on "sports," or large deviations of characters from parent to offspring. Rather, he notes the following two interrelated (and much more complex and significant) problems with Darwin's theory as expressed in the *Origin.* First, how is variation distributed? If the distribution is continuous, then we must use statistics to describe it. If, on the other hand, it is not a continuous, populational sort of variation, but rather individual and isolated instances, these instances must be measured, and the odds of some particular variation being eliminated by chance must be determined. Second, what is the method of transmission of characters to offspring? If offspring carry a mixture of the characters of their parents, as Darwin and most others assumed, how can the problem of regression to the mean be avoided? Gayon notes that this constitutes an "impressive list of problems" for future theorists:

> Using modern vocabulary, they can be summarized as follows: continuity or discontinuity of hereditary variation, blending or particulate inheritance, the effect

> of the mating system, the quantitative nature of 'advantage' and the interaction between selection and samping effects. (Gayon, 1998, p. 97)

Gayon describes this as a "dilemma" for Darwin. Natural selection is sometimes described in the *Origin* as "a principle of the progressive accumulation of variations that appear in an isolated manner in individuals.… But Darwin also presented natural selection as acting on 'infinitesimally small' variation, in other words, on continuous variation" (Gayon, 1998, p. 97). When Darwin refers to "individual variation," then, he thus elides the difference between these two ways of presenting selection. Depew and Weber argue that Darwin's own response to this problem is highly unsatisfactory. In the last two editions of the *Origin,* all he does to respond to this charge is to posit the existence of more continuous variation and fewer "sports" – as Depew and Weber note, "by fiat," changing the singular nouns referring to variation to plurals (1995, p. 196).

Given that Darwin's response here was so unsatisfying, what was to be done about the dilemma that Jenkin raised? The long-term solution, of course, was the rejection of the blending model of inheritance. But this would have to wait for the "rediscovery" of Mendel's paper and the birth of genetics, almost thirty years after Darwin's death (Druery and Bateson, 1901).

In the meanwhile, defenders of Darwin's theory sought refuge in statistics – taking the 'continuous variation' horn of the dilemma, and establishing how natural selection could work in a gradualist, statistical manner. The most prominent such early defender was Darwin's cousin Francis Galton. The publication of Darwin's *Origin* sparked in him a deep interest in breeding – particularly in eugenics and the heredity of human intelligence and other abilities. To that end, Galton found two things unsettling about the about the trouble with blending

Figure 2.1: Karl Pearson (left) and Francis Galton (right), taken in 1909, two years before Galton's death.

inheritance. First was the potential undermining of his cousin's theory of evolution by natural selection, which Galton had described in a letter to Darwin as engendering "a feeling that one rarely experiences after boyish days, of having been initiated into an entirely new province of knowledge which, nevertheless, connects itself with other things in a thousand ways" (Galton and Darwin, 1859). But second, and more importantly, were the eugenic implications of the blending argument. Assume, as the eugenics movement invariably did, that we have an accurate idea of which traits are superior and which families are in possession of them: respectively, intelligence, culture, and the other values of the Victorian British aristocracy, and the few upper-class families (such as the Wedgwood-Galton-Darwins) who clearly deserved them. The British upper-crust, therefore, constitutes the far right-hand tail of the normal distribution of variation in the human population. Here enters Jenkin's problem. For unless heredity and variation work in precisely the right way, it remains possible that the eugenic program is a failure before it begins: that *even with* the aid of severe eugenic programs, we will still be unable to preserve these characters within the families that are entitled to them.

Thus was the problem of blending inheritance doubly magnified for Galton. How did he propose to resolve it? He began with a radically different view of the way in which inheritance operates. Relatively early during his study of heredity, Galton shifted to a population-based, statistical view of the transmission of characters from parents to offspring. As Porter notes (1986, p. 136), Galton began with the work of Quetelet. To use Quetelet's framework, however, required that we claim that there is a substantial degree of statistical similarity between the characters in the population of parents and those in the population of offspring. How could we do this in the biological context? Not by biological experiment or mathematical derivation, surprisingly, but by a *social* analogy: the transmission of hereditary characters is

analogous to "indiscriminate conscription: thus, if a large army be drawn from the provinces of a country by a general conscription, its constitution, according to the laws of chance, will reflect with surprising precision, the qualities of the population whence it was taken" (Galton, 1872, p. 397). Thus was the bridge between Quetelet's social statistics and the world of biological inheritance built.

Galton used this perspective to develop a view of particulate inheritance on which many small heritable factors – some "latent" and some "developed" or "patent" in the adult – combine and compete for a small number of "places" within the offspring. The closest metaphor we can create for such inheritance, Galton writes, is this. Consider "an urn containing a great number of balls, marked in various ways, and a handful to be drawn out of them at random as a sample: this sample would represent the person of a parent [his or her developed characters]." Then we mix another, similarly sized urn in with the first, representing the contribution of the other parent, and draw out a second sample. "There can be no nearer connexion justly conceived to subsist between the parent and child than between the two samples" (Galton, 1872, p. 400).

The very foundations of heredity, therefore, can now be considered statistically – as a vast, population-level urn-drawing experiment. From here, we can turn toward mathematizing the relationship between parent and offspring. Gayon argues that

> right from the outset, Galton viewed Darwin's 'individual differences' in terms of language and images taken directly from Quetelet's social statistics. Beginning with his 1865 text on 'hereditary talent and character', it is clear that, for Galton, 'heredity' was not to be considered only in terms of the individual, but rather in terms of the family and the race.… (Gayon, 1998, p. 117)

Heredity thus was, from the time of Galton's first articles on the subject, best dealt with at a population-based, statistical level. Radick (reporting Weldon's view of Galton) describes

Galton's use of "the throwing of dice as a model for relations between generations" as having "culminated in the famous Galton-Pearson law of ancestral heredity" (Radick, 2011, p. 133), to which we now turn.

The primary mathematical contribution to evolutionary theory made by Galton himself, the law of ancestral heredity describes the extent to which the contribution of heritable characters in ancestors influences the characters of offspring – "the integration of *all* hereditary phenomena in a single conceptual framework or expression," in the words of Gayon (1998, p. 132). In *Natural Inheritance,* Galton describes the law as follows:

> [T]he influence, pure and simple, of the Mid-Parent [the average of the mother and father] may be taken as ½ and that of the Mid-Grand-Parent [the average of all four grandparents] as ¼, and so on. Consequently the influence of the individual Parent would be ¼, and of the individual Grand-Parent 1/16, and so on. It would however, be hazardous on the present slender basis, to extend this sequence with confidence to more distant generations. (Galton, 1889, p. 136)

To see what Galton is attempting to do here, consider the characters of an offspring. We know that there is a strong force of regression to the mean, so the interesting question becomes: at what fidelity are the characters of parents (and earlier ancestors) transmitted to their offspring? Galton first determined empirically that the coefficient of correlation between sons and 'mid-parents' was ⅔. However, this correlation includes not only characters from the parents themselves, but also some from the grandparents passed on to the parents and then the offspring – we have to "factor out" this grandparental contribution if we want to determine the "pure" contribution of the parent. By two separate estimations (one assuming a constant diminution of transmission in all generations and one assuming a diminution that increases over time), Galton arrives at the value of ½ for the mid-parent contribution.

Galton's technical conclusions aside, we can clearly see the tools and methods of statistics deeply embedded in his work. The value upon which the entire derivation of the

law of ancestral heredity rests, the mid-parent to offspring correlation of ⅔, was determined empirically via regression on measurements of height, and Galton sought to confirm it via statistical measurement of moth populations, human eye-color, artistic talent, disease, and so forth. Gayon, describing the impact of Galton's statistical work on the study of evolution, notes that

> Galton's statistical studies gave a quantitative and functional interpretation to concepts which in Darwin's writings remained undeveloped and even mysterious: 'individual differences', 'heredity', 'variability' and 'reversion'. (Gayon, 1998, p. 131)

Galton has, indisputably, made evolution a statistical discipline, and statistics proved to be here to stay. We thus very clearly find in the work of Galton the first historical event for which we have been searching – the theory of evolution is now a statistical one.

But note the depth of the use of statistics – Galton's concern with eugenics and breeding means that we don't find this statistical viewpoint throughout the foundations of evolutionary theory, but only in heredity. Galton at points gestures at a statistical view of natural selection (e.g., Galton, 1877b, p. 533), but not in anything like a robust or empirically grounded way. The thorough integration of statistics into the rest of evolutionary theory would be executed by Pearson and Weldon, to whom we will return below.

2.3.1. Galton on Chance

What about the role of objective chance in Galton's theorizing? He is nearly silent on this issue, but we can divine two conclusions. First, return to Galton's discussion of his statistical theory of heredity. Galton sees both the transmission of elements to offspring and the development of organisms as complex but necessarily strictly Newtonian or mechanistic

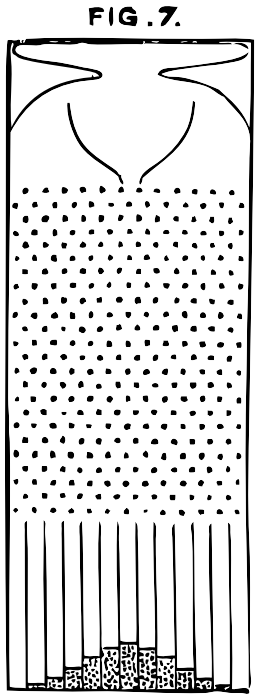

Figure 2.2: Galton's quincunx device used to demonstrate the normal curve.

causal processes. He describes "segregation" as a straightforward process of competition (which Radick (2011) has likened to natural selection), saying that "for each place [in an organism's set of developed characters] there have been many unsuccessful but qualified competitors" (Galton, 1872, p. 395). On development, he says that if we had sufficient information, "statistical experiences would no doubt enable us to predict the average value of the form into which they would become developed … but the individual variation of each case would of course be great, owing to the *large number of variable influences* concerned in the process of development" (Galton, 1872, p. 396, emphasis added). This sounds much like Quetelet's view of social statistics as the result of the aggregation of a myriad small, non-statistical causes.

Second, we can consider Galton's famous use of the quincunx device. Consider the outcome of dropping a handful of shot into the top of the device depicted in figure 2.2. The

shot falls through the series of pins set in the board, and collects at the bottom in a series of bins. The shot will, Galton notes, pile up in these bins in precisely the distribution described by the normal curve (shown at the bottom of the device). Importantly for us, consider Galton's description of how the device approximates the law of errors:

> The principle on which the action of the apparatus depends is, that a number of small and independent accidents befall each shot in its career. In rare cases, a long run of luck continues to favour the course of a particular shot towards either outside place, but in the large majority of instances the number of accidents that cause Deviation to the right, balance in a greater or less degree those that cause Deviation to the left. […] This illustrates and explains why mediocrity is so common. (Galton, 1889, pp. 64–65)

Setting aside the moral overtone present in Galton's invocation of 'mediocrity', here we have another invocation of precisely the same sense of chance we saw expressed by Darwin. The law of errors is nothing more or less than the record of a very large number of small, deterministic causes acting on the same system over time – and it is merely our ignorance, or our inability to follow those "runs of luck," that makes the use of statistics necessary.

But the full tale of Galton's view of chance must be slightly more complicated than this. For the analogy between the quincunx and the "large number of variable influences" in heredity is not a perfect one. The most common way of describing the correct understanding of the law of errors in Galton's day comes from Quetelet, who, Galton was right to note, believed that the "essence" of the law of errors "is that differences should be wholly due to the collective actions of a host of independent *petty* influences in various combinations, as was represented by the teeth of the harrow [in the quincunx]" (Galton, 1877a, p. 512). But this cannot be the full explanation for the facts of heredity. "[A]lthough characteristics of plants and animals conform to the law [of errors]," he argues, "the reason of their doing so is as yet totally unexplained," because the processes of heredity "are not petty influences, but very

important ones" (Galton, 1877a, p. 512). Thus, we are forced to conclude "that the processes of heredity must work harmoniously with the law of deviation, and be themselves in some sense conformable to it" (Galton, 1877a, p. 512). While Galton does not therefore believe the statistical account of heredity is a *direct* analogue of the behavior of the quincunx, we must explain the fact that the various non-statistical and decidedly *non*-petty processes of heredity are "conformable" to statistical explanation. Whatever the account of this coincidence, we clearly have no cause for inferring anything other than an interpretation of chance as unpredictability for Galton. The laws derived for the various processes of heredity, he argues, "may never be exactly correct in any one case, but at the same time they will always be approximately true and always serviceable for explanation" (Galton, 1877b, p. 532). It is clear that it is merely our ignorance of the precise details of these processes that makes higher-level statistical laws necessary and particularly "serviceable." If we are looking for the first invocation of objective chance, we are not to find it in the work of Galton.

2.4. Pearson and Weldon: Minor Characters?

Following the narrative of Depew and Weber, we would now turn our focus forward to the work of Sewall Wright, where we would find the first instance of chancy evolutionary processes in his shifting-balance theory. We will not here, however, evaluate this second historical event – Depew and Weber may well be quite correct that the first instance of objectively chancy theories is found in the work of Wright. Rather, I want to advance a much shorter increment, to two of Galton's students – Karl Pearson and W.F.R. Weldon.

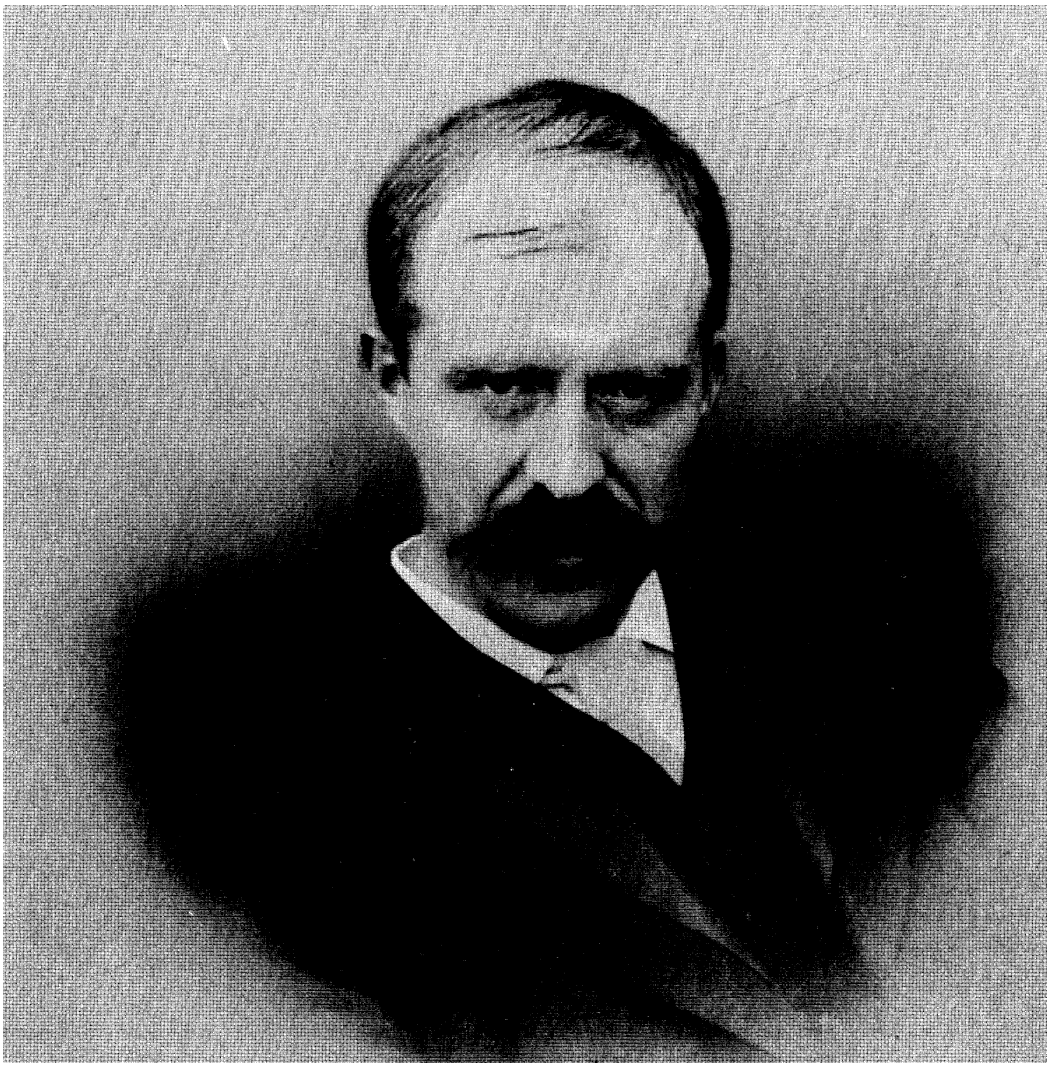

Figure 2.3: Walter Frank Raphael Weldon.

Pearson, whose life has been masterfully detailed by Porter (2004), was a particularly interesting character. He studied at Cambridge, and after having received his degree with Mathematical Honors, departed for Germany, becoming deeply affected by the Romantic tradition and publishing and lecturing in German history.[11] He developed an intense interest in socialist politics as well as women's rights. Finally, upon returning to England, he was appointed chair of Applied Mathematics and Mechanics at University College, London, where he was chiefly occupied with teaching mathematics to students of engineering. His work there, which included teaching geometry and drawing, would have a substantial influence on the significant visual aspect of his later work in statistics.

His completion of William Kingdon Clifford's *Common Sense of the Exact Sciences* (1885) provided an early glimpse of Pearson's philosophy of the physical sciences. Broadly positivist in nature – that is, emphasizing the importance of mathematical formulae in the development of scientific knowledge and espousing a strict form of empiricism – he would go on to develop this philosophy of science in his widely known *magnum opus,* the *Grammar of Science,* first published in 1892 and later revised and expanded (with more material on evolution) in 1900.

W.F.R. Weldon, known as Raphael, was born in 1860.[12] He attended University College and later King's College, studying biology under Lankester and Balfour. After finishing there, he worked at both the Naples Zoological Station and Cambridge, finally being appointed as Lecturer in Invertebrate Morphology at University College, London in 1884. He became

11. This brief biography follows that presented by Pearson's son in Pearson (1936).

12. This brief biography follows Pearson's memorial of Weldon (Pearson, 1906). No scholarly biography of Weldon has as yet been prepared.

quite active in the Marine Biological Laboratory at Plymouth after its construction in 1888, eventually running several large-scale experiments there.

In 1890, Weldon was appointed to the Jodrell Professorship of Zoology at University College, London, and Weldon and Pearson quickly formed a friendship. Pearson had been driven to the study of biology by reading Galton's *Natural Inheritance,* which had been published the year before (Pearson, 1936, pp. 210–211), and early in 1890, Weldon had published his first work applying statistics to biology (Weldon, 1890). The mathematics in Weldon's paper had been prepared under the direct tutelage of Galton himself, who was sent the paper as a referee (Pearson, 1906, p. 17). A bit later, in November 1891, Pearson delivered the Gresham College Lectures in Geometry. He would go on to deliver thirty lectures in this series on the subject of chance and statistics – in particular, focusing on visual aids and graphical representations of various kinds to make the material accessible to a broad student audience (Porter, 2004, pp. 235–236). As time went on, however, evolution featured ever more strongly in the lecture content. We can see, Porter notes, "a vision … taking form, even as he wrote his lectures, that evolution by natural selection could be comprehended statistically" (Porter, 2004, p. 238).

By this point, then, the collaboration between Pearson and Weldon was off to the races. Nearly all of Weldon's papers from 1893 until his untimely death from pneumonia in 1906 involved statistical collaboration with Pearson, and Pearson would publish a series of some twelve papers titled "Mathematical Contributions to the Theory of Evolution," describing various applications of statistical methodology to the evolutionary process. With Weldon's death and the increasingly hostile climate of the battle between the biometricians and Mendelians, Pearson would largely abandon the study of biology after 1906, taking

over the directorship of Galton's Eugenics Laboratory at University College (Pearson, 1936; Magnello, 1999a,b).

In the remainder of this chapter, we will consider how Pearson and Weldon fare as regards the two historical events that have formed our framework here: do they utilize a statistical theory of evolution, and do they consider this statistical theory to be undergirded by an objective notion of chance extant in the world? We will discover that, on Depew and Weber's view, Pearson and Weldon are relatively garden-variety: their situation with respect to our two focal questions is only a slight advance on their mentor Galton's.

As far as the statistical nature of biological theory, they play an important but relatively minor role. Whereas Galton, as we saw in the last section, deployed statistical notions primarily within the study of heredity, Pearson and Weldon brought statistics to the study of variation, heredity, inheritance, correlation, and natural and sexual selection – the entire breadth of evolutionary study. While Galton, as we mentioned above, viewed *portions* of evolutionary theory statistically, the *entirety* of evolution was to be studied statistically for Pearson and Weldon. Pearson, for example, in the second edition of his *Grammar of Science,* claims that only the growth of the statistical picture of evolution had made it possible to provide a "precise definition of fundamental biological concepts" (1900, p. 372). We thus have the introduction of a thoroughgoing statistical methodology in portions of evolutionary theory where Galton's use of statistics and empirical data had only been cursory. Pearson and Weldon's combination of mathematics and experiment was exceptionally productive.

What about their views on the role of an objective notion of chance in biology? I will consider in turn the two places where Pearson and Weldon invest the most time in explicating the philosophical underpinnings of their view of chance: the two editions of

Pearson's *Grammar of Science,* and a lecture by Weldon entitled "Inheritance in Animals and Plants."[13]

2.4.1. Pearson's *Grammar of Science*

The first edition of Pearson's magnum philosophical opus, his *Grammar of Science,* was published in 1892. Though the two men were unknown to each other, we can recognize what we would now call a "Machian" view of physics as much of the motivation for Pearson's *Grammar* – indeed, Mach would write to Pearson in 1897, plaintively noting "how useful would it have been for me to know back in 1872 that I didn't stand *alone* in my efforts."[14] Pearson focuses extensively on the usefulness of science for the economy of thought, denigrates the speculative use of 'metaphysics' in science, and extensively praises an austere form of empiricism.

Because of this Machian or positivist bent in Pearson's work, he did not believe, nor could he consistently have believed, that our scientific theories somehow latch onto objective chance in the world. Objective, reified chance is an inhabitant of the realm of things-in-themselves, which Pearson barred from his philosophy. Indeed, Pearson claims that the object of the *Grammar* is to argue "that science is in reality a classification and analysis of the contents of the mind; and the scientific method consists in drawing just comparisons and inferences from stored sense-impresses and the conceptions based upon them" (Pearson,

13. This lecture by Weldon is discussed exceedingly rarely in the literature; see Radick (2011) and Pence (2011, reprinted here as the next chapter).

14. "Wie werthvoll wäre es mir gewesen schon 1872 zu wissen, dass ich mit meinen Bestrebungen nicht *allein* stehe." Ernst Mach to Karl Pearson, Jul. 12, 1897, published in Thiele (1969, p. 537).

1892, p. 63).

Such a strict empiricism seems to give Pearson two options for an entirely subjective usage of probabilities and chances, and he avails himself of both. First, we have the use of chance as credence. In what must be a very early expression of probabilism in epistemology, Pearson writes that

> We ought first to notice that the use of the word *belief* in our language is changing: formerly it denoted something taken as definite and certain on the basis of some external authority; now it has grown rather to denote credit given to a statement on a more or less sufficient balancing of probabilities. (Pearson, 1892, p. 71)

Second, and much more importantly for our purposes, Pearson offers an ignorance interpretation of the scientific use of probabilities, just as Galton and Darwin had before him. In a section titled "The Bases of Laplace's Theory lie in an Experience as to Ignorance" (Pearson, 1892, p. 171), he argues that the underlying justification behind the use of probabilistic claims in science is an equiprobability assumption, and this equiprobability assumption is justified as the best course of action in the face of ignorance: "In our ignorance we ought to consider before experience that nature may consist of all routines, all anomalies, or a mixture of the two in any proportion whatever, and that all such are equiprobable" (Pearson, 1892, p. 172). He goes on to offer an extensive justification of why our past experience with situations of incomplete information does indeed justify the use of equiprobability as a canon of legitimate inference.

Interestingly, however, Pearson fails entirely to explore these philosophical concerns about chance in the context of the biological sciences. As of the first edition of the *Grammar* in 1892, Pearson avows ignorance of biology – he spends the chapter entitled "Life" exploring the relationship between biology and physics, the varying hypotheses for the generation of life from non-life, a brief overview of natural selection, and the Spencerian application of

selection to the history of man, as these are more suited to "the range of the present work [and] the power of its author" (Pearson, 1892, p. 388).

In 1900, however, Pearson publishes a second edition of the *Grammar,* spurred in large part by his (by this point extensive) work with Weldon on biological problems. He notes in the second-edition preface that "the progress in this [i.e., the biological] direction during the last few years enables me to define several of these conceptions much more accurately than was possible in 1892" (Pearson, 1900, pp. viii–ix). The second edition contains two new chapters specifically dedicated to evolution – the first to variation and selection, the second to reproduction and inheritance. But these chapters, as well, do not consider the role of 'chance' in biology: rather, the focus is on the fact that we are finally endowed, thanks to statistics, with the ability to demonstrate evolution's action quantitatively. For example, after discussing the various types of selection that have been proposed, Pearson writes that "before we can accept [any cause of progressive change] as a factor we must not only have shown its plausibility, but if possible have demonstrated its quantitative validity" (Pearson, 1900, p. 380). And the "chancy" nature of biological processes is to be demonstrated empirically, not considered philosophically. As Pearson himself makes the point (referring to some collected data on the height of Englishmen):

> [J]ust as all boys of eleven years of age of the same height do not grow into a group of boys of the same height at thirteen, but into an array of definite variability, so all the progeny of an individual of organ or character *a*, form an array with definite variability, nor is the type of this array, that is, its mean, identical with *a*, but with an organ whose quantitative value is nearer to the modal value *c* than *a*. These, the main features of inheritance, are well established as we shall see later. (Pearson, 1900, p. 422)

The fact of the statistical nature of heredity here is established by empirical observation and taken as a given, not a matter for theoretical or philosophical reflection. We thus have no

evidence to support any conclusion other than that Pearson viewed 'chance' in biology just as he did 'chance' more generally: as ignorance, precisely in line with both Darwin and Galton before him.

2.4.2. Weldon's "Inheritance in Animals and Plants"

In 1906, a lecture by Weldon on the topic of inheritance (delivered the previous year) was published in a volume of *Lectures on the Method of Science.* While the bulk of the lecture is relatively uninteresting (if not downright confusing), Weldon begins by defending the use of statistical methods in the life sciences. While we will return to this text in more detail in the next chapter, a brief outline here will be useful for contextualizing Weldon's views on objective chance.

In physics, Weldon argues, statistics serves two common uses. First, it averages over errors in measurement. Weldon gives the example of the determination of the latitude of the Radcliffe Observatory – even though highly skilled workers are responsible for its measurement, the values obtained fall into a range, in terms of the observatory's position, of about thirty-four yards (Weldon, 1906, p. 86). One function of taking a single number and declaring it *the* latitude of the observatory, then, is to average over small errors in these various measurements. Of course, this use of the "law of error" – and the fact that the example is one drawn from observational astronomy – has its own historical explanation. The law of errors had already received extensive use in the astronomical realm; the term "correlation" was first used in French to refer to the systematic error introduced by inaccuracy in the

relative positions of two observatories observing the same phenomenon (Porter, 1986, p. 273). Weldon is, therefore, situating his defense of statistics in one of the first traditions in science to make extensive use of statistical inference.

Returning to Weldon's discussion, we also obviously have a purpose for determining the latitude of the observatory – it is a number that we wish to utilize in calculations. Weldon notes that even though it may be the case that there *is no single value* for the latitude of the observatory (due to, for example, changes over time in the position of the equator), we still are pragmatically required to ascribe the latitude a single value. That is, statistics can be useful for "attributing to the latitude of the Radcliffe telescope a constancy it does not really possess" (Weldon, 1906, p. 88).

The ideal, then, in the physical sciences, is to be able to use a method which can separate these two sorts of discrepancy – which can describe all of the results thusfar obtained, discarding human error in measurement *without* discarding the genuine variability in the data. The fine statistical variation in the latitude of the telescope, Weldon claims, does indeed have a physical source, which would have been missed if it was thought to be the result of human error: "the slight changes in the position of the Earth's axis, by which the latitude of a telescope is changed, not by a change in the position of the telescope but by a shifting of the Equator,—these changes have only been discovered by examining all the records of experience, and refusing to replace discordant observations by an imaginary uniformity" (Weldon, 1906, p. 93).

So much for the ideal in the physical sciences. Should we apply statistics to biology in precisely the same way? No, Weldon argues. The variation in the biological case is too great. "[I]f I tell you," he writes, "that Englishmen are 5 feet 7½ inches high, you remember

your father who is five feet ten, and your cousin who is over six feet, and you think I am talking nonsense" (Weldon, 1906, p. 94). The kind of simplifying use of statistics deployed in the physical sciences doesn't work in biology. Rather, we need a way to capture *all* of the variation in biological systems – we need to collect and preserve statistical data in its entirety, in order to come up with a complete description of our observations.[15]

Is this, then, evidence for the existence of objective chance on Weldon's view? No – he, as well, affirms a straightforward interpretation of chance as subjective unpredictability. "All experience, which we are obliged to deal with statistically, is experience of results which depend upon a great number of complicated conditions, so many and so difficult to observe that we cannot tell in any given case what their effect will be" (Weldon, 1906, p. 97). Weldon, again, follows Darwin, Galton, and Pearson in adopting a notion of chance grounded entirely in ignorance. The introduction of an objective notion of chance in evolution is not to be found in the work of either Pearson or Weldon.

2.5. Moving Away from the Standard History

The explanation given by the standard history of the early development of chance in evolution is relatively straightforward. We begin with Darwin, who develops a non-statistical theory of non-chancy biological systems. Galton, endeavoring to respond to the troubles of blending inheritance, statisticalizes the theory of heredity. Pearson and Weldon expand this usage of statistics to selection itself, making them minor innovators. Neither Galton nor

15. More on this motivation of Weldon's in the next chapter.

his students discard Darwin's ignorance interpretation of chance in the objective biological world – this was Sewall Wright's doing, introducing objective chance in the context of his shifting balance theory.

If we consider merely the two events brought out in Depew and Weber's analysis, it is not obvious why Pearson and Weldon are even worthy of mention at all, much less of a systematic development of their views on chance. On the contrary, the case of Pearson and Weldon, I argue, is an excellent example for use in teasing apart more thoroughly the history of the introduction of chance into evolutionary theory. Most worryingly, it is difficult to search for the first use of objective chance at all in Pearson and Weldon's case, for two reasons. First, as with all examinations of the metaphysical entailments of biological theories, we are hampered by biologists' uncertain attitude toward the metaphysical or ontological claims of their theories (see, e.g., Waters, 2011). Second, despite the fact that some early work on chance in the late nineteenth and early twentieth centuries did make room for the possibility of indeterminism, we have little evidence that 'chance' in an objective, reified sense was a concept even considered by any of the authors whose work we have considered here.[16]

These two factors conspire, then, to make it extremely unlikely that the actual practice of biologists like Pearson and Weldon will be particularly amenable to analysis in terms of Depew and Weber's search for objective chance.[17] And it seems likely, as well, that if we do not consider a question to which the work of Pearson and Weldon responds, any

16. For various objective uses of chance *prior* to the introduction of quantum mechanics, see, for example, Stöltzner (2008) on the Exner school in physics, Beatty (1984) for a brief mention of the relationship between Darwin and Peirce's tychism, or Dale (1999, p. 399) for John Venn's frequentist theory of probability.

17. Some of these issues, such as the concern about the uneasy nature of the metaphysical entailments of biological theories, remain troublesome even for contemporary work in the philosophy of biology, a problem which will be discussed in detail in chapter 4.

deeper understanding of the conceptual underpinnings of their biological theories will be uninformed, if not incorrect.

The instance of Pearson and Weldon is then, I claim, precisely the historical impetus we need to develop a new way of understanding the role of chance in the early development of evolution. A schism that developed between the two men, often unremarked-upon in the historical literature, reveals that they were engaged in a serious, long-standing debate not over the question of a reified, objective sense of 'chance', but over the question of *the relationship between biological systems and the statistical theories used to describe them.* It is clear that the two historical events of the Depew and Weber view do not let us capture this distinction, so it is beholden upon us to find a way of framing the issues that allows us to recognize and understand it. That will be the project of the next chapter.

CHAPTER 3

THE STATISTICAL PHILOSOPHIES OF W.F.R. WELDON AND KARL PEARSON

As I argued at the close of the last chapter, the case of W.F.R. Weldon and Karl Pearson provides the perfect historical impetus for the development of a novel way of understanding the question of chance in evolutionary theory. It is the goal of this chapter, then, to provide a detailed case study of the work of Pearson and Weldon, motivating the novel question which I introduced above and preparing us to develop it in more detail in the second half of the dissertation.

I have suggested at the end of the last chapter that we ought to approach the work of Pearson and Weldon by considering the relationship between the statistical theories these authors deploy and the biological systems that these theories describe. What is it about this question that makes it amenable to analysis in this important historical case, while the standard history's questions are less applicable? The answer lies in yet another historical contingency.

When Weldon and Pearson began to promote their statistical picture of biological theory – and in particular, when they began to defend it against the attacks of those whose focus lay in heredity and experimentation rather than selection and statistics – they were pressed quite forcefully on their motivations. Why would one begin to approach the biological

world with the tools of statistics? This is a topic familiar from the history of statistics, and two answers traditionally present themselves.[1] First, one might use statistics to simplify vastly complicated data, reducing it to the mean in order to construct a picture of "the average man." This position is all but synonymous with the name of Adolphe Quetelet, who coined the very phrase *l'homme moyen* (Porter, 1986, p. 52). As Hacking describes it, Quetelet began with the normal curve, previously derived as either an error curve or the limit-distribution of the result of games like coin-tossing, and he "applied the same curve to biological and social phenomena where the mean is not a real quantity at all, or rather: *he transformed the mean into a real quantity*" (1990, p. 107, original emphasis). This shift created not a real individual, but rather "a 'real' feature of a population" (1990, p. 108). Quetelet then uses his average man to "represent this [population] by height [or some other character], and in relation to which all other men of the same nation must be considered as offering deviations that are more or less large" (Quetelet in 1844, quoted in Hacking, 1990, p. 105).

On the other hand, one might use statistics to attempt to model diversity, to study a statistical distribution with the intent of capturing outliers. Galton, as Hacking tells the story, is a paradigm of this motivation for statistical study. Galton concerns himself, again on Hacking's picture, with "those who deviate widely from the mean, either in excess or deficiency" (Galton in 1877, quoted in Hacking, 1990, p. 180).[2] Hacking calls this a "fundamental transition in the conception of statistical laws," a shift toward Galton's "fascination with the exceptional, the

---

1. For example, in Porter (1986), Hacking (1990), or even Igo (2007).

2. It is notable that this picture of Galton is up for debate – I thank an anonymous reviewer for pointing out that Galton's work on composite portraits (e.g., Galton, 1879) looks much like Quetelet's use of "averaging."

very opposite of Quetelet's preoccupation with mediocre averages" (1990, p. 181).[3]

What does all this discussion of motivation mean for our question about the relationship between theory and world? The standard history would read both Pearson and Weldon as traditional Galtonians. Pearson, of course, took over leadership of Galton's Eugenics Laboratory (Magnello, 1999a,b) and wrote a laudatory, three-volume biography of Galton (Pearson, 1914, 1924, 1930). Weldon, to the extent that he is ever considered independently of Pearson, is squarely placed in the same camp, having published his first statistical-biological article under the direct mathematical guidance of Galton (Weldon, 1890). This gives them both, and the biometrical school in turn, a very obvious place within the history of statistics and biology.

I wish to argue for two related critiques of this view. First, if we look at Weldon's philosophy and motivation on its own, independent from that of Pearson, we can, despite their mutual connection to Galton, see an important difference between the two men with respect to their motivation for engaging in statistical practice. Pearson views statistics as part of a project consistent with his broader positivist philosophy of science – statistics is an appropriate tool to bring to biological systems in order to simplify them and reduce them to their underlying mathematical laws. Weldon, on the other hand, appears more focused on the preservation of diversity, arguing that only statistics allows us to take account of the real variability present in the biological world.

Second, and more importantly for our purposes, this distinction between these two scientists is drawn out and understood best in terms of the new question I presented in the last

3. A more metaphysical and less "historicized" version of this thesis is discussed in illuminating detail by Sober (1980).

chapter. For both of these motivations are grounded in views of how our statistical theories are related to the processes they describe. While I do not pretend to claim that a sufficient level of philosophical sophistication is present in either author to elucidate these views in detail, Pearson's view points toward a picture of statistical theories as acausal, anti-realist descriptions of the biological world. Weldon's view, on the other hand, seems to indicate a rich, causal view of biological theory, describing the very structure of biological systems.

I will begin by attempting to lay out a new view of Pearson's motivation for engaging in statistics, consonant with his philosophy of science, his prescriptions on methodology, and the conclusions of recent biographical work. I will then consider a much-neglected debate between Weldon, Pearson, and a few of their opponents. We find here our first evidence of the distance between Weldon and Pearson – a philosophical disagreement that one would not expect if both men merely innovated trivially on the work of Galton. I will then turn to developing a new conception of Weldon's motivation for engaging in statistics, grounded in a broader reading of Weldon's own philosophy of science, reconstructed in particular from the few sources in which Weldon self-consciously reflects on questions of philosophy. Weldon's view of science brought him to statistics by a profoundly different route than the positivism of Pearson.

There is a substantial body of literature on the history of biometry, particularly on the contentious debates between the biometricians and various proponents of discontinuous (and later, Mendelian) evolution, including William Bateson.[4] Weldon's work, however, has

4. Notably, I am forced to pass over the importance of this debate to the sociology of science; see MacKenzie (1978, 1979, 1981) and Norton (1978) for the sociological perspective, as well as Roll-Hansen (1980) and Olby (1989) for qualifications. See also Kim (1994) for a helpful dissection of the various classes of participants in these debates. Finally, see Cock and Forsdyke (2008) for a biography of William Bateson.

generally been seen only within Pearson's shadow.[5] I hope, in the end, to demonstrate as well that the lack of study of his thought is much to be regretted: Weldon's philosophy of science, and his reasons for adopting the biometrical method, are far more interesting than the usual stories would lead us to believe, and can direct us to insights not just about Weldon himself, but about Pearson, the general history of biological theory, and even the role of chance within it.

3.1. Pearson and Statistics

In addition to being a pioneer in statistics, Pearson was a profound philosopher of science in his own right, and was intensely reflective about his methodology and motivations. His philosophy of the physical sciences in particular, as expressed in his completion of W.K. Clifford's *Common Sense of the Exact Sciences* and his own *Grammar of Science,* was extensively developed, and, while formulated independently from the views of Ernst Mach (with whom Pearson corresponded only late in his career),[6] bears much resemblance to Mach's positivism.[7]

Gayon offers us a helpful place to begin by condensing Pearson's philosophy of science into three broadly positivist tenets: (1) science rests ultimately only on *phenomena;* (2)

5. For Weldon, the best biographical source is still Pearson's obituary (1906), though see, for example, Radick (2011), with significant insight into Weldon's work. Pearson's life is extensively detailed in Porter (2004).

6. See Thiele (1969) for their correspondence.

7. While several authors, such as Alexander (1964), Kevles (1985), and Plutynski (2006) argue that Pearson is best seen as a "Machian," Porter's recent biography places such a causal connection between Mach and Pearson in substantial doubt. See Porter (2004).

scientific laws *economize* our thought regarding these phenomena (by reducing them to mathematical formulae); and (3) science must not engage in *metaphysical speculation* (Gayon, 2007). Biometry can be readily seen to exemplify all three of these basic principles.

First, we have the phenomenological basis of science. Biometry consists crucially in the search for *empirical* trends in *observed data.* The extent to which this was adopted as a central claim in biometrical methodology can be seen as early as 1893, in the first paper produced from the collaboration of Pearson and Weldon. In it, Weldon claims that statistical investigation is "the only legitimate basis for speculations" regarding evolutionary theory: the study of phenomena is the only appropriate method in biology (Weldon, 1893, p. 329).

Second, we may turn to the economization of thought by mathematics. Pearson seems to adopt this unequivocally, equating the concepts of formula, law, and cause – all natural laws are merely mathematical formulas, and to describe the causes at work in a system *just is* to describe the laws (or formulas) governing it. Most directly, he says in the *Grammar of Science* that the last step of the scientific method is "the discovery by aid of the disciplined imagination of a brief statement or *formula,* which in a few words resumes the whole range of facts. Such a formula … is termed a *scientific law.* The object served by the discovery of such laws is the economy of thought" (Pearson, 1892, p. 93). Further evidence for this view may be found throughout his other work on biometry. In one of Pearson's many "Mathematical Contributions" articles, he mentions, regarding fertility, that "if it be correlated with any inherited character … then we have a source of progressive change, a *vera causa* of evolution" (Pearson et al., 1899, p. 258). This cause is to be investigated, not merely by asserting the existence of a correlation, but by determining the precise mathematical law which relates the quantities at issue (Pearson et al., 1899, p. 267). Pearson is noticeably silent about what

would constitute the appropriate mathematical laws for biology, but it might be inferred, on the basis of his enthusiasm for his version of Galton's Law of Ancestral Heredity, that this was the sort of thing he had in mind: a law which could tell us the expected deviation of an offspring from the generation mean based on the characteristics of its parents, grandparents, and so on.[8]

Such claims abound in Pearson's *Grammar of Science.* Commenting on the concept of "laws of nature," he says that

> law in the scientific sense only describes in mental shorthand the sequences of our perceptions. It does not explain *why* those perceptions have a certain order, nor *why* that order repeats itself; the law discovered by science introduces no element of necessity into the sequence of our sense-impressions; it merely gives a concise statement of *how* changes are taking place. (Pearson, 1892, p. 136)

This view of laws supports the understanding of science as economizing our thought from, as it were, another direction – by claiming that natural law, the supposedly basic explanation for the necessary connections holding within nature, cannot perform the role demanded of it by traditional ideas of causality.

Importantly, Pearson's view of causation creates a high bar for science – we must know quite a bit about the system under investigation in order to construct relationships of the sort that he demanded. In a paper read at the end of 1895 and published in the *Transactions of the Royal Society* for 1896, Pearson seems skeptical that biological causes can be found, given the current level of knowledge: "The causes in any individual case of inheritance are far too complex to admit of exact treatment; and up to the present the classification of the circumstances under which greater or less degrees of correlation … may be expected has

8. See the last chapter for an overview, or Froggatt and Nevin (1971) for more precise information on the form and development of the law of ancestral heredity.

made but little progress" (Pearson, 1896b, p. 255). That is, the complexity of biological systems makes the project of delineating their formal structure with precision incredibly difficult, and the completion of such a project has, in Pearson's view, been far from successful.

One more example may be cited. In the second edition of the *Grammar of Science,* published in 1900, Pearson adds the following (my emphasis):

> In the last chapter we freely used the words 'evolution' and 'selection' as if they had current common values. Now this is very far from being the case, and it is accordingly desirable to give to these terms and to other subsidiary terms definite and consistent meanings. It is only within the last few years, however, *with the growth of a quantitative theory of evolution,* that precise definition of fundamental biological concepts *has become possible.* (Pearson, 1900, p. 372, emphasis added)

It is worthy of note that in the intervening years between 1895 and 1900, Pearson seems to have become substantially more optimistic about the odds for success of a "quantitative theory of evolution." Pearson sees the introduction of biometrical methods as the only way by which we can expose the true scientific, lawlike, or causal (all three identical for Pearson) foundations of biological concepts. This position might seem odd, until we consider that such a grounding for biology consists of a description of the mathematical dependence of phenomena on one another. In this light, Pearson's philosophy of science appears broadly unified.

This focus on statistical/causal laws was also noticed by Pearson's son, who, in his two-part obituary for his father, mentions that, given the tenor of the nascent biometrical method as espoused in the first (1892) edition of the *Grammar of Science,* this process was all but inevitable:

> Looking back it is easy to follow where these trends of thought led, almost at once, in action: to an interest in Galton's Law of Ancestral Heredity; to a more accurate statement of this Law, involving the development of the theory

> of multiple correlation; to the testing of its adequacy as a descriptive formula by an extensive collection and analysis of data on inheritance.… (Pearson, 1936, pp. 216–217)

In other words, the very essence of the biometrical school, for Pearson, led almost inexorably to the utilization of an entirely functional notion of cause – the attempt to flesh out descriptive, mathematical laws which can summarize extensive amounts of data.

Finally, we may turn to the third positivist tenet underlying Pearson's philosophy of science, the avoidance of "metaphysical speculation." Arid theorizing about the material basis of heredity or the precise physiological or causal significance of observational results, Pearson argues, will do nothing but damage the progress of the science. Empirical grounding is the way to avoid mere blind guessing, as Weldon, collaborating with Pearson, insisted in 1895:

> These [statistical results] are all the data which are necessary, in order to determine the direction and rate of evolution; and they may be obtained without introducing any theory of the physiological function of the organs investigated. The advantage of eliminating from the problem of evolution ideas which must often, from the nature of the case, rest chiefly upon guess-work, need hardly be insisted upon. (Weldon, 1895b, p. 379)

This claim rings strongly of both a grounding in phenomena and a reticence to engage in metaphysical speculation unwarranted by available data. Even more striking is Pearson's complaint, expressed in his extended 1896 article on panmixia (i.e., random mating, or, for Pearson, the effect of completely random interbreeding *without* the influence of natural selection), that the current lack of progress in biology is

> largely owing to a certain prevalence of almost metaphysical speculation as to the causes of heredity, which have usurped the place of that careful collection and elaborate experiment by which alone sufficient data might have been accumulated, with a view to ultimately narrowing and specialising the circumstances under which correlation was measured. (Pearson, 1896b, p. 255)

When we look at Pearson's considered philosophy of science, then, it is no wonder that he found himself attracted to the biometrical methodology. Kevles describes Pearson as being drawn to biology because it was "rife with speculative concepts … that purported to explain vital phenomena yet were beyond operational test. He found [the biometrical] program appealing because of its positivist determination to deal only with directly observable quantities" (Kevles, 1985, p. 29). And a further conclusion can be drawn. Pearson's work, throughout his revisions of the *Grammar of Science,* remained emphatic about the usefulness of science for the *economy of thought.* The complexity of organisms is undeniable, as is our relative inability to specify with any true precision their internal workings. Biological data is thus a vast, tangled web of observations – on various characteristics, of different organisms, at different times, in different environments. We need the statistical method in biology so that we can *simplify our way out of this tangle:* only through statistics can we hope to offer economized laws of nature, which can encapsulate this data in a comprehensible manner. E.S. Pearson, writing about his father's reasons for leaving the study of evolution, claimed that "in the growing complexity of the Mendelian hypothesis," Pearson "could not see those simple descriptive formulae which held so important a place in his conception of scientific law" (Pearson, 1936, p. 241).

Statistics is thus useful for establishing the *vera causa* of evolution, as Pearson describes it, not only because it lets us capture the outliers in some particular statistical distribution, but more importantly because we can *then* continue onward, simplifying and economizing our thought regarding biological systems, distilling those biological processes into mathematical, functional laws – the only things which Pearson would recognize as "causes." This, then, is Pearson's motivation for engaging in the biometrical program – the

reduction of biological complexity to simple, functional laws, phrased in terms of statistics.

Two clarifications of Pearson's philosophy should be raised here. First, we have substantial evidence that Pearson's philosophy of science was not *merely* positivist. It was, as detailed especially by Porter (2004), a strange amalgam of influences, some of which are positivist, some based in German idealism, and some grounded in Pearson's political views. None of these other philosophical inclinations, however, do a better job than positivism of explaining Pearson's motivation for engaging in statistics. Fundamentally, positivism is clearly a *sufficient* explanation for Pearson's use of statistics – statistics clearly does economize our thought in the way Pearson valued. Moreover, we have much direct documentary evidence that Pearson viewed statistics in a profoundly positivist manner – it is positivism to which he returns time and again throughout the very papers in which he elaborates the biometrical methodology, and it is positivism to which his son ascribes his reluctance to engage in Mendelian genetics. Such evidence is simply lacking with respect to any other explanation for the motivation behind Pearson's statistical project.

Second, it is certainly true that Pearson is often very concerned with the preservation of outliers, complex distributions, and so forth: Pearson can hardly be said to want to simplify or reduce all statistical distributions to mean values, as, for example, did Quetelet.[9] Pearson worked extensively with non-normal distributions, in what could be argued was an attempt to preserve their structure, or their variability (e.g., Pearson, 1894b, 1895). In a review harshly titled "Dilettantism in Statistics," Pearson rails against an investigator who reduces his data to a simple normal curve, discarding the distribution's important skewness, which Pearson

9. I thank Theodore Porter for encouraging me to review this side of Pearson's thought.

claims constitutes the data's only important characteristic (Pearson, 1894a). Further, Stigler even reports that Weldon became angry with Pearson as a collaborator because he refused to take Weldon's suggestion to remove the outliers from his curves (Stigler, 1986, pp. 337–338).[10]

In evaluating this aspect of Pearson's work, we must be careful to separate two distinct features of Pearson's thought. On the one hand, we have his methodological prescription that we must always preserve outlying points, skew distributions, and so forth. On the other hand, we have the issue of Pearson's motivation for engaging in the statistical method in the first place. On the first point, I have no wish to argue that Pearson's statistical methodology was as simple as Quetelet's – such a claim is obviously ludicrous when applied to a mathematician as sophisticated as Pearson. But importantly, his technical methodology is consistent with several *motivations* for statistical practice – positivism foremost among them.

Turning to Pearson's harsh criticism of reduction to a normal curve, the "simple descriptive formulae" that Pearson's son described as so central to his father's view of scientific law should, to borrow an old cliché, be as simple as possible, but no simpler. I think we see in Pearson's critical review an instance of Pearson attacking oversimplification – a perfectly acceptable critique even on positivist grounds. At the same time, the reason that we engage in statistical work in the first place is because of its exceptional ability to provide us with the descriptive formulae that positivism places at the center of scientific research. Pearson and Weldon may have in fact disagreed about whether a given data point was (biologically) legitimate or not, but this disagreement fails to speak to the two men's fundamental motivation behind the practice of statistics.

---

10. In fact, Weldon is lamenting in this letter that Pearson is often more concerned with applying complicated statistical analyses than adhering to biological accuracy.

It is to Weldon's motivation that we should now turn, beginning by attempting to separate Weldon's view of science from that of Pearson – pulling Weldon out from under Pearson's gargantuan shadow.

3.2. Weldon and the *Nature* Debate

On the traditional reading of the relationship between Pearson and Weldon, we would expect the two men to view statistics in precisely the same way. As the story usually goes, Weldon is the empirically minded biologist who approaches Pearson when he feels his experimental problems might be helped by statistical methods. Beyond this point, Pearson and Weldon are deemed to be all but philosophically, methodologically, and motivationally identical.[11]

It is understandable that this is the accepted reading of their relationship. Extracting a distinct view of Weldon's thought is a difficult enterprise for several reasons. No comprehensive biography of Weldon has yet been prepared, and he was a strict naturalist of the highest order – his published articles rarely stray from relatively straightforward reporting of the biological data which he devoted his entire (and, sadly, too-short) career to collecting. Weldon, unlike Pearson, very rarely stopped to consider the philosophical and motivational grounding of his own methods.

I will explore three sources in order to separate Weldon's motivation for engaging

11. Froggatt and Nevin (1971, pp. 3–4) describes them both as drawn to the same problems by the same reading of Galton's *Natural Inheritance.* Sloan (2000, p. 1071) and Norton (1978, p. 4) have a similar reading of their early relationship, though Sloan complicates Weldon's later development.

in the biometrical program from Pearson's. First, in this section, I will consider a debate which occurred in the correspondence pages of *Nature* between Weldon, Pearson, Joseph T. Cunningham, and E. Ray Lankester.[12] In the next section, I will examine two other sources – a lecture which Weldon published in a volume on the methods of science and the first paper Weldon wrote with Galton, before his collaboration with Pearson began.

The *Nature* debate is yet another chapter in the long and storied argument between the biometricians and their Batesonian (and later Mendelian) opponents. We should begin by setting the stage.

3.2.1. The Opposition

The level of acrimony between the opponents and supporters of biometry around 1900 is indeed legendary. It would take a monograph to describe this conflict in detail, but a little context is useful here. We begin with the publication of Bateson's *Materials for the Study of Evolution* in 1894 (Provine, 1971; Sloan, 2000, p. 1074). Bateson worried, as Cock argues, about the twin problems of the usefulness of small variations and the difficulty of preserving variation over time (Cock, 1973, p. 8). These issues coalesced a community of scientists concerned with, first and foremost, describing the mechanism of heredity.

We can see many levels of disagreement between, on one side, Bateson, his allies, and even, as we will see below, scientists as diverse as E. Ray Lankester (a British Haeckelian and

12. In general, these "paradigm articulators" (to use the phrase of Kim (1994)) are not well known, excepting Lankester (see Lester, 1995). See Ankeny (2000) and Tabery (2004) for further study of two other important, smaller players in these debates.

dyed-in-the-wool defender of Darwin) and J.T. Cunningham (a highly influential British neo-Lamarckian) – and, on the other, the early biometrical school of Pearson and Weldon. First, the biometrical method was highly technical. The life sciences had engaged in research for centuries without the aid of complex mathematics, and many practitioners saw no need for it now.

Secondly, the Batesonian group was convinced that the sort of variation that would resolve their problems would be *discontinuous.* Cock and Forsdyke (2008, part V) argue that if one issue can be said to have motivated William Bateson throughout his career, it is the conviction that there was something qualitatively different about the discontinuous variation responsible for the generation of new species. The biometricians, in contrast, were committed gradualists – staunch defenders of an orthodox Darwinism (Froggatt and Nevin, 1971, p. 10).

Further, Pearson and Weldon, as I noted in section 3.1, explicitly deemphasized in their biometrical methodology the discovery of the physical mechanism of heredity. The Batesonians, therefore, failed to see how Pearson and Weldon's statistical methods could even be *relevant* to the study of evolution. Observational work ought to intend, as that in Bateson's *Materials* did, to test and explain theories of heredity like theirs. These theories are not the sorts of things *even subject to* investigation using Pearson's unnecessarily complicated tools. A later, though representative, statement of the objection can be found in a critique by Bateson of one of Pearson's later works:

> …much of the statistical work produced by Professor Pearson and his followers has, I believe, gone wide of its mark, if that aim is the elucidation of Evolution. More fitly might this work be described as "Mathematical Contributions to a Theory of Normality." […] By the one word *Variation* we are attempting to express a great diversity of phenomena in their essence distinct though merging insensibly with each other. The attempt to treat or study [these phenomena] as similar [i.e., by using advanced statistics like Pearson's] is leading to utter

confusion in the study of evolution. (Bateson, 1901, pp. 203–204)[13]

Looking at the statistics which so interested Pearson, Bateson claims, smoothes over precisely the sorts of differences we are concerned with capturing in the study of evolutionary variation.

This dispute was further complicated by some preexisting bitterness between Weldon and Bateson – Weldon had written an unfavorable review of Bateson's *Materials,* and Weldon and Bateson had argued at length in the correspondence pages of *Nature* in 1895 about an issue concerning the *Cineraria,* a genus of small, shrub-like flowering plants (Cock, 1973, p. 8). In the same year, Weldon had spearheaded the organization of the Evolution Committee of the Royal Society as a haven for biometrical work. After Bateson and his allies roundly criticized Weldon's article on the evolution of crabs (about which more later), Galton pressured Weldon to place Bateson on the committee. Bateson promptly took over and stacked the committee, causing Pearson and Weldon to finally resign in 1900 (Pearson, 1936, p. 228; Froggatt and Nevin, 1971, p. 9). Shortly thereafter, the hostile climate for the biometricians spurred the founding of the journal *Biometrika,* intended to be a place for them to publish their works without interference (Weldon et al., 1901). At this point, the dispute between Bateson and Weldon had become so bitter that Weldon called it "paltry and dirty beyond measure" (quoted in Magnello, 1998, p. 72).

Lastly, we have the role of Mendel. After the "rediscovery" of Mendel's paper and the publication of its translation in the *Journal of the Royal Horticultural Society* (Druery and Bateson, 1901), the Batesonians eagerly picked up Mendel's banner, in large part because they felt his theory would be a highly useful way to approach their concerns in both heredity and

13. "Mathematical Contributions to the Theory of Evolution" was the title of a series of more than a dozen papers Pearson wrote during his collaboration with Weldon.

breeding (Darden, 1977; Olby, 1987).[14] Despite attempts by some at the time to synthesize the work of the biometricians and the Mendelians, Mendelian genetics rapidly became the front line in this controversy – and the Mendelians rapidly won converts.[15] From 1900 until 1906, the story for the biometricians is one of a steady loss of allies, as attempts were made to discover how the Law of Ancestral Heredity, a central biometrical principle which Pearson had extended from Galton's original formulation, might be related to Mendelian inheritance (Pearson, 1898, 1904; Weldon, 1902). Weldon's death in 1906 precipitated Pearson's retirement from the study of evolution, and he would attend only one meeting of the British Association after 1904 (Pearson, 1936, p. 231). Until the early synthetic work of Fisher (1918, 1922), leading to the later contributions of authors like Wright and Dobzhansky, the Mendelians carried the day.

3.2.2. The Debate Begins

I want to narrow the focus, however, to one particular debate between Pearson, Weldon, and two opponents that took place in the letters to *Nature* in 1895 and 1896. This exchange has been discussed before: Bowler (1992, p. 4) cites it as evidence that the Darwinians during the "eclipse of Darwinism" weren't able "to maintain a unified front" against their opponents.[16] But this (admittedly accurate) portrayal conceals a very interesting aspect of

14. See, however, Pearson (1908), Magnello (2004), and Porter (2005) for the complexity of the biometricians' response to Mendel. For the impact of Mendel in other fields, see Roll-Hansen (2000).

15. On these near-syntheses, see Tabery (2004) and Morrison (2002).

16. It is also briefly mentioned by Olby (1989) and Plutynski (2006).

the exchange: the insight it brings into the relationship between Pearson and Weldon on deep, philosophical points. First, a brief discussion of the context in which the discussion took place.

In March of 1895, Weldon published a summary of his seminal paper on the statistical analysis of measurements of the crab *Carcinus maenas.*[17] Weldon had collected extensive data on several morphological quantities of interest – one of which, "frontal breadth" (a relatively unimpressive morphological characteristic of these crabs), he claimed could be shown to be under selective pressure. Pearson's influence on the paper was extensive, as the amount of statistical work required to demonstrate the influence of selection was massive. First, one had to normalize for the simple *growth* of the crabs over their lifespan (a profoundly difficult statistical feat), and discard data for obviously wounded or malformed crabs. Once this was done, Pearson believed he had arrived at a *slightly* non-normal distribution for frontal breadth. This distribution curve could be factorized (using a new method which Pearson had just developed) into the superposition of *two* normal curves. The crucial claim was that this superposition provided evidence that the population itself was bimodal – that is, that natural selection had *split* the population into two sub-groups which were evolving away from one another. Weldon hypothesized that the selective pressure at work was due to the turbidity of the water at various places in the crabs' environment.

It is not surprising, given the tenuous nature of these inferences, that controversy soon developed. The first encounter on the *Nature* correspondence pages occurred when the botanist William T. Thiselton-Dyer submitted a letter commenting on Weldon's paper

17. The original is Weldon (1895b); the summary printed in *Nature* is Weldon (1895a).

(Thiselton-Dyer, 1895).[18] Thiselton-Dyer believed that the statistical method could be used to shed light on the "stability problem" – the tendency of a "mean specific form" to be preserved in a population. This was quite a live question in the biological community in 1895 – Galton would present a request to the Entomological Society just three weeks after Thiselton-Dyer's letter was published, asking "those who have had experience in breeding" for data bearing on "a theoretical question of much importance; namely, the part played in Evolution by 'organic stability'" (Galton, 1895, p. 155). Galton's proposal of the Law of Regression had attempted to formalize the observation that, as he put it, "offspring [plants] did *not* tend to resemble their parent seeds in size, but to be always more mediocre than they" (Galton, 1886, p. 246). It was a solution to this problem which Thiselton-Dyer believed he had spotted in Weldon's work.

But despite his optimism that statistical methods might be used to solve problems of regression, Thiselton-Dyer was more skeptical when it came to Weldon's methodological claims:

> I am not sure that I quite understand Prof. Weldon when he says that "the statistical method is the only one at present obvious by which [the Darwinian] hypothesis can be experimentally checked." In the first place, I should myself hardly call it experimental at all. In the next place … in the important cases where evolution is actually taking place, the mathematical analysis appears to me to be beset with very great difficulties. (Thiselton-Dyer, 1895, p. 461)

Thiselton-Dyer, it seems, entirely missed the point – confirmed by the fact that *both* the biometricians and their opposition cite him as an opponent. He managed to read into Weldon a concern – namely, the demonstration of the law of regression to the mean – in which Weldon had little interest, and flatly dismissed as unintelligible the point that Weldon was actually trying to make. (As it turns out, the statistical method *is* highly amenable to the explication

18. This letter is also known for sparking the debate over the origin of the cultivated *Cineraria,* mentioned earlier (Froggatt and Nevin, 1971, p. 9).

of regression to the mean – Pearson would publish a paper doing precisely that the following year (Pearson, 1896b, pp. 306*ff*).)

Nonetheless, the floodgates had opened – spurred, no doubt, by Thiselton-Dyer citing one of the most controversial statements in Weldon's paper – and the parties to the debate quickly formed: J.T. Cunningham and E. Ray Lankester on one side, and Weldon on the other.[19]

Joseph T. Cunningham was a marine biologist and zoologist at the Marine Research Station at Granton, whom Bowler has called "an important but by no means typical British Lamarckian" (Bowler, 1992, p. 89). He was deeply engaged in the battle against Weismann; his obituary read that "he remained to the last one of the most eminent of the neo-Lamarckians" (Mudge, 1935, p. 42). E. Ray Lankester, on the other hand, "was one of the giants of late-nineteenth-century British science," and had positioned himself as "a champion of the Darwinian selection theory against Lamarckism" (Lester, 1995, pp. 1, 87). But he was very interested in the German or Haeckelian version of the "problem of variation," espousing a theory of "correlated variation," according to which variation "is limited by the already selected and emphasized characteristics of the group. Every part … varies in accordance with the constitutional tendency of the organism, which may be called its ancestral bias, or group bias."[20] With figures as influential as these lined up against Weldon, and given the confusing content of the string of *Nature* letters, a reconstruction of the play-by-play is bound to be useful here.

---

19. Some of the more interesting letters include Cunningham (1895), Lankester (1896a,b), Pearson (1896a), and Weldon (1896a,b).

20. From the Lankester papers, privately held, quoted in Lester (1995, p. 89).

The first letter following Thiselton-Dyer's was from Cunningham (1895), who brought with him an exceptional dose of methodological vitriol. After claiming that all Weldon had done was to show that some *future* demonstration of natural selection *might* be possible using statistical methods, he railed that:

> Prof. Weldon says that if we know that a given deviation from the mean is associated with a greater or less percentage of death-rate, we do not require to know how the increase or decrease of death-rate is brought about, and all ideas of functional adaptation become unnecessary. This may be his own state of mind on the subject, but I venture to state that it is not Darwinism, and that he cannot shut others out from the most interesting and most important fields of biology in this way. (Cunningham, 1895, p. 510)

Not only, then, are statistical methods good for delivering us little more than a promissory note on future results, Weldon's methodological prescriptions actually hinder the advancement of biological science.

After this outburst, the debate fell silent for a little more than a year. In June of 1896, Alfred Russel Wallace presented a paper to the Linnean Society regarding the existence and utility of the "specific character," or the set of characteristics that separate a species from the other members of its genus. Lankester wrote to *Nature,* ostensibly to comment on the views of Wallace. But his intent was clearly otherwise: he said outright that his "chief object in writing this letter is to draw attention to the views of Prof. Weldon" (Lankester, 1896b, p. 245).

As Lankester tells the story in his letter, during the discussion at the Linnean Society after the reading of Wallace's paper, he had argued for the importance of his "correlation of variation." Weldon declared Lankester's theory entirely irrelevant, because (again, as reported by Lankester) given a case of two characters, both of which are positively correlated with favorable selection, it is "absolutely impossible to separate these two correlated phenomena. The coloured skin is as much a cause of the survival of the dark man as is the germ-destroying

property of his blood" (Lankester, 1896b, p. 245).

Lankester was dumbfounded. "I was not prepared," he laments, "for an empty wrangle in regard to the proper uses or improper uses of the word 'cause'" (Lankester, 1896b, pp. 245–246). Lankester immediately proceeded to a philosophical debate over Weldon's use of causation. He claimed that Weldon

> has deliberately departed from the simple statement which his observations warranted, viz., that such-and-such a proportion of frontal measurement accompanies survival, and has unwarrantably (that is to say unreasonably) proceeded to speak of the "effect" of this frontal proportion, to declare it to be a *cause* of survival, to estimate the "advantage" and "disadvantage" of this same proportion, and finally to maintain that its "*importance*" may be estimated without troubling ourselves to inquire how it operates, *or whether indeed it operates at all.* (Lankester, 1896b, p. 246, original emphasis)

I have quoted this passage at some length to give an idea of the philosophical level on which this debate took place. Importantly, we see a shift from the position of Cunningham to that of Lankester. For Cunningham, the problem is merely about the methodological claims of biometry: statistical investigation of correlation is fine, he seems to say, but it cannot constitute a replacement for comparative physiology and the investigation of functional adaptations (a point with which, we will see, Weldon actually agrees). Lankester, on the other hand, seems to indict Weldon for *philosophical* mistakes. Weldon's *underlying philosophy of science* is inadequate if it leads him to think that the discovery of correlation is sufficient to determine causal influence. Lankester's letter closes by asserting that biometrical methods "appear to me not merely inadequate, but in so far as they involve perversion of the meaning of accepted terms and a deliberate rejection of the method of inquiry by hypothesis and verification, injurious to the progress of knowledge" (Lankester, 1896b, p. 246).

Weldon responded to Lankester's letter with one of his own, and if there were any doubt that the argument had become genuinely philosophical by this point, his response

should remove it. Weldon quoted, at length, Hume's definition, from the *Enquiry,* of cause as constant conjunction, and challenged Lankester as to whether he had the audacity to disagree with Hume (Weldon, 1896a, p. 294). Weldon stated he only ever intended to discuss cause under Hume's definition "or in Kant's extension of it [!]; but Prof. Lankester seems to go beyond it" (Weldon, 1896a, p. 294).

Lankester rushed to the defense of the Humean acceptability of his method, claiming that he merely desired, given the existence of two features correlated with some positive outcome, to engage in a process of hypothesis and experiment in order to determine the "true order and relation" of "a complex group of related phenomena" (Lankester, 1896a, p. 366).

Weldon, perhaps finally realizing in full detail the proposition with which he was disagreeing, beat a hasty retreat in August of 1896 (Weldon, 1896b). He explained that he was "far from rejecting the method of imaginative hypothesis and subsequent experiment and observation." "A complete knowledge," he wrote, "of the processes associated with this relation between frontal breadth and death-rate is a thing of very great interest, and I believe, as firmly as Prof. Lankester, that every effort should be made to attain to it" (Weldon, 1896b, p. 413). Though such a theory is quite hard to obtain (even, perhaps, impossible) due to the complexity of the interrelations of the organs of any organism, it must nonetheless be sought. Weldon would indeed undertake a series of experiments attempting to determine the influence on the crabs of the amount of silt in their water in 1897 and 1898, concluding that "a narrow frontal breadth renders one part of the process of filtration of water more efficient than it is in crabs of greater frontal breadth" (Weldon, 1898, p. 901; Pearson, 1906, pp. 26–27).

It is worth pausing here to note that this letter and Weldon's subsequent experiments, despite their appearances, do not necessarily constitute a *motivational* retreat on Weldon's

part. For both the claims that hypothesis and experiment are a good way to guide our future statistical research, and that more detailed knowledge of the correlations at issue in a given biological system is a desirable thing, are fully consistent with Weldon's belief that the *correlations* are the complete and sufficient endpoint of biological research. But more on this in the next section.

Surprisingly, Pearson entered the debate at this point, with a response *countering* Weldon (Pearson, 1896a). Pearson began by arguing that the statistical analysis in Weldon's 1895 paper was far too simplistic to constitute a genuine verification of natural selection. Given the complex nature of the assumptions regarding the general growth of the crabs that had to be made, Pearson declares it "very improbable" that the true growth curve was found. Further, he writes that "when the law of the growth of crabs has been accurately ascertained … I am convinced that it will require much more complex analysis than that of the *Report* to ascertain whether a selective death-rate does or does not exist" (Pearson, 1896a, p. 460). Then, Pearson turned to the defense of a concept of causation as *more* than mere Humean event-correlation. Even if the data in Weldon's 1895 paper were perfect, Pearson claimed, it would still not be enough to show that frontal-breadth *is the cause of* death-rate. "Very probably it may be, but the demonstration is not logically complete, or at any rate a definition of cause has been adopted which does not appear of much utility to science" (Pearson, 1896a, pp. 460–461). As an example, he turns to a discussion of artificial selection in cows, describing what contemporary philosophy of biology would recognize as the difference between selection *of* and selection *for* – both establish correlations, only one (selection *for*) involves causation. Without much more research, he claims, "it seems to me that a link is really missing in the chain of demonstration" (Pearson, 1896a, p. 461). After Pearson's input,

the discussion tapered off.

On the traditional reading of the relationship between Pearson and Weldon, we should by no means expect a disagreement here, nor should we expect a disagreement of this kind. Pearson and Weldon, it is alleged, shared precisely the same view of science and the role of statistics. Yet it is on these fundamental issues that Pearson and Weldon disagree here. Let us consider a few more sources.

3.3. Weldon on Statistics

In 1906, the year of his death, Weldon contributed a piece on "Inheritance in Animals and Plants" to a collection of lectures on scientific method (Weldon, 1906). Its breadth and central concerns are quite exciting – Weldon offers a sustained defense of the use of statistical methods in science, both generally and with particular emphasis on the biological sciences.

Why would statistical methods require a defense? According to Weldon, they inherently require a compromise – a methodological value judgment by a group of scientific practitioners. "Men measure a certain thing," he writes, "and find that up to a certain point their measurements agree with each other, and their experience is uniform; but beyond that point [i.e., in the very fine details], their experience is contradictory" (Weldon, 1906, p. 88). We can use statistics to smooth over these contradictory results (e.g., by averaging), but at a price – we must decide "how far the variability of the actual experience depends on imperfect observation, and how far it is a true record of differences in the thing measured" (Weldon, 1906, p. 88). What we truly desire – in Weldon's words, the "ideal description of

every experience, the description which alone makes further progress possible" (Weldon, 1906, p. 93) – is the correct description of all the observed results, without having neglected any inconsistencies whatsoever. Physics and chemistry possess successful, general, formal methods because they have "succeeded in confining the limits within which these inconsistencies occur, so that the proportion of the whole experience affected by them is very small. But biologists have not yet advanced so far as this: the margin of uncertainty in their experience is so large that they are obliged to take account of it in every statement that they make" (Weldon, 1906, p. 93). That is, physics and chemistry have advanced theoretically to the point that these disciplines are confident that they can sort any discrepancies in their standard processes of measurement into clear, well-marked types: either they are a result of individual failings of experimenter or apparatus, or they are an indication that the fundamental theory needs to be refined.

Biology, however, is a different story. We do not yet know enough about the underlying structure of biological systems to know what constitutes important or unimportant variation. Weldon offers an example:

> If we want to make a statement about the stature of Englishmen, we must find a way of describing our whole experience; we must find some simple way of describing our whole experience, so that we can easily remember and communicate to others how many men of any given height we find among a thousand Englishmen. We must give up the attempt to replace our experiences by a simple average value and try to describe the whole series of results our observation has yielded. (Weldon, 1906, p. 94)

The gist is this: physicists have more knowledge regarding their experimental systems, and work tirelessly at searching for sources of error. All in all, they are in a much better position than biologists to render themselves confident that any variation within a set of measurements is due to operator error. They can then take the mean, and (given their precautions) assume

that, for all practical purposes, they have arrived at the correct answer. Such a claim is consistent with Weldon's paradigm for statistics in physics being its use to correct for errors in observational astronomy (Porter, 1986, pp. 95–96). In biology, we cannot yet do this – due to the peculiarities of biological systems (especially their complexity), and our lack of knowledge, we can almost never simplify our observations in this way. We must therefore strive to preserve *all their details,* in a manner which still may be communicated – and, as you may have guessed, statistics is how we accomplish such a feat.

We should now have enough of a grip on Weldon's thought to try to articulate his view of biological method. Statistics is essential to the investigation of biological systems, due to their vast complexity and our lack of comprehension of their fine-structure. However, statistics is *not* essential for the reason we ascribed to Pearson – the primary motivation for using statistics is not to simplify or economize data for limited knowers, though this is certainly part of its merit. Its main purpose is rather to *permit us to retain* (as much as is possible) *the complexity of the biological world.* Any overzealous act of simplification (in Weldon's example, by doing things like substituting averages for statistical distributions) is equivalent to claiming that any remaining variation in our data is an artifact, not a feature of the outside world. We must use statistics (and keep the entirety of the statistical distribution intact) in order to hold on to the range of observational data which nature gives us.

We can see this same statistical philosophy in Weldon's earlier works. In his first publication to use statistical methodology, Weldon, with the help of Francis Galton, endeavored to draw conclusions from some data on shrimp (Weldon, 1890). By the standards of his later work with Pearson, his methods are horribly primitive – they simply consist in measuring the data's absolute deviation from a normal curve provided by Galton. The stated goal of the

paper is to "determine the degree of accuracy with which [the] adjustment [of a local variety to its environment] is effected, and the law which governs the occurrence of deviations from the average" (Weldon, 1890, p. 445). In other words, since no character of any organism is *perfectly* adapted to its environment, and assuming that the mean value constitutes the (selective) optimum, to what extent do organisms deviate from this optimum? Galton had proposed a curve (the "law of error," known today as the normal distribution) as an answer to this question, extracted from his empirical studies, and Weldon aimed to test this proposed solution in his shrimp.

Interestingly, however, the conclusions of the paper are broader than this. Weldon collected his measurements, and noted that, indeed, all the characters he measured were normally distributed – with, however, different means and standard deviations in the three environments from which he collected the shrimp. But he also considers as confirmed a stronger proposal of Galton's, namely, that even though selective pressures "must vary in intensity in different places," "the frequency with which the observed deviations from the average occur is in all three cases expressed by a curve of error" (Weldon, 1890, p. 451). That is, though natural selection can alter the mean and standard deviation of a given character, the characters in *any* environment, under *any* selective pressure, will remain normally distributed.

We cannot infer from this paper anything like the statistical and methodological sophistication which would later come to the biometrical school after many years of developing its tools and techniques. But we can see, sixteen years before publication of the "Inheritance" lecture, the first glimmers of Weldon's statistical philosophy of science. The variations present in his shrimp, he writes, "depend not only on the variability of the individuals themselves (which is possibly nearly alike in all races), but also on the selective action of the surrounding

conditions" (Weldon, 1890, p. 451). Clearly, these causes are too complex to allow our exact treatment. And we cannot rely on the mean values either, "for I am aware of no case in which the individuals composing any race of animals – however small and isolated the area in which they live, however uniform the conditions which obtain throughout that area – have been shown to resemble one another *exactly* in any character" (Weldon, 1890, p. 445). To accept the mean value is to discard important information about the population. In order to preserve necessary variation in the data, then, we must focus on the *normal distribution itself* – it is only at this level that we are permitted to draw conclusions about the population. Weldon's view here may lack the precision of his later formulations, but we see all of its important features.

Finally, we should return to Weldon's contribution to the *Nature* debate in 1895 and 1896. With this understanding of Weldon's statistical philosophy, we can easily see why he would adopt a "Humean" definition of cause. It is not an empiricist worry about the legitimacy of the imposition of a necessary structure on nature (*à la* Hume) that drives Weldon to such a position – rather, it is his recognition that the structure of biological systems is far too intricate to make such a determination possible. Correlation (which Weldon seems to think is similar to, or at least acceptable on the standards of, Humean "constant conjunction") is the only type of connection that can be drawn between biological systems of the kind Weldon was interested in investigating *without discarding some of the essential features of these systems.* To single out one cause and one effect is to commit precisely the same fallacy as substituting the mean for the statistical distribution – it is to unjustifiably decide that some aspect of one's data (whether variability or a consistent correlation) is *unimportant.*

It is also easy to understand Weldon's response to Lankester's persistent probing. For

Weldon clearly thought that investigation into particular facets of biological systems was worth undertaking – he did so himself. And such investigation is fully compatible with his philosophy as I have laid it out here: there is nothing dishonest about attempting to understand more fully the detailed nature of organisms – as long as such knowledge does not come with the implication that some features of organisms are to be privileged at the expense of others. To use Lankester's example, "the coloured skin" and "the germ-destroying property of his blood" both *cause* "the survival of the dark man" on Weldon's view because they are both elements of the broad, complex picture of such a human being (Lankester, 1896b, p. 245). Any attempt to choose one of these *over* the other must necessarily discard vitally important information, unless and until the correlation itself breaks down (e.g., when we detect "coloured" individuals *without* the property of exceptional survival "in malarial regions").

Recall from the beginning of this chapter that Pearson's engagement with biometry is best interpreted in the light of his positivist philosophy of science. Statistical correlation, for Pearson, is one step in the process of determining the precise mathematical laws or causes underlying a given biological system. We can now see clear evidence that Weldon, on the contrary, enters statistical practice out of a much more broadly empiricist concern for the preservation of variation in biological systems. For this is precisely what Weldon means when he speaks of "describing our whole experience" – statistics is the enterprise that lets us preserve and study the full range of biological phenomena. It is a grave methodological error to attempt to simplify away – to attempt to *economize* – this data, even when we might think we have good reasons for ignoring certain correlations or variability. Furthermore, why would we need to do so when we have the tools of biometry available to us, which allow us to

study biological systems in the full array of their natural variety?

3.4. Weldon, Pearson, and Chance in Evolution

With a novel conception of Weldon's philosophy of science in hand, we can turn to reevaluating the relationship between Pearson and Weldon. While their overall methodologies were all but identical, as were their ideas of valuable data and good experimental process, this agreement masks the fact that their motivations for engaging in statistics were interestingly divergent. Pearson, I have argued, views statistics as a tool for positivist simplification, while Weldon sees it as essential for the preservation of variation.

On such a view, the tension between Pearson and Weldon concerning the notion of causation (as reflected in the *Nature* debate) becomes manifest – and fully explicable. Pearson has adopted a positivist view of causation – we want to examine biological systems until we can reduce their behavior into a series of simple mathematical laws. When we know these laws, we will have the only thing which might pass for "causal" knowledge in biology. Weldon, on the other hand, has adopted what we might call a statistical view of causation – the only way in which we may accurately claim "causal" knowledge of a system, without destructive simplification, is to point to correlations within the system as a whole. Such correlations, on Pearson's philosophy, would constitute a very weak sort of causation – they would be clearly *necessary* for a causal link between two features, but far from *sufficient.* That is, a correlation is a form of functional relationship of the variety Pearson recognized, but an unacceptably *weak* one. The sharpening of these simple correlations into true *laws* must be one of the

projects of a positivist biometry. It is obvious, then, that Pearson would decry Weldon's view as "a definition of cause … which does not appear of much utility to science" (Pearson, 1896a, pp. 460–461).

Pearson, then, seems to engage in the biometrical program on broadly positivist philosophical grounds. He is moved to introduce statistics based on an emphasis on simplification and economization of thought, a central positivist tenet, best interpreted in light of his *Grammar of Science* and other positivist writings.

Weldon, on the other hand, holds a profoundly different motivation, a more traditionally empiricist reliance on the diversity of biological phenomena. His rich form of empiricism impelled him to approach the great diversity of biological observation as a necessary and even beautiful feature of life.[21] Far from it being our duty to take this diversity and simplify it by using statistical tools, Weldon claimed that "it is the first business of a scientific man to describe some portion of human experience as exactly as possible. It does not matter in the least what kind of experience he chooses to collect; his first business is to describe it" (Weldon, 1906, p. 81).

This much is clear from the historical record. More speculatively, however, we may consider the impact that these well-described motivations have on our study of the role of chance in evolutionary theory. Despite all the philosophical work these authors put into describing their use of statistics, we can see that there is precious little engagement here with the second question set out by the standard history in the last chapter – the presence or absence of a reified, objective notion of "chance" in the evolutionary process.

---

21. Radick comes to effectively the same conclusion regarding Weldon in his (2005) and (2011).

On the other hand, I argue that it is clear that these motivations for engaging in statistical practice express coherent and interesting philosophical positions. Pearson's simplifying, positivist philosophy of science indicates an acausal, anti-realist view of the relationship between biological theory and world. After all, the statistical theories in which Weldon claimed to find causation were merely poor approximations for the sort of genuine causes that Pearson hoped to find (causes that, arguably, can be found only in lower-level sciences than biology). Weldon, on the other hand, located causation – a causal summary of the myriad elements of biological variability – in the very statistical theories themselves. This distinction, then, seems little like that expected on the standard history described in the previous chapter. Rather, it fits more comfortably within the new question I introduced in the previous chapter: that of the relationship between our statistical biological theories and biological processes in the world.

I do not intend the argument for this last, speculative point to be air-tight. Rather, the overall case presented in these first three chapters is intended to support the claim that the literature on the role of chance in evolution – both contemporarily and historically – is often deeply confused. Two clear reasons for this have emerged. First, we have simple conflation of the various notions of "chance" itself. And second, we often ask questions of biological theory and practice that seem to be quite difficult to analyze in real-world cases.

More importantly, I have gestured at a way out of this confusion. I argue that if we focus on the relationship between our statistical theories and the biological processes that those theories aim to describe, we will consider a question with both more historical *and* philosophical relevance. Before turning to the new framework I will develop in the final two chapters, however, we have one more piece of ground-clearing to pursue.

CHAPTER 4

A CRITICAL REEVALUATION OF THE CAUSALIST/STATISTICALIST DEBATE

The first three chapters have thus endeavored to make the case that we have been engaging with the wrong question in our work on chance in evolutionary theory. Rather than an obsession with a reified sense of "objective" chance present in the world, we should, I have argued on the basis of both historical and contemporary examples, focus instead on the relationship between our statistical biological theories and the processes in the world those theories aim to describe. This picture, I claim, best helps us avoid a problematic conflation between several distinct notions of chance, and also lets us understand the work of important historical actors involved in the introduction of chance in evolutionary theory. Now, however, it is time for more positive payoff. A focus on the relationship between statistical theories and biological processes can also, as we will now see in the following chapters, help us interpret the central concepts of evolution by natural selection.

Philosophers of biology familiar with the landscape of the current literature, however, will immediately note that a debate over precisely this question has raged in the literature for the last decade: the debate over whether or not natural selection and genetic drift are *causal processes,* already mentioned several times in the first few chapters. This issue, in turn, has touched on several other core features of evolutionary theory. Are we right to describe

selection, drift, mutation, migration, and so on as analogous to Newtonian forces, each of which drives a population in a given direction? If we are, which of these should be considered part of the "inertial" state (analogous to Newton's first law), and which should be considered "special" forces (analogous to Newtonian gravitation and described by the second law)? What is the role of fitness in natural selection? If selection is causal, is fitness causal as well? Or is fitness merely a non-causal property of organisms, tallied for biological convenience?

Broadly, two positions in this debate have solidified. The "causalist" picture, canonically stated by Sober (1984), considers selection and drift to be causal processes, which in many cases can profitably be compared with Newtonian forces. The "statisticalist" interpretation, which was inaugurated by Walsh et al. (2002) and Matthen and Ariew (2002), disagrees. On this interpretation, selection, drift, and fitness are all non-causal. Rather, they are simply useful statistical summaries of the genuinely causal events that occur at other (individual, genetic, biochemical) ontological levels. The force analogy, the statisticalists argue, breaks down, particularly in the case of genetic drift.

A profitable way to characterize these two positions in more detail is, in fact, to begin with what *isn't* at stake. Both sides grant an identical conception of the individual-level causal account of the biological world. Individuals live and die, they are predated, they mutate, they give birth to other individuals. All of these are causal events, all are known to occur, and all are, in some sense or another, "responsible for" the higher-level phenomena that we call natural selection and genetic drift.

This commonality is best illustrated by Matthen and Ariew (2009), a recent statisticalist work. They introduce the tried-and-true example of Kettlewell's moths, and then wholeheartedly endorse the proposition that "[*v*]*ariation in camouflage* causes evolutionary

change (of the moth population)." They go on to say that "[n]obody today doubts this result, least of all us" (Matthen and Ariew, 2009, p. 203). Shapiro and Sober, on the causalist side, hold that the appropriate supervenience base for natural selection includes some subset of "the causal processes that impinge on individual organisms" (Shapiro and Sober, 2007, p. 251). It is clear, then, that both sides suppose that happenings in the lives and deaths of individual organisms are causal, and that these in some way result in evolutionary change.

Second, it is of course granted that natural selection and genetic drift (and, hence, evolutionary change in general) are theories that demand probabilistic explanations – that is, they are expressible mathematically only as a set of statistical trends that exist within populations of organisms. Of the causalists, Hodge's (1987) position perhaps makes this most clear – selection and drift *just are* varieties of sampling, and this sampling is best described by the statistical formulation of traditional population genetics.

This, it would seem, is a fairly robust picture of the biological world. We have agreement on both the underlying causes that (in some way or another) constitute the evolutionary process, and on the broad structure of the statistical framework that must be used to describe population change over time. What, then, remains at stake, after this description of the world is agreed upon? I will argue that three main issues are left unresolved. Each of these issues, in turn, is emblematic of serious problems within this debate – circularity, irrelevance to the philosophy of biology, and the neglect of a significant and unresolved gap in our interpretation of evolutionary theory.

First, causalists and statisticalists disagree on the definition of a causal process, the topic of section 4.1. After all, we are not merely asking about whether or not natural selection *exists* – rather, we are asking whether or not it is a very particular *kind* of thing: a causal

process. In order to resolve this question, we must clearly have laid out the necessary and sufficient conditions for being a causal process. By and large, the causalist/statisticalist debate has not done this. More importantly, we will see that we *require* metaphysical evidence, independent of this debate in the philosophy of biology, to settle this question. Without it, our failure to specify our notion of causal processes results in *a serious risk of begging the question* – the concept of 'causal process' one assumes has a very large impact on the tenability of each of the causalist and statisticalist positions.

The second primary issue at stake in the causalist/statisticalist debate is a core metaphysical concern: in which circumstances do the organism- and biochemical-level causes agreed upon by both sides compose to produce causal processes of natural selection, genetic drift, and so on, active at "higher" ontological levels? I will call this "the composition question" (or CQ).[1] I will argue in section 4.2 that this question is not one deriving primarily from the philosophy of biology, but rather is a general metaphysical problem. If discussion of it is to be fruitful in the biological context, then it is important for the composition question to be demonstrably *relevant* to the philosophy of biology. There are two ways, I argue, in which this relevance could be demonstrated. First, natural selection and genetic drift might constitute a revealing set of case studies for the investigation of CQ. I will argue in section 4.3 that this is not the case – for two reasons, selection and drift constitute a poor place to study the composition question. Second is the converse: the metaphysical study of CQ might be a profitable way to approach the interpretation of selection and drift. I will argue in section 4.4 that this, as well, is doubtful. While there has been some encouraging progress in this direction, it

1. It has also been termed an instance of "epiphenomenalism" by Shapiro and Sober (2007), about which more later.

is unlikely that any general metaphysical work on CQ will increase our understanding of selection and drift. The composition question is, therefore, *broadly irrelevant to the philosophy of biology.*

Finally, we move to the third live issue in the debate, and the one which, I argue, has genuine biological significance. Causalists and statisticalists broadly differ on their definitions of just what natural selection and genetic drift *are.* As we will see in section 4.6, there is a divide between process and product definitions of selection and drift, one which is commonly thought to parallel the rift between the causalist and statisticalist positions. This identification, however, is incorrect. I will argue that whether one adopts a process or product definition of selection or drift has nothing whatsoever to do with one's position on the causal potency of these evolutionary factors, indicating that this is a deeper issue in our understanding of evolution. Finally, in section 4.7, I will argue that while these two classes of definitions are quite different intensionally, their extensions – how they actually classify particular cases as selection or drift – are very nearly identical. The difference between them is thus quite subtle, and the instances in which they come apart from one another constitute an interesting and challenging place for investigating which of the various intensions are correct. We cannot hope to resolve the causalist/statisticalist debate without *settling this issue in the interpretation of evolutionary theory.*

Given these three interconnected problems, it is the corner cases that separate our definitions of selection and drift on which we should focus in order to increase our understanding of the interpretation of evolution. The issues of the composition question and the definition of a causal process, while interesting to philosophers of biology, are not any more readily resolved by dealing with them in the context of evolutionary theory. Rather, they

ought best be left within the confines of the metaphysics room. The one issue that is genuinely a problem in the philosophy *of biology* – the disparate definitions of natural selection and genetic drift – has often received short shrift in this debate, and it is this crucial, and uniquely biological, issue on which we should center our argumentative efforts.

As we have already seen, the implications of these central disagreements can be found in many areas of biological theory – the question of "forces," the analysis of the concept of fitness, and so on. My focus here will remain narrow, however – the problems I discuss can be described without leaving the heart of the causalist/statisticalist debate. Tracing out the impact of the troubled nature of these arguments into these farther-removed fields is another project entirely.

4.1. What is a Causal Process?

One fundamental issue in the causalist/statisticalist debate has, surprisingly, been almost entirely absent from the literature: if we are considering whether or not these individual-level events constitute causal processes, we must agree on the concept of a causal process.

In the context of the general metaphysics of science, a few authors have offered process-based accounts of causation. Perhaps most famously, Wesley Salmon (1984; 1994) accounted for causality in terms of causal processes which serve as the transmitters of causal influence. While causal interactions, defined in terms of "causal forks," are more fundamental than processes on this view, both are part of the "basic" causal furniture of the world. What is a causal process, for Salmon? He begins by defining the notion of a process *simpliciter,* which

he claims "can reasonably be regarded as a primitive concept that can be made sufficiently clear in terms of examples and informal descriptions" (Salmon, 1994, p. 297). The primary goal of a theory of causation is to separate these processes into two varieties: genuinely causal processes and mere pseudo-processes. For Salmon, this separation arises from the fact that only genuinely causal processes may plausibly be said to *transmit* a physical quantity through time.

Phil Dowe (1992; 1995), beginning with a handful of worries about Salmon's project, went on to develop a relatively similar "conserved quantity" theory of causation. For Dowe, we begin with the definition of a conserved quantity – "any quantity universally conserved according to current scientific theories" (Dowe, 1992, p. 210). A causal process, then, "is a world line of an object which manifests a conserved quantity," where an object is "anything found in the ontology of science (such as particles, waves or fields), or common sense" (Dowe, 1992, p. 210). As with Salmon, processes and interactions are the basic causal apparatuses of the world, and for Dowe it is the specification of quantity conservation that suffices to separate causal processes from pseudo-processes.

Anjan Chakravartty (2007) offers a realist account of causal processes that is subtly different from both of these. As a realist, Chakravartty can countenance a fully fledged notion of *causal properties,* as well as an associated account of *dispositions* – both of which might introduce Humean problems for the staunchly empiricist Salmon and Dowe. For Chakravartty, then, the existence of causal processes is a straightforward implication of the interaction of objects that bear causal properties:

> Dispositions borne by objects in virtue of their properties are continuously manifested in accordance with the presence and absence of other objects and properties. Objects with causal properties are thus in a continuous state of

> causal interaction, a state in which relations between causal properties obtain. (Chakravartty, 2007, p. 108)

Once we have specified what a causal property is, a causal process is just what happens when properties like these continuously interact. This, then, is a definition of 'process' even more minimal than those offered by Salmon and Dowe.

In all three of these cases, we can see that we have a fundamental, simple notion of "process" – worldlines through spacetime for Salmon and Dowe, and property interactions for Chakravartty – and a criterion that separates genuinely causal processes from mere pseudo-processes. But what definition we adopt (whether one of these or another), and how we spell out its details in the biological case, will have drastic implications for whether or not natural selection and genetic drift qualify as causal processes.

A handful of authors in the causalist/statisticalist debate have at least gestured at what they take a causal process to be, and they have laid out vastly disparate notions. On the causalist side, Millstein has been the most vocal on the definition of a process. She even explicitly defines the term at one point, saying that "by 'process' I mean a series of physical states occurring over time, whereas by 'outcome' I mean the effect, or ending state at a particular point in time, of that process" (Millstein, 2006b, p. 679). This is, however, only a definition of a process *simpliciter,* leaving unspecified the criterion (if there is one) that would separate causal from pseudo-processes. We lack sufficient material here (and elsewhere within Millstein's work) to know whether any such series of physical states might be counted as a *causal* process. If any such series of states is indeed a causal process, we might, for example, be able to speak of causal processes that are very loosely causally connected, like a causal process of urbanization in twentieth-century America. Alternatively, Millstein might

invoke something like Salmon or Dowe's criteria for causal processes, though how exactly to spell out these theories (which are defined in general in terms of quantities from physical theory) at the biological level would constitute a significant challenge.[2]

On the statisticalist side, Matthen and Ariew offer a much more restrictive definition of process. They speak of natural selection as failing to be a causal process, because "it is not a process in which the earlier events cause the latter" (2002, p. 79). Although this view is minimally spelled out, Matthen has argued elsewhere (Matthen, unpublished) that a causal process must have either "an overarching cause that is responsible for each event in the sequence, and the order in which they occur," or be such that each event in the sequence is causally responsible for future events in the sequence.[3] (Whether this is partial causal responsibility or complete causal responsibility is left unspecified.) On this alternative definition, then, there may well be no such thing as the causal process of urbanization in twentieth-century America, if this process is the result of a constellation of causes, none of which can integrate the various states in a sufficiently tightly-knit manner.

As already mentioned, we must also consider how to parse these definitions in the particular cases of selection and drift.[4] Sometimes, for example, Millstein argues that selection and drift, *qua* processes, can operate at the same time. In one article, she presents an example

---

2. Millstein has claimed in a new paper (Millstein, 2013) that she intends to refer to Salmon's account of causal processes, an argument that I have yet to consider.

3. Matthen has suggested to me (pers. comm.) that the definition of causal process intended in the earlier Matthen and Ariew (2002) is Salmon's, though this leaves unspecified, as above, the way in which we should apply Salmon's criterion to the biological case. It also leaves unspecified the entities which are engaged in these processes – whether they are individuals, populations, etc. It is further unclear how the published remarks in Matthen and Ariew (2002) can be reconciled with Salmon's view (Millstein, for example, claims they cannot).

4. This task will be made more difficult by the fact that our very definitions of selection and drift are unsettled, to which we will return in section 4.6.

regarding snails, and claims that it and cases like it "may be analyzed as being cases of natural selection together with random drift." She continues, saying that in this example

> there is a process in which the physical difference between the snails is causally irrelevant to the difference in their reproductive success.... Thus, the process of random drift is occurring in the population. However, there also is a process in which the heritable physical difference between the snails *is* causally relevant to the difference in their reproductive success.... So, a discriminate parent sampling process, natural selection, is occurring in the population as well. (Millstein, 2002, p. 44)

Here, then, we have multiple causal processes which may simultaneously occur within a population. Different concatenations of physical states may result in a process which either does or does not consider physical differences between organisms in the population to be relevant. Those that do are selective processes, those that do not are drift processes.

Elsewhere, though, Millstein seems to claim that drift and selection are mutually exclusive states of affairs:

> I, on the other hand, would label the deviations from expectations as drift if and only if there were no physical differences that were causally relevant to the change in frequency. If there were causally relevant physical differences, then I would attribute the change to selection. (Millstein, 2005, p. 172)[5]

On this view, then, there exists just one process – a sampling process – and this process can either take into account no physical differences (in which case it is a drift process), or it can take into account some physical differences (in which case it is a selective process). Selection and drift cannot happen simultaneously.[6]

Both the difficulty of applying these definitions of causal processes to the biological world and the gaps between the various notions we've seen here serve as reminders of the

---

5. It is notable that Millstein is here talking about how to label the *outcomes* of the evolutionary process, some particular set of "deviations from expectations." This makes comparing this passage with the above passage difficult. We will return to this process/product distinction below.

6. Millstein was strongly criticized for this position by Brandon (2005).

obvious truth that whether or not selection and drift qualify as causal processes will depend crucially on what the necessary and sufficient conditions are for something to constitute a causal process. But the extent to which the debate can turn on the definition of a causal process is often not fully appreciated. If a causal process *just is* a series of physical events (a set, temporally ordered concatenation, or the like), then since the individual-level events themselves are causal, there is a clear and trivial sense in which the process is causal as well. Selection and drift would straightforwardly then be causal processes – taking as the *only* required assumption the shared claim between causalists and statisticalists alike that individual-level events are causal. On the other hand, if we require an overarching, unifying cause to be present in order to qualify as a causal process, or if we require that there be some degree of causal responsibility of earlier events for later events, then whether or not selection and drift are causal processes may much more clearly turn on the debate over the composition question.

Given, then, that our definition of a causal process goes a long way toward resolving the debate between causalists and statisticalists outright, it is clear that we require metaphysical evidence independent of this debate in order to decide on the correct interpretation of causal processes. Otherwise, we are merely begging the question – one notion of a causal process is much more amenable to the existence of causal processes of selection and drift, and the other is much more conducive to a statisticalist interpretation. Without some independent method of establishing the concept of a causal process, we run a serious risk of circular argument.[7]

---

7. One might even think that the correct conception of a causal process is *prior* to understanding instance-causation; this seems to be an implication of Chakravartty's discussion of causal processes, for example (2007, p. 108).

4.2. The Metaphysics of Multi-Level Causation

Whether or not natural selection and genetic drift are genuine *causal processes* is, supposedly, the central issue in this debate. But we've already seen that the precise definition of causal processes hasn't been settled – and we thus run a serious risk of circularity. Further, when we look at the specific arguments offered by both sides, we will see that they primarily concern not the specifics of the case of drift and selection, but the circumstances, in general, in which a set of lower-level causal events constitutes a higher-level causal process, which I have termed the *composition question.* In this and the next two sections, I will argue that CQ is, by and large, not a problem that is relevant for the philosophy of biology. Let's consider three examples of causalist/statisticalist arguments.

Individual-level causal selection is a category mistake. Several causalist readings of natural selection place the causal action at the level of individual organisms (e.g., Hodge, 1987; Bouchard and Rosenberg, 2004). Walsh et al. (2002) offer an argument to the effect that this move is a category mistake. Selection, they argue, is a sorting process. And sorting "is not a force that causes a coin to fall head or tail" – rather, the sorting outcome "is explained and predicted by appeal to some statistical property [of the sequence], an *average* of individual propensities" (Walsh et al., 2002, p. 463). We therefore commit a category mistake if we attempt to "point to one or other of these individual-level causal processes and say *this* is a cause of the error and *this* one is a cause of sorting" (Walsh et al., 2002, p. 465).[8]

8. Millstein (2006a) makes the same argument using instead as a premise that biologists always define selection at the population level, therefore selection at the individual level should be ignored. I am skeptical of this argument, because I am doubtful that attempts to "read off" ontology from biological theories will succeed

The causal interpretation fails to understand population subdivisions. Consider a coin-tossing example in which two coins are tossed 50 times each (Walsh, 2007, p. 293).[9] The tosses are performed ten at a time by two experimenters – Walsh provides simulated data for such an experiment. Now, Walsh writes, "[t]here are at least three different, equally legitimate ways to describe this process" (2007, p. 293): (1) a single series of coin tosses, 49 heads and 51 tails; (2) two series of 50 tosses, 20/30 and 29/21; and (3) 10 series of 10 tosses, ranging from 5/5 to 7/3. If drift is just the analogue of sampling error in this experiment, then, Walsh argues, the causal interpretation "is committed to the claim that drift-the-cause is strong in the aggregate of 10 sequences of 10 tosses," as well as the claim that "[i]n the single sequence of 100 tosses, however, drift is not very strong at all. But these are not two populations; *they are different ways of describing the same population*" (Walsh, 2007, p. 296, original emphasis). Thus a causal notion of drift is both strong and not-strong in the same population, a supposed contradiction.

Northcott (2010, pp. 459–460) provides a causalist rebuttal to this argument – a case in which subdivisions are non-additive in the same way as genetic drift, but which is, at least intuitively, nonetheless causal. He has us imagine 100 slaves pushing a large rock in service of the pharaoh. The pharaoh is impressed by uniformity of effort and angered by lack of uniformity – so we may plausibly say that the variance of the 100 pushes *causes* the pharaoh's anger. This cause is non-additive in precisely the way that genetic drift is – the pharaoh should be angrier (on the average, at least) at smaller groups of workers than he should be at

(see later in this section for an argument for this skepticism).

9. It is interesting to note (in connection with my claim below that these arguments are non-biological) that while several initial expressions of this argument were phrased in biological terms (Walsh et al., 2002; Walsh, 2004; Pigliucci and Kaplan, 2006), the most recent (and clearest) version comes from Walsh (2007), in which he uses the example of a coin-tossing experiment.

the group as a whole. But this doesn't necessarily entail that the variance somehow *fails to cause* the pharaoh's anger.

Selection is a spontaneous statistical tendency, not a force. Matthen and Ariew (2009) lay out and critique a particular picture of the causal interpretation, one which they claim derives from Sober (1984). On this view, natural selection exists as an intervening step in the causal diagram of any evolutionary process – variation in traits causes selection, which in turn causes the births and deaths of animals responsible for evolutionary change. Such a "tertium quid," they argue, is improbable. Consider a series of tosses of a coin biased toward heads. They claim that "[t]he bias implies that it is probable, in a series of tosses of this coin, that heads will come up more often than tails. No process of 'toss selection' is needed for this result" (Matthen and Ariew, 2009, p. 206). The addition of a causally potent natural selection (on this reading of the causal interpretation) is therefore superfluous.

What's the point in discussing these three arguments? One feature of them, I think, stands out – or should. None of them are phrased with any reference to actual biological instances of drift and selection. In this and the other arguments relating to the causal efficacy of selection and drift, we see apple carts (Walsh, 2007), Newtonian gravitation, centers of mass (Matthen and Ariew, 2009), pharaoh's laborers (Northcott, 2010), scatter plots, smoking, heart disease, painkillers (Walsh, 2010), race cars (Ramsey, 2013c), and (many, many) coin flips. This seems to be good prima facie evidence that we are dealing with a general question about the metaphysics of multi-level causation, not a specific problem in the philosophy of biology.

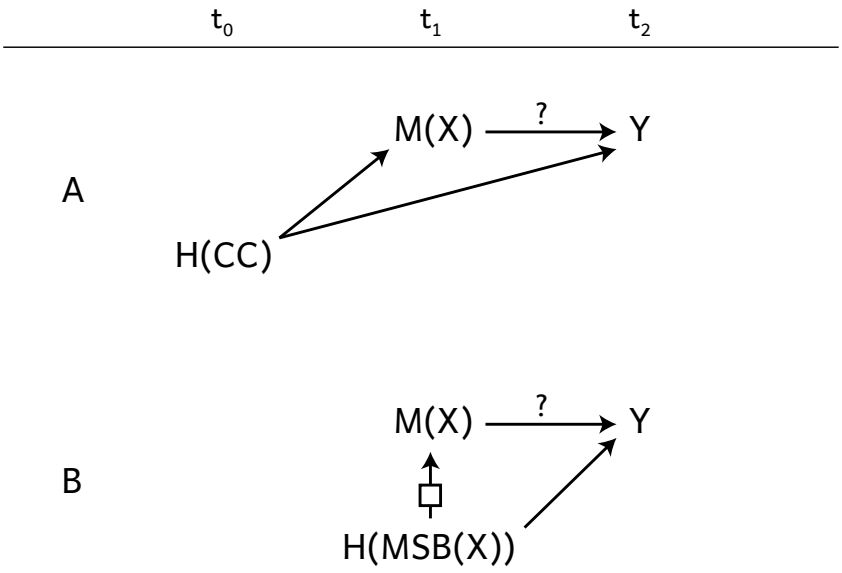

Figure 4.1: Two ways to test for the presence of a cause. $H(X)$ represents holding condition X fixed, $M(X)$ represents manipulating X. CC are the common causes, and $MSB(X)$ is the micro-level supervenience base of the macro-level phenomenon X. The boxed arrow holds with nomological necessity, and the question-marked arrow is the causal relation to be studied.

The clearest elucidation of this issue at the general, metaphysical level comes from Shapiro and Sober (2007). They consider both the causalist/statisticalist debate and the debate over epiphenomenalism of mental properties in the philosophy of mind – arguing that the question of mental epiphenomenalism and the composition question are two sides of the same coin. In essence, they claim, both of these boil down to a dispute over how to appropriately perform a manipulation to test whether or not a macro-level cause is present (see figure 4.1).

In both cases, we want to know whether or not manipulating the macro-level event X (at time t_1) will result in a change in the macro-level event Y (at time t_2). The question, Shapiro and Sober argue, is what we hold fixed in order to bring this manipulation about in the appropriate way. Two obvious possibilities present themselves. First, we could hold fixed

the prior (at some earlier time t_0) common causes of X and Y (called CC in the figure) – this is "option A." Alternatively, we could hold fixed the micro-level supervenience base of X (called $MSB(X)$ in the figure) – this is "option B." They argue that adopting option B amounts to evaluating what would happen in a nomologically impossible situation: considering whether or not Y would occur if a nomologically-connected sufficient condition for X occurred (namely, $MSB(X)$), but X did not. Option A is thus the correct choice. The statisticalists, they hope to show, have argued on the basis of option-B analyses of selection and drift, while causalists have argued (correctly) from option-A analyses.

Whatever the merits of this particular argument about manipulability turn out to be, Shapiro and Sober have correctly noted that the composition question in the biological case is a metaphysical question about causation and supervenience in multiple-level systems, to be analyzed by appeals to our best theories of what counts as a cause in general.[10] Of course, in and of itself, this is no problem at all. For our evaluation of the causalist/statisticalist debate, however, one question is now particularly important: is this general debate over the composition question truly relevant to the philosophy of biology?

There are two ways, it seems, in which the marriage of CQ and the debate over selection and drift could be a particularly natural one. On the one hand, it might be the case that natural selection and genetic drift constitute a particularly fertile place in which to study CQ itself – that is, selection and drift might serve as an example that could help us resolve this broader metaphysical problem. On the other hand, it could be the case that taking our general metaphysical work on CQ and applying it to the specific case of selection and drift is

10. They also draw an analogy between the composition question in general and the particular case of epiphenomenalism in the philosophy of mind, an analogy which I will later critique.

useful for our understanding of the evolutionary process – that is, that CQ can serve as a lens for considering the way in which we ought to understand selection and drift. In the next two sections, I will argue that neither of these possibilities pans out. In the end, the composition question is not a relevant issue for the philosophy of biology.

4.3. Examining the Composition Question with Selection and Drift

Let's consider first the use of selection and drift as a case study for the composition question. I claim that evolutionary theory does not serve as a particularly illuminating example for CQ, for two reasons. First, as we will see in section 4.6, the very definitions of natural selection and genetic drift are themselves disputed – and much of the literature that attempts to sort out this thorny causal issue establishes neither what selection and drift *are,* nor how the various possible definitions on offer relate to one another. Without clarity at this fundamental level, it is quite difficult to see how selection and drift could profitably contribute to the debate over CQ. In particular, the differing definitions of selection and drift result in differing causal structures for the evolutionary process, a confusion which renders it that much more difficult to make inferences about the composition question from the example of evolutionary theory.

Second, and more importantly, there is a significant problem with the connection between biological theory and general issues in the metaphysics of science.[11] Philosophers of biology who work at the highly general or foundational level at which the causalist/statisticalist

11. For example, while I lack the space to pursue the claim here, much of what I argue in this and the next section applies equally well to the exploration of the realism/anti-realism debate in the biological context.

debate is situated often – of necessity – employ particularly simple models of selection and drift. Only in the context of these austere biological models may we clearly partition the influences of selection and drift, forming tractable thought experiments on which the contributions of the various components of the evolutionary process can be considered. But we all are aware that the real world is much more complex than this. The precise interplay of selection, drift, mutation, and migration required to offer a plausible explanation of even a relatively uncomplicated Mendelian trait like blood type in humans is incredibly hard to describe in detail.[12] As another example, the standard recurrence equations used to provide predictions for how trait fitness affects a population when selection alone is at work are derived in the context of Hardy-Weinberg equilibrium. But populations at Hardy-Weinberg equilibrium are highly idiosyncratic – in fact, the standard response to evolutionary stasis is to invoke powerful stabilizing selection, as noted by McShea and Brandon (2010, p. 120). It seems, therefore, that giving a complete, biologically accurate account of selection and drift will make it substantially *more difficult* to tease out valuable insights for general issues such as CQ. It thus appears all the more likely that selection and drift do not provide us a beneficial context in which to investigate the composition question.

Two responses to this conclusion are obvious. First, one might argue that the intent of using simplified biological models, which I have argued is a problem for the study of the composition question, is precisely to avoid the two concerns that I mention here – to (roughly) standardize our working notions of selection and drift, and to make them tractable examples for metaphysical theorizing. After all, this is a relatively common argumentative

12. I thank Mark Jordan for the example as well as for pushing me to consider the impact of biological complexity on the debate.

tactic in metaphysics: we locate a simply stated yet relevantly challenging example with bearing on a metaphysical problem, and we explore how various intuitive analyses of this example might alter our general metaphysical perspective. The simplicity is a necessary condition for gaining intuitive traction on the case.

There are several problems with such a reply, however. As we will see in section 4.6, even in the simple cases usually presented, we in fact have failed to sufficiently standardize our definitions of selection and drift. Further, the complexity of contemporary work on causation makes it all the more unlikely that genuine progress on the composition question can be made by appeal to a few "intuitive" biological examples. It is not "intuitively clear," for example, even in these simple cases, whether natural selection satisfies the criteria of a manipulationist, counterfactual, or a mechanism-based account of causation, nor should we expect it to be. Finally, intuition can be an unsteady guide in many evolutionary cases. As mentioned above in the contexts of Hardy-Weinberg equilibrium and blood type, cases which seem simple on one axis (population stasis or simple Mendelian heredity, respectively) can frequently be profoundly complex along other axes, and complex in ways that produce a very opaque causal structure, making it difficult to draw inferences from biological examples to metaphysical conclusions.

As a second response to the problems raised here, one might look toward biological practice. If the use of simplified biological examples is not an effective way to examine the composition question, perhaps we could turn the other way – attempting to "read off" the correct answers to our metaphysical worries from the use of biological theories in real-world examples from the scientific literature. One strand within the debate has in fact attempted to work in this manner, remaining biologically informed and working out the way in which

selection and drift are distinguished (and, it is argued, used *as causes*) in evolutionary biology. Most prominent here is the work of Millstein, who has used examples from the montane willow beetle (Millstein, 2006a), land snails (Millstein, 2008), and the early development of the theory of genetic drift by Fisher and Wright (Millstein et al., 2009). The effort to bring the debate back to a biological context is long overdue, and crucial to the relevance of our philosophical work. When applied to the composition question (and other metaphysical issues), however, I am skeptical of its chances for success. Millstein herself is all too cognizant of the limitations of this approach when the questions at issue are as abstract as these. She notes that her examples can only show that viewing "natural selection [as] a causal process *is consistent with at least some biological practice*" (Millstein, 2006a, p. 637, emphasis added), or that "Wright and Fisher had good reason to think that there were biological processes in nature that *needed to be represented in drift models* in order to better track evolutionary changes in a population" (Millstein et al., 2009, p. 8). In both statements, we see a clear recognition of the difficulty of making metaphysical inferences from the practice of working biologists.

And, I claim, there is a well-grounded philosophical reason for the existence of this difficulty. Several authors, perhaps most persuasively Waters (2011), have made the point that many practicing biologists are what we might call "toolbox theorists." Such scientists "seek true theoretical accounts that best address particular interests" in specific examples (Waters, 2011, p. 233). Later, he elaborates:

> Perhaps, given the messiness of the world, the ideal theory turns out to be more like a toolbox than a fundamental framework. Perhaps what is sometimes thought to be a single concept, such as fitness, consists of a family of concepts, each useful for theorizing about different aspects, parts, or scales of entangled causal processes. (Waters, 2011, p. 240)

Setting aside the accuracy of this view as a broader prescription for science or philosophy as a whole, we have ample evidence that practicing biologists do often view their theoretical enterprise in this way. Picking up Waters's invocation of fitness, for example, the biologist Gerdien de Jong has reviewed the use of 'fitness' (in *only* the restricted numerical sense familiar from population genetics) in the biological literature, and covers dozens of fitness measures which can be categorized in several ways, and which feature in several different models of the selective process. In her conclusion, however, she asks the following provocative question:

> What is fitness supposed to measure? Some sort of overall performance or quality of design or aptness for life, or general adaptedness? *This seems to be a discussion that is standing outside the practice of evolutionary biology.* What we usually ask is how such measures … are interrelated, and how they relate to the change in phenotypic traits. Given knowledge of the life history *of the population,* the causes and values of the fitness components, *these relations can be spelled out in specific models. We need not ask whether expected time to extinction or genotypic weight is the proper fitness measure.* What we would like to know is how they are related in a mechanistic model for a specified situation. (de Jong, 1994, p. 18, emphasis added)

I have quoted extensively from de Jong not in an attempt to establish her view as the proper understanding of fitness, but rather because her survey provides a particularly clear demonstration of toolbox theorizing. The question of the "correct" model of fitness, she argues, lies entirely outside the province of evolutionary biology. Rather, we take a particular system, with a particular model which specifies the components of fitness of interest. This model and system, then, provide a context in which we study the relationship of the various fitness measures we find in our toolbox. If this is how practicing biologists view (or, at least, how one influential practicing biologist views) fundamental theoretical concepts like fitness, selection, and drift, then it is a mistake to turn to the biological literature in an attempt to resolve metaphysical questions such as these factors' status as causal processes.

We thus have two compelling reasons that selection and drift make for a bad environment in which to investigate CQ. First, as we will see below, the definitions of selection and drift themselves are unsettled. And second, the simplistic models of selection and drift we use in this context drive a large wedge between our understanding of CQ in abstract cases (like that of figure 4.1) and the real biological world. Two ways of responding to these issues – using either simplified biological models or real-world case studies – have also been shown to be problematic. Selection and drift are not helpful in advancing our understanding of the composition question in general.

4.4. Examining Selection and Drift with the Composition Question

Perhaps, then, the relevance of the composition question for the philosophy of biology points the other way: general work in the metaphysics of science might be profitably applied to the case of evolutionary theory in order to sharpen or enhance our understanding of natural selection and genetic drift. Several papers have offered us an example of how this debate might look given significant interaction with work on causation in metaphysics and general philosophy of science. Within the last decade, we've seen a surge of new interest in the development of theories of causation robust enough to explain causal inferences in the special sciences – particularly relevant examples include the "manipulationist" definition of causation (Woodward, 2000, 2003; Woodward and Hitchcock, 2003; Hitchcock and Woodward, 2003; Woodward, 2010) and contemporary work on "mechanisms" (Machamer et al., 2000; Barros, 2008; Fazekas and Kertész, 2011).

Several authors in the causalist/statisticalist debate have endeavored to utilize these results to clarify the causal status of natural selection and genetic drift. Reisman and Forber (2005; Forber and Reisman, 2007) have considered whether or not the manipulationist account of causation would deem selection and drift to be causes; they conclude that it would. The point is, unsurprisingly, controversial – Lewens (2010) and Walsh (2010) argue that Reisman and Forber misapply the manipulationist causal criterion, while Shapiro and Sober (2007) and Abrams (2007) agree with Reisman and Forber's analysis.[13] While the current state of this particular line of argument may be unsettled and unsatisfying, it demonstrates an emerging recognition that the correct answer to CQ cannot either be constructed from simple examples or read off of the theories of selection and drift.

There are, however, several reasons to be skeptical that even this reimagined version of the composition question project will succeed. First, it is important not to overstate the level of unanimity in the literature on causation. There is most certainly no "received view" which may be readily applied to the biological case. The position with perhaps the most traction of late – Woodward and Hitchcock's manipulationist approach – is still highly controversial, and as noted above, the way in which it should be applied to the biological case has been hotly contested. Second, the problems related to definitions of selection and drift and the complexity of the biological world raised in the last section will still apply to this approach. General work on causation cannot tell us much without settled, biologically realistic definitions of selection and drift. Finally, one might be worried about the prospects for the resolution of CQ in general. It might not admit of a univocal, general answer – and

13. See also Gildenhuys (2009) for a manipulationist analysis of genetic drift. Millstein (2006a) offers brief arguments that selection satisfies *all* of the current contenders for an analysis of causation.

even if it does, it is possible that whatever answer is provided by our best metaphysics cannot be applied to the particular instance of natural selection and genetic drift.[14] If one of these more pessimistic scenarios is the case, then the failure of the philosophy-of-biology approach combined with the failure of the metaphysical approach would, it seems, force us to conclude that the composition question as applied to selection and drift is genuinely unanswerable. The prospects for the success of this project, then, should be considered an open question.

4.5. A Few Objections

The last three sections, then, have argued that the composition question is simply not relevant for the practice of the philosophy of biology. Several objections are obvious at this point, and I should stop to consider them before continuing. First, I want to forestall one particular misreading of my argument. Setting aside the caveats raised above, I do not wish to claim that CQ in general – describing the circumstances under which lower-level causal events compose to produce higher-level causal events – is uninteresting or unanswerable. Recent work on causation is quite promising, and could plausibly be extended and refined to provide us with a resolution to CQ. I further don't have any reason yet to conclude that the composition question *in the specific biological context* is vacuous or unanswerable. It seems to me plausible that there is a unique and correct answer to whether or not the individual-level events that are responsible for natural selection form a higher-level causal process. But at the moment, the concerns raised in the preceding sections make it unlikely that the philosopher

14. I thank Dan Hicks for bringing this point to my attention.

of biology has anything substantive to contribute to this debate – evidence, I claim, that CQ is best left within the confines of the metaphysics room.

Second, one might read my argument as an injunction against the analysis of the causal structure of theories in general. This, as well, is a claim I do not intend to make. The analysis of the causal structure of scientific theories is a part of the stock-in-trade of the philosophy of science, and much positive and fruitful work has been done in this vein for many different theories. It is, rather, a consequence of several particular features of the landscape of *this* debate that make evolutionary theory a special case – in particular, the complexity of the biological world and the difficulty of reading ontological commitments off of biological theories, combined with the general difficulty of CQ as a whole. These facts conspire to make our inferences about the metaphysical structure of evolutionary theory particularly tenuous. And the most obvious way to avoid the unpalatable consequences of this is the solution that I advocate – the analysis of the composition question within the context of the metaphysics of science.

Finally, Shapiro and Sober have argued that the composition question, both in general and in the philosophy of biology, is roughly parallel with the issue of the epiphenomenalism of mental properties in the philosophy of mind. One might, however, be worried that this analogy is defective.[15] Consider the general argument regarding epiphenomenalism presented by Shapiro and Sober and diagrammed in figure 4.1. In the philosophy-of-mind case, we have a set of brain states at the "micro" level, and a set of mental properties at the "macro" level. One way of interpreting the problem of mental properties is not as a causal issue, but

15. Thanks to Anjan Chakravartty for pointing out this objection, and to Mohan Matthen for encouraging me to drive a larger wedge between these two problems.

as an issue of determinates vs. determinables. Is there some determinable mental property, "anger" (say), of which the neuronal configurations are determinates or instances? Phrased in this way, the debate over epiphenomenalism becomes clearly metaphysical – the question of whether or not determinables "really exist" is almost certainly one with an *a priori* answer.

If this is the correct reading of mental epiphenomenalism, then there seems to be an important disanalogy between this and the composition question in biology. The important question for us as regards selection and drift is not whether or not there "really is" natural selection – selection is not a universal for which we are in search of instantiations, and which we may interpret in a realist or anti-realist manner. Rather, the question is whether or not the individual-level events form a *very particular kind of* macro-level thing: a *causal process*. This returns us to the trouble described in section 4.1: the debate has not settled on a clear definition of a causal process. And the trouble with unsettled definitions, as we will now see, goes deeper than this: we also are not operating with a single definition of either natural selection or genetic drift.

4.6. Selection and Drift: Process or Product?

Finally, we turn to the issue which, I argue, evinces a genuine *biological* difference between the causalist and statisticalist positions: our varied definitions of natural selection and genetic drift. First, in this section, we will see that there exists a wide array of definitions of selection and drift, which can be categorized readily based upon whether they define drift and selection as processes or products. Importantly, while these two categories are highly

correlated with the two sides in the causalist/statisticalist debate, the mapping is not perfect, and each category is perfectly compatible with a causal or statistical interpretation of selection and drift. Then, in the next section, I will argue that while these definitions *intensionally* entail different conceptual structures for the evolutionary process, their *extensions* are much less distinct than is often supposed. Answering the question of which intension is correct, then, will be best accomplished by looking at cases where the extensions come apart – where there is a genuine conceptual difficulty in understanding how the various components of the evolutionary process are to be identified and individuated. These cases, then, ought to constitute our focus in coming to better understand the structure of evolving systems.

We should begin by drawing a distinction between two categories of selection and drift definitions: what I call "process" and "product" concepts.[16] Though it is difficult to speak for all authors in the discussion (and I will elaborate on these definitions in the next section), we might initially characterize their positions as follows:

> *selection (process):* a process of sampling that discriminates between individual organisms based on differences in fitness (or merely any physical or causally-relevant differences whatsoever)
>
> *selection (product):* a change (result, outcome) which is predicted (or explained, or both) by differences in the fitnesses of traits
>
> *drift (process):* a process of sampling that does *not* discriminate between individual organisms
>
> *drift (product):* a change (result, outcome) which is *not* predicted (or explained, or both) by differences in the fitnesses of traits

The process view of selection tends to be most often associated with the causalists. Beatty defines natural selection as "a sampling process that discriminates, in particular, on the basis of *fitness* differences" (1984, p. 190). Hodge defines it as "what is occurring when

16. This distinction has also been drawn out in detail by (among others) Matthen and Ariew (2009) and by Millstein (2002), who describes it as a difference between process and outcome notions of selection and drift.

and only when there is the nonfortuitous differential reproduction of hereditary variants" (1987, p. 251). Bouchard and Rosenberg describe it as "a contingent causal process in which individual fitness differences are the causes and subsequent population differences are the results" (2004, p. 710). Millstein argues that natural selection is a process that acts on populations of organisms, on the basis of "differences in abilities to survive and reproduce" (2006a, p. 643).[17]

Much the same is true of the process view of genetic drift. Beatty defines drift as "indiscriminate" sampling, on which "any physical differences between the entities in question are irrelevant to whether or not they are sampled" (1984, p. 189). Millstein describes it as "an indiscriminate sampling process whereby physical differences between organisms are causally irrelevant to differences in reproductive success" (2006a, p. 640). Slightly differently, Bouchard and Rosenberg define drift as occurring in cases where "the initial conditions in the divergence [are] rare, improbable, and unrepresentative of the whole population of initial conditions" (2004, p. 708).[18]

The product view of natural selection, on the other hand, is associated with the statisticalists. Walsh has defined selection several slightly different ways, always consistent with a product-based view – as "the consequence [elsewhere, an 'effect'] of the differential rates of distinct causal processes occurring within individuals" (2000, p. 141), as that which "explains the changes in the structure of a population by appeal to differences in trait fitness" (2003, p. 289), and as "a change in population structure predicted and explained by variation

17. For another instance of this view, see Filler (2009, pp. 774–775).

18. For more instances of this process view of drift, see Hodge (1987), Millstein et al. (2009, p. 1) and Filler (2009, pp. 774–775).

in trait fitness" (2004, p. 351). Walsh, Lewens, and Ariew describe selection as what occurs "only when the relative frequency of trait types changes in a population as a consequence of differences in the average fitnesses of individuals in different trait-classes" (2002, p. 464). Matthen and Ariew describe it as "a statistical trend emerging from events that occur in these 'substrates' [chromosomes, reproductive systems, body plans, developmental sequences, etc.]" (2002, p. 68).[19]

And again, for genetic drift, we see the product view often adopted by statisticalists. Walsh defines drift as "a change in the structure of a population that is not predicted or explained by trait fitnesses" (2004, p. 351), and elsewhere claims that "what it is for a change in relative trait frequencies to constitute selection (or drift) is merely for it to be susceptible to a certain kind of statistical description" (2007, p. 282). Matthen argues that genetic drift is best described as the "uncertainty associated with" the "spread of possible outcomes" in a finite series of trials of a specific sort of probabilistic process (2009, p. 484).

As a general rule, then, we can see that causalists argue from process definitions of drift and selection, while statisticalists argue from product definitions. Why not, then, just refer to these as causal and statistical notions of selection and drift?

As it turns out, things are more complicated than they first seem, which necessitates a shift in terminology. First, one prominent definition of genetic drift adopted by some causalist authors is clearly a product-based notion. Most noteworthy among these is Brandon, who has several times defined drift as "any deviation from the expected result due to sampling error" (2005, p. 158). Brandon is, however, most certainly a causalist (see, e.g., Brandon (2006,

19. And yet more: see Brunnander (2006, p. 245), Walsh (2007, p. 282), Matthen (2009, p. 484), and Matthen and Ariew (2009, p. 222).

2010); McShea and Brandon (2010)). Referring to product-drift as "statistical," then, would muddy the waters.

But more importantly, whether one accepts process or product definitions of selection and drift is *independent* of whether or not one takes them to be causal. It is worthwhile to defend this claim in some detail, as the relatively strong alignment of causal/statistical with process/product may make it rather surprising.[20]

Consider, as an exemplar of the process view of selection and drift, Hodge's indiscriminate and discriminate sampling processes. Indiscriminate sampling can certainly be non-causal. Imagine shuffling a deck of cards and then spreading them out on a table in a line. After the fact, we choose the first five or last five cards, and call this a "sampling" event. Nothing about the card selection process causally impinged on the cards. Indeed, the fact that a central element of the process was entirely subjective seems to provide paradigmatic evidence for its being "non-causal" in this sense (a point that is often made by the statisticalists in the context of genetic drift). And discriminate sampling could be non-causal as well – we simply choose instead the first five red cards, or the last five black cards, and we have discriminate, non-causal (again, subjective) sampling.

Of course, these processes could also be causal. Spread out the same deck of cards, then fire a very large dart with a suction cup on the end at the line. Say the suction cup isn't perfectly sticky, so it grabs only 80 percent of the cards it touches. Now we have two causal processes that have, together, resulted in indiscriminate sampling – the firing of the dart gun, clearly causal, and the precise physical details of the stickiness of the suction cup, also

20. Some language in Walsh (2007) seems to indicate that he may have noticed this feature of these two classes of definitions, though I cannot say for certain.

causal. We can make the sampling discriminate by making some of the cards stickier than others – these sticky cards are more likely to be picked up by the suction cup, so they will be overrepresented in our random sample.[21]

For the product view of natural selection, we need only consider the outcomes from the cases already considered. In the non-causal, discriminate case above, if we choose the first five red cards, then their being chosen is (partially) explained by (and would have, before-the-fact, been predicted by) the fact that the cards are red (the higher "trait fitness" of red against black). This, then, is a non-causal product view of selection. And the example of the sticky cards and dart gun is causal product-selection – the cards that are chosen are again (partially) explained and predicted by the fact that some of them are sticky (the higher "trait fitness" of sticky against non-sticky).

As a token example of the product view of genetic drift, the "sampling error" definition of drift proposed by Brandon and Walsh is also amenable to both causal and non-causal analysis. We have already seen a causal example – the 80% success rate of our suction cup counts as a causal sort of sampling error. But it could be non-causal as well, as both Walsh and Brandon note. Let's say we spread out our deck of cards and want to select one third of them. Of course, we can't successfully grab 17 1/3 cards, so we will by necessity be off by a fraction of a card. This seems to be a paradigm case of non-causal sampling error – the "error" in our sample is merely a result of mathematics.[22]

---

21. This example is similar to one involving colored balls in an urn developed by Brandon and Carson (1996, pp. 321–325), though they take this case as evidence that selection and drift are both causal and "stochastic" in their sense of the term. The claim that drift must be causal because of this example is refuted by the analysis here. Brandon and Carson's notion of stochasticity was already problematized in chapter 1.

22. Again, this case is very similar to one that Brandon and Carson (1996) describe, although the conclusion they draw from it – that genetic drift is "inevitable" in certain types of populations – only follows if

It is thus clear that the distinction between process and product notions of genetic drift and natural selection has little to do with whether drift and selection are causal or statistical. So what *can* we say about these two sets of definitions?

4.7. Are Process and Product Definitions Related?

It is clear that the process and product definitions of selection and drift entail substantially distinct views of the conceptual structure of evolutionary theory. But I will argue in this section for an unconventional, if perhaps unsurprising, thesis: that these two sets of definitions *extensionally* differ much less than is often thought.[23] The interesting cases, I claim – the cases where we might actually resolve the question of which *intension* is the correct one – lie at the fringes where these two classes of definitions come apart. In these instances, we have genuine conceptual difficulty in delineating the structure of the evolutionary process – and here is ripe work for philosophers of biology.

One strange feature of the argument in this section should be noted at the outset. We are focusing here on the definitions of selection and drift deployed by both process and product advocates in the debate – *not* on whether or not selection and drift should be conceived of as causes, for reasons already noted. Thus, I will, for the moment, drop all reference to the causal or non-causal character of these processes or outcomes.

A second issue, made particularly acute by section 4.1, has to be resolved as well. What

one adopts their product definition of genetic drift. A process notion of genetic drift, as we have seen, might or might not be operating in such an instance.

23. Unsurprising, I claim, because these are after all intended to be multiple definitions of *the same* concept – we should expect their extensions to be roughly equivalent. The structure of this equivalence, however, has not been carefully explored in the literature.

will my working definition of 'process' be in the following? Despite the risk of circularity discussed above, we must take *some* stance in order to be able to analyze these conceptions of selection and drift. The risk, however, is entirely avoided if we recall a point made above: whether selection and drift are processes or products in the sense at work in this discussion has nothing to do with whether or not they are causal. For our purposes here, then, we do not require a definition of a *causal* process, merely a definition of a process *simpliciter,* like that deployed initially by Salmon. For that purpose, I will use a particularly minimal definition – by "process" in the remainder of this discussion I intend nothing more than a temporally ordered sequence of states of a system. To reiterate, I do *not* mean to imply that any such process is causal – to do that would be to beg the question against the statisticalist. It may merely be a pseudo-process, a convenient way of understanding the behavior of a system, with no causal import whatsoever. Importantly, though, we will see that even on this weak definition of process, we still generate the kinds of interesting corner cases that I will describe below.

My approach, then, will be to consider two different definitions of natural selection and four different definitions of genetic drift.[24] (See table 4.1 for the particular definitions of drift and selection used in the following.) We will take them pairwise, seeing in what circumstances selection (or drift) in the first sense is extensionally equivalent to selection (or drift) in the other, and vice versa.

---

24. Considerations of space prevent me from engaging the NINPIC view of Gildenhuys (2009), which endeavors to specify quite precisely just the sorts of causes that ought to count as genetic drift. I predict many of the arguments I apply to GD^{SP}_{proc} would apply to the NINPIC view as well.

TABLE 4.1

DEFINITIONS OF SELECTION AND DRIFT

Abbrev.	Process / Product	Definition	Citations
NS_{proc}	process	A sampling process which discriminates on the basis of fitness differences	Hodge (1987); Beatty (1984); Millstein (2006a)
NS_{prod}	product	Population outcomes predicted/explained by trait fitness differences	Walsh (2004, 2007); Matthen (2009)
GD^{SP}_{proc}	process	A sampling process which does *not* discriminate on the basis of fitness differences	Hodge (1987); Beatty (1984); Millstein (2006a)
GD^{IC}_{proc}	process	Initial conditions as unrepresentative of the possibility space	Bouchard and Rosenberg (2004)
GD^{TF}_{prod}	product	Population outcomes *not* predicted/explained by trait fitness differences	Walsh (2004, 2007); Matthen (2009)
GD^{SE}_{prod}	product	Any deviation from the expected result due to sampling error	Brandon (2005, 2006)

NS_{proc} vs. NS_{prod}. Natural selection is the easiest, as it in fact has been defined in only one way by process and product theorists, respectively.[25] On the process side, natural selection is (NS_{proc}) a process that discriminates between organisms based upon differences in their fitness values. On the product side, it is (NS_{prod}) any change in a population that can be explained by variation in fitness.

It is immediately obvious that these two definitions can quite peacefully coexist in effectively all biological circumstances. Any process that discriminates between individuals based on fitness differences is highly likely to produce outcomes that require for their explanation reference to those very fitness differences – the process definition, in most cases, entails the product definition. And if we have some change in a population that may be explained by variation in fitness, then there must have been *some* process (again, in the minimal sense assumed for this section) which brought that difference about – and that process, given that fitness features in the explanation, must take fitness differences into account. (Note, to reiterate a point made earlier, that I do *not* mean to imply that this process would or would not be causal.) The product definition thus usually entails the process definition. For natural selection, therefore, the process and product definitions are, in almost all instances, extensionally equivalent.

GD^{SP}_{proc} vs. GD^{TF}_{prod}. A tradition in process definitions of genetic drift, beginning with Hodge (1987), defines drift (analogous to the definition of selection above) as (GD^{SP}_{proc}, for "Sampling Process") an indiscriminate sampling process – sampling that does not discriminate between

25. Ramsey (2013a) gestures at another view, which has yet to be described in detail; it is a selective counterpart to GD^{IC}_{proc}, intended to be interpreted causally. I lack the space to engage fully with this view, but I hazard it will be equivalent to NS_{proc}, in much the same manner as I demonstrate below for GD^{IC}_{proc} and GD^{SP}_{proc}.

organisms based on fitness differences. And the most common product definition of drift (again, analogous to selection) defines drift as (GD^{TF}_{prod}, for "Trait Fitness") any change in a population that cannot be predicted or explained by variation in fitness.

Once again, similar to the case of selection, it is clear that these two definitions are extensionally equivalent in almost all cases – processes that do not discriminate based on fitness will not be explained by reference to fitness differences. But the precise details are made more complicated by GD^{TF}_{prod}, primarily because it requires us to demonstrate a negative – that the change at work *cannot* be explained by reference to trait fitness.

First, as Millstein points out, "an indiscriminate sampling process can produce what looks like a directed outcome (mimicking the most likely outcomes of a discriminate sampling process)" (2005, p. 172). That is, it is possible that a population which is actually undergoing no selection whatsoever could produce the same result that would have been expected were offspring produced according to fitness differences. In this case, we might well infer on the basis of NS_{prod} that this change could be explained by fitness differences (and was selective), when in fact an *in*discriminate sampling process was responsible for it. Walsh has precisely this characteristic in mind when he notes that "[i]t is an unfortunate consequence of [the conjunction of NS_{prod} and GD^{SP}_{proc}] that drift-the-process [i.e., GD^{SP}_{proc}] causes selection-the-effect [i.e., NS_{prod}]" (Walsh, 2010, p. 154).[26]

Second, consider what it would take to prove in general that GD^{TF}_{prod} is extensionally equivalent to GD^{SP}_{proc}. We take as a premise that a population has undergone a change which cannot be explained or predicted by reference to fitness differences. But how can we really

26. Notably, no author in this debate simultaneously accepts both NS_{prod} and GD^{SP}_{proc}, so it is not clear who is the target of Walsh's argument.

infer further that the process responsible for that change did not discriminate at all with respect to fitness differences? Perhaps the process has fitness differences as a core feature, but combines these with a substantial stochastic element, altering the results significantly enough that the best *explanation* of the change no longer requires any reference to those fitness differences. Explanations have a subjective element that causal processes do not. In such a case, we would have drift in the sense of GD^{TF}_{prod} without GD^{SP}_{proc}, a failure of extensional equivalence.

It is precisely these cases which deserve further scrutiny, and to which I will return later. But it is clear that for the vast majority of biologically plausible processes, GD^{SP}_{proc} and GD^{TF}_{prod} provide the same demarcation of selection and drift. Any process that in fact does not discriminate with regard to fitness differences should in most cases not require fitness differences in the explanation of its outcomes – GD^{SP}_{proc} thus usually entails GD^{TF}_{prod}. And any change in a population which is not explained by fitness differences at least gives us a prima facie reason to believe that the process responsible for that change, in turn, does not rely on fitness differences – GD^{TF}_{prod} does seem to, at least weakly, provide evidence for the presence of drift in the sense of GD^{SP}_{proc} (with acknowledgment of the caveats above).

GD^{SP}_{proc} vs. GD^{SE}_{prod}. A second product definition of genetic drift is popular in the literature, primarily due to Brandon – genetic drift is (GD^{SE}_{prod}, for "Sampling Error") any deviation from the expected outcome in a population as a result of sampling error. In general, the varieties of sampling error implicated in this definition are broad: the sampling of gametes to form offspring, the sampling of parents in the current generation to form the offspring in the next generation (sampling that could result from an event like a forest fire or disease), sampling of

survivors of a population bottleneck, sampling of individuals in a geographic isolation event, and so forth.[27] What is the relationship between this definition and GD^{SP}_{proc}?

The question here is a bit thorny. When we consider the specific biological processes that most authors include under the heading of the sampling error mentioned in GD^{SE}_{prod}, we can see that the vast majority of these satisfy the indiscriminate sampling criterion present in GD^{SP}_{proc}. Beatty, for example, writes that these processes

> are importantly similar agents of change – both involving elements of randomness, and "randomness" in the same sense in both cases. According to this sense of "randomness", sampling from a population is random when each member of the population has the same chance of being sampled. (Beatty, 1984, p. 190)

But this is not sufficient to demonstrate that GD^{SE}_{prod} entails GD^{SP}_{proc}. For recall that GD^{SE}_{prod} defines drift as any *deviation from the expected result* as a result of sampling error. It is thus not sufficient, on GD^{SE}_{prod}, for a process like this to be *acting* – it must further be the case that the result in some particular instance *differs* from the result we would expect given the action of natural selection alone.

To put the difficulty differently, if we have the circumstances supposed by GD^{SE}_{prod} – a population deviation from expected results due to sampling error – then it is clearly the case that we may infer GD^{SP}_{proc}. All the various processes that would generate this deviation from expectation, as it is described by the authors that defend it, satisfy the indiscriminate sampling criterion of GD^{SP}_{proc}.

But what about the opposite inference? Here we meet with the problem. For we require the existence of a process like that supposed by GD^{SP}_{proc}, and, in addition, we require that process to drive the population to an outcome *other* than that predicted on the basis of

27. See Beatty (1992) for a helpful categorization of the various meanings of genetic drift in the biological literature.

selection alone. We mentioned with regard to GD^{TF}_{prod} that there are many cases in which an indiscriminate sampling process produces the same result as would have been predicted on the basis of selection, and that this could cause difficulty for product definitions of drift and selection. In this case, too, such situations *definitionally* fail to qualify as genetic drift – if the outcome is that predicted by selection, then there *is no* genetic drift according to GD^{SE}_{prod}. Again, we are forced to conclude that GD^{SP}_{proc} provides only a weak inference to GD^{SE}_{prod}, and we are pointed toward a set of biological instances where the connection between the two breaks down.

GD^{IC}_{proc} and GD^{SP}_{proc}. Finally, we have one more process notion of genetic drift to consider, due to Bouchard and Rosenberg (2004). On this definition, genetic drift (GD^{IC}_{proc}, for "Initial Conditions") occurs whenever the initial state of a population, prior to an evolutionary change, constitutes an unrepresentative sample of the space of all possible initial conditions. It is worth noting that one might dispute whether this is genuinely a "process" notion of genetic drift. It indeed does not describe features of the *processes* responsible for drift, but rather a feature of the *relationship* between the population and those processes – a claim about where the members of the population stand within the population's space of possible initial conditions (which is a property of the process).

My decision to describe this as a process notion of genetic drift, however, is due to the analysis that follows. In most cases, we will see that GD^{IC}_{proc} is extensionally equivalent to GD^{SP}_{proc} – if a population changes as a result of the unrepresentativeness of its initial conditions, then the process responsible for that change must have been indiscriminate in the sense of GD^{SP}_{proc}.

Let's begin by clarifying what Bouchard and Rosenberg mean by "unrepresentative." They must not mean that the population is undergoing genetic drift if it is unrepresentatively distributed with respect to *fitness* (or, to make the metaphor more picturesque, unrepresentatively distributed *on the fitness landscape*). For a population may be distributed in any given way on the fitness landscape and nonetheless be under only the influence of selection. Consider, for example, an infinite population, undergoing no "drift-like" processes whatsoever. This population is situated on an odd fitness landscape – the entire landscape slopes upward in the direction of some particular allele, excepting a tiny bump – a local optimum – at an otherwise very low fitness value (see figure 4.2). Now situate the infinite population near, but not quite at, the peak of this local optimum. Its selective force will be much smaller – due, to return to the visual metaphor, to the smaller gradient of the fitness landscape near this peak – than a population at any other location on the fitness landscape. It is thus undergoing a peculiar sort of change due to the unrepresentative nature of its distribution across the fitness landscape. But there is no drift here, by hypothesis – Bouchard and Rosenberg must not mean unrepresentativeness in this sense.

Indeed, Bouchard and Rosenberg's extensive focus on ecology leads us to believe that the correct sort of unrepresentative distribution must be unrepresentative *ecological* distribution – or, to take the useful framework deployed by Pence and Ramsey (2013), the unequal distribution of the population in the space of *possible life histories* of organisms.[28]

Why would this sort of unequal distribution result in something looking like genetic drift? Consider what a discriminate and indiscriminate sampling process would each look

28. See the next chapter for more on the notion of possible life histories.

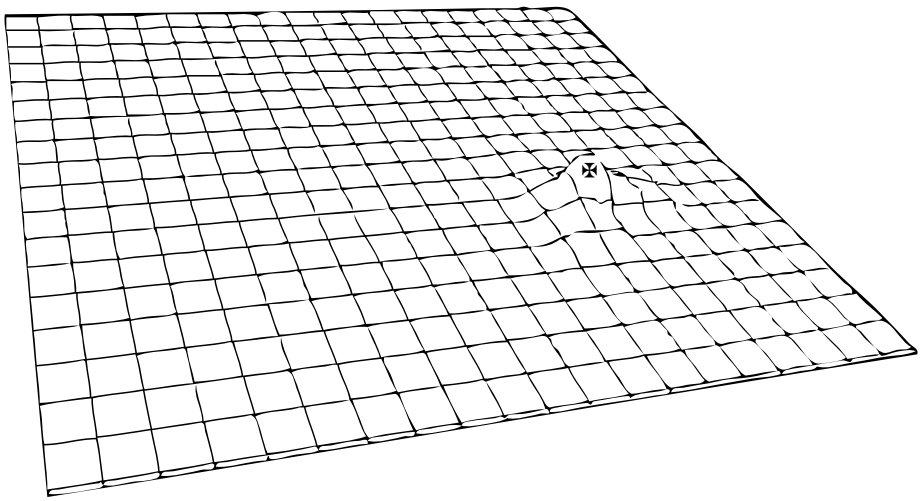

Figure 4.2: A fitness landscape with an "unrepresentative" population. While almost all the landscape slopes upward with a fairly drastic gradient, a population near the peak of the small local optimum, denoted by ✠, will be subject to a much smaller selective force, one which is not representative of the average selection pressure.

like when described in terms of their action on populations in this space of possible life histories. Say we have two different types of organisms. These types are very similar to one another, and inhabit the same environment. The vast majority of their possible life histories, therefore, will be identical. Type X, however, outperforms type Y in one particular sort of interaction (a certain kind of food gathering, say), that occurs in around 10% of possible life histories. On the characterization of GD^{SP}_{proc}, therefore, almost all the processes that these organisms are subjected to are indiscriminate – they almost all result in the same possible life history trajectory for both X- and Y-type organisms. Only the one food-gathering process discriminates between the two types. So much for the distinction between selection and drift on GD^{SP}_{proc}.

What about GD^{IC}_{proc}? Consider what happens if the population is evenly distributed

across the space of possible life histories.[29] For all of the *indiscriminate* processes, the X-type organisms experience 'favorable' outcomes at precisely the same rate as Y-type organisms – because the distribution is representative – and similarly with 'unfavorable' outcomes. The only difference in outcome that could result in this case, therefore, would be due to the *discriminate* processes. Selection alone is acting, and not drift. And selection and drift on GD^{IC}_{proc} thus align in precisely the same way that GD^{SP}_{proc} would predict.

Now turn to an unrepresentative population distribution (by definition, a finite population). In this case, we may, as a result merely of this random distribution, wind up with different outcomes for X- and Y-type organisms. Say, for example, that the X-type organisms are placed in a particularly 'unfavorable' region of outcomes with respect to indiscriminate sampling processes – a food-poor region of the environment, or an area with many predators. The Y-type organisms, on the other hand, might be placed in a 'favorable' region of outcomes – rich food resources and plentiful cover. In this case, we would have a population-level change (the Y organisms outperforming the X, *contra* the prediction of selection) that was a result of the organisms' differing exposure to indiscriminate sampling processes. Both selection and drift may be acting on the population,[30] but the distinction between selection and drift drawn by GD^{IC}_{proc}, again, exactly mirrors that expected on GD^{SP}_{proc}.

Assuming, then, that my ecological reading of Bouchard and Rosenberg's GD^{IC}_{proc} is correct, it turns out to be extensionally equivalent to GD^{SP}_{proc}. We can thus say, in turn, that

---

29. Of course, a perfectly representative sample only occurs in an infinite population (as spatial distribution is an element of possible lives, and the possible number of spatial distributions is infinite), another way of expressing the claim that only infinite populations can be entirely free of genetic drift.

30. Since it was left unspecified whether or not the various life histories engaged in the feeding behavior that manifests a selective pressure.

GD^{IC}_{proc} is, with the caveats mentioned above, roughly extensionally equivalent to GD^{TF}_{prod} and GD^{SE}_{prod}.

Given the most common process and product definitions of selection and drift, therefore, we can see that in many plausible cases they turn out to be very close to extensionally equivalent. Further, our attention was repeatedly drawn to an interesting set of biological cases on which the conceptual parsing of the various components of evolution is quite difficult: the handful of cases where these extensions are in fact disjoint. Such instances are ripe for philosophical work, which seems to be the most likely place to decide the issue of which of these definitions of drift and selection are truly intensionally appropriate.

The causalist/statisticalist debate has, however, by and large not recognized this troubling issue in the interpretation of evolutionary theory. On the one hand, authors like Millstein (2006a; 2008) have argued on the basis of the *extensions* of these selection and drift concepts that one or another position in the causalist/statisticalist debate must be correct. In many cases, we can see based on the analysis here that this will be a fruitless endeavor, as the two sets of definitions are often extensionally equivalent. On the other hand, writers such as Walsh (2010) have far too easily slipped back and forth between process and product definitions of selection and drift. While their extensions only rarely differ, they *do* differ, and their intensions imply drastically different conceptual structures for evolution. We cannot expect to resolve the causalist/statisticalist debate without some fundamental agreement on the correct definitions of selection and drift.

4.8. Conclusions

We thus stand at a troubling point for the causalist/statisticalist debate. First and foremost, it runs a substantial risk of begging the question, as the notion of a causal process must be settled by independent metaphysical evidence. Second, it has by and large focused on the issue of the composition question, which, as we have seen, is neither illuminated by nor illuminates our understanding of selection and drift. The prospects for resolving the composition question within the philosophy of biology community (and without the aid of work from metaphysics or general philosophy of science), I have argued, are bleak at best.

And this focus on metaphysical issues has come at the expense of an important, neglected interpretive issue in how we understand biological theory – our varying definitions of natural selection and genetic drift. These definitions are nearly extensionally equivalent, resulting in a substantial challenge to authors who wish to explore which of the dramatically different conceptual pictures implied by them is correct. Our attention was drawn to a set of interesting and difficult biological cases where these definitions do in fact have disjoint extensions. It is these cases where work in the philosophy of biology might hope to genuinely advance our understanding of selection and drift – and might help us resolve which of these differing conceptual structures for evolutionary theory is the correct one.

What does this mean for the causalist/statisticalist debate overall? It certainly doesn't mean that it has been fruitless. While it has required some non-trivial reconstruction to see that this is so, the debate has illuminated some difficult issues about the definitions of selection and drift, and brought some interesting contemporary research on causation to

bear on biological systems.

But it does mean that the persistent focus on the issues of the composition question and causal processes within the philosophy of biology has been misguided. All evidence shows that these questions won't (or even shouldn't) be answered by a clearer understanding of evolution, but by a clearer understanding of causation and inter-level causal relationships – an understanding that is not made any more accessible by considering the issue using selection and drift as a lens.

We therefore owe it to ourselves as philosophers *of biology* to focus on the genuine biological problem which this debate has foregrounded: what is the best way to interpret selection and drift? The complex interplay of the similar extensions and radically different intensions of these two sets of definitions has made the problem difficult both to see and to solve. It is here, however, that we might hope to genuinely increase our understanding of evolutionary theory. To that end, we will turn in the next chapter to the development of a new model of individual fitness – one which can provide us a new and especially fruitful framework for approaching the study of chance in evolutionary theory. Finally, in chapter 6, we will see how we can use this new model of fitness to gain insight into some of the most difficult questions surrounding the role of chance in natural selection and genetic drift.

CHAPTER 5

A NEW FOUNDATION FOR THE PROPENSITY INTERPRETATION OF FITNESS

In this chapter, I will present a novel model of the fitness of individual organisms, developed with Grant Ramsey. How can our conception of organismic fitness help us understand the role of chance in evolution? The full explication of this connection forms the topic of the next chapter. For now, recall from the first and fourth chapters that individual fitness – in particular, the propensity interpretation of fitness – is intimately related to many of the issues raised in the causalist/statisticalist debate. If the propensity interpretation (or something like it) is right, then the probabilities which it proposes require interpretation of their own, of the sort discussed in chapter 1. If the propensity interpretation fails to pan out, then we require yet another account of how the probabilities found in fitness are to be understood. But let's not move too quickly: what *is* the propensity interpretation?

The propensity interpretation of fitness (PIF) was introduced in the late 1970s with two articles, one by Brandon (1978) and the other by Finsen (née Mills) and Beatty (1979). These papers, among other things, provided a solution to what has since come to be known as the "tautology problem": If fitness is defined in terms of actual evolutionary outcomes, fitness cannot then causally explain these outcomes – the "survival of the fittest" reduces to "the survival of those that survive." They proposed that fitness (or, in Brandon's terminology,

'adaptedness') is a probabilistic propensity to produce offspring. Each organism, that is, is taken to have a disposition to produce certain numbers of offspring, with differing probabilities associated with each possible offspring number. If the fitness of an individual organism is traceable to this propensity, then it is clear that the tautology problem is solved: it is not vacuous to say that the fittest organisms will tend to survive more often than their less fit counterparts – just as a sturdier glass will tend to break less often than a fragile one.

5.1. Mathematical Models and Counterexamples

Fitness, however, fills more roles than merely the prevention of tautology. Most models of evolutionary change employ fitness as a scalar numerical value, comparable between organisms. In addition to providing a rank-ordering of the organisms in a population – which can justify claims like "a is fitter than b" – these fitness values are utilized by models such as those in population genetics to predict the future evolutionary trajectory of a given population.

The PIF, then, has traditionally been presented alongside a mathematical model which can serve to translate this probability distribution into a single, privileged measure on the distribution. The primary such mathematical model of the PIF, introduced by Brandon, has defined fitness as an organism's expected or average number of offspring, weighted by the associated probabilities. He described this formally as

$$A(O, E) = \sum P\left(Q_i^{OE}\right) \cdot Q_i^{OE}, \tag{5.1}$$

where each Q_i^{OE} is a possible number of offspring and $P(Q_i^{OE})$ is the probability of that number of offspring being realized. As mentioned above, Brandon used the term 'adaptedness' for fitness and $A(O, E)$ should thus be read as "the adaptedness of organism O in environment E." This is the standard model of the PIF, and it is shared across most major presentations of the PIF, including those of Mills and Beatty (1979), Beatty and Finsen (1989), Brandon (1990), and Sober (2001).

5.1.1. The "Generality Problem"

When the PIF was introduced, Brandon also saw that it could be useful for the solution of another problem in the philosophy of biology, one which we will call here the "generality problem." One area of work in the philosophy of biology has endeavored to theorize at a very abstract level about the process of evolution by natural selection – we might, consistent with similar terminology in the philosophy of physics, call this the study of the 'foundations' of evolutionary theory. Explanations of evolution at this level do not focus on particular episodes of natural selection, but rather on what it is that is common to *every* instance of natural selection, across every environment, system of heredity, unit of selection, and so forth where natural selection might be instantiated.

Returning to Brandon, in the same paper in which the PIF is introduced, he introduces the following as a "law of nature" (which he would later call the Principle of Natural Selection (Brandon, 1990, p. 11)):

> If *a* is better adapted than *b* in environment *E*, then (probably) *a* will have more (sufficiently similar) offspring than *b* in *E*. (Brandon, 1978, p. 187)

He argues that it is this law, or something like it, that is presupposed by all general explanations of natural selection, including the three traditional Lewontin conditions for evolution by natural selection: variation, heritability, and fitness differences (Lewontin, 1970).

It is thus crucial to the understanding of philosophical work such as this that we provide a definition of what Brandon calls "better adapted" – which many other authors, including Bouchard and Rosenberg (2004), call "fitter." That is, we need a notion of fitness that is capable of serving in the phrase "*a* is fitter than *b*" in *every* evolutionary system in which a statement like Brandon's Principle of Natural Selection holds true. This is what we will call the generality problem: the problem of producing a notion of fitness which may be correctly applied in explanations that encompass all possible evolutionary systems.

The generality problem does not only appear in the context of these unrestricted, general principles of natural selection, however. Indeed, were this its only use, it would be a matter of debate whether such a notion of fitness were really necessary at all, as some authors have constructed frameworks for natural selection on which a PNS like that deployed by Brandon is not necessary.[1] The reason why the generality problem is important is that it is central to the debate over the causal structure of natural selection, fitness, and genetic drift – the debate between the "statisticalists" and "causalists."

When, for example, Abrams proposes as an "elaboration of the PIF" that "if there are individual-level probabilities which are in some sense causal, natural selection and drift are causal in the same sense" (2007, p. 670), and Walsh argues in response that "fitness distribution explains but does not cause population change" (2010, p. 168), these authors are

1. The most prominent recent approach which discards the PNS is the spatial framework of Godfrey-Smith (2009).

not arguing over the causal forces present in some particular biological population. Rather, they are asking us to consider what the appropriate interpretation of evolutionary theory is, again, in every circumstance in which it applies. Does fitness reflect a causal property in all cases? Or is it merely a method of bookkeeping, a subjective tally of objective organismic lives and deaths? Again, it seems that a prerequisite for this debate is to find a concept of fitness and a mathematical model of that concept which apply in all cases – to find, that is, a solution to the generality problem.

And the generality problem is not exclusive to the philosophical domain. When Lewontin proposes his "three principles" for describing evolution, the second of these is that "[d]ifferent phenotypes have different rates of survival and reproduction in different environments (differential fitness)" (Lewontin, 1970, p. 1). This invocation of "differential fitness" is not relativized to any specific biological population, or even any specific model of fitness. Later work in mathematical biology has attempted to rigorize this notion of general fitness – Metz et al. argue that "the biomathematical literature of the last 10–20 years reflects the implicit acceptance of a common evolutionary framework, the core idea being that there exists a unique general fitness measure that concisely summarizes the overall time course of potential invasions by initially rare mutant phenotypes" (1992, p. 198). The model of fitness that we will propose here, as we will see later, extends this work of Metz et al. to the level of the fitness of individuals and draws out its philosophical implications, unifying the philosophical and biological approaches to the generality problem.

One obvious objection to this project is that it is not clear that we require a general concept of fitness in order to solve the generality problem. Several authors have endeavored to construct theories of natural selection that make no reference to fitness. Prominent among

these is the work of Millstein, who notes that "selection requires differences in abilities to survive and reproduce." She goes on to say that some would term these "'fitness differences' – I avoid the term to prevent entanglement in disputes over the definition of fitness" (2006a, p. 643). Elsewhere, she defines selection in general as a "discriminate sampling process whereby physical differences between organisms are causally relevant to differences in reproductive success" (Millstein, 2006a, p. 640), substituting *causally relevant physical* differences in place of *fitness* differences. If such a definition of selection will suffice, why do we need to produce a general model of fitness at all?

We claim that this picture of selection is perfectly reasonable, but only insofar as it smuggles in an implicit reference to a concept of fitness. It's clear that not just *any* physical difference, or even any physical difference that's causally connected to survival and reproduction, will suffice for being counted as taking part in selection. Each individual mammal, for example, has a unique pattern of hair follicles, and if hair is causally relevant to survival in some species, then *a fortiori* the pattern of individual hairs is as well. But it does not therefore follow that there is a selective difference between each pair of individuals that is *due to* their follicle pattern difference. We thus need some way to cash out selection in terms of *relevant* physical differences between organisms. The causal connection of a physical difference with survival and reproduction works as a basic criterion of relevance, but (as we have seen) seems to occasionally include irrelevant features. Further, the numerical quantification of selection – surely an important task for biology – requires a notion that can differentiate just *how relevant* a given physical difference is to selection. The elaboration of this measure (quantifying how relevant a given physical difference is to an individual's reproductive success) will, we argue, just consist in the elaboration of a model of individual fitness.

Finally, one more approach to the generality problem should be mentioned here. The "statisticalist interpretation," as the position of Walsh, Ariew, Matthen, Lewens, and others has come to be known, attempts to solve this problem by replacing the PIF with an interpretation of fitness as a set of facts about the statistical distribution of evolutionary outcomes. Walsh, for example, states that "what it is for a change in relative trait frequencies to constitute selection (or drift) is merely for it to be susceptible to a certain kind of statistical description" (2007, p. 282).[2] As we will now see, one of the primary reasons for adopting such a position is that there does not exist an account of the PIF that is free of counterexamples. By producing such an account here, we therefore substantially weaken a key motivation for adopting the statisticalist position along with their solution to the generality problem.

5.1.2. Counterexamples to the PIF

In addition to offering their own solution, the statisticalists (as well as some reluctant yet honest propensity theorists) have offered several counterexamples that purport to demonstrate that the PIF is *not* in fact suitable as an answer to the generality problem. It is notable, as an aside, that were the PIF not taken to offer us a solution to the generality problem, a counterexample to it that showed a particular population or set of populations to which it did not apply would be neither surprising nor germane. Biologists model the fitnesses of organisms in specific kinds of populations (with a particular genetic system, population size, etc.) in different ways throughout the biological literature. But given that the PIF does claim

2. See Matthen and Ariew (2002); Walsh et al. (2002); Krimbas (2004); Brunnander (2006); Walsh (2007); Ariew and Ernst (2009); Walsh (2010) for other works in the statisticalist program.

to offer a solution to the generality problem, these putative counterexamples are taken to constitute a significant problem for the PIF.

Before turning to these counterexamples, however, we would like to highlight one key distinction that is frequently overlooked. The distinction is that between the propensity interpretation of fitness as a philosophical position or an "interpretation" in a broad sense – the claim that the fitness of an organism is traceable to the probability distribution over its possible numbers of offspring produced – and the mathematical model which reduces that propensity to a scalar value (which we will henceforth call a "model of" the propensity), as expressed, for example, by Brandon's equation (5.1). Importantly, all the counterexamples raised against the PIF are counterexamples *to the mathematical model* – they demonstrate that the fitness of an organism can change without the fitness value determined by equation (5.1) changing. This point will be vital when we turn to describing the different ways in which we can respond to these counterexamples. The three most devastating such issues that have been raised are the *moments problem,* the *delayed selection problem,* and the *timing of offspring problem.* Let's consider them briefly in turn.

The moments problem. Brandon's equation computes the weighted average of *O*'s possible offspring. But individual fitness, it turns out, is sensitive not just to the average number of possible offspring, but also to higher moments of the possible offspring distribution, like variance, skew, and so forth. As Beatty and Finsen (1989) and others have pointed out, for example, if two organisms have the same average number of possible offspring, but one of them has a higher variance in possible offspring, the higher-variance organism will be less fit, ceteris paribus.

In light of this realization, Brandon modified his equation, introducing a correction factor intended to compensate for the effects of the higher moments (Brandon, 1990):

$$A^*(O,E) = \sum P\left(Q_i^{OE}\right) \cdot Q_i^{OE} - f\left(E, \sigma^2\right). \tag{5.2}$$

Brandon has the new element in the equation, $f(E, \sigma^2)$, "denote some function of the variance in offspring number for a given type, σ^2, and of the pattern of variation" (1990, p. 20).

This new formalization of the PIF treats it not as a single equation, but as an equation schema describing a "family" of models: The exact nature of the propensity in a given case can only be specified once the details of the population are determined. This solution to the problem of the higher moments has two shortcomings. First, it is not obvious that, if this is the best model of the PIF, the PIF still offers a solution to the generality problem. Brandon has traded a single equation for an infinite disjunction of equations, and it is not clear that such an infinite disjunction can, for example, tell us anything about (or feature in) the general causal structure of natural selection and genetic drift. Second, Brandon's proposed solution does not achieve the desired generality: the following counterexamples emerged, which show that there is more to fitness than expected number of offspring corrected for the effects of the higher moments.

The delayed selection problem. The Q_i^{OE} are possible numbers of *offspring*. But there are many biological situations in which offspring production is not a good correlate for fitness. Consider the classic case of the grandchildless mutation found in some species of the fruit fly *Drosophila* (Crow and Kimura, 1956). This mutation has no effect on the number of offspring

produced, but it causes all of an organism's offspring to be sterile – that is, it has a major effect on grandoffspring production. This is a counterexample to any measure of fitness founded solely on offspring production. We might attempt to solve this problem simply by modifying the PIF to be based on an expected number of grandoffspring instead of offspring. This would solve the problem of the grandchildless gene, but other species have mutations that end in sterility not one or two, but dozens of generations later (Ahmed and Hodgkin, 2000). This is not so readily fixable, since even if the PIF were based on the current maximal number of generations necessary for all extant species, future evolution may increase or decrease the required number of generations. Brandon, it seems, would have to add a second correction factor to the infinite disjunction of equations, making its suitability as a solution to the generality problem even more dubious.

Timing of reproduction. If two organisms have the same expected number of offspring, but one is disposed to reproduce earlier, then it will be fitter, ceteris paribus. This, too, is an effective counterexample to equation (5.2), since timing of reproduction is a component of fitness but can change independently of expected offspring number. Brandon, again, could add a third correction factor, one for the timing of offspring production. But this will only make it a poorer solution to the generality problem.

5.1.3. The Need for a New Model

There are thus ample counterexamples to the mathematical models of the PIF expressed by equations (5.1) and (5.2). Three different ways of responding to these counterexam-

ples are obvious. We could (i) abandon the PIF and adopt another solution to the generality problem, (ii) argue that the PIF can be defended without a corresponding mathematical model, or (iii) look for a more robust, counterexample-free model. If we choose the first option, the statisticalist interpretation is the most obvious replacement for the PIF in this context. It solves the generality problem by abstracting over all causal details of the biological case at hand, and describing only the evolutionary outcomes in terms of their statistical distribution.

Although the statistical response is not without merit, it has considerable shortcomings. Defining selection and drift merely in terms of their population-level outcomes runs the risk of obscuring the distinction between selection and drift, as well as making trouble for the traditional ways in which biologists understand these differing contributors to evolutionary change (Millstein, 2002; Brandon and Ramsey, 2007; Millstein et al., 2009). As mentioned in the last chapter, the relationship between the various accounts of causation on offer and evolutionary theory seems to produce prima facie evidence that drift and selection are in fact causal, *contra* the statisticalist position (Reisman and Forber, 2005; Millstein, 2006a; Forber and Reisman, 2007; Northcott, 2010). Perhaps most worryingly, there seem to be instances where the statisticalist interpretation simply gets the empirical data wrong – claiming that selection and drift cannot be distinguished or that selection is not acting on a population, when in fact the opposite is true (Brandon and Ramsey, 2007; Millstein, 2008). These claims are all the subject of intense argument (see, e.g., Lewens, 2010; Walsh, 2010), but it is worth our while to investigate ways in which a defender of the PIF could salvage the PIF's basic insights.

In order to understand and evaluate the other two possible responses to the counterex-

amples, we must begin by returning to the important distinction we made above between the PIF and the mathematical model of this propensity. The counterexamples just described serve as counterexamples specifically to Brandon's mathematical models of the PIF as expected number of offspring (possibly with a correction factor). In order, then, for them to serve as counterexamples to the PIF itself, two additional premises are required: (i) the formulation of a successful mathematical model of the PIF is *required* for the project to go through, and (ii) Brandon's original equation is either *the only* or *the best possible mathematical model* of the PIF. It is open to the defender of the PIF to reject either of these latent premises.

Perhaps the simplest way to reject the first premise would be to abandon the search for a mathematical model of the PIF entirely. We would then focus on the correctness of the PIF as a philosophical understanding of fitness, ignoring, or leaving to the biologists, the matter of determining the precise mathematical details of how this interpretation of fitness might be formalized. Alternatively, we could reject the second premise, and resume the search for a new mathematical foundation for the propensity interpretation of fitness. It is this latter approach, we believe, which stands the best chance of solving the generality problem while saving the possibility of a causal interpretation of fitness, natural selection, and genetic drift.

We must, however, defend this choice. Why is it that the correct response for a defender of the PIF is to continue the search for a mathematical model? Might we not best interpret many of the arguments of the statisticalists as proving to us that such a search is likely to be fruitless? We claim that it is not. First and foremost, if a counterexample-free model of the PIF can be developed, this implies the tenability of the PIF-as-interpretation. While the lack of a model does not imply the incoherence of the PIF, a counterexample-free model shows

that the PIF can be modeled in a clear and explicit manner. And further, the development of models of the propensity interpretation that are connected with biological practice can form a bridge between philosophical theory and scientific practice. If we can craft a model of the PIF that connects it with current work in biology, then the PIF – which otherwise may seem esoteric and non-biological – can be shown to be directly tied to contemporary evolutionary theory.

We argue that the lack of recognition of the complexity of the biological world has been one of the key mistakes made by defenders of the PIF. Beginning with expected number of offspring one generation into the future and then adding a host of correction factors has been, as we have seen, an ultimately fruitless path. We will instead discard this formulation and begin afresh, offering a model that does indeed avoid these counterexamples, and connects directly with the much more sophisticated mathematical models arising from cutting-edge mathematical biology. Our model can thus serve as the new foundation for the PIF. Further, and more importantly for our purposes here, we will see in the next chapter that it can also serve as an incredibly useful way to approach the two other most important instances of chance in evolutionary theory: chance in natural selection and genetic drift.

5.2. A New Formalization

Let us begin by carefully considering the structure of the PIF itself – that which any mathematical model of it is intending to capture. (Henceforth, when we refer to "the PIF," we intend to denote the PIF-as-philosophical-interpretation, not any particular model thereof.)

We will begin with the sketch offered by Ramsey's "block fitness" (2006), attempting to provide it with some mathematical rigor. This formal structure, at this point in the argument, is not intended to capture any specific mathematical model of the PIF – that is, we are not yet offering our own model of the PIF. Rather, the following is meant to give us a vocabulary in which any mathematical model of the PIF might be phrased (including, as we will see later, Brandon's original equation (5.1), among others). The formalism offered here will be quite general – containing far more terms than are required by most models of the PIF. This would allow one to formalize models of any complexity by simply ignoring terms for elements not required in a given model.

Consider an individual organism o, in a given environment E with a given genome G.[3] Over time, o may produce a population of offspring, o_i^1. And these, in turn, may produce offspring o_i^2. We say *may* produce because, for each organism (in each generation), there is a set of possible reproductive outcomes for that organism's life. These ways that organisms might live (or "possible lives"[4] for short) might include dying early of malnutrition, being preyed upon as a juvenile, or living to maturity and producing many offspring. These possible lives reflect not just the overall reproductive output of an organism (as used in Brandon's equation (5.1)), but many other features besides, such as the timing of offspring production. It is thus clear that mathemtical models of the PIF can, if we wish, rely on many more theoretical

3. We do not intend a particularly restrictive definition of either "environment" or "genome" here. Genome, for example, should be taken to include all heritable factors passed on from parent to offspring, including epigenetic and behavioral transmission. These two factors will be characterized more fully in the next chapter.

4. No modal ontology should be read into these "possible lives." In fact, we believe – consistent with the propensity interpretation – that these possible lives are best understood as the manifestation of a probabilistic dispositional property over time, a propensity.

resources than are utilized by standard formalizations like equations (5.1) and (5.2).

Let us return to the further development of our theoretical vocabulary. Combining these possible lives over generational time leads to a set of "possible daughter populations" of o.[5] There are many such possible sets, each containing all the descendants that o might produce in some set of circumstances. Call each of these possible daughter populations ω_i, and call the totality of such possible daughter populations (the set containing all of them) our sample space Ω.[6]

Now that we have our set of possible daughter populations, we need a way of tracking how probable these various possibilities are. To do this, we define a σ-algebra and probability measure, $\mathcal{F}$ and Pr, over Ω, in the traditional way. The details of this operation need not concern us here – this is the traditional mathematical formalization that lets us assign probabilities to the elements of our sample space. The probability $\Pr(\omega_i)$ assigned to each possible daughter population is simply the probability that ω_i will be the actual daughter population of o.[7]

This set of possible daughter populations along with their associated probabilities clearly gives a very thorough picture of the "success" of o. But we have yet to offer a mathematical model of the PIF – merely a very precise, perhaps unnecessarily large, mathematical

5. We should note here that these are not "populations" in any sense familiar from population genetics or ecology. A more appropriate term might be "lineage," but we wish to avoid confusion with several current theories of "lineage fitness" (see below).

6. Equivalently, one could define Ω as the space of functions $\omega : \mathbb{R} \to \mathbb{R}^n$, which take a time $t \in [0, \infty)$ to a 'state vector' consisting of some fine number of real-valued degrees of freedom. The 'possible daughter population' formulation, however, is more perspicuous, so we will use it for the remainder of the discussion.

7. Note that the sample space Ω is uncountable, necessitating that we integrate with respect to the probability measure Pr. See the last optional section 5.6 for information on the size of Ω, as well as a proof that Ω is well-behaved.

vocabulary in which many various mathematical models of the PIF might be phrased. These raw sets of possibilities and their associated probabilities cannot, for example, be directly compared to produce a fitness rank-ordering. The task of constructing our novel mathematical model of the PIF from these elements, then, is the aim of the remainder of this section. We are searching for the measure of individual fitness that can be extracted from this expansive set of theoretical resources which is maximally sensitive to the features of the raw sets of possibilities, and therefore as free as possible from the sorts of counterexamples articulated in section 5.1.2.

As we discussed in the introduction, the traditional way of turning the raw sets of possibilities and their associated probabilities into scalar values, expressed in the original formulation of the propensity interpretation (Brandon, 1978; Mills and Beatty, 1979), is the following. First, define a function $\phi(\omega, t)$ which computes the size of some particular possible daughter population ω at time t. Then, fixing T as the time one generation into the future, we define individual fitness as

$$F = \int_{\omega \in \Omega} \Pr(\omega) \cdot \phi(\omega, T)\, \mathrm{d}\omega. \tag{5.3}$$

This is simply the continuous analogue of the weighted average: We multiply the probability of each possible life by the size of the daughter population for that outcome one generation in the future, and then integrate to obtain the expected daughter-population size in the next generation. This is precisely equivalent to the result obtained by the traditional propensity interpretation of fitness, giving the same numerical results as equation (5.1) – that is, this is just Brandon's original formulation of the PIF expressed in our new theoretical vocabulary.

As mentioned, however, this formulation is subject to many problems. First and foremost, we need to remove the reliance of equation (5.3) on T, and hence on the daughter-population size only one generation into the future, in order to resolve the delayed selection problem. A first attempt at removing this time-dependence might lead us to compute something like the limit $F_\infty = \lim_{T\to\infty} F$, computing individual fitness in the "infinite long run." This would assuredly accomplish our goal of capturing all causal influences which might impact the future fate of an organism within a population.

There is, however, no guarantee that this infinite limit converges, is finite, connects with biological practice, or in any way tracks other measures of individual fitness. First and foremost, it seems quite likely that in all cases, $F_\infty = 0$. If every possible daughter population of o goes extinct in the infinite long run, then for every $\omega \in \Omega, \lim_{t\to\infty} \phi(\omega, t) = 0$, and thus $F_\infty = 0$. On the other hand, if some possible daughter populations do not go extinct, then it is possible that the population dynamics at infinite time are so chaotic that the limit in F_∞ does not converge to a stable value.

How can we, then, produce a long-run measure of individual fitness from Ω? As it turns out, the problem of determining this function is equivalent to a well-studied issue in demographics and mathematical biology – the asymptotic behavior of sequences of random, non-negative matrices. Results in this theoretical arena (following Tuljapurkar and Orzack, 1980; Tuljapurkar, 1989; Caswell, 1989; Tuljapurkar, 1990) can guarantee the existence of a limit much like F_∞. (The details of this derivation may be found in the last optional section 5.6.) These results allow us to define individual fitness instead as

$$F = \exp\left(\lim_{t\to\infty} \frac{1}{t} \int_{\omega\in\Omega} \Pr(\omega) \cdot \ln(\phi(\omega, t))\, \mathrm{d}\omega\right). \tag{5.4}$$

Before we consider the peculiarities of this new model (such as its infinite limit, logarithms/exponents, and factor of $1/t$), let's examine how it solves the problems of extinction and chaotic future dynamics. We must enforce three assumptions on the possible daughter populations and their associated probabilities in order to guarantee that the limit in equation (5.4) converges (Tuljapurkar, 1990, p. 25): (i) demographic weak ergodicity, (ii) that a random and stationary process generates the Pr-function, and (iii) that the logarithmic moment of the growth rate is bounded. For the sake of brevity, we will pass over the detailed mathematical characterization of these assumptions here.[8] In short, demographic weak ergodicity assures that there exists some non-zero probability of the population's survival at all times t, getting us around the extinction problem mentioned above. While this assumption is biologically unrealistic, it is quite common in mathematical demography, and can be dealt with either by describing extinction as a threshold (that is, a population "goes extinct" when its size falls to less than some small value n), or by introducing some random environmental sampling variation into the model (see Keiding, 1975). A guarantee that the Pr-function is generated by a stationary random process assures that chaotic population dynamics are not permitted. (The boundedness of the logarithmic moment of the growth rate is of purely technical interest.) In general, however, we defer here to Cohen, who states that "under reasonable conditions, which are likely to be satisfied in demographic applications, the stochastic process and the matrices in Leslie form [the population analogue of our individual daughter populations and probabilities] are such that the limits in [equation (5.4)] exist" (Cohen, 1979, p. 164).

8. For demographic weak ergodicity, the reader may consult Seneta (1981, pp. 80–91), Tuljapurkar (1990, p. 17), Tuljapurkar and Orzack (1980, pp. 319–20), and Cohen and Heyde (1985, p. 123). For a discussion of the stationarity and ergodicity requirements on the random process, consult Tuljapurkar and Orzack (1980).

These three conditions do, however, have biological significance for our model. Most importantly, they imply that the selective pressures at work are *density-independent,* and that the population dynamics are *non-chaotic.* These are the most substantial limitations of our model, and because of this, equation (5.4) clearly cannot provide the PIF with a complete answer to the generality problem. However, the necessary mathematical work to generalize our derivation to cases of non-static environments and density dependence, as well as chaotic population dynamics, has been published within a research program known as adaptive dynamics, to which we will return shortly. We omit it in order to simplify our derivation, as it relies on a hefty theoretical apparatus which considerations of length and accessibility prevent us from presenting here.[9] Our equation (5.4) is the density-independent, non-chaotic limit of this more sophisticated work, and thus, given these restrictions, is equivalent to this more general model. Further, and perhaps most importantly, all the counterexamples that have been offered in the philosophical literature to the traditional mathematical model of the PIF (equation (5.3)) are resolved by equation (5.4).

Let us take stock. We have explicated the PIF itself via a very extensive picture of the success of an organism *o*, by considering all the possible daughter populations to which it might give rise, and we have then modeled this propensity by defining a function, equation (5.4), that successfully encapsulates the behavior of this sample space in the infinite limit.

But what does this function actually represent? Does it correspond to any other

9. The interested reader is referred to Rand et al. (1994); Grant (1997); Caswell et al. (2004); Caswell and Takada (2004); Caswell (2009). Many of the most important conclusions for individual fitness follow directly from Theorem 1 of Rand et al. (1994, p. 271). See Benton and Grant (2000) for a comparison of various approaches to fitness in population genetics and adaptive dynamics.

known models of individual fitness, and if so, in what circumstances? And what should we make of its peculiar mathematical features?

5.2.1. The New Model and Biological Theory

Although equation (5.4) was derived via a reflection on the philosophical thesis of the PIF, we will now show that the same mathematical formula has been independently derived within the field of mathematical biology. In fact, in adaptive dynamics, a variation of this equation is argued to be the optimal predictor of the fates of populations. We therefore suggest (and will attempt to demonstrate in the remainder of this section) that equation (5.4) is deeply connected to biological theory and practice, and that with certain simplifying assumptions, one can derive from equaton (5.4) many standard fitness measures. Thus, despite the fact that it may seem less connected to biological theory and practice than the standard formulation of the PIF articulated by Brandon (1978), equation (5.4) is much more closely connected with contemporary biological theory and practice, and is sensitive to the advances in mathematical work that have occurred in the decades since Brandon first published his attempt at providing a model of the PIF.

To explore the relationship between equation (5.4) and other (recently developed as well as traditional) models of fitness, let's begin with two different models that are precisely equal to equation (5.4). First, equation (5.4) is equal to a model of fitness known in mathematical biology as Tuljapurkar's a, replacing the population-level quantities in Tuljapurkar's original function with equation (5.4)'s individual-level quantities. Indeed, many of

his results have been crucial to the derivation of our model here (Tuljapurkar, 1989, 1990).[10] Tuljapurkar's intent in creating this model of fitness was to produce an expanded notion of lifetime reproductive success (LRS) which can apply to the general case of environments that vary stochastically over time. Notably, Tuljapurkar's a has proven to be a successful measure of fitness in the wild – Cohen and Heyde (1985), for example, use it to study the growth of breeding populations of striped bass in the Potomac River. They determined that it provides a highly accurate determination of average growth rate (while noisy data make estimation of other parameters like variance more difficult).

This model has also appeared as one of the fundamental quantities in the research program known as adaptive dynamics, which we briefly mentioned above. Adaptive dynamics endeavors to produce a highly general notion of fitness applicable in many ecological contexts, based on two considerations: (i) the modeling of populations in variable environments, and (ii) the determination of fitness on the basis of invasion, consistent with much work throughout evolutionary ecology. In one of the seminal articles of adaptive dynamics, titled "How should we define 'fitness' for general ecological scenarios?", Metz et al. note that "the long-run growth rate" of a population in their framework "can be defined as the limiting value, as (time) T approaches infinity, of the quantity $T^{-1}\{\ln|\mathbf{N}(T)| - \ln|\mathbf{N}(0)|\}$" (1992, p. 198). This quantity, again, is precisely equivalent to our model.[11] As we discussed earlier, work in this field has also produced substantially more sophisticated models which can be used to take account of both chaotic population dynamics and density-dependent selection. Equation (5.4)

10. Technically, $\ln(F) = a$; see section 5.6 for details.

11. We omit the derivation here, as it is almost precisely identical – though phrased in a different theoretical vocabulary – to the derivation of Tuljapurkar's a presented in section 5.6.

is the density-independent, non-chaotic limit of these more complex models.[12]

To connect equation (5.4) to futher biological models, we must introduce some simplifying assumptions. First, if we assume that multigenerational effects are absent from the population, we may approximate equation (5.4) by taking its value at time T, one generation into the future. Assuming that $T \approx 1$,[13] we then remove the limit and factor of $1/t$ from equation (5.4), resulting in

$$F \approx \exp\left(\int_{\omega\in\Omega} \Pr(\omega) \cdot \ln(\phi(\omega, T))\, \mathrm{d}\omega\right). \tag{5.5}$$

To further simplify equation (5.5), we should note that it has the form of a geometric mean. That is, the geometric mean of a function $f(x)$ applied to some sample space X is defined as

$$\mathrm{GM}(f(x), X) = \exp\left(\int_{x\in X} \Pr(x) \cdot \ln(f(x))\, \mathrm{d}x\right),$$

and equation (5.5) thus states that $F \approx \mathrm{GM}(\phi(\omega, T), \Omega)$.

Notably, the use of a geometric mean in fitness models in biology is by no means a new concept (see Lewontin and Cohen, 1969; Boyce and Perrins, 1987; Sober, 2001; Simons, 2002; Lee and Doughty, 2003). In particular, however, given the prominence of this work in the philosophical community, geometric mean fitness is discussed extensively by Gillespie (1977), in his summary of his earlier work. Specific translation of his work into our arena is

12. Therefore, with arguments very similar to the ones provided here, these more complex models also reduce to our model, Brandon's model of the PIF, and the other standard biological fitness measures we discuss in the following.

13. We also must assume that the function $\int_{\omega\in\Omega} \Pr(\omega) \cdot \ln(\phi(\omega, t))\, \mathrm{d}\omega$ is continuous at T, so that $\lim_{t\to T}\left[\int_{\omega\in\Omega} \Pr(\omega) \cdot \ln(\phi(\omega, t))\, \mathrm{d}\omega\right] = \int_{\omega\in\Omega} \Pr(\omega) \cdot \ln(\phi(\omega, T))\, \mathrm{d}\omega$. This should be true in effectively all circumstances, particularly since we have already guaranteed non-chaotic population dynamics.

difficult, as he is discussing the fitness of traits, rather than the fitness of individual organisms (see below for more discussion of this point). He notes that in general, however, when we have stochasticity resulting from "temporal fluctuations in the environment, for example, the best measure of fitness turns out to be the geometric mean of the offspring number, averaged over time" (Gillespie, 1977, p. 1011).[14] He then establishes several results using the series expansion of the geometric mean, to which we will now turn.

We know from standard results in statistics that the geometric mean can be expressed as

$$\ln(\mathrm{GM}(X)) = \ln\left(\overline{X}\right) - \frac{1}{2\overline{X}^2} M_2\left(\overline{X}\right) + \frac{1}{3\overline{X}^3} M_3\left(\overline{X}\right) - \cdots,$$

where $\overline{X}$ is the arithmetic mean of the distribution X, and $M_i\left(\overline{X}\right)$ is the i-th central moment of the distribution X (i.e., its variance, skewness, kurtosis, and so on, for $i = 2, 3, 4, \ldots$) (Jean and Helms, 1983). We can thus see that when the higher moments of a distribution are negligible (i.e., when effects of variance, skew, etc. can be neglected), $\mathrm{GM}(X) \approx \overline{X}$. In these cases, we may thus consider individual fitness as though it were defined in terms of an arithmetic mean. We can therefore approximate equation (5.5) by an arithmetic mean, which gives us

$$F \approx \int_{\omega \in \Omega} \Pr(\omega) \cdot \phi(\omega, T)\, \mathrm{d}\omega. \tag{5.6}$$

Equation (5.6), then, is just equation (5.3): we have reduced our new formulation to that of the traditional model of the PIF with the aid of two relatively plausible simplifying assumptions. (Of course, in the presence of the conditions described in the various counterexamples to

14. Gillespie probably has something like predictive accuracy in mind when he invokes the "best measure of fitness," but these details need not concern us here.

equation (5.3) described above, these simplifying assumptions do not hold.)

Several short-term measures of fitness are particularly common in the biological literature, as they are easy to estimate and can be derived from readily available empirical data. The first is the net reproductive rate (or ratio), R_0, which is a common measure of the single-generation reproductive output of a population (Murray and Gårding, 1984; Murray, 1990). Lifetime reproductive success (LRS) is the individual analogue of this population measure. It is well-known from the literature on the original model of the PIF that equation (5.6) is equal to the LRS, and thus equation (5.4) reduces to the LRS. Another quite common biological fitness measure is the Malthusian parameter, the growth rate of a population given an exponential growth model. Given the LRS, we may derive the (individual analogue of the) Malthusian parameter as well: $r = T^{-1}\ln(F)$ (Charlesworth, 1970; Denniston, 1978; Charlesworth, 1980; Murray, 1990).

To recap, then, our new model of the PIF as described by equation (5.4) is precisely equivalent to several advanced models of fitness, and with two plausible simplifying assumptions – that (i) multigenerational effects are absent, and (ii) effects of the higher moments are negligible – can be reduced to the original model of the PIF as well as to the most commonly used biological fitness measures (R_0, LRS, and the Malthusian parameter). Our model thus dovetails very tightly with contemporary work on the measurement of fitness in biology.

5.3. Possible Objections to *F*

Now that we have seen that equation (5.4) is closely connected with biological theory and practice, let's pause to consider some of the ways philosophers might object to the model. We will, in the following section, then turn to the question of whether it successfully dodges the counterexamples discussed in section 5.1.2. Our formulation is clearly a *multi-generational* or *long-run* measure of fitness. And this long-run measure of fitness makes extensive reference to not merely the organism itself, but to the organism and all its possible daughter populations – that is, our definition of fitness depends crucially on the organism's *lineage.* Both long-run and lineage fitness models have been challenged in the past, and in this section we will show that our model is not undermined by these challenges.

There is nothing novel about the concept of long-run fitness. In the philosophical literature, the varying time-scales required in definitions of fitness have been discussed extensively by Ramsey (2006) and Abrams (2009c). In the biological literature, two prior models of very-long-term fitness have been proposed: Thoday's definition of fitness as "the probability that … a unit of evolution will survive for a given long period of time, such as 10^8 years" (1953, p. 98), and Cooper's definition of fitness as "expected time to extinction" (1984).

Our new model is clearly not equivalent to either of these models and, we will argue, neither of them can serve as a satisfactory answer to the generality problem in the context of the PIF. To see why, let us first consider Thoday (1953). Thoday's model simply sets a large future time (10^8 years) and calculates the probability of the survival of a given organism's lineage to that point. The primary objection to a model like Thoday's is simply that this

time-frame is entirely arbitrary. Why should daughter-population events at time $10^8 - 1$ years be included in an organism's fitness, but events at time $10^8 + 1$ years be ignored? It seems that no philosophically defensible answer to this question can be found. It is for this reason that we have attempted to include all possible future causal influences in the scope of our model. Further, should evolutionary dynamics be chaotic (a possibility introduced in the last section), there is no guarantee that the probability to which Thoday refers will even have a definite value.

Cooper's (1984) definition of fitness as expected time to extinction suffers from a different, equally fatal flaw. The precise numerical expected time to extinction of a given organism seems to have only a very loose relationship to other more commonly used models of fitness. Expected time to extinction may well be correlated with individual fitness, and one could obviously derive expected time to extinction from the theoretical resources offered by our set of possible daughter populations. But for expected time to extinction to serve as a fundamental model of fitness, this derivation would need to be reversible – one would need to be able to derive other standard models of fitness (such as those mentioned at the end of the last section) *from* expected time to extinction, which is impossible. ETE thus solves the generality problem at the price of alienating the PIF from all other work on the concept of fitness in both philosophy and biology. Given that our model does not suffer from this flaw, we believe this price is too high.

Our model, despite being long-term, is multigenerational in a different manner than either of the models of Thoday and Cooper. Its time-scale is infinite, encompassing all possible future causal influences on organisms. At the same time, it still reduces to other common short-term models of fitness. It thus effectively avoids the problems that Thoday

and Cooper fall prey to.

Finally, we will consider several other objections that have been raised against both long-run and individual models of fitness.

Objection 1: Natural selection is short-term. Brandon argues, referring to Thoday and Cooper's long-term notions of fitness, that "they fail to explain how the process of natural selection can be sensitive to differences in long-term probabilities of surviving offspring." Selection, he notes, is "the differential reproduction of phenotypes that is due to the differential adaptedness of those phenotypes to a common environment. Evolution results from this process if the phenotypic differences are heritable. How could this process be sensitive to long-term probability (i.e., over many generations) of surviving offspring?" (Brandon, 1990, p. 25). Long-term notions of fitness are, that is, irrelevant to the process of natural selection – since selection is short-term, individual fitness must be short-term as well.

In response to this charge, Sober notes, quite correctly, that "the fact that selection occurs one generation at a time does not mean that it is wrong to define a quantity that describes a trait's long-term expected fate" (Sober, 2001, p. 313).[15] To consider a similar example, just because the half-life of a particular sample of uranium may be defined in terms of its expected behavior thousands of years into the future, we need not say that this somehow means that radioactive decay does not "act on" the uranium "as it currently is." Radioactive decay has no more foresight than natural selection – and both may make reference to future

15. Similarly, Abrams (2009b, p. 751) argues that "since probabilities of long-term effects can be derived from probabilities of short-term effects, the former are simply mathematical properties of causes acting in the short term."

(or even *possible* future[16]) events.

One might reply that the behavior of radioactive decay is somehow more "regular" or "predictable" than the behavior of biological organisms, and that therefore this analogy fails.[17] But this isn't the relevant feature of the analogy (if, in fact, the behavior of uranium *is* any more "regular" than that of organisms, which is itself not obvious). The objection as argued by Brandon seems to claim that it is a *category mistake* to include future events in the definition of fitness, as selection acts only in the present. A half-life certainly does make reference to future events for its definition, and radioactive decay certainly acts only in the present. This facet of the analogy thus clearly holds. This objection does not, therefore, challenge the coherence of our model.

More importantly, however, this objection misunderstands the purpose of our model. We are attempting to produce a model of the PIF that forms a successful solution to the generality problem. Solutions to the generality problem require casting an expansive net, including the extension of our model of fitness to the long term. Other models of fitness will of course be useful in other pragmatic contexts – the generality problem is often (or even usually) far from our minds when we work on evolutionary systems. And our demonstration in the last section that our model reduces, in the short term (and given several other limiting assumptions), to several well-known biological models of fitness can give us hope that this new model of the PIF can both provide a solution to the generality problem and ground a theoretically unified picture of fitness.

---

16. If there had been only one molecule of uranium in the entire universe, that molecule would still have the same half-life as the uranium which we know, even though this half-life could be defined in terms of other possible (not actually existing) atoms of uranium.

17. We thank an anonymous reviewer for pointing out this possible reply.

Objection 2: Descendants are only minimally related to ancestors. Another objection to long-term fitness is offered by Ramsey. He writes that since, over time (for sexually reproducing species), the organisms constituting the daughter population of some organism are related less and less to that organism (a factor of 1/2 in the first generation, 1/4 in the next, and so on), the long-term descendants of two organisms in a population may well be roughly identically related to each of those ancestors. It is therefore a mistake to think that the fitness of those two organisms somehow depends on the characteristics of those descendants.

Two responses can be made to this objection. First, Ramsey is at this point concerned with the *operationalization* of his notion of fitness. It is true that as we move from the short term to the long term, we begin to consider organisms that may be quite different than the original organism we intend to study. And it also may be true that the precise fitness difference between two organisms hinges upon an organism that we cannot measure, for it lies many generations into the future. But we have already seen that in many relevant cases, the long-term notion of fitness we describe here reduces to easily measurable short-term fitness measures. Of course, these require simplifying assumptions, and will therefore occasionally produce the wrong answers. In this case as in many others, we must assess the common tradeoff between complex, accurate biological models that are difficult to measure and simpler, less accurate models that are more empirically tractable.

Second, another way of reading this objection would confuse an accurate point about trait fitness with an inaccurate claim about individual fitness. When we are considering the fitness of traits, Ariew and Lewontin (among others) remind us that "the rate of reproduction *by* a genetic type is not the same as the rate of reproduction *of* a genetic type" (2004, p. 352). In measuring the long-term fitness of a trait, then, the genetic relatedness of those offspring

with their ancestors is, in fact, a crucial point. If we lose track of this relatedness, we run the risk of inflating the fitness of ancestral types (a mistake which Ariew and Lewontin ascribe to Fisher).

In the case of individual fitness, on the other hand, such a worry does not arise. Individual fitness, as we have explicated the picture provided by the propensity interpretation, is concerned with the size of the daughter population of a given individual. Membership in a given daughter population does not come in degrees – one is either a descendant of a given organism or one is not. If a given future organism is a descendant, then it is counted by our model.[18] While the "coefficient of relatedness" of some arbitrary, distant descendant to two members of the ancestral population may be both similar and very small, this gives us no reason to think that the fitnesses of these two ancestors will be equivalently similar. Effects of path dependence on the daughter populations in the intervening generations, for example, may well be quite significant.

Objection 3: Evolutionary time-scale is pragmatically determined. Sober claims that we should wish to remain agnostic over whether we should choose long-term or short-term measures of fitness. "Long-term fitness," he writes, "is a coherent concept that may be useful in the context of certain problems; however, its coherence and desirability do not undermine the concept of short-term fitness" (Sober, 2001, p. 313). In general, we agree. As noted above, we have introduced our model with the intent of cementing the place of the PIF (and along with it, the causal interpretation of fitness, selection, and drift) as one possible solution to the

18. Of course, there are manifold issues concerning individuality in evolutionary biology, and it thus may not always be clear when an "organism" is to count as a "descendant" (Godfrey-Smith, 2009). Our model, however, suffers from this problem no more acutely than any other of the many theories in evolutionary biology that require the counting of individuals.

generality problem. But short-term predictions, as we have seen above in our discussion of the relationship of this model to other biological definitions of fitness, can readily be derived from this model under various plausible sets of limiting assumptions. We thus are fully entitled to utilize these short-term predictions when they are needed. And as we have already noted, the solution to the generality problem – an active issue in the philosophy of biology – requires the model to take the long-term view.

Objection 4: Long-term fitness is lineage fitness. We claim to be setting out a model of *individual* or *organismic* fitness. Our derivation, however, seems to traffic only in *lineages* or *daughter populations.* Isn't this model, then, in fact a model not of individual but of lineage fitness?

We should begin by noting that we do not intend to disparage the usefulness of lineage fitness. Jost, who defines lineage fitness as "the number of descendants weighted by their degree of relatedness with the ancestor of the lineage" (2003, p. 331), has made a persuasive argument for lineage fitness and used it effectively to analyze the emergence of altruism (though see Okasha (2006) for a criticism of this concept).

It is true, however, that the objection that equation (5.4) represents lineage fitness instead of organismic fitness would defeat our claim to be offering an improved model of the propensity interpretation of fitness. But reading our model as a model of lineage fitness misunderstands our work. Clearly, the possible future lineages of an organism are one of the *determinants* of its fitness. But this is not equivalent to the claim that lineages are the *bearers* of fitness. Our model, that is, is a model *of* individual fitness that *depends on* lineage characteristics.

Objection 5: The theory of evolution by natural selection fundamentally concerns trait fitness, not individual fitness. Our model is, as we have stated, a model of the fitness of individual organisms. Organismic fitness plays important roles in parts of ecology and evolutionary biology, and is the concept of fitness underlying the PIF. On the other hand, much of the active work of biologists, as well as many of the arguments of philosophers, relies instead upon trait fitness. Further, many counterexamples have been raised against propensity-based models of trait fitness, including critiques by Ariew and Lewontin (2004)[19] and Krimbas (2004).[20] One might be concerned, then, both that our model fails to respond adequately to these other counterexamples present in the fitness literature, and that our model fails to offer an account of the fitness of traits.

Trait fitness, however, is commonly understood in two different ways. First, we have trait fitness as the average fitness of all individual organisms that bear a given trait (Sober, 2001; Walsh et al., 2002; Walsh, 2003). Second, we have trait fitness as a prediction of future trait prevalence – the quantity that lets us predict the frequency of a trait in the next generation given its current frequency (Matthen and Ariew, 2002; Walsh, 2003). If the first of these two definitions is adopted, then trait fitness is straightforwardly parasitic on individual fitness, and a model of individual fitness must be provided in order to make sense of the fitness of traits. If the second defintion is adopted, however, then we are dealing with quite a different quantity than the one modeled here. Trait fitness in this second sense relies on individual

19. Ariew and Lewontin intend their critique to be targeted at "a scalar quantity … which then predicts changes in the representation of types" (2004, p. 350).

20. Though he speaks occasionally of the fitness of individuals, Krimbas's main concern is "the absolute or Darwinian fitness of a certain genetic constitution of individuals of the same species in a population" (2004, p. 190), clearly a notion of trait fitness.

fitness as well, but also includes factors like heritability. Thus, under either of the standard ways of understanding trait fitness, individual fitness is in some sense foundational: trait fitness values are either directly derived from individual fitness values, or individual fitness values are a component of trait fitness. Because of this, we are justified in simply providing a model of individual fitness as the foundational concept in the PIF.

5.4. Response to Counterexamples

At the beginning of the chapter, we examined three counterexamples to the traditional model of the PIF as expected number of offspring. In this section, we will show that our new model of the PIF as represented by equation (5.4) does not fall prey to these counterexamples.

5.4.1. Timing of Reproduction

If two organisms, O_1 and O_2, have the same expected number of offspring, but O_1 is disposed to reproduce earlier than O_2, then it will be fitter, ceteris paribus. The reason for this is that O_1 (and presumably its offspring, if the trait responsible for its different reproductive behavior is heritable) will have a shorter generation time. If both O_1 and O_2 have two offspring, but O_1 has them twice as early in life, then O_1 will have twice as many descendants as O_2 when O_2 finishes reproducing. If this is not a dramatic fitness difference, then nothing is. Yet, if we measure the fitness of O_1 and O_2 with equation (5.2), we arrive at the same value. Thus, even the revised version of the original PIF does not take this fitness difference into account.

The reason for this is that the equation is a function of possible lifetime reproductive success. That is, it merely tallies the reproductive event outcomes of entire life histories without being sensitive to other important properties of these life histories, such as the temporal arrangement of reproductive events.

Equation (5.4), however, solves this problem, by using time instead of number of generations to determine fitness. It is important, however, that the right time scale is used. If O_1 reproduces at the age of five and O_2 at ten, then defining fitness as the expected number of offspring at one year would mistakenly compute both of their fitness values to be zero. The time scale has to be at least as large as the longest generation time for any individual whose fitness is being compared, but it will regularly need to go far beyond a single generation. The fact that equation (5.4) is based on an infinite limit clearly provides us with a sufficiently large time scale to account for all possible variability in reproductive timing.

5.4.2. Delayed Selection

Not all offspring are created equal. Some will share the fertility of their parents, while others will be infertile. The grandchildless (and other, longer-term) mutations mentioned at the beginning of the chapter, which have been found in *Drosophila* and other taxa, show that a single-generation time scale is not sufficient to capture fitness. One might, however, think that the grandchildless mutation does not necessitate a super-generational time scale. Instead, one could attempt to preserve single-generational fitness by simply discounting offspring based on their fertility – infertile offspring would not be counted as offspring. This proposal,

we argue, merely smuggles in a multigenerational time scale. Consider, for example, why it is that we must not count *O*'s infertile offspring. The correct answer to this, it seems, is that they will not give *O* any *grand*-offspring. Thus, the choice to discount the fitness of infertile offspring *relies* on fitness being multigenerational. Furthermore, both fertility and viability affect fitness, and both come in degrees. Thus, not only would infertile individuals have to be excluded, each offspring would have to be assigned a weight, depending on its relative fertility and viability. But what is the assignment of such a weight, but an assignment of fitness values to the offspring? There are thus two reasons why the offspring-weighting proposal fails: First, it defines fitness in terms of fitness. Second, it defines the fitness of the offspring in terms of the grand-offspring, and so on, making it in fact multigenerational.

Fortunately, equation (5.4) is sensitive to these and similar mutations. Determining the fitness of individuals via the size of their daughter populations "at infinity" allows us to incorporate the fact that not all descendants are created equal. Our fitness model takes account of the unlimited variability of fertility and viability, since the long-term descendant pool is sensitive to these factors.

5.4.3. The Moments Problem

As we saw in the introduction, fitness is a function of not only the first moment (i.e., arithmetic mean) of the distribution of possible offspring, but the higher moments as well (variance, skew, etc.). This problem was not recognized in the original formulation of the propensity interpretation of fitness, but was later recognized and solutions were offered.

The solution offered by Brandon (1990), as we saw (equation (5.2)), was to add a correction factor that would discount the fitness of individuals based on the structure of the distribution of possible offspring. But this solution was unsatisfying, since it merely offered "some" unspecified function of this distribution.

Our model, unlike previous models, solves the moments problem by proposing a concrete, specific function, equation (5.4). This function solves the moments problem by virtue of its long time scale. To see why this is the case, consider a simple example of two organisms, O_1 and O_2, which reproduce asexually and clonally, and have discrete generations. Each generation, the first organism and its descendants have $\Pr(\phi(\omega, T) = 1) = 1$ (that is, the probability of having 1 offspring each generation is 1), and the second organism and its descendants have $\Pr(\phi(\omega, T) = 0) = 0.5$ and $\Pr(\phi(\omega, T) = 2) = 0.5$ (equal probability of having either 0 or 2 offspring). Both O_1 and O_2 have the same expected number of offspring after one generation (namely, one), but O_1 is fitter. To see why, consider what you would expect the daughter population size of O_1 to be as $t \to \infty$. Since O_1 has a probability of 1 of having 1 offspring, the daughter population size will be one. But O_2, on the other hand, will not fare so well. The probability of the population going extinct will approach 1 as $t \to \infty$, since extinction is all but guaranteed in the long run: in each passing generation, each of O_2's descendants runs a 50% risk of lineage extinction. Taking the long view, as does equation (5.4), correctly evinces the superior fitness of O_1. We lack the space to provide more examples of higher moments here (such as those presented in Beatty and Finsen (1989) and Abrams (2009b)), but they reveal themselves in the long term as well.

5.5. Conclusion

There has long been a perception that the PIF is in dire straits. As we have seen, it has from its earliest days purported to offer us a solution to the generality problem, yet the mathematical model of it commonly offered is subject to a host of counterexamples. Many or even most real-world biological populations are subject to one of the difficulties described, making the PIF appear to be an inadequate answer to the generality problem – and thus not able to support the causal interpretation of fitness, selection, and drift against its statisticalist opposition. Propensity theorists have tried to save the PIF by amending its standard mathematical model with correction factors. These amended versions have saved the PIF from counterexamples, however, only by rendering it a poor solution to the generality problem. Furthermore, as the statisticalists have eagerly pointed out, these proposed correction factors do not seem to dodge all of the counterexamples. That is, there still has not been a model of the PIF offered in the literature robust enough to withstand the full array of philosophical difficulties with prior models.

As was argued at the beginning of the chapter, however, all these counterexamples are problems not with the propensity interpretation of fitness itself, but with the various mathematical models of it which have been proposed. The opponents of the PIF have done nothing by way of arguing that the extant mathematical models of the PIF are either the only or the best ways to formalize this interpretation of fitness. We considered three possible ways for a proponent of the PIF to respond. One could (i) jump ship and embrace the statisticalist interpretation of fitness, selection, and drift (or something like it), (ii) abandon the search for

a mathematical model of the PIF entirely, or (iii) craft a model of the PIF that can offer it a new mathematical foundation. With equation (5.4), we have accomplished the third option. We have constructed a model that retains the PIF's purported ability to solve the generality problem without being subject to the counterexamples that have been proposed against it.

We have not, notably, argued directly against the statisticalist position itself. As noted above, the statisticalist position is not without difficulty, and has been criticized on many fronts (Millstein, 2002; Brandon and Ramsey, 2007; Millstein et al., 2009; Northcott, 2010; Otsuka et al., 2011). One of the seemingly compelling arguments in its favor, however, is that it is not subject to the sorts of counterexamples that undermine the PIF. Since we have provided a model of the PIF that avoids these counterexamples (and, hence, have demonstrated that there is no reason to think that the PIF cannot solve the generality problem), one of the main reasons for adopting the statisticalist position has been called into question. The statisticalists, we suggest, would need to attempt to form counterexamples to this new model, not the older formulation provided by Brandon (1978).

While this by no means resolves the debates over the role of fitness in evolutionary theory – one could, even taking our model into account, still reject the notion of fitness as a causal property – we hope that the presentation of a mathematical model that resists the now-common counterexamples to the PIF will allow for a more sophisticated debate between causalists and statisticalists.

This account of individual fitness provides us with a dramatically different way of understanding just why it is that fitness is inherently "chancy." Fitness, as we have seen, relies essentially on the set of possible lives that a particular organism might live, and how likely it is that each of those lives might actually occur. It is time, then, to see how this approach to

chance at the level of fitness can illuminate the two other most important loci for chance in evolutionary theory. I will turn in the last chapter to analyzing the connection between the model presented here and the roles of chance in natural selection and genetic drift.

5.6. Optional: Mathematical Derivation of Fitness Model

Note: This section is optional. It provides the details of the mathematical derivation of equation (5.4)*, and was published as an appendix to Pence and Ramsey (2013).*

Begin with the set Ω of possible daughter populations of some individual o, as defined above. We noted that we define a σ-algebra and probability measure, $\mathcal{F}$ and Pr, over Ω, in the usual way. In order to do so, however, we must demonstrate that Ω is sufficiently well-behaved that a probability measure may be defined over it – one might plausibly think that this set is simply too large to be suitable as a sample space. Consider our provisional definition (from footnote 6) of Ω as the space of functions from $\mathbb{R} \mapsto \mathbb{R}^n$. This set has cardinality $|\Omega| = 2^{2^{\aleph_0}}$, which makes it impossible to establish a standard probability measure over Ω in the normal manner – by exhibiting an isomorphism between Ω and either $\mathbb{R}$ or $[0, 1]$.

However, as demonstrated by Nelson (1959), we are able to define a standard σ-algebra $\mathcal{F}$ and a Borel probability measure Pr over certain subsets of this larger set. Namely, we can get what we need if we restrict our attention to (1) continuous functions ω, (2) functions ω with only point discontinuities, or (3) functions ω with only jump discontinuities.

Though we would like to remain neutral on how the n degrees of freedom available in the state vectors at each time t might be parameterized, it doesn't appear out of the question

to assume that the ω functions will have at worst jump discontinuities. If this is the case, then the proofs in Nelson (1959) demonstrate that a working probability measure can be reasonably defined over Ω.

With the sample space suitably defined, we may now derive our model. As we did above, define $\phi(t, \omega)$ as the function that takes a particular point in the sample space ω and a time t to the number of o's progeny living at time t on that outcome.

Now return to the problem of constructing a long-term picture of organismic fitness. We cannot evaluate the ϕ-function at time $t = \infty$, as we will obtain (at least potentially) infinite values. We thus need to define some function of the ϕ-values which will converge as t goes to ∞. We know, however, as stipulated above, that the ϕ-values are generated by a stationary random process, that demographic weak ergodicity holds, and that the logarithmic moment of vital rates is bounded. From this we may conclude (Tuljapurkar, 1989, pp. 233–234) that the following limit exists:

$$a = \lim_{t \to \infty} \frac{1}{t} \cdot E_w[\ln(\phi(t))], \tag{5.7}$$

with E_w an expectation value weighted by the probabilities given by our σ-algebra as defined above, and removing the parameter ω from the ϕ-function when it appears inside a mathematical expectation. If we take the exponent of both sides of equation (5.7), we arrive at

$$F = \exp\left(\lim_{t \to \infty} \frac{1}{t} \cdot E_w[\ln(\phi(t))]\right). \tag{5.8}$$

This value F is, then, the value of fitness in our model: equation (5.8) is equivalent to equa-

tion (5.4). It is precisely equal to the exponential of Tuljapurkar's a (that is, equation (5.7) just *is* Tuljapurkar's a), and therefore roughly equal (under simplifying assumptions) to the net reproductive rate and related to the Malthusian Parameter ($r = \ln(F)/T$, with T the generation time).

CHAPTER 6

UNDERSTANDING CHANCE IN FITNESS, SELECTION, AND DRIFT

In this last chapter, I want to explore what the the model developed in the previous chapter can tell us about the role of chance in evolutionary theory. How do we account for chance in fitness, natural selection, and genetic drift, if this picture is correct? In particular, invoking the new driving question I presented in chapter 3, what does this model say about the relationship between our statistical biological theories and the processes in the world those theories aim to describe? This is, of course, a vast topic, and I lack the space to pursue all of its contours in full detail – the discussion here serves, in part, as an exploration of how I would like to pursue the research undertaken in this dissertation in the future.

If we look over the model of fitness constructed in the last chapter, we see that "chance" enters in at only one point: the assignment of probabilities to the various possible lives that an organism might lead. This is a tantalizingly simple view of the relationship between statistical theories and biological processes. We ground the entirety of our use of probability in evolution in the simple observation that there are different ways that organisms' lives might go, and these are made more or less probable by a vast array of characteristics of the organisms and their environments.

The question that we have to answer in this chapter, therefore, is this: what is the

relationship between the probabilities found within our model of individual fitness and those found in natural selection and genetic drift? How, that is, do we spell out this theory/world relationship for each of the primary components of evolutionary theory discussed by philosophers of biology?

We can start by breaking up this larger question into two smaller foci for analysis. First, we should look at the interpretation of these probabilities themselves. In virtue of what do they take the values that they do, or on what do they depend? What kind of probabilities are they, and what sort of chance do they exemplify, according to the taxonomy laid out in chapter 1? Second, we should consider how these probabilities are manifested in genetic drift and natural selection – and in the process of answering this question, we must lay out exactly how selection and drift are to be described in terms of the model.

What we will see, surprisingly, is that from this single source of probability – the distribution across the possible lives of an individual organism – we can account for a wide array of observed "chancy" phenomena in evolution. The chance we know to be present in fitness, natural selection, and genetic drift, even that thought to be best understood at the population rather than the individual level, can all be shown to derive from this single source. This unified picture of the grounding of chance in evolutionary theory – of the connection between statistical biological theories and the biological world – will, I hope, be highly useful for future work in the philosophy of biology.

6.1. Probabilities in the Pence-Ramsey Model

Let's begin, in this section, by considering where and how probabilities enter into the model described in the last chapter (which I will henceforth refer to as the Pence-Ramsey model), and the appropriate empirical grounding and metaphysical interpretation of these probabilities. The first and most obvious question is this: what fixes their values? On which empirical facts do they rest?

6.1.1. What Fixes Their Values?

As mentioned in the introduction, each of the probabilities $\Pr(\omega)$ that appear in the model describe how likely it is that some particular possible life is the life actually experienced by a given organism.[1] It's clear that these probabilities are influenced by a whole host of properties of organisms and their environments. We captured this in the formal definition of the model with the two parameters for "genotype" and "environment," G and E, but noted that an extremely expansive reading of these two variables would be required to make sense of the various factors that govern the probability values. The "genotype" parameter should be taken to include all the various properties of the organism itself that influence how it lives its possible lives. In addition to the genotype itself, the probability values will also be altered by any other, non-genetic methods of inheritance, such as epigenetic transmission (Jablonka and Raz, 2009; Shea et al., 2011), niche construction (Odling-Smee et al., 2003;

1. Technically, since the probability space is uncountable, we can only assign probability values to events, or sets v of possible lives, but this detail isn't particularly important for us here.

Krakauer et al., 2009), the influence of parental behavior on offspring traits (Bateson, 2004), or cultural transmission (Feldman and Cavalli-Sforza, 1976; Boyd and Richerson, 2005; Laland et al., 2010). Since it is the phenotype, not the genotype, that governs an organism's interaction with the environment, the elusive "G-P map," the connection between genotypes and phenotypes, will also have a profound influence on these probability values (Pigliucci, 2010). Finally, evolutionary developmental biology has dramatically increased our estimation of the impact of development on the evolutionary process (Gilbert et al., 1996; West-Eberhard, 2005; Pigliucci, 2009; Olson, 2012), and development enters our model by virtue of its importance for the G-P map.

The "environment," as well, casts a wide net. It first includes all the features traditionally described as parts of an organism's environment – the specification of which is already a longstanding, difficult, and broadly unsolved problem in the philosophy of biology (Brandon, 1990; Abrams, 2009c). It also includes all conspecific organisms, including an organism's various possible mates. This point will be important when we come to a discussion of genetic drift, so it is worth elaborating in more detail. In the sense of the term familiar from population genetics and ecology, there is no obvious notion of a "population" present anywhere in the Pence-Ramsey model – we have only individual organisms and their possible *daughter* populations, none of which, as noted in the last chapter, are actually extant groups of conspecifics in the usual sense. Conspecific organisms, in the vocabulary of the model, appear as external features, elements of the environment with which the organism must interact. A difference in population membership, in this model, expresses itself as a difference between environmental interactions in various classes of possible lives.

On the basis of this extensive catalog of facts on which these probabilities depend,

one might be worried that they include *too many* features and are for this reason somehow ill-defined. In order to answer this worry, we should turn to our next task. In the last chapter, we claimed that these probabilities are consistent with a *propensity* interpretation of fitness. What are propensities, and can these probabilities be interpreted in this way?

6.1.2. Are They Really Propensities?

Before we can turn to propensities, we should begin with a quick account of dispositional properties in general. What is a disposition? There are several ways of drawing the distinction between dispositional and non-dispositional (or categorical) properties, but a common one centers on the idea that some properties describe not how objects *are,* but how they *will act* under certain conditions. Consider the property of solubility. A piece of salt is soluble not because it *is currently* dissolving in water, but rather because it *would* dissolve in water *if* appropriately situated with respect to some water – and this claim is true even if this particular piece of salt never is actually exposed to any water whatsoever.

To formalize this, we can minimally define dispositions as follows (following Fara, 2005, p. 70). A property P is a dispositional property iff it is such that an object N which bears it has some intrinsic property or properties such that it Ms when some set of conditions C holds. We call the variable M the manifestation of the property, the variable C the manifestation conditions, and we then say that P is the property of being disposed to M when C. As an example, a piece of salt (N) is soluble (P) iff it has some intrinsic property or properties in virtue of which it dissolves (M) when it is placed in contact with a solvent (C). The

manifestation of solubility is dissolving, and the manifestation conditions consist in being placed in water (or some other solvent, which is not itself saturated with salt, etc.). While this definition, sometimes called the "conditional analysis" of dispositions, is the most natural way to describe them, it is unfortunately beset with problems that must be resolved in order to provide a satisfactory and complete account.[2] Thankfully, however, this brief description will suffice for our purposes here.

What, then, is a propensity? A propensity (following Mellor, 1971, pp. 66–70) is just a dispositional property in which the manifestation – the property M – is a probability distribution. A coin, for example, has a propensity to be biased 60% in favor of heads and 40% tails if and only if it is disposed to express such a probability distribution when it is flipped. Again, note that the coin could have this property (i.e., this statement could be true of the coin) even if it is minted and immediately melted down without ever having once been flipped, since the property is dispositional. Propensities, on Mellor's view, are properties of individual objects, and, like all dispositions, are expressed only within certain ranges of manifestation conditions. As with dispositions, problems abound, and this is just one definition of many. Popper (1959), for example, would have propensities be properties of the *experimental setup* rather than of the individual objects themselves. On this interpretation, then, our example would be a dispositional property of some sort of coin tossing *experiment,* and the manifestation conditions would just be the setting up and executing of this particular type of experiment. These details, however, don't matter for us. What counts is that, in virtue of these propensities, the probability distributions at issue emerge as an interaction

2. See, for example, chapter 3 of Mumford (2003).

between the intrinsic properties of individual objects and some set of external manifestation conditions (however we might choose to divide these conceptually), and they are distributions over a specific range of values for some variable (such as a range of flipping velocities).

With a basic account of propensities under our belt, let's return to the probabilities $\Pr(\omega)$ from the Pence-Ramsey model of fitness. If these are to be the manifestation of some propensity, how are we to describe it? The original propensity interpretation of fitness (PIF) considered an individual organism's *propensity to have offspring,* which gives rise to the single-generation distribution of probability over the possible numbers of offspring that an organism might have. This propensity is clearly inappropriate for the new model – the probabilities described by the new model are fixed by *multi*-generational facts, which cannot be grounded in the disposition to have one generation of offspring. One terminological feature of propensities also stands out. They are, as a rule, propensities to attain some sort of *outcome* – the propensity to *land heads,* or to *have offspring.* The probability distribution, then, describes how likely this outcome really is in various circumstances, or for various particular specifications of the outcome conditions. This rules out a few candidates for which propensity might govern the $\Pr(\omega)$. For example, we wouldn't want to say that it was a propensity to *live out a particular possible life* – this doesn't seem like a suitable kind of outcome, but rather a process in which a particular organism and its offspring are engaged.

The best conclusion, then, seems to be that the propensity governing the $\Pr(\omega)$ is a propensity *to produce a daughter population.*[3] We would thus say that the probabilities $\Pr(\omega)$ are the manifestations of a propensity – the propensity to produce a daughter population

3. Again, for the peculiar definition of "daughter population" described in the last chapter.

(measured by the complex function introduced in the last chapter), and are manifested when an organism and its descendants reproduce.

Are these probabilities somehow too complex to be interpreted as propensities? I argue that they are not – they are no more complex, that is, than garden-variety examples of dispositions and propensities. Consider, for example, a dispositional account of fragility. The physical facts responsible for the fragility of an object are almost innumerable – relational properties that specify its shape, and many atomic-level properties specifying the hardness, bond strength, and deformability of the object. As regards fragility, these myriad intrinsic properties manifest themselves as either breaking or not breaking, in the presence of specific sorts of manifestation conditions (namely, being dropped in different ways). The mere claim that many facts determine whether or not an object will break when dropped doesn't seem to offer us any reason to think that a dispositional account of 'fragility' is ill-defined.

Turn, then, from a dispositional view of fragility to a propensity view of the probability of breakage of an object into a certain number of pieces given certain kinds of impacts. If the disposition of fragility is well-formed, then it's plausible to think that the propensity to break into pieces when dropped, which manifests given a particular drop as a probability distribution specifying the chance of breakage into different numbers of pieces, is also well-formed. In the case of the propensity to break into pieces, just like the propensity account of fitness, we simply have a property whose description summarizes a rather large amount of information. The "probability of breaking into X pieces given Y impact" seems to look much like the "probability of producing X daughter population given Y possible life" Ramsey (2013a).

Another objection to considering fitness as a propensity might rest on a different way

of drawing the distinction between categorical and dispositional properties. In the case of solubility and other clearly dispositional properties, one might argue, part of our evidence that the property at issue is dispositional is that the manifestation conditions could, clearly, fail to obtain – a particular piece of salt might never be exposed to any water. This does not seem to be the case for fitness, however. As the objection goes, fitness comes with an *unavoidable* set of manifestation conditions – a particular organism, even if it produces *no* offspring, is still producing *some* kind of daughter population. The manifestation conditions for the disposition we have described seem to merely be "living out a life."[4] It might thus be concluded, the objector claims, that this constitutes evidence that we have not a dispositional, but a categorical property.

We can respond, however, by noting that there do in fact exist many dispositions that are continually manifested by objects, and which those objects cannot fail to manifest in some way. Gravitation, for example, could be taken to describe a disposition of objects to accelerate toward one another in proportion to their masses and the square of the distance between them. An object with this disposition *always* manifests it – even if there are no other masses in the universe or it is currently experiencing no acceleration whatsoever. The reason that gravitation needs to be dispositional is that it describes how an object *will* react to different external conditions – just like fitness does.

Finally, we have one more interesting metaphysical question about these probabilities to consider. What are the *bearers* of these probabilities? By calling the model a definition of individual fitness, I have claimed that they are properties of individual organisms. Exactly

4. The same objection would, notably, apply to the traditional PIF. Since having no offspring is still an outcome as regards having offspring, it seems that organisms cannot avoid the manifestation conditions of the propensity to produce offspring.

what the bearers of these propensities are, however, will depend at least in part how we understand the notion of a propensity in general. (Recall that the conceptions of Popper and Mellor regarding the appropriate conception of the bearers of propensities.)

Setting the larger question of the correct account of propensities aside (as it lies well outside the scope of my project here), there is one feature of this fitness propensity in particular that seems to cast doubt on the claim that it can really be borne by individual organisms. Our account of dispositional properties above required that they hold in virtue of the *intrinsic properties* of the objects that bear them.[5] Don't the probabilities in this model of fitness reference too many factors that are external to the individual organism? For example, this dispositional account of fitness will make reference to whether or not a particular organism happens upon a food source, clearly not an intrinsic property of the organism itself. Does this imply that individual fitness is not a legitimate dispositional property? I argue that it does not, for three reasons. First, we must be cautious to separate the way in which we describe dispositions – which clearly makes reference to the external manifestation conditions – from the properties in virtue of which the disposition manifests, which may not in fact make such reference. Second, many of the properties that fall under the heading of the "genotype" above clearly are intrinsic properties of the organism – any properties of the organism that it obtains via genetic, cultural, behavioral, or epigenetic modes of transmission are intrinsic properties, among others.

But we must, finally, consider the influence of external factors on the organism – and these will appear under both our headings of the "genotype" (e.g., development) and, more

5. Several authors have considered the possibility of *extrinsic* dispositions, another possible response to this objection that I will not consider here. See McKitrick (2003); Contessa (2012).

obviously, the "environment." To account for these external influences, we should return to the structure of propensities. Consider the example of the propensity of a coin to fall heads or tails. This propensity manifests itself whenever the coin is "flipped." But what do we mean by flipped? Assuming that we're talking about an experimental setup that describes normal, human coin-flippers, we must mean that there is some *range* of values of height from the floor, initial linear and angular momentum, and other relevant physical parameters that makes individual trials count as a successful flip. For example, we might say that if the coin fails to tumble through the air a few times (failing to turn over, like a tossed pizza), or is flipped merely an inch above the floor (only inverting a few times), we have an "illegal" flip. The probability distribution manifested by this propensity just is a 60% chance of landing heads *across* the broad, fuzzy range of "legal" coin-flips. A similar story applies to a propensity account of the probability of breaking when dropped. The relationship between the shape and hardness of the object and these same characteristics of the surface that it impacts are vital components of the odds that an object will break – and these are expressed in this propensity by the looseness of our definition of "dropped."[6]

It's precisely here where we can provide an account of the impact of external factors on the probabilities described by this model of fitness. The propensity to produce a daughter population is manifested whenever organisms (and their descendants) live their lives, succeeding or failing to produce offspring. Just as in the cases of "flipped" and "dropped," these probabilities are a distribution over a whole range of various possible interactions with the external environment that the organism might experience – sets of interactions that all

6. Notably, this fuzziness is why Popper wanted propensities to be a property of the experimental setup, taken to include the process that generates this range of circumstances.

qualify as "living their lives." Once again, it seems that these probabilities, when analyzed carefully, look much like those that are ascribed in virtue of traditional propensities.

6.1.3. What Sort of 'Chance'?

Finally, to conclude our analysis of these probabilities, we should return to the taxonomy that I developed in the first chapter. What is the correct way to understand the *type of 'chance'* embodied by these probabilities? In that chapter, I described unpredictability as the inability of a given agent, with a given epistemic state, to predict the precise outcome of a given process. We can begin, then, by quickly dismissing the idea that these probabilities might be the result of unpredictability in this sense. They arise as a result of the interaction between organisms and environments, features which are surely observer-independent.

The next question, then, is whether we should consider them to be objective chance or a result of probabilistic causal processes. Recall that this question is equivalent to asking whether or not, if we conditionalized on enough information, the values of these probabilities would "collapse" – that is, would take on the values one or zero. Getting to the correct answer here is much more difficult, and the search encompasses several distinct questions. One is clearly empirical: what is the impact of quantum mechanics on biological systems? There are most certainly instances of quantum indeterminism that can "percolate up" to the level of biological phenomena (Stamos, 2001; Davies, 2004): perhaps most impressive among these is a model of avian detection of magnetic North which relies upon careful maintenance of entanglement and superposition, possibly even for longer periods of time than can be

successfully executed in controlled laboratory experiments (Maeda et al., 2008; Gauger et al., 2011). In these cases, it is plausible that the probabilities of organisms' leading particular possible lives will have a non-trivial, non-eliminable component of quantum randomness, and would thus be non-collapsible, *objective chance* in the sense of chapter 1.

Many other causal influences on organisms, though, will be roughly deterministic (Rosenberg, 2001). These will include the bulk of things traditionally studied by evolutionary biologists – predation, migration, and so forth. Because they are essentially deterministic processes, these are collapsible probabilities, as their values could become extremal given enough additional information regarding initial conditions. This means that some of the probabilistic influences on the fitness probabilities are going to reflect *probabilistic causal processes* (again, as in chapter 1).

The precise mix of collapsible and non-collapsible probability in evolution is, therefore, a matter for further empirical research. What does this mean, though, for the overall probability values $\Pr(\omega)$ and the propensity which manifests them? What is the correct way to interpret a propensity which is grounded, at least in part, in further probabilistic properties found at lower ontological levels (i.e., physics or chemistry), and hence which seems to be determined by, or encompass the effects of, 'chance' in the senses of both objective chance *and* probabilistic causal processes?

On the one hand, the precise definitions of collapsible and non-collapsible probabilities seem as though a single jot of objective chance makes any probability influenced by that objective chance, in turn, another instance of objective chance. That is, if a probability X is at all determined by some non-collapsible probability Y, no amount of conditionalization could result in X always precisely equaling one or zero – there is no way to eliminate the influence

of the non-collapsible Y. This makes X, too, non-collapsible, and yet another instance of objective chance.

But, on the other hand, is this really the right way to evaluate the status of these probabilities? This analysis seems to overstate the role of objective chance. For example, consider a case where some probability which begins with a value of 0.6 will collapse, given a complete set of information, to a tiny value, such as 10^{-8}. It seems incorrect to say, in a case like this, that our initial value was *only* (or even *mainly*) a matter of objective chance – collapsible probabilistic influences were responsible, it seems, for the lion's share of the probability's value.

Finally, with regard to the probabilities of individual fitness in particular, we have the problem – both in general and in every real-world case – that we cannot specify the precise relationship between the components of fitness (that is, the fundamental, organism-level instances of objective chance and probabilistic causal processes) and the overall probabilities that feature in fitness itself. Of course, the issue of whether or not such epistemic factors should be taken into account in our modeling of (or metaphysical conclusions regarding) biological systems is a matter of some debate (see, e.g., Wimsatt, 2007; Richardson, 2007). But if we should indeed take these factors into account, this could suggest another approach to the interpretation of the probabilities in this model of fitness. Perhaps the correct course of action in this case is to analyze the role of chance as it arises in fitness *without* thinking about these inter-level connections at all. Perhaps, that is, the inter-level connections are *irrelevant* to the correct analysis of chance in fitness, and the fitness probabilities should be interpreted as *sui generis.*

All these factors combine to make it exceedingly unclear what the correct account of

these fitness probabilities will turn out to be, and I do not yet have a compelling resolution to all these various problems. I am sympathetic to the idea that the final way to understand this propensity that I considered is correct – that the inter-level relationships I have described (and problematized) here are in fact irrelevant to the interpretation of these probabilities. I lack a fully fledged argument for this conclusion, however; I can only claim that I suspect this pluralistic approach will provide the most fruitful and useful analysis of chance in biological theory – one that can have an impact on biological practice itself, rather than taking the "relevant" loci for the analysis of chance in evolution to be chemical, physical, or even metaphysical. I hope to clarify such an analysis in future work.

6.2. Chance in Natural Selection

Let's now turn to the second task laid out in the introduction. What does this picture of the probabilities in individual fitness say about the role of chance in natural selection and genetic drift? Importantly, before we can figure out the role of chance in natural selection according to this new model, we first have to consider how selection is to be defined in our theoretical vocabulary. Unsurprisingly, of course, how much we're able to say about natural selection depends on how one defines selection in the first place – returning to chapter 4, whether we adopt a process or product definition of selection.

Let's begin with the product definition. On this view, selection just is population change that's tied to differences in fitness. What, then, is responsible for those differences here? As discussed above, many facts are responsible for setting the probabilities of the various

possible lives, and these will then be, in turn (though much more mediately), "responsible for" population change due to natural selection.

This, however, isn't very informative. Given the wide variety of facts implicated in fixing the values of these probabilities, it's no real surprise that some set of these facts is mediately responsible for natural selection in the product sense. This is, though, all we should expect to be able to say about selection on a product definition. Defining selection as population outcomes that correlate with differences in fitness doesn't say anything about the processes that bring that change about.

Can we see anything more interesting, then, if we move to a process-based definition of natural selection? Recall that on this view, natural selection just is the set of all causal processes that discriminate between organisms on the basis of fitness differences. How would these causal processes wind up being represented in the vocabulary of the Pence-Ramsey fitness model?

We have an interesting problem of explanatory inversion here. For discriminate causal processes, on this definition of selection, are those that act differently on different organisms, on the basis *of their fitness values* – that is, the differences in fitness are taken to be fundamental, and selection is described in terms of them. But now, we're investigating the converse scenario: taking selection as fundamental, can we describe these causal processes which constitute selection in terms of our model of fitness?

One initial thought might be to describe natural selection as (NS1) the full suite of things that set our possible-life probabilities. But this clearly can't be correct. The set of factors that govern the values of these probabilities seems to contain things that should be considered the basis for (or causes of) genetic drift, mutation, and other components of evolution, not

natural selection. For example, mutation rates (and the environmental features that govern their values) influence the likely size of an organism's future lineage, and therefore are going to feature in the set of things that fix the values of the possible-life probability values.

If this is wrong, could we instead define natural selection as (NS2) the set of factors that are directly responsible for a difference in fitness value between two organisms? This also can't be right, for we can have instances of selective change without differences in individual fitness values (Walsh, 2007). Consider the response of two organisms to two different environmental influences. Say that organism *A* outperforms *B* in one scenario, and *B* outperforms *A* in the other. When we evaluate individual fitness across all possible lives, these influences happen to offset, and *A* and *B* have identical individual fitness values. As a matter of chance, however – that is, as a result of the particular possible life that each organism actually experiences – say that organism *A* outperforms organism *B*, and it is due precisely to this difference in response to environmental influences described above. It seems that we would, indeed, want to call this natural selection – on the process view of selection, for example, *A* has outperformed *B* in virtue of a causal process that discriminated between the two organisms. But there is no corresponding individual fitness difference, and thus no influences to point to as "responsible" in the sense of (NS2). Of course, there was a difference in the fitness of *traits* at work here – the traits behind the differing responses to the two environmental effects.[7] But we're not dealing with trait fitness, nor is it yet clear how we could extract a concept of trait fitness from the Pence-Ramsey model.

Neither (NS1) nor (NS2), then, seems to be a viable candidate for a definition of

7. Things are made even more complex by the existence of cases that are precisely inverse to this one – where we want to say that a selective event has taken place, and we have a difference in individual fitness but fail to have a difference in trait fitness (Ramsey, 2013c).

natural selection. Return to the way in which we cashed out "discriminate" above. Another way to think of these discriminate causal processes would be to try to establish some sort of equivalence relation on the set of possible lives, and then say that a causal process is discriminate when the outcomes for two different organisms differ in what are, on this equivalence relation, "the same" possible lives. Assuming the equivalence relation were defined in the appropriate way, this could capture the case we described above – organisms *A* and *B* have different numbers of offspring when they are exposed to "the same" sequence of particular environmental conditions. We might then try to define natural selection as (NS3) processes which cause different organisms to experience different outcomes in "the same" possible lives, for some such equivalence relation.

There are two problems with this definition, however, that make it difficult to see how we could adopt it. First, how could we possibly cash out this equivalence relation? We clearly don't mean identical possible lives, as it is impossible that any two organisms could possibly lead *strictly identical* possible lives (since, for example, position in spacetime is clearly part of a possible life). It seems that specifying this equivalence relation in any nontrivial way will be incredibly difficult. Second, the same seems to be true of the notion of "different outcomes" which (NS3) invokes. Precisely the reason that we have introduced this complex measure of fitness is because defining the correct notion of an "outcome" is exceedingly complicated.

Sorting out natural selection on the Pence-Ramsey model, therefore, is a complex and open problem. If we adopt an outcome based view of natural selection, then the model doesn't seem to give us much of a handle on the influences responsible for selection. The same, it turns out, is true for a process view of selection.

One might respond, of course, that this is to be expected – as I noted above, we usually

attempt to define selection in terms of differences in fitness, not the other way around. But we saw in our exploration of (NS2) that this project, as well, may be met with some difficulties. As was discussed in chapter 4, the definition of selection is a recurring and unresolved problem with implications throughout the philosophy of biology, and this new model of fitness makes things all the more problematic. I hope to be able to return to this issue in the future.

Returning to the main theme of this chapter, is there anything we can say about the role of chance in natural selection, even without being able to fully resolve the problem of its definition?

The main insight regarding chance that we must take into account in the case of natural selection dates back to Darwin – all selection can guarantee is that the fitter organisms will *tend to* out-reproduce the less fit. And accounting for the presence of that "tend to" only requires one of the most straightforward implications of the propensity interpretation of fitness that we detailed in the last chapter: the actual reproductive output of organisms isn't necessarily connected to their fitness values.

And once again, this insight derives from the single source of probabilities in the Pence-Ramsey model: that there are different possible ways an organism's life might go. To the limited extent that chance is present in natural selection – however we might wind up treating selection in this model – that instance of chance can be readily accounted for.

6.3. Chance in Genetic Drift

What about genetic drift? The case here is the most difficult of all. Most importantly, drift is almost always considered to be the result of some population-level cause or sampling process, especially since its effect is modulated by population size. This is particularly problematic for the model developed in the last chapter – as mentioned above, there is no obvious way to extract a concept of a "population" from the model, or even from the broader theoretical vocabulary we developed there. Notable as well is the fact that we don't need just any population size to get theories of drift off the ground. As argued extensively by Otsuka et al. (2011), the population size which modulates the effect of drift is the *effective* population size – a very particular sort of population size, the definition and measurement of which is one of the central difficulties in the application of the models of population genetics to real-world, empirical cases. This makes it all the more doubtful that we would be able to extract the correct value of effective population size from the framework of the Pence-Ramsey model.

This way of analyzing the situation, however, takes as a premise that since drift is modulated by effective population size, we must describe it in terms of processes that act on populations (Millstein, 2006a). But this isn't the only way in which we might describe the structure of drift. Another option for explaining a population-level effect like drift is to look for the set of individual-level properties in virtue of which genetic drift occurs in populations. There may, of course, be no well-defined set of these properties, or the properties might be connected with drift effects in only a very complex way – in these cases, we would then say

that drift is an emergent phenomenon at the population level, and attempts to explain drift from the individual level would be relatively fruitless. If, however, we could deploy this sort of individual-level analysis of drift, it might be possible to describe drift in terms provided by the model developed in the last chapter. It's therefore worthwhile to pursue that goal here.

To get a handle on what features of individuals might possibly ground genetic drift, we can start by considering one of drift's most commonly cited characteristics. Infinite populations, it is commonly stated, entirely lack all drift effects. In the terms of the Pence-Ramsey model, is there some feature of infinite populations, grounded in individual-level facts, that can explain this difference between finite and infinite populations? For one, infinite populations form a *perfectly representative sample* of the space of possible lives.[8] In effectively all empirical cases, as well, it's clear that *only* an infinite population could form such a sample.[9] Might we find within this fact an individual-level grounding for genetic drift?

A good way to try to tease this out, then, would be to consider a toy model of a finite population that happened to be perfectly representative of its space of possible lives. Would this population experience genetic drift, or not? Imagine a population composed of two types of individuals. One type has only one life history, on which it has one offspring, and then dies. Another type has two possible life histories. On one of these, it leaves behind no offspring, and on another, it produces two. One each of these offspring will live out each of the two possible life histories. Imagine, then, a population that begins with four organisms –

---

8. This is, of course, not quite correct, as an infinite sample from an infinite population does not *have* to be representative – consider the sample of all even numbers from the natural numbers. It is, however, the case that the probability of a sample's being representative goes to 1 as the sample size goes to infinity.

9. For example, as mentioned in chapter 4, since spatial trajectory is going to distinguish possible lives, almost all real-world examples of the space of possible lives will be infinite, and thus only an infinite sample can be perfectly representative of an infinite space.

two of the first type, and two of the second type, one living out each life history. The outcome of this case is clearly deterministic. There will be two organisms of the first type at every generation, and two organisms of the second type at every generation. And this (finite!) population experiences no genetic drift.

What's more, it seems that this case will generalize. Any time we have a population that is *perfectly* representative of the space of possible lives, there will be no genetic drift. In all such cases, population change will, definitionally, proceed precisely according to individual fitness values. All population change, therefore, will be attributable to natural selection.[10] Of course, as mentioned above, we will never find a real-world population in which this perfect representativeness condition holds: populations must either be infinite or have bizarre, constrained sets of possible lives as in our last example to perfectly sample the space.

What is the moral of this example? For one thing, Ramsey (2013a) has argued that this entails that there is a difference between *drift* and *the response to drift,* paralleling the distinction commonly drawn in the case of selection. A finite population that happens to be perfectly representative might have *the conditions necessary for drift,* but actually experience no drift-like population change. In our example, for instance, the second organism type with two life histories is highly amenable to drift-like change. Were there a *non*-representative sampling of possible lives, many more of the zero-offspring organisms might chance to survive instead of two-offspring organisms (or vice versa), producing a dramatic change attributable to genetic drift. In Ramsey's terminology, this makes those individuals have a high *driftability* value (the quantity which measures the extent to which an individual might

10. We might also have mutation or migration (or other evolutionary forces), depending on the details of the example, but there will clearly be no genetic drift.

experience drift, analogous to the role of fitness in natural selection), but might nonetheless not experience any population change as a result of drift.

But this point is orthogonal to our discussion of chance. What makes genetic drift chancy on this picture of how drift works – where do the probabilities in genetic drift come from? The key difference between cases with and without drift that we described above had to do with the extent to which the distribution of organisms throughout the space of possible lives is representative of the distribution overall. If the population is distributed in a highly representative way, it will very likely experience little change as a result of drift. If it is distributed unrepresentatively, then it is likely to experience much change as a result of genetic drift.

There are, unfortunately, a variety of ways to quantify the extent to which a sample drawn from a distribution is representative. I consider it another open research question to pursue which of these measures will reproduce the equations of the various models of genetic drift familiar from population genetics, such as the Wright-Fisher model. But even without being able to specify the mathematics of the connection precisely, this picture lets us explain two interesting features of genetic drift.

First, we can see precisely where we get the appearance of a population-level effect. Clearly, any measure of whether or not a sample is statistically representative is going to introduce a sample size parameter, and is going to be a property *of* the sample. But we can further see how genetic drift could be describable as a population-level effect while being grounded in the individual-level properties that we have already described.[11] The peculiar

11. For more on this connection, see Ramsey (2013a).

status of drift as a property of populations that is at the same time deeply connected to a wide array of disparate individual-level features becomes clear on this view.

Second, let's return to the discussion of definitions of drift from chapter 4. Given the view developed here, we can see what's right about (a few of) the process and the product pictures of genetic drift.

1. The description of drift as "sampling error" – what I called GD^{SE}_{prod} – is at least in part correct. The extent to which a sample of a possible-life distribution is representative is a variety of sampling error, though a variety different than those discussed by any of the authors that advocate GD^{SE}_{prod}.

2. The ecological interpretation that I developed of the "initial conditions" view of Bouchard and Rosenberg (2004), which we called GD^{IC}_{proc}, turns out to be *very close* to being correct – the only difference between their view and ours is that our "possible lives" are multigenerational, for the reasons described in the last chapter.

3. And, while it isn't possible to offer a full analysis of this point here, it seems plausible that the set of causes at the individual level that are singled out by the "indiscriminate sampling process" definition of genetic drift (GD^{SP}_{proc}) are just those causes that increase the variance in the distribution of possible lives, making it more likely that a sample of a given size will be less representative.

We also can begin to understand why the borderline cases of genetic drift that we described at the close of chapter 4 are so strange. These cases, arguably, influence not just the variance of the distribution of possible lives (increasing or decreasing the expected amount of change

from genetic drift), but also change the fitness values themselves, giving them elements of both drift-based and selection-based changes.

6.4. Conclusions

The last few points here have been fairly sketchy and underdeveloped. We can see, however, that this model of fitness constitutes an incredibly fertile arena for discussing the question I proposed in chapter 3. Our "chancy" biological theories – at least, the theories of natural selection and genetic drift[12] – are connected to the biological world through one single source of "chanciness," the distribution of probability values over the set of possible lives that an organism might live.

And these, in turn, are grounded in the simple insight of the PIF-as-interpretation developed in the last chapter: that an individual organism has myriad different ways that its life might go. In constructing biological theory, we quantify that probability, and we proceed to split the various influences causally responsible for it and the effects that arise from it. Some of these we call natural selection, and some we call drift. But the recognition that all of the various aspects of "chance" in these theories actually derive from a single source is a profoundly important insight – making it far easier to understand the true roots of chance in evolution.

This brings to a close the wide arc of my project here. We began by exploring the

12. It is notable, of course, that I lack the space in this dissertation to pursue how this connection might be extended to a whole host of other "chancy" components of evolution, including mutation, macroevolutionary contingency, and the use of stochastic models in ecology. This, as well, is something I hope to pursue in future work.

contemporary status of the literature on chance in evolution in chapter 1, noting that much of that literature rests on a persistent conflation between many distinct senses of 'chance' and related concepts. In the second and third chapters, we saw that literature on the historical introduction of chance into evolutionary theory also frequently misunderstands the aim and work of the scientists who were the first to comprehensively introduce statistics and probability into evolution. This historical case provided the impetus for the development of a new way to think about chance in evolution: the relationship between our statistical theories and the biological world.

In the fourth chapter, I considered one common way of approaching precisely this question in the philosophical literature – the causalist/statisticalist debate – and found that it, too, was lacking. Finally, in the fifth chapter, I proposed a novel way of understanding individual fitness – one that, as we have seen in this last chapter, can offer us a new and especially simple way of responding to the charge presented in chapter 3. The relationship between our statistical theories and the biological world is simpler than we might have thought. Much of the influence of chance in evolution passes through a single, unified theoretical point: the distribution of probabilities over the different ways that an organism's life might go.

While this chapter has managed to raise as many questions as it answers, I hope it is clear that I have developed here a novel and interesting way to try to approach one of the most difficult points in the philosophical interpretation of probability, statistics, and evolution. Exploring the role of chance in evolution brings us right to the heart of Darwin's theory and the hundred and fifty years of succeeding scholarship in biology. A sophisticated understanding of this connection can provide a fruitful lens through which to observe and

appreciate the "tangled bank" of beautifully complex facts explained by evolutionary theory.

BIBLIOGRAPHY

Abrams, M. 2007. *How do natural selection and random drift interact?* **Philosophy of Science**, 74(5):666–679. doi: 10.1086/525612.

———. 2009a. *Fitness "kinematics": biological function, altruism, and organism-environment development*. **Biology and Philosophy**, 24(4):487–504. doi: 10.1007/s10539-009-9153-2.

———. 2009b. *The unity of fitness*. **Philosophy of Science**, 76(5):750–761. doi: 10.1086/605788.

———. 2009c. *What determines biological fitness? The problem of the reference environment*. **Synthese**, 166(1):21–40. doi: 10.1007/s11229-007-9255-9.

———. 2012. *Mechanistic probability*. **Synthese**, 187(2):343–375. doi: 10.1007/s11229-010-9830-3.

———. 2013. *Populations and pigeons: prosaic pluralism about evolutionary causes*. **Studies in History and Philosophy of Biological and Biomedical Sciences**, 44:294–301. doi: 10.1016/j.shpsc.2013.06.005.

Ahmed, S. and J. Hodgkin. 2000. *MRT-2 checkpoint protein is required for germline immortality and telomere replication in* C. elegans. **Nature**, 403(6766):159–164. doi: 10.1038/35003120.

Alexander, P. 1964. *The philosophy of science, 1850–1910*. In **A critical history of Western philosophy**, page 402–425. The Free Press of Glencoe, London.

Ankeny, R. 2000. *Marvelling at the marvel: The supposed conversion of A.D. Darbishire to Mendelism*. **Journal of the History of Biology**, 33(2):315–347. doi: 10.1023/A:1004750216919.

Ariew, A. and Z. Ernst. 2009. *What fitness can't be*. **Erkenntnis**, 71(3):289–301. doi: 10.1007/s10670-009-9183-9.

Ariew, A. and R.C. Lewontin. 2004. *The confusions of fitness*. **British Journal for the Philosophy of Science**, 55(2):347–363. doi: 10.1093/bjps/55.2.347.

Barros, D.B. 2008. *Natural selection as a mechanism*. **Philosophy of Science**, 75(3):306–322. doi: 10.1086/593075.

Bateson, P. 2004. *The active role of behaviour in evolution*. **Biology and Philosophy**, 19(2): 283–298. doi: 10.1023/B:BIPH.0000024468.12161.83.

Bateson, W. 1901. *Heredity, differentiation, and other conceptions of biology: A consideration of Professor Karl Pearson's paper 'On the principle of homotyposis'*. **Proceedings of the Royal Society of London**, 69(451-458):193–205. doi: 10.1098/rspl.1901.0099.

Beatty, J.H. 1984. *Chance and natural selection*. **Philosophy of Science**, 51:183–211. doi: 10.1086/289159.

———. 1992. *Random drift*. In Keller, E.F. and E.A. Lloyd, editors, **Keywords in evolutionary biology**, page 273–281. Harvard University Press, Cambridge, MA.

———. 2006a. *Chance variation: Darwin on orchids*. **Philosophy of Science**, 73(5):629–641. doi: 10.1086/518332.

———. 2006b. *Replaying life's tape*. **Journal of Philosophy**, 103(7):336.

Beatty, J.H. and S. Finsen. 1989. *Rethinking the propensity interpretation of fitness: A peek inside Pandora's box*. In Ruse, M., editor, **What the philosophy of biology is: Essays for David Hull**, page 18–30. Kluwer Academic Publishers, Dordrecht.

Benton, T.G. and A. Grant. 2000. *Evolutionary fitness in ecology: Comparing measures of fitness in stochastic, density-dependent environments*. **Evolutionary Ecology Research**, 2:769–789.

Bouchard, F. and A. Rosenberg. 2004. *Fitness, probability and the principles of natural selection*. **British Journal for the Philosophy of Science**, 55(4):693–712. doi: 10.1093/bjps/55.4.693.

Bowler, P.J. 1992. *The eclipse of Darwinism: anti-Darwinian evolution theories in the decades around 1900*. Johns Hopkins University Press, Baltimore, MD.

Boyce, M.S. and C.M. Perrins. 1987. *Optimizing great tit clutch size in a fluctuating environment*. **Ecology**, 68(1):142–153. doi: 10.2307/1938814.

Boyd, R. and P.J. Richerson. 2005. *The origin and evolution of cultures*. Oxford University Press, Oxford.

Brandon, R.N. 1978. *Adaptation and evolutionary theory*. **Studies in History and Philosophy of Science**, 9(3):181–206. doi: 10.1016/0039-3681(78)90005-5.

———. 1990. *Adaptation and environment*. Princeton University Press, Princeton, NJ.

———. 2005. *The difference between selection and drift: A reply to Millstein*. **Biology and Philosophy**, 20(1):153–170. doi: 10.1007/s10539-004-1070-9.

———. 2006. *The principle of drift: biology's first law*. **Journal of Philosophy**, 103(7): 319–335.

———. 2010. *A non-Newtonian model of evolution: the ZFEL view*. **Philosophy of Science**, 77(5):702–715. doi: 10.1086/656901.

Brandon, R.N. and S. Carson. 1996. *The indeterministic character of evolutionary theory: No "No hidden variables proof" but no room for determinism either*. **Philosophy of Science**, 63(3):315–337. doi: 10.1086/289915.

Brandon, R.N. and G. Ramsey. 2007. *What's wrong with the emergentist statistical interpretation of natural selection and random drift?* In Hull, D.L. and M. Ruse, editors, **The Cambridge companion to the philosophy of biology**, pages 66–84. Cambridge University Press, Cambridge.

Brunnander, B. 2006. *What is natural selection?* **Biology and Philosophy**, 22(2):231–246. doi: 10.1007/s10539-005-9008-4.

Caswell, H. 1989. *Matrix population models*. Sinauer Associates, Sunderland, MA.

———. 2009. *Sensitivity and elasticity of density-dependent population models*. **Journal of Difference Equations and Applications**, 15(4):349–369. doi: 10.1080/10236190802282669.

Caswell, H. and T. Takada. 2004. *Elasticity analysis of density-dependent matrix population models: the invasion exponent and its substitutes*. **Theoretical Population Biology**, 65(4):401–411. doi: 10.1016/j.tpb.2003.09.007.

Caswell, H., T. Takada, and C.M. Hunter. 2004. *Sensitivity analysis of equilibrium in density-dependent matrix population models*. **Ecology Letters**, 7(5):380–387. doi: 10.1111/j.1461-0248.2004.00595.x.

Chakravartty, A. 2007. *A metaphysics for scientific realism: knowing the unobservable*. Cambridge University Press, Cambridge.

Charlesworth, B. 1970. *Selection in populations with overlapping generations. I. The use of Malthusian parameters in population genetics*. **Theoretical Population Biology**, 1(3): 352–370. doi: 10.1016/0040-5809(70)90051-1.

———. 1980. *Evolution in age-structured populations*. Cambridge University Press, Cambridge.

Clifford, W.K. and K. Pearson. 1885. *The common sense of the exact sciences*. D. Appleton and Company, New York.

Cock, A.G. 1973. *William Bateson, Mendelism and biometry*. **Journal of the History of Biology**, 6(1):1–36. doi: 10.1007/BF00137297.

Cock, A.G. and D.R. Forsdyke. 2008. *Treasure your exceptions: the science and life of William Bateson*. Springer, New York.

Cohen, J.E. 1979. *Comparative statics and stochastic dynamics of age-structured populations*. **Theoretical Population Biology**, 16(2):159–171. doi: 10.1016/0040-5809(79)90011-X.

Cohen, J.E. and C.C. Heyde. 1985. *Confidence intervals for demographic projections based on products of random matrices*. **Theoretical Population Biology**, 27:120–153. doi: 10.1016/0040-5809(85)90007-3.

Contessa, G. 2012. *Do extrinsic dispositions need extrinsic causal bases?* **Philosophy and Phenomenological Research**, 84(3):622–638. doi: 10.1111/j.1933-1592.2010.00435.x.

Conway Morris, S. 2009. *The predictability of evolution: glimpses into a post-Darwinian world*. **Die Naturwissenschaften**, 96(11):1313–1337. doi: 10.1007/s00114-009-0607-9.

Cooper, W.S. 1984. *Expected time to extinction and the concept of fundamental fitness*. **Journal of Theoretical Biology**, 107(4):603–629. doi: 10.1016/S0022-5193(84)80135-6.

Crow, J. and M. Kimura. 1956. *Some genetic problems in natural populations*. **Proceedings of the third Berkeley symposium on mathematical statistics and probability**, 4: 1–22.

Cunningham, J.T. 1895. *The statistical investigation of evolution [letter of Mar. 28, 1895]*. **Nature**, 51(1326):510. doi: 10.1038/051510a0.

Dale, A.I. 1999. *A history of inverse probability: from Thomas Bayes to Karl Pearson*. Springer, New York.

Darden, L. 1977. *William Bateson and the promise of Mendelism*. **Journal of the History of Biology**, 10(1):87–106. doi: 10.1007/BF00126096.

Darwin, C. 1837. *Notebook B: [Transmutation of species (1837–1838)]. CUL-DAR121*. Rookmaker, K., editor. Darwin Online, URL: http://darwin-online.org.uk/.

———. 1838a. *Notebook M: [Metaphysics on morals and speculations on expression (1838)]. CUL-DAR125*. Rookmaker, K., editor. Darwin Online, URL: http://darwin-online.org.uk/.

———. 1838b. *Notebook C: [Transmutation of species (1838.02–1838.07)]. CUL-DAR122*. Rookmaker, K., editor. Darwin Online, URL: http://darwin-online.org.uk/.

———. 1838c. *Notebook N: [Metaphysics and expression (1838–1839)]. CUL-DAR126*. Rookmaker, K., editor. Darwin Online, URL: http://darwin-online.org.uk/.

———. 1838d. *Notebook E: [Transmutation of species (10.1838–7.1839)]. CUL-DAR124*. Rookmaker, K., editor. Darwin Online, URL: http://darwin-online.org.uk/.

———. 1838e. *Old & useless notes about the moral sense & some metaphysical points [(1838–1840)]. CUL-DAR91.4-55*. Barrett, P.H., editor. Darwin Online, URL: http://darwin-online.org.uk/.

———. 1861. *On the origin of species*. John Murray, London, 3rd edition.

———. 1871. *The descent of man and selection in relation to sex*. John Murray, London.

———. 1875. *The variation of animals and plants under domestication*. John Murray, London, 2nd edition.

———. 1876. *On the origin of species*. John Murray, London, 6th, corr. edition.

———. 1909. *The foundations of the Origin of Species: two essays written in 1842 and 1844*. Darwin, F., editor. Cambridge University Press, Cambridge.

Davies, P.C.W. 2004. *Does quantum mechanics play a non-trivial role in life?* **BioSystems**, 78 (1-3):69–79. doi: 10.1016/j.biosystems.2004.07.001.

de Jong, G. 1994. *The fitness of fitness concepts and the description of natural selection*. **Quarterly Review of Biology**, 69(1):3–29. doi: 10.1086/418431.

Denniston, C. 1978. *An incorrect definition of fitness revisited*. **Annals of Human Genetics**, 42:77–85. doi: 10.1111/j.1469-1809.1978.tb00932.x.

Depew, D.J. and B.H. Weber. 1995. *Darwinism evolving: systems dynamics and the genealogy of natural selection*. Bradford Books, Cambridge, MA.

Desjardins, E.C. 2011. *Historicity and experimental evolution*. **Biology and Philosophy**, 26: 339–364. doi: 10.1007/s10539-011-9256-4.

Diaconis, P. 1998. *A place for philosophy? The rise of modeling in statistical science*. **Quarterly of Applied Mathematics**, 56(4):797–805.

Dietrich, M.R. 2006. *Three perspectives on neutrality and drift in molecular evolution*. **Philosophy of Science**, 73(5):666–677. doi: 10.1086/518521.

Dowe, P. 1992. *Wesley Salmon's process theory of causality and the conserved quantity theory*. **Philosophy of Science**, 59(2):195–216. doi: 10.1086/289662.

———. 1995. *Causality and conserved quantities: a reply to Salmon*. **Philosophy of Science**, 62(2):321–333. doi: 10.1086/289859.

Druery, C.T. and W. Bateson. 1901. *Experiments in plant hybridization [translation of Mendel, J. G., 1865, Versuche uber Pflanzenhybriden]*. **Journal of the Royal Horticultural Society**, 26:1–32.

Eagle, A. 2013. *Chance versus randomness*. In Zalta, E.N., editor, **The Stanford Encyclopedia of Philosophy (Spring 2013 Edition)**. URL: http://plato.stanford.edu/archives/spr2013/entries/chance-randomness/.

Fara, M. 2005. *Dispositions and habituals*. **Noûs**, 39(1):43–82. doi: 10.1111/j.0029-4624.2005.00493.x.

Fazekas, P. and G. Kertész. 2011. *Causation at different levels: tracking the commitments of mechanistic explanations*. **Biology and Philosophy**, 26:365–383. doi: 10.1007/s10539-011-9247-5.

Feldman, M.W. and L.L. Cavalli-Sforza. 1976. *Cultural and biological evolutionary processes, selection for a trait under complex transmission*. **Theoretical Population Biology**, 9 (2):238–259. doi: 10.1016/0040-5809(76)90047-2.

Filler, J. 2009. *Newtonian forces and evolutionary biology: a problem and solution for extending the force interpretation*. **Philosophy of Science**, 76:774–783. doi: 10.1086/605799.

Fisher, R.A. 1918. *The correlation between relatives on the supposition of Mendelian inheritance*. **Philosophical Transactions of the Royal Society of Edinburgh**, 52:399–433.

———. 1922. *Darwinian evolution by mutations*. **Eugenics Review**, 14:31–34.

Forber, P. and K. Reisman. 2007. *Can there be stochastic evolutionary causes?* **Philosophy of Science**, 74(5):616–627. doi: 10.1086/525608.

Froggatt, P. and N.C. Nevin. 1971. *The 'law of ancestral heredity' and the Mendelian-ancestrian controversy in England, 1889-1906*. **Journal of Medical Genetics**, 8 (1):1–36. doi: 10.1136/jmg.8.1.1.

Galton, F. 1872. *On blood-relationship*. **Proceedings of the Royal Society of London**, 20: 394–402. doi: 10.1098/rspl.1871.0075.

———. 1877a. *Typical laws of heredity. II*. **Nature**, 15(389):512–514. doi: 10.1038/015512b0.

———. 1877b. *Typical laws of heredity. III*. **Nature**, 15(390):532–533. doi: 10.1038/015532a0.

———. 1879. *Composite portraits, made by combining those of many different persons into a single resultant figure*. **Journal of the Anthropological Institute of Great Britain and Ireland**, 8:132–144. doi: 10.2307/2841021.

———. 1886. *Regression towards mediocrity in hereditary stature*. **Journal of the Anthropological Institute of Great Britain and Ireland**, 15:246–263. doi: 10.2307/2841583.

———. 1889. *Natural inheritance*. Macmillan, London.

———. 1895. *Questions bearing on specific stability*. **Transactions of the Royal Entomological Society of London**, 43(1):155–157. doi: 10.1111/j.1365-2311.1895.tb01667.x.

Galton, F. and C. Darwin. 1859. *Darwin Correspondence Project Database. Letter 2573 – Galton, Francis to Darwin [9 Dec. 1859]*. URL: http://www.darwinproject.ac.uk/entry-2573.

Gauger, E.M., E. Rieper, J.J.L. Morton, S.C. Benjamin, and V. Vedral. 2011. *Sustained quantum coherence and entanglement in the avian compass*. **Physical Review Letters**, 106(4):040503. doi: 10.1103/PhysRevLett.106.040503.

Gayon, J. 1998. *Darwinism's struggle for survival: heredity and the hypothesis of natural selection*. Cambridge University Press, Cambridge.

———. 2007. *Karl Pearson: les enjeux du phenomenalisme dans les sciences biologiques*. In **Conceptions de la science: hier, aujourd'hui, demain: hommage à Marjorie Grene**, page 305–334. Ousia, Brussels.

Gilbert, S.F., J.M. Opitz, and R.A. Raff. 1996. *Resynthesizing evolutionary and developmental biology*. **Developmental Biology**, 173(2):357–372. doi: 10.1006/dbio.1996.0032.

Gildenhuys, P. 2009. *An explication of the causal dimension of drift*. **British Journal for the Philosophy of Science**, 60(3):521–555. doi: 10.1093/bjps/axp019.

Gillespie, J.H. 1977. *Natural selection for variances in offspring numbers: a new evolutionary principle*. **American Naturalist**, 111(981):1010–1014. doi: 10.1086/283230.

Glymour, B. 2001. *Selection, indeterminism, and evolutionary theory*. **Philosophy of Science**, 68(4):518–535. doi: 10.1086/392940.

Glynn, L. 2010. *Deterministic chance*. **British Journal for the Philosophy of Science**, 61(1): 51–80. doi: 10.1093/bjps/axp020.

Godfrey-Smith, P. 2009. *Darwinian populations and natural selection*. Oxford University Press, Oxford.

Grant, A. 1997. *Selection pressures on vital rates in density-dependent populations*. **Proceedings of the Royal Society of London B: Biological Sciences**, 264(1380):303–306. doi: 10.1098/rspb.1997.0043.

Graves, L., B.L. Horan, and A. Rosenberg. 1999. *Is indeterminism the source of the statistical character of evolutionary theory?* **Philosophy of Science**, 66(1):140–157. doi: 10.1086/392680.

Hacking, I. 1990. *The taming of chance*. Cambridge University Press, Cambridge.

Harman, G. 1973. *Thought*. Princeton University Press, Princeton, NJ.

Hilts, V.L. 1978. *Aliis exterendum, or, the origins of the Statistical Society of London*. **Isis**, 69 (1):21–43. doi: 10.1086/351931.

Hitchcock, C. and J. Woodward. 2003. *Explanatory generalizations, part II: Plumbing explanatory depth*. **Noûs**, 37(2):181–199. doi: 10.1111/1468-0068.00435.

Hodge, M.J.S. 1977. *The structure and strategy of Darwin's 'long argument'*. **British Journal for the History of Science**, 10(3):237–246. doi: 10.1017/S0007087400015685.

———. 1985. *Darwin as a lifelong generation theorist*. In Kohn, D., editor, **The Darwinian heritage: a centennial retrospect**, page 207–243. Princeton University Press, Princeton, NJ.

———. 1987. *Natural selection as a causal, empirical, and probabilistic theory*. In Krüger, L., G. Gigerenzer, and M.S. Morgan, editors, **The probabilistic revolution, volume 2: ideas in the sciences**, page 233–270. Bradford Books, Cambridge, MA.

———. 1989. *Darwin's theory and Darwin's argument*. In Ruse, M., editor, **What the philosophy of biology is: Essays for David Hull**, page 136–182. Kluwer Academic Publishers, Dordrecht.

———. 1992a. *Darwin's argument in the Origin*. **Philosophy of Science**, 59(3):461–464. doi: 10.1086/289682.

———. 1992b. *Biology and philosophy (including ideology): a study of Fisher and Wright*. In Sarkar, S., editor, **The founders of evolutionary genetics**, page 231–293. Kluwer Academic Publishers, Dordrecht.

———. 2000. *Knowing about evolution: Darwin and his theory of natural selection*. In Creath, R. and J. Maienschein, editors, **Biology and epistemology**, page 27–47. Cambridge University Press, Cambridge.

Hodge, M.J.S. and G. Radick, editors. 2009. *The Cambridge companion to Darwin*. Cambridge University Press, Cambridge, 2nd edition.

Hull, D.L. 1973. *Darwin and his critics: The reception of Darwin's theory of evolution by the scientific community*. Harvard University Press, Cambridge, MA.

———. 2009. *Darwin's science and Victorian philosophy of science*. In Hodge, M.J.S. and G. Radick, editors, **The Cambridge companion to Darwin**, page 173–196. Cambridge University Press, Cambridge, 2nd edition.

Igo, S.E. 2007. *The averaged American: surveys, citizens, and the making of a mass public*. Harvard University Press, Cambridge, MA.

Ismael, J. 2009. *Probability in deterministic physics*. **Journal of Philosophy**, 106(2):89–108.

———. 2011. *A modest proposal about chance*. **Journal of Philosophy**, 58(8):416–442.

Jablonka, E. and G. Raz. 2009. *Transgenerational epigenetic inheritance: prevalence, mechanisms, and implications for the study of heredity and evolution*. **Quarterly Review of Biology**, 84(2):131–176. doi: 10.1086/598822.

Jean, W.H. and B.P. Helms. 1983. *Geometric mean approximations*. **Journal of Financial and Quantitative Analysis**, 18(3):287–293. doi: 10.2307/2330720.

Jeffrey, R.C. 2004. *Subjective probability: the real thing*. Cambridge University Press, Cambridge.

Jenkin, F. 1867. *[Review of] The Origin of Species*. **North British Review**, 46:277–318.

Jost, J. 2003. *On the notion of fitness, or: the selfish ancestor*. **Theory in Biosciences**, 121(4): 331–350. doi: 10.1078/1431-7613-00065.

Keiding, N. 1975. *Extinction and exponential growth in random environments*. **Theoretical Population Biology**, 8(1):49–63. doi: 10.1016/0040-5809(75)90038-6.

Kevles, D.J. 1985. *In the name of eugenics: genetics and the uses of human heredity*. University of California Press, Berkeley.

Kim, K. 1994. *Explaining scientific consensus: the case of Mendelian Genetics*. The Guilford Press, New York.

Kolmogorov, A.N. and V.A. Uspenskii. 1987. *Algorithms and randomness*. **Theory of Probability and its Applications**, 32(3):389–412. doi: 10.1137/1132060.

Krakauer, D.C., K.M. Page, and D.H. Erwin. 2009. *Diversity, dilemmas, and monopolies of niche construction*. **American Naturalist**, 173(1):26–40. doi: 10.1086/593707.

Krimbas, C.B. 2004. *On fitness*. **Biology and Philosophy**, 19(2):185–203. doi: 10.1023/B:BIPH.0000024402.80835.a7.

Laland, K.N., J. Odling-Smee, and S. Myles. 2010. *How culture shaped the human genome: bringing genetics and the human sciences together*. **Nature Reviews Genetics**, 11(2): 137–148. doi: 10.1038/nrg2734.

Lankester, E.R. 1896a. *The utility of specific characters [letter of Aug. 20, 1896]*. **Nature**, 54 (1399):365–366. doi: 10.1038/054365b0.

———. 1896b. *Are specific characters useful? [letter of Jul. 16, 1896]*. **Nature**, 54(1394): 245–246. doi: 10.1038/054245c0.

Laughlin, R.B., D. Pines, J. Schmalian, B.P. Stojkovic, and P. Wolynes. 2000. *The middle way*. **Proceedings of the National Academy of Sciences**, 97(1):32–7.

Lee, M.S.Y. and P. Doughty. 2003. *The geometric meaning of macroevolution*. **Trends in Ecology & Evolution**, 18(6):263–266. doi: 10.1016/S0169-5347(03)00103-4.

Lennox, J.G. 2005. *Darwin's methodological evolution*. **Journal of the History of Biology**, 38:85–99. doi: 10.1007/s10739-004-6511-4.

Lenormand, T., D. Roze, and F. Rousset. 2009. *Stochasticity in evolution*. **Trends in Ecology & Evolution**, 24(3):157–165. doi: 10.1016/j.tree.2008.09.014.

Lester, J. 1995. *E. Ray Lankester and the making of modern British biology*. Bowler, P.J., editor. British Society for the History of Science, Oxford.

Lewens, T. 2009. *The Origin and philosophy*. In Ruse, M. and R.J. Richards, editors, **The Cambridge companion to the "Origin of Species"**, page 314–332. Cambridge University Press, Cambridge.

———. 2010. *The natures of selection*. **British Journal for the Philosophy of Science**, 61 (2):313–333. doi: 10.1093/bjps/axp041.

Lewontin, R.C. 1970. *The units of selection*. **Annual Review of Ecology and Systematics**, 1 (1):1–18. doi: 10.1146/annurev.es.01.110170.000245.

Lewontin, R.C. and D. Cohen. 1969. *On population growth in a randomly varying environment*. **Proceedings of the National Academy of Sciences**, 62(4):1056–1060. doi: 10.1073/pnas.62.4.1056.

Machamer, P., L. Darden, and C.F. Craver. 2000. *Thinking about mechanisms*. **Philosophy of Science**, 67(1):1–25. doi: 10.1086/392759.

MacKenzie, D.A. 1978. *Statistical theory and social interests: a case-study*. **Social Studies of Science**, 8(1):35–83. doi: 10.1177/030631277800800102.

———. 1979. *Karl Pearson and the professional middle class*. **Annals of Science**, 36(2): 125–143. doi: 10.1080/00033797900200441.

———. 1981. *Statistics in Britain, 1865–1930: the social construction of scientific knowledge*. Edinburgh University Press, Edinburgh.

Maeda, K., K.B. Henbest, F. Cintolesi, I. Kuprov, C.T. Rodgers, P.A. Liddell, D. Gust, C.R. Timmel, and P.J. Hore. 2008. *Chemical compass model of avian magnetoreception*. **Nature**, 453(7193):387–390. doi: 10.1038/nature06834.

Magnello, M.E. 1998. *Karl Pearson's mathematization of inheritance: from ancestral heredity to Mendelian genetics*. **Annals of Science**, 55(1):35–94. doi: 10.1080/00033799800200111.

———. 1999a. *The non-correlation of biometrics and eugenics: rival forms of laboratory work in Karl Pearson's career at University College London, part 1*. **History of Science**, 37: 79–106.

———. 1999b. *The non-correlation of biometrics and eugenics: rival forms of laboratory work in Karl Pearson's career at University College London, part 2*. **History of Science**, 37: 123–150.

———. 2004. *The reception of Mendelism by the biometricians and the early Mendelians (1899–1909)*. In Keynes, M., A.W.F. Edwards, and R. Peel, editors, **A century of Mendelism in human genetics**, page 17–30. CRC Press, Boca Raton, FL.

Manier, E. 1978. *The young Darwin and his cultural circle*. D. Riedel Publishing Company, Dordrecht.

Marsaglia, G. 1995. *The Diehard battery of tests for randomness*. URL: http://www.stat.fsu.edu/pub/diehard/.

Martin-Löf, P. 1966. *The definition of random sequences*. **Information and Control**, 9: 602–619. doi: 10.1016/S0019-9958(66)80018-9.

———. 1969. *Algorithms and randomness*. **Review of the International Statistical Institute**, 37(3):265–272. doi: 10.2307/1402117.

Matthen, M. 2009. *Drift and "statistically abstractive explanation"*. **Philosophy of Science**, 76(4):464–487. doi: 10.1086/648063.

———. unpublished. *Commentary on Grant Ramsey, 'Drift as an individual-level process,' delivered at APA Pacific 2011*.

Matthen, M. and A. Ariew. 2002. *Two ways of thinking about fitness and natural selection*. **Journal of Philosophy**, 99(2):55–83.

———. 2009. *Selection and causation*. **Philosophy of Science**, 76(2):201–224. doi: 10.1086/648102.

McKitrick, J. 2003. *A case for extrinsic dispositions*. **Australasian Journal of Philosophy**, 81 (2):155–174. doi: 10.1080/713659629.

McShea, D.W. and R.N. Brandon. 2010. *Biology's first law: the tendency for diversity and complexity to increase in evolutionary systems*. University of Chicago Press, Chicago.

Mellor, D.H. 1971. *The matter of chance*. Cambridge University Press, Cambridge.

Merlin, F. 2010. *Evolutionary chance mutation: a defense of the Modern Synthesis' consensus view*. **Philosophy and Theory in Biology**, 2:e103.

Metz, J.A.J., R.M. Nisbet, and S.A.H. Geritz. 1992. *How should we define 'fitness' for general ecological scenarios?* **Trends in Ecology & Evolution**, 7(6):198–202. doi: 10.1016/0169-5347(92)90073-K.

Mills, S.K. and J.H. Beatty. 1979. *The propensity interpretation of fitness*. **Philosophy of Science**, 46(2):263–286. doi: 10.1086/288865.

Millstein, R.L. 2000. *Chance and macroevolution*. **Philosophy of Science**, 67:603–624. doi: 10.1086/392857.

———. 2002. *Are random drift and natural selection conceptually distinct?* **Biology and Philosophy**, 17:33–53. doi: 10.1023/A:1012990800358.

———. 2005. *Selection vs. drift: a response to Brandon's reply*. **Biology and Philosophy**, 20 (1):171–175. doi: 10.1007/s10539-004-6047-1.

———. 2006a. *Natural selection as a population-level causal process*. **British Journal for the Philosophy of Science**, 57(4):627–653. doi: 10.1093/bjps/axl025.

———. 2006b. *Discussion of "Four case studies on chance in evolution": philosophical themes and questions*. **Philosophy of Science**, 73(5):678–687. doi: 10.1086/518522.

———. 2008. *Distinguishing drift and selection empirically: "The Great Snail Debate" of the 1950s*. **Journal of the History of Biology**, 41(2):339–367. doi: 10.1007/s10739-007-9145-5.

———. 2011. *Chances and causes in evolutionary biology: How many chances become one chance*. In Illari, P.M., F. Russo, and J. Williamson, editors, **Causality in the sciences**, page 425–444. Oxford University Press, Oxford.

———. 2013. *Natural selection and causal productivity*. In Chao, H., S. Chen, and R.L. Millstein, editors, **Mechanism and causality in biology and economics**. Springer, New York.

Millstein, R.L., R.A. Skipper, and M.R. Dietrich. 2009. *(Mis)interpreting mathematical models: drift as a physical process*. **Philosophy and Theory in Biology**, 1:e002.

Morrison, M. 2002. *Modelling populations: Pearson and Fisher on Mendelism and biometry*. **British Journal for the Philosophy of Science**, 53(1):39–68. doi: 10.1093/bjps/53.1.39.

Mudge, G.P. 1935. *J. T. Cunningham, M.A. [obituary]*. **British Medical Journal**, 2(3887):42. doi: 10.1136/bmj.2.3887.42.

Mumford, S. 2003. *Dispositions*. Oxford University Press, Oxford.

Murray, B.G. 1990. *Population dynamics, genetic change, and the measurement of fitness*. **Oikos**, 59(2):189–199. doi: 10.2307/3545534.

Murray, B.G. and L. Gårding. 1984. *On the meaning of parameter x of Lotka's discrete equations*. **Oikos**, 42(3):323–326. doi: 10.2307/3544400.

Nelson, E. 1959. *Regular probability measures on function space*. **Annals of Mathematics**, 69 (3):630–643. doi: 10.2307/1970027.

Northcott, R. 2010. *Walsh on causes and evolution*. **Philosophy of Science**, 77(3):457–467. doi: 10.1086/652965.

Norton, B.J. 1978. *Karl Pearson and statistics: the social origins of scientific innovation*. **Social Studies of Science**, 8(1):3–34. doi: 10.1177/030631277800800101.

Odling-Smee, F.J., K.N. Laland, and M.W. Feldman. 2003. *Niche construction: the neglected process in evolution*. Princeton University Press, Princeton, NJ.

Okasha, S. 2006. *Evolution and the levels of selection*. Clarendon Press, Oxford.

Olby, R. 1987. *William Bateson's introduction of Mendelism to England: a reassessment*. **British Journal for the History of Science**, 20(4):399–420. doi: 10.1017/S0007087400024201.

———. 1989. *The dimensions of scientific controversy: the biometric-Mendelian debate*. **British Journal for the History of Science**, 22(3):299–320. doi: 10.1017/S0007087400026170.

Olson, M.E. 2012. *The developmental renaissance in adaptationism*. **Trends in Ecology & Evolution**, 27(5):278–287. doi: 10.1016/j.tree.2011.12.005.

Otsuka, J., T. Turner, C. Allen, and E.A. Lloyd. 2011. *Why the causal view of fitness survives*. **Philosophy of Science**, 78(2):209–224. doi: 10.1086/659219.

Pearson, E.S. 1936. *Karl Pearson: an appreciation of some aspects of his life and work. Part I: 1857–1906*. **Biometrika**, 28(3/4):193–257. doi: 10.2307/2333951.

Pearson, K. 1892. *The grammar of science*. Walter Scott, London, 1st edition.

———. 1894a. *Dilettantism and statistics: [Review of] The Growth of St. Louis Children*. **Nature**, 51(1311):145–146. doi: 10.1038/051145a0.

———. 1894b. *Contributions to the mathematical theory of evolution*. **Philosophical Transactions of the Royal Society of London A**, 185:71–110. doi: 10.1098/rsta.1894.0003.

———. 1895. *Contributions to the mathematical theory of evolution. II. Skew variation in homogeneous material*. **Philosophical Transactions of the Royal Society of London A**, 186:343–414. doi: 10.1098/rsta.1895.0010.

———. 1896a. *The utility of specific characters [letter of Sep. 17, 1896]*. **Nature**, 54(1403): 460–461. doi: 10.1038/054460c0.

———. 1896b. *Mathematical contributions to the theory of evolution. III. Regression, heredity, and panmixia*. **Philosophical Transactions of the Royal Society of London A**, 187: 253–318. doi: 10.1098/rsta.1896.0007.

———. 1898. *Mathematical contributions to the theory of evolution. On the law of ancestral heredity*. **Proceedings of the Royal Society of London**, 62:386–412. doi: 10.1098/rspl.1897.0128.

———. 1900. *The grammar of science*. Adam and Charles Black, London, 2nd edition.

———. 1904. *Mathematical contributions to the theory of evolution. XII. On a generalized theory of alternative inheritance, with special reference to Mendel's laws*. **Philosophical Transactions of the Royal Society of London A**, 203:53–86. doi: 10.1098/rsta.1904.0015.

———. 1906. *Walter Frank Raphael Weldon. 1860–1906*. **Biometrika**, 5(1/2):1–52. doi: 10.1093/biomet/5.1-2.1.

———. 1908. *On a mathematical theory of determinantal inheritance, from suggestions and notes of the late W. F. R. Weldon*. **Biometrika**, 6(1):80–93. doi: 10.1093/biomet/6.1.80.

———. 1914. *The life, letters, and labours of Francis Galton, volume 1: birth 1822 to marriage 1853*. Cambridge University Press, Cambridge.

———. 1924. *The life, letters, and labours of Francis Galton, volume 2: researches of middle life*. Cambridge University Press, Cambridge.

———. 1930. *The life, letters, and labours of Francis Galton, volume 3A: correlation, personal identification, and eugenics*. Cambridge University Press, Cambridge.

Pearson, K., A. Lee, and L. Bramley-Moore. 1899. *Mathematical contributions to the theory of evolution. VI. Genetic (reproductive) selection: Inheritance of fertility in man, and of fecundity in thoroughbred racehorses*. **Philosophical Transactions of the Royal Society of London A**, 192:257–330. doi: 10.1098/rsta.1899.0006.

Pence, C.H. 2011. *"Describing our whole experience": The statistical philosophies of W. F. R. Weldon and Karl Pearson*. **Studies in History and Philosophy of Biological and Biomedical Sciences**, 42(4):475–485. doi: 10.1016/j.shpsc.2011.07.011.

———. forthcoming. *The early history of chance in evolution*. **Studies in History and Philosophy of Science**.

Pence, C.H. and G. Ramsey. 2013. *A new foundation for the propensity interpretation of fitness*. **British Journal for the Philosophy of Science**, 64(4):851–881. doi: 10.1093/bjps/axs037.

Pfeifer, J. 2005. *Why selection and drift might be distinct*. **Philosophy of Science**, 72(5): 1135–1145. doi: 10.1086/508122.

Pigliucci, M. 2009. *An extended synthesis for evolutionary biology*. **Annals of the New York Academy of Sciences**, 1168:218–228. doi: 10.1111/j.1749-6632.2009.04578.x.

———. 2010. *Okasha's evolution and the levels of selection: toward a broader conception of theoretical biology*. **Biology and Philosophy**, 25(3):405–415. doi: 10.1007/s10539-010-9197-3.

Pigliucci, M. and J.M. Kaplan. 2006. *Making sense of evolution: the conceptual foundations of evolutionary theory*. University of Chicago Press, Chicago.

Plutynski, A. 2006. *What was Fisher's fundamental theorem of natural selection and what was it for?* **Studies in History and Philosophy of Biological and Biomedical Sciences**, 37(1):59–82. doi: 10.1016/j.shpsc.2005.12.004.

Popper, K. 1959. *The propensity interpretation of probability*. **British Journal for the Philosophy of Science**, 10(37):25–42. doi: 10.1093/bjps/X.37.25.

Porter, T.M. 1986. *The rise of statistical thinking, 1820–1900*. Princeton University Press, Princeton, NJ.

———. 2004. *Karl Pearson: the scientific life in a statistical age*. Princeton University Press, Princeton, NJ.

———. 2005. *The biometric sense of heredity: statistics, pangenesis and positivism*. In Müller-Wille, S. and H. Rheinberger, editors, **A cultural history of heredity III: 19th and early 20th centuries**, page 31–42. Max Planck Institute for the History of Science, Berlin.

Provine, W.B. 1971. *The origins of theoretical population genetics*. Princeton University Press, Princeton, NJ.

Radick, G. 2005. *Other histories, other biologies*. **Royal Institute of Philosophy Supplement**, 56(3-4):21–47. doi: 10.1017/S135824610505602X.

———. 2011. *Physics in the Galtonian sciences of heredity*. **Studies in History and Philosophy of Biological and Biomedical Sciences**, 42(2):129–138. doi: 10.1016/j.shpsc.2010.11.019.

Ramsey, G. 2006. *Block fitness*. **Studies in History and Philosophy of Biological and Biomedical Sciences**, 37(3):484–498. doi: 10.1016/j.shpsc.2006.06.009.

———. 2013a. *Driftability*. **Synthese**, 190(17):3909–3928. doi: 10.1007/s11229-012-0232-6.

———. 2013b. *Can fitness differences be a cause of evolution?* **Philosophy and Theory in Biology**, 5:e401.

———. 2013c. *Organisms, traits, and population subdivisions: two arguments against the causal conception of fitness?* **British Journal for the Philosophy of Science**, 64: 589–608. doi: 10.1093/bjps/axs010.

Rand, D.A., H.B. Wilson, and J.M. McGlade. 1994. *Dynamics and evolution: evolutionarily stable attractors, invasion exponents and phenotype dynamics*. **Philosophical Transactions of the Royal Society of London B: Biological Sciences**, 343(1035):261–283. doi: 10.1098/rstb.1994.0025.

Reisman, K. and P. Forber. 2005. *Manipulation and the causes of evolution*. **Philosophy of Science**, 72:1113–1123. doi: 10.1086/508120.

Richards, R.J. 1992. *The meaning of evolution: the morphological construction and ideological reconstruction of Darwin's theory*. University of Chicago Press, Chicago.

———. 2009. *Darwin's theory of natural selection and its moral purpose*. In Ruse, M. and R.J. Richards, editors, **The Cambridge companion to the "Origin of Species"**, page 47–66. Cambridge University Press, Cambridge.

Richardson, R.C. 2006. *Chance and the patterns of drift: a natural experiment*. **Philosophy of Science**, 73:642–654. doi: 10.1086/518519.

———. 2007. *[Review of] Re-engineering philosophy for limited beings: piecewise approximations to reality*. **Notre Dame Philosophical Reviews**.

Roll-Hansen, N. 1980. *The controversy between biometricians and Mendelians: A test-case for the sociology of scientific knowledge*. **Social Science Information**, 19(3):501–517. doi: 10.1177/053901848001900302.

———. 2000. *Theory and practice: the impact of Mendelism on agriculture*. **Comptes Rendus de l'Académie des Sciences: Series III: Sciences de la Vie**, 323:1107–1116. doi: 10.1016/S0764-4469(00)01259-2.

Rosenberg, A. 2001. *How is biological explanation possible?* **British Journal for the Philosophy of Science**, 52:735–760. doi: 10.1093/bjps/52.4.735.

Rosenthal, J. 2010. *The natural-range conception of probability*. In Ernst, G. and A. Hüttemann, editors, **Time, chance, and reduction: Philosophical aspects of statistical mechanics**, page 71–90. Cambridge University Press, Cambridge.

Rutherford, H.W. 1908. *Catalogue of the library of Charles Darwin now in the Botany School, Cambridge*. Cambridge University Press, Cambridge.

Salmon, W.C. 1984. *Scientific explanation and the causal structure of the world*. Princeton University Press, Princeton, NJ.

———. 1994. *Causality without counterfactuals*. **Philosophy of Science**, 61(2):297–312. doi: 10.1086/289801.

Seneta, E. 1981. *Non-negative matrices and Markov chains*. Springer, New York.

Shapiro, L. and E. Sober. 2007. *Epiphenomenalism – the dos and the don'ts*. In Wolters, G. and P. Machamer, editors, **Thinking about causes: from Greek philosophy to modern physics**, page 235–264. University of Pittsburgh Press, Pittsburgh, PA.

Shea, N., I. Pen, and T. Uller. 2011. *Three epigenetic information channels and their different roles in evolution*. **Journal of Evolutionary Biology**, 24(6):1178–1187. doi: 10.1111/j.1420-9101.2011.02235.x.

Simons, A.M. 2002. *The continuity of microevolution and macroevolution*. **Journal of Evolutionary Biology**, 15(5):688–701. doi: 10.1046/j.1420-9101.2002.00437.x.

Sloan, P.R. 2000. *Mach's phenomenalism and the British reception of Mendelism*. **Comptes Rendus de l'Académie des Sciences: Series III: Sciences de la Vie**, 323:1069–1079. doi: 10.1016/S0764-4469(00)01255-5.

Sober, E. 1980. *Evolution, population thinking, and essentialism*. **Philosophy of Science**, 47 (3):350–383. doi: 10.1086/288942.

———. 1984. *The nature of selection*. The MIT Press, Cambridge, MA.

———. 2001. *The two faces of fitness*. In Singh, R.S., editor, **Thinking about evolution: historical, philosophical, and political perspectives**, page 309–321. The MIT Press, Cambridge, MA.

———. 2010. *Evolutionary theory and the reality of macro probabilities*. In Eells, E. and J.H. Fetzer, editors, **The place of probability in science: in honor of Ellery Eells (1953–2006)**, number 284 in Boston Studies in the Philosophy of Science, page 133–162. Springer, Dordrecht.

Stamos, D.N. 2001. *Quantum indeterminism and evolutionary biology*. **Philosophy of Science**, 68(2):164–184. doi: 10.1086/392871.

Stephens, C. 2004. *Selection, drift, and the "forces" of evolution*. **Philosophy of Science**, 71 (4):550–570. doi: 10.1086/423751.

Sterrett, S.G. 2002. *Darwin's analogy between artificial and natural selection: how does it go?* **Studies in History and Philosophy of Biological and Biomedical Sciences**, 33(1): 151–168. doi: 10.1016/S1369-8486(01)00039-5.

Stigler, S.M. 1986. *The history of statistics: the measurement of uncertainty before 1900*. Belknap, Cambridge, MA.

Strevens, M. 2011. *Probability out of determinism*. In Beisbart, C. and S. Hartmann, editors, **Probabilities in physics**, page 339–364. Oxford University Press, Oxford.

———. 2013. *Tychomancy: inferring probability from causal structure*. Harvard University Press, Cambridge, MA.

Stöltzner, M. 2008. *The causality debates of the interwar years and their preconditions: revisiting the Forman thesis from a broader perspective*. In Joas, C., C. Lehner, and J. Renn, editors, **HQ-1: conference on the history of quantum physics**, page 113–126. Max Planck Institute for the History of Science, Berlin.

Tabery, J.G. 2004. *The "evolutionary synthesis" of George Udny Yule*. **Journal of the History of Biology**, 37(1):73–101. doi: 10.1023/B:HIST.0000020390.75208.ac.

Thiele, J. 1969. *Karl Pearson, Ernst Mach, John B. Stallo: Briefe aus den Jahren 1897 bis 1904*. **Isis**, 60(4):535–542. doi: 10.1086/350540.

Thiselton-Dyer, W.T. 1895. *Variation and specific stability [letter of Mar. 14, 1895]*. **Nature**, 51 (1324):459–461. doi: 10.1038/051459c0.

Thoday, J.M. 1953. *Components of fitness*. **Symposia of the Society for Experimental Biology**, 7:96–113.

Tuljapurkar, S. 1989. *An uncertain life: demography in random environments*. **Theoretical Population Biology**, 35:227–294. doi: 10.1016/0040-5809(89)90001-4.

———. 1990. *Population dynamics in variable environments*. Springer-Verlag, New York.

Tuljapurkar, S.D. and S.H. Orzack. 1980. *Population dynamics in variable environments I. Long-run growth rates and extinction*. **Theoretical Population Biology**, 18(3): 314–342. doi: 10.1016/0040-5809(80)90057-X.

Walsh, D.M. 2000. *Chasing shadows: natural selection and adaptation*. **Studies in History and Philosophy of Biological and Biomedical Sciences**, 31(1):135–153. doi: 10.1016/S1369-8486(99)00041-2.

———. 2003. *Fit and diversity: explaining adaptive evolution*. **Philosophy of Science**, 70 (2):280–301. doi: 10.1086/375468.

———. 2004. *Bookkeeping or metaphysics? The units of selection debate*. **Synthese**, 138(3): 337–361. doi: 10.1023/B:SYNT.0000016426.73707.92.

———. 2007. *The pomp of superfluous causes: the interpretation of evolutionary theory*. **Philosophy of Science**, 74(3):281–303. doi: 10.1086/520777.

———. 2010. *Not a sure thing: fitness, probability, and causation*. **Philosophy of Science**, 77(2):147–171. doi: 10.1086/651320.

———. 2013. *Descriptions and models: some responses to Abrams*. **Studies in History and Philosophy of Biological and Biomedical Sciences**, 44:302–308. doi: 10.1016/j.shpsc.2013.06.004.

Walsh, D.M., T. Lewens, and A. Ariew. 2002. *The trials of life: natural selection and random drift*. **Philosophy of Science**, 69(3):429–446. doi: 10.1086/342454.

Waters, C.K. 2009. *The arguments in the Origin of Species*. In Hodge, M.J.S. and G. Radick, editors, **The Cambridge companion to Darwin**, page 120–143. Cambridge University Press, Cambridge, 2nd edition.

———. 2011. *Okasha's unintended argument for toolbox theorizing*. **Philosophy and Phenomenological Research**, 82(1):232–240. doi: 10.1111/j.1933-1592.2010.00472.x.

Weldon, W.F.R. 1890. *The variations occurring in certain decapod Crustacea. — I. Cragnon vulgaris*. **Proceedings of the Royal Society of London**, 47:445–453. doi: 10.1098/rspl.1889.0105.

———. 1893. *On certain correlated variations in Carcinus maenas*. **Proceedings of the Royal Society of London**, 54:318–329. doi: 10.1098/rspl.1893.0078.

———. 1895a. *Variation in animals and plants*. **Nature**, 51(1323):449–450. doi: 10.1038/051449a0.

———. 1895b. *An attempt to measure the death-rate due to the selective destruction of Carcinus maenas with respect to a particular dimension*. **Proceedings of the Royal Society of London**, 57:360–379. doi: 10.1098/rspl.1894.0165.

———. 1896a. *The utility of specific characters [letter of Jul. 30, 1896]*. **Nature**, 54(1396): 294–295. doi: 10.1038/054294a0.

———. 1896b. *Utility of specific characters [letter of Sep. 3, 1896]*. **Nature**, 54(1401):413–413. doi: 10.1038/054413a0.

———. 1898. *Address of the President of Section D (zoology)*. **Report of the British Association for the Advancement of Science**, 68:887–902.

———. 1902. *Mendel's laws of alternative inheritance in peas*. **Biometrika**, 1(2):228–254. doi: 10.1093/biomet/1.2.228.

———. 1906. *Inheritance in animals and plants*. In Strong, T.B., editor, **Lectures on the method of science**, page 81–109. Clarendon Press, Oxford.

Weldon, W.F.R., K. Pearson, and C.B. Davenport. 1901. *Editorial: the scope of Biometrika*. **Biometrika**, 1(1):1–2. doi: 10.1093/biomet/1.1.1.

West-Eberhard, M.J. 2005. *Developmental plasticity and the origin of species differences*. **Proceedings of the National Academy of Sciences**, 102(2):6543–6549. doi: 10.1073/pnas.0501844102.

Wimsatt, W.C. 2007. *Re-engineering philosophy for limited beings: piecewise approximations to reality*. Harvard University Press, Cambridge, MA.

Woodward, J. 2000. *Explanation and invariance in the special sciences*. **British Journal for the Philosophy of Science**, 51(2):197–254. doi: 10.1093/bjps/51.2.197.

———. 2003. *Making things happen*. Oxford University Press, Oxford.

———. 2010. *Causation in biology: stability, specificity, and the choice of levels of explanation*. **Biology and Philosophy**, 25(3):287–318. doi: 10.1007/s10539-010-9200-z.

Woodward, J. and C. Hitchcock. 2003. *Explanatory generalizations, part I: A counterfactual account*. **Noûs**, 37(1):1–24. doi: 10.1111/1468-0068.00426.