ABSTRACT

Title of Dissertation: MECHANISM AND CHANCE: TOWARD AN ACCOUNT OF STOCHASTIC MECHANISM FOR THE LIFE SCIENCES

Lane Thomas DesAutels, Doctor of Philosophy, 2014

Dissertation directed by: Professor Lindley Darden
Department of Philosophy

In this dissertation, my aim is to develop some important new resources for explaining probabilistic phenomena in the life sciences. In short, *I undertake to articulate and defend a novel account of stochastic mechanism for grounding probabilistic generalizations in the life sciences*.

To do this, I first offer some brief remarks on the concept of mechanism in the history of philosophical thought. I then lay out some examples of probabilistic phenomena in biology for which an account of stochastic mechanism seems explanatorily necessary and useful: synaptic transmission in the brain, protein synthesis, DNA replication, evolution by natural selection, and Mendelian inheritance. Next, I carefully examine the concept of regularity as it applies to mechanisms—building on a recent taxonomy of the ways mechanisms may (or may not) be thought to behave regularly. I then employ this taxonomy to sort out a recent debate in the philosophy of biology: is natural selection regular enough to count as a mechanism? I argue that, by paying attention to the forgoing taxonomy, natural selection can be seen to meet the regularity requirement just fine. I then turn my attention to the question of how we should understand the

chance we ascribe to stochastic mechanisms. To do this, I form a list of desiderata that any account of stochastic mechanism must meet. I then explore how mechanisms fit with several of the going philosophical accounts of chance: subjectivism, frequentism (both actual and hypothetical), Lewisian best-systems, and propensity. I argue that neither subjectivism, frequentism, nor best-system-style accounts of chance will meet all of the proposed desiderata, but some version of propensity theory can. Borrowing from recent propensity accounts of biological fitness and drift, I then go on to explore the prospects for developing a propensity interpretation of stochastic mechanism (PrISM) according to which propensities are (i) metaphysically analyzable and operationally quantifiable via a function of probability-weighted ways a mechanism might fire and (ii) not causally efficacious but nonetheless explanatorily useful. By appealing to recent analyses of deterministic and emergent chance, I argue further that this analysis need not be vulnerable to the threat of metaphysical determinism.

MECHANISM AND CHANCE: TOWARD AN ACCOUNT OF STOCHASTIC MECHANSIM FOR THE LIFE SCIENCE

By

Lane DesAutels

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park, in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2014

Advisory Committee:

Professor Lindley Darden (Chair)
Professor Peter Carruthers
Professor Mathias Frisch
Associate Professor Jessica Pfeifer
Assistant Professor Aidan Lyon

Table of Contents

Chapter 1. Introduction

Abstract: The two main goals of this introductory Chapter are (1) to supply some motivation for why my reader should care about this dissertation, and (2) to provide some historical, conceptual, and lexicographical framework for what is to follow.

1.1 A Brief Introduction to the idea of Mechanism in the History of Philosophical Thought

For quite a long time, the heart beat was deeply puzzling. It wasn't until William Harvey (1628) discovered that the beating heart produces a continuous circulation of blood through the interconnected vasculature at the extremities of the body that an explanation of it could be given. Harvey discovered that the beating heart was a central part of the *mechanism* for blood circulation in animals. French researchers Francois Jacob and Jacques Monod discovered messenger RNA (1961): the missing link between DNA and protein; they found a key part of the protein synthesis *mechanism*. In order to explain puzzling phenomena in the living world, life scientists often search to find and describe underlying mechanisms.

Recently, much work in the philosophy of science has been devoted to understanding what exactly it is that scientists look for when they search for mechanisms. One now widely accepted philosophical characterization of mechanism was put forward in Machamer, Darden, and Craver's seminal paper "Thinking about Mechanisms" (MDC 2000).

> MDC: Mechanisms are entities and activities organized such that they are productive of regular changes from start or set-up to finish or termination conditions. (MDC 2000, 3)[1]

On this characterization of mechanism, the beating heart clearly qualifies as a mechanism. It is composed of entities (aorta, ventricles, arteries, and so on) and activities (beating, pumping etc..) that are organized to produce regular changes (blood circulation) from the start of an animal's life to its end.

[1] Similar characterizations have been put forward by Glennan (1996, 2002) and Bechtel & Abrahamsen (2005), but MDC's characterization has received the most attention in the literature—so theirs is the one on which I focus my discussion.

As elegant and straight-forward as it seems, however, the MDC characterization of mechanism raises some difficult questions. One such question has to do with how to appeal to mechanisms, a concept traditionally associated with regular, machine-like, deterministic behavior, to explain *probabilistic* phenomena. This question becomes especially important once it is recognized that the life sciences are rife with probabilistic generalizations. Mendel discovered that the chance of a hybrid green and yellow pea plant to produce yellow peas in the (what we now call the) F_2 generation is .75. In neuroscience, the release of neurotransmitters can fail to result in the successful initiation of electrical activity in a particular postsynaptic neuron up to 90% of the time[2]. Evolution by natural selection is subject to the whims of genetic mutation—where the evolutionary consequences of genetic mutation are conceptualized in terms of the chance (per unit of time) a gene has of changing from one state to another. A question of significant import to philosophers of science is: what makes these statements true? What in the world, if anything, grounds these probabilistic facts?

In this dissertation, my aim is to develop some important new resources for explaining probabilistic phenomena in the life sciences. In short, *I undertake to articulate and defend a novel account of stochastic mechanism for grounding probabilistic generalizations in the life sciences*.

Before beginning this difficult task, some context and framework needs building. The job of this introductory chapter is to motivate this project by providing some background-giving historical context to the discussion of mechanism—a concept that once played a dominant role in our explanations of the natural world, but one that fell dramatically out of favor in the past century. I then briefly apprise the reader of the recent resurgence of mechanisms in recent philosophy of science, some of the reasons why mechanisms have resurfaced in contemporary

[2] Kandel et al. 2013, 270

scientific thought, and the way that several leading advocates of mechanistic philosophy of science have explicitly formulated what mechanisms are and how they function in scientific explanations. I then begin to draw attention to an area of this recent mechanisms literature that needs more development: how are we to appeal to mechanisms, a concept commonly associated with regular, deterministic operation and outputs, to explain probabilistic phenomena in the life sciences? To show why this question is important, I formulate a few examples of probabilistic phenomena of interest to proponents of mechanistic explanation in biology. Having these examples on hand will help as I refer back to them in order to show (1) that extant accounts of mechanism are under-equipped to explain these probabilistic phenomena, and (2) that the account of stochastic mechanism I develop can. This chapter concludes by laying out the central task of this dissertation, some key terms and distinctions on which I will rely, as well as the goals and thesis of the dissertation as a whole.

1.2 A Brief History of Mechanism

Why should you, or indeed *anyone*, read a dissertation about mechanisms? Hasn't the concept of mechanism fallen away like Aristotle's hylomorphic soul or Ptolemy's crystalline spheres? What use could such an old-fashioned atavism of early thought possibly be to us now?

To begin to answer these questions, it is worth briefly looking back at the varied role of the idea of mechanism in the history of philosophical thought—if only to distinguish a contemporary understanding of mechanism from the many ways it has been understood throughout history.

At first thought, the word 'mechanism' seems easy enough to define. A mechanism, many of us would imagine, is simply a structure of moving parts that perform some function. A clock is a mechanism; a tractor is a mechanism; so is particle accelerator. In the history of

philosophy, however, the term ‘mechanism’ came to represent a great deal more than this simple notion.

As with most (if not all) important philosophical concepts, the origin of ‘mechanism’ traces back to the ancient Greeks. The Greeks’ idea of mechanism derived from the machines they created to do work, work that they conceived of as being opposed to natural forces (Bechtel 2006, 42). Nature, on the ancient Greek view, did not operate mechanistically. This is apparent in the ancient text *Mechanica*:

> Nature often operates contrary to human interest, for she always follows the same course without deviation, whereas human interest is always changing. When, therefore, we have to do something contrary to nature, the difficulty of it causes us perplexity and art has to be called to our aid. The kind of art which helps us in such perplexities we call Mechanical Skill. (Mechanica 847a14f as quoted by Whewell 1837, 94)

This view of mechanisms began to change with Aristotle. Rather than mechanisms operating in the natural world, Aristotelians envisaged nature as composed of entities made up of distinct forms—each with a specific *telos* (goal). Natural entities, on the Aristotelian (and later scholastic) worldview, are of certain substantial forms each with their own directed behaviors. Rocks and other Earthly entities fall downward because their natural goal is to be part of the Earth. Fire rises up because its goal is to reach the heavens. On the Aristotelian/Scholastic worldview, objects move almost as if they *intend* to reach their natural goal.

This worldview began to change dramatically in the 17th century at which time the teleological Aristotelian/Scholastic worldview ultimately gave way to the *mechanical worldview*. The distinguishing feature of the mechanical worldview, according to Richard Westfall, was to conceive of nature as “a huge machine” such that the role of natural philosophy was to “explain the hidden mechanisms behind phenomena” (Westfall 1971, 1). The beginning of this change is exemplified by the development of Galileo’s mechanics. According to Galileo, the same

mechanics explain mundane Earthly bodies as apply to celestial bodies like the sun and moon. With his three laws of falling bodies, Galileo single-handedly succeeded in undermining the Aristotelian theory of motion[3]. And it was this idea—that the same few laws of motion succeed in explaining the movements of earthly and celestial bodies alike—that ushered in the new mechanistic worldview, a worldview according to which the universe is composed of a collection of moving parts all working together under the same basic forces, a worldview that saw the working parts of the universe like the gears of a clock.

The early-modern mechanistic worldview reached its culmination with Descartes. With the Scholastics squarely in aim, Descartes undertook to replace the "mentally" influenced depiction of physical qualities in Scholastic natural philosophy with a theory that requires only extensional properties (e.g., size, shape, and motion) to describe the manifest order of the natural world. He writes, "I have described this earth and indeed the whole universe as if it were a machine: I have considered only the various shapes and movements of its parts" (Cottingham et al. 1985, 188), On Descartes's mechanical view, the movement of one object necessarily moved another, because he did not believe there to be any empty space. Like a jar full of tiny ball bearings, there is no moving one without the movement of many others. Descartes not only conceived of the universe in mechanistic terms, he also conceived of living things this way. He writes,

> I should like you to consider that these functions (including passion, memory, and imagination) follow from the mere arrangement of the machine's organs every bit as naturally as the movements of a clock or other automaton follow from the arrangement of its counter-weights and wheels." (ibid, 108)

Modulo only the disembodied soul, the function of which for Descartes was surprisingly limited, human bodies were fully mechanical arrangements of interconnected parts.

[3] Cf., Sharratt 1994, p. 203

A generation later, Boyle coined the term "mechanical philosophy" (Boyle 1666)[4]. On Boyle's view the explanation of natural phenomena should make reference only to particles of matter (which he called *corpuscles*), their motion, and their interaction under the governance of a few basic laws of nature. Because of its insistence that reference to particles of matter, their motion, and their interaction was sufficient to fully explain any given natural phenomenon, the mechanistic worldview became associated with metaphysical determinism. Just as the working gears of a clock determine that it must strike 3:00 an hour after striking 2:00, all phenomena in the mechanistic universe must be completely determined, past, present or future. As Laplace famously claimed,

> We may regard the present state of the universe as the effect of the past and the cause of the future. An intellect which at any given moment knew all of the forces that animate nature and the mutual positions of the beings that compose it, if this intellect were vast enough to submit the data to analysis, could condense into a single formula the movement of the greatest bodies of the universe and that of the lightest atom; for such an intellect nothing could be uncertain and the future just like the past would be present before its eyes. (Laplace 1814)

The mechanistic worldview dominated the early-modern period of philosophy and science. It was not until the beginning of the 20th century that the mechanistic worldview began to fall out of favor, and with it, the predominant position of mechanisms in philosophical thought. The reason for its demise was, in large part, a growing suspicion that matter—at least on its smallest scale—does not behave like the gears of a clock.

The beginning of the end of the classical mechanistic worldview came with the development of the theory of quantum mechanics. Early quantum mechanics discovered that, while elementary particles behave fairly predictably in many experiments, they become highly unpredictable in others—such as when attempting to measure individual particle trajectories

[4] Although, it is worth pointing out that there are many historiographical disputes as to its precise origins. See Garber (2013, 3-4) for an excellent discussion of this.

through a simple physical apparatus. As a result of this discovery, new models were needed to explain this behavior, models that violated several of the central tenets of classical, deterministic physics. Nature no longer seemed to fit into a clock-like, mechanical worldview. Instead, scientists and philosophers had to come to terms with the fact that nature had to be described in terms that were deeply and essentially probabilistic.

1.3 The Return of Mechanistic Explanation in Biology

> "At least in biology, most scientists see their work as explaining types of phenomena by discovering mechanisms, rather than explaining theories by deriving them from or reducing them to other theories, and *this* is seen by them as reduction, or as integrally tied to it." (Wimsatt 1976, 671)

Why, after its spectacular fall from grace at the hands of the quantum mechanics, have mechanisms returned to favor among philosophers of science interested in biology? A brief look at the history of scientific explanation will help to supply an answer to this.

The modern discussion of explanation begins in earnest with the Deductive-Nomological (D-N) model developed by Hempel (1942, 1965) and advocated by many others (Popper 1959, Braithwaite 1953, Gardiner 1959, Nagel 1961). On the D-N model of explanation, what is required for a description of a phenomenon (an explanandum) to be explained by a set of sentences (an explanans) is: 1) the explanandum must follow deductively from the explanans; 2) the sentences that constitute the explanans must be true; and 3) the explanans must contain at least one "law of nature".

There are many well-known problems with the D-N model of explanation. For one thing, there are intuitive cases that show that the satisfaction of the D-N model is neither *required* for a good explanation, nor does it *guarantee* one. Suppose, for example, that I witness my dog knock a cup off the coffee table in a fit of exuberant tail-wagging. There is a sense in which the

explanation of why this event occurred does not require reference to the initial position of the cup, and the relevant laws of nature. A simple, "…because my dog hit it with her tail" would suffice. Furthermore, there seem to be instances when explanations seem bad even when the D-N model is satisfied. Woodward describes a particularly devastating one (owing originally to Bromberger 1966):

> There are many cases in which a derivation of an explanandum *E* from a law *L* and initial conditions *I* seems explanatory but a "backward" derivation of *I* from *E* and the same law *L* does not seem explanatory, even though the latter, like the former, appears to meet the criteria for successful *DN* explanation. For example, one can derive the length *s* of the shadow cast by a flagpole from the height *h* of the pole and the angle θ of the sun above the horizon and laws about the rectilinear propagation of light. This derivation meets the *DN* criteria and seems explanatory. On the other hand, a derivation of *h* from *s* and θ and the same laws also meets the *DN* criteria but does not seem explanatory. Examples like this suggest that at least some explanations possess directional or asymmetric features to which the *DN* model is insensitive. (Woodward 2009)

What this example illustrates is that there is a deeply problematic *asymmetry* to D-N explanation. The D-N model appears to explain when we derive the length of the shadow from the height of the pole and the position of the sun, but it seems not to when we derive the height of the pole from the length of the shadow and the position of the sun—this despite the fact that both instances meet the requirements set forth by the D-N model.

Many philosophers of science have taken this explanatory asymmetry as an indication that the D-N model leaves something important out of scientific explanation. Salmon argues forcefully that the missing ingredient is *causation*. The reason we find that the height of the pole explains the shadow length (and not vice versa) is because the pole-height plays an important *causal role* in the length of the shadow it casts and not the other way around. As such, Salmon (1989) famously concludes that to explain, at least in the context of science, necessarily involves locating an event in "a causal nexus".

Asymmetry problems aside, D-N explanation does not seem to cohere well with the special sciences for additional reasons. The main reason is its requirement that all explanations make reference to at least one *law of nature*. Here is an argument, implicitly supported in the literature, for why the generalizations made in the special sciences *should not* be construed as laws:

> (1) Generalizations made in the special sciences are contingent.
>
> (2) Whatever else a law is, it must be more than contingently true.
>
> (3) Therefore, generalizations made in the special sciences should not be construed as laws (Smart 1963; Schiffer 1991; Beatty 1995; Earman and Roberts 1999; Woodward 2000, 2001, 2002, and 2003).

Some have disagreed with this conclusion, maintaining that these differences merely show that we need to have two categories of laws: *ceteris paribus* laws and *strict laws*, where the special sciences should only be seen to have the former (Fodor 1991; Hausman 1992; Pietroski and Rey 1995; Lange 2000, 2002). Still other philosophers have dismissed *ceteris paribus* formulations of laws as deeply problematic, but have nevertheless maintained that certain of the generalizations made in the special sciences should count as laws. Adherents to this last sort of view have either proposed that these laws need not be more than contingently true (Mitchell 1997, 2000, 2002, 2003), or they have attempted to reformulate their generalizations so that they are not contingent at all (Sober 1997, Elgin 2006). The result of all of this debating about the prospects of laws in the special sciences is that there appears to be no real consensus on the matter.

Perhaps biologists don't need laws *of biology* in order to do D-N explanation. Perhaps biologists can do all the explaining they need by appeal, say, to the fundamental laws of physics together with initial conditions and do their explanatory derivation that way. Unfortunately, there are well known problems here too. For one, there are some who argue that there are no

fundamental laws of physics (e.g., Cartwright 1983). And second, among those who allow for fundamental physical laws, there is very little agreement among philosophers as to what is required of a generalization for it to qualify as a law of nature. Do they need to be necessary relations between universals (a la Armstrong 1983)? Are laws no more than regularly occurring conjunctions of events (a la Hume 1739)? Do laws of nature have to be exceptionless? Support counterfactuals? Do laws get confirmed by their own instances (a la Goodman 1947)? Each of these positions has its proponents, and its critics. And little progress, if any, has been made in reaching agreement.

Even if a philosophical consensus on laws of nature *could* be accomplished, there is a remaining difficulty with D-N explanation: that it simply doesn't deliver satisfying answers to many of our explanation-seeking questions. This is because D-N explanation, at its base, is a kind of *subsumption.* That is, it takes an explanandum event and purports to explain it by showing that it can be subsumed under some law of nature. But, if we stop and think about it, there is an important sense in which this kind of explanatory subsumption strategy does not deliver a satisfactory explanation. When I ask a why some event *x* occurred, should I be satisfied by the explanation that *x* occurred because it is the sort of event that gets subsumed under a class of events that we already know always occurs? Put another way, there is a sense in which D-N explanation merely tells us that *x* occurred because *x has to occur*. But that doesn't so much tell me *why* x occurred as that *x* is the sort of event that some law (whatever that is) tells us must occur.

At the very least—given all of this consternation—*it would be nice if we could do explanation in the special sciences without laws.* That is precisely what mechanisms portend to do.

So what is mechanistic explanation? In short, the idea is this:

> ME: Some description of a phenomenon (explanandum) is explained by some set of sentences (explanans), just in case the explanans is a description of the mechanism that produces this phenomenon.

Because 'produces' in this context is a causal notion, I take mechanistic explanation to be a kind of causal explanation. What is a causal explanation? I follow Woodward's view that the distinguishing feature of causal explanations is that "They show how what is explained depends on other, distinct factors, where the dependence in question has to do with some relationship that holds as a matter of empirical fact rather than for logical or conceptual reasons" (Woodward 2003, 4-5). Mechanistic explanation does just this. It explains some phenomenon by showing how it depends on some other, distinct factor: the *mechanism* responsible for its production. So a mechanistic explanation is a causal explanation—but one of a certain sort. And not only does ME satisfy Salmon's requirement that good scientific explanations be causal, it does so without any appeal to laws of nature. That is a good thing.

Furthermore, the manifest fact is that life scientists, as a matter of actual practice, have done (and continue to do) incredibly successful science by searching for and describing mechanisms. A keyword search for articles appearing in the journal *Nature* (and its subsidiary journals) turned up 3,380 articles with the word 'mechanism' in their title published between 2000-2014. Machamer, Darden, and Craver's 2000 article "Thinking about Mechanisms" is the most cited paper published in the journal *Philosophy of Science* over the last three years. Since its publication, its total current number of citations is 1062.

1.4 Taking stock

This brief glance at the some of the history of the role that the idea of mechanism has played in the history of philosophy and explanation serves to draw attention to a fundamental tension

which will motivate the rest of the dissertation. On the one hand, the mechanical worldview has been associated with the thesis that any phenomenon can be entirely explained by reference to its material constituents, their motion, and some basic physical laws. However, on the other hand, the post-mechanical worldview is committed to the thesis that some phenomena are fundamentally random, and their behavior cannot be explained merely by reference to the arrangement of their material constituents, their motion, and some basic physical laws. My hope in what follows is to begin to forge a middle path, a path on which we can maintain the benefits of understanding, at least some of, the living world in mechanistic terms while at the same time admitting that much of the world we seek to explain behaves in a probabilistic fashion. Before tackling this project, however, a bit more framework is needed.

1.5 Current Characterizations of Mechanisms

In section 1.3, we saw some of the history of mechanistic thought as well as some of the reasons for its resurgence in the life sciences. Now we can begin to talk more specifically about how the idea of mechanism is understood in contemporary philosophy of science. There are several current philosophical characterizations of mechanism. Because much of the discussion that follows will draw on various aspects of these characterizations, it will be helpful to list them at the outset. They are as follows.

> *Glennan 1*: A mechanism underlying a behavior is a complex system which produces that behavior by the interaction of a number of parts according to direct causal laws (Glennan 1996, 52).
>
> *Glennan 2*: A mechanism for a behavior is a complex system that produces that behavior by the interaction of a number of parts, where the interactions among parts can be characterized by direct, invariant, change relating generalizations. (Glennan 2002, S344)
>
> MDC: Mechanisms are entities and activities organized such that they are productive of regular changes from start or set-up to finish or termination conditions (MDC 2000, 3).

> *B&A*: A mechanism is a structure performing a function in virtue of its components parts, component operations, and their organization. The orchestrated functioning of the mechanism is responsible for one or more phenomena. (Bechtel and Abrahamsen 2005, 47).
>
> *Cartwright*: A nomological machine is a stable enough arrangement of components whose features acting in consort give rise to (relatively) stable input/output relations (Cartwright 2009, p. 8).

As is evident by the forgoing list, there are many prominent philosophers of science interested in how we should philosophically characterize mechanisms. However, as we will see in the coming Chapters, there are many interesting questions that arise from these attempted characterizations.

1.6 Some Mechanism Lexicography

As with any area of intellectual area of discourse, the philosophical literature on mechanisms carries with it its own jargon—some of it somewhat idiosyncratic. As such, it will also be helpful to offer a brief lexicon of some of the central terms at play in the mechanisms literature.

> *Activities*: what goes on in a mechanism, or as MDC describes them, activities are the producers of change (e.g., *binding* between neurotransmitters or *bonding* between a DNA base and its complementary base) (MDC, 3)
>
> *Entities*: the things in a mechanism that engage in activities (e.g., neurotransmitters, strands of DNA) (Ibid, 3). MDC's entities, for our purposes, may be thought of as analogous to Glennan's and B&A's notion of *parts.*
>
> *Set-up conditions*: background conditions that, when satisfied, allow for the possibility of a mechanism's start-up conditions to obtain.
>
> *Start-up conditions*: conditions that, when satisfied, initiate the operation of a mechanism (a.k.a. its firing).
>
> *Triggering*: that which occurs upon the instantiation of a mechanism's start-up conditions.
>
> *Firing*: a general way of saying the mechanism has begun operation. (may also refer to the whole operation span of a mechanism).
>
> *Termination conditions*: conditions that, when satisfied, indicate that the mechanism is finished operating or has achieved its function.

Behavior: may refer to that which entities do internal to the mechanism (in which case it is analogous to MDC's activities), or it may refer to the outcomes produced by a mechanism (in which case it is analogous to MDC's 'phenomenon').

Function: may refer to roles played by entities and activities or that which the mechanism as a whole is set up to do.

Phenomenon: the observable stuff in the world that a mechanism description purports to explain.

1.7 Some Probability Lexicography

Since my primary topic will be to understand the mechanistic explanation of probabilistic phenomena in biology by appeal to, what I will call, 'stochastic mechanisms', we need to have some understanding at the outset of what is meant by *probabilistic phenomena* and *stochasticity* as well as some of the concepts closely surrounding it. So more lexicography is needed. Many of the terms defined below will get more careful definitions later in the dissertation. But the ones I offer here, I hope, will suffice for getting us off and running.

Time: a set of points (call them T points) ordered linearly.

History: a temporal path through some state space.

Deterministic History: a history that, for any given T point, has only one possible continuation.[5]

Indeterministic History: a history that, for at least one of its T points, has more than one possible continuation.

Deterministic Phenomenon: an observable state of affairs with a deterministic history, i.e., a state of affairs whose preconditions together with the laws of nature were sufficient to bring about its occurrence.

Metaphysical Global Determinism: the view that the actual world is entirely a deterministic one.

[5] Here, I am happy to understand possibility in terms of possible worlds or in terms of nomological possibility in the actual world.

Deterministic World: A deterministic world is one whose entire history supervenes on the world's laws of nature with the complete state of the world at any given time (Earman 1986).

Metaphysical Indeterminism: the view that at least some of the world's histories have been non-deterministic.

Indeterministic Phenomenon: an observable state of affairs which has a non-deterministic history.

Probabilities: a measure of the likelihood that an event will occur.

Probabilistic Phenomenon: an observable state of affairs that is best described by appeal to [non-extremal] probabilities. (note: this may be because it is indeterministic, but it may also be that it is deterministic and too complex to be described as such).

Probabilistic Behavior: the activity of some observable state of affairs that is best described by appeal to [non-extremal] probabilities.

Stochasticity: the property instantiated by a state of affairs by virtue of the fact that it behaves probabilistically.

Objective Chance: an understanding of probabilities that meets a set of commonly accepted platitudes regarding its relationship to several related concepts: credence, possibility, future, lawfulness, intrinsicness, and lawfulness (to be laid out in detail in Chapter 6)

Probabilistic Statement: a sentence that makes reference to a probability.

Having now specified preliminary definitions for several of philosophy of probability-related terms that will arise in many of our discussions to come, I should pause to point out a few troubling ambiguities.

Sometimes non-philosophers use the above terminology in ways that are different from what I have just stipulated. For example, biologists will often use the term 'stochastic' to be synonymous with 'unpredictable'. For reasons I'll give in Chapter 4, I do not want to use the term this way. 'Stochastic' also gets used a lot to describe a particular kind of scientific model, e.g., one that does not make determinate predictions. Stochastic models of this sort are, no doubt, interesting and useful in science. But again, this is not how I will be using the term. For me,

stochasticity is a metaphysical property, one that is instantiated by actual states of affairs (e.g., mechanisms). So when I use the term 'stochastic mechanism' I will be referring to actual, existing structures that, for various reasons to be specified in the following Chapter, are disposed to behave irregularly.

But before we get to the project of trying to develop a more detailed understanding of the ways in which a mechanism might be thought to behave irregularly, it will be helpful for me to lay out some of the primary scientific examples to which I'll be referring in the upcoming Chapters.

1.8 Examples of Probabilistic Phenomena Targeted for Mechanistic Explanation

Now that much of the lexicographical framework for both our upcoming discussion of mechanisms and probability has been laid out, we can begin to explore some science. In what follows, I offer brief descriptions of five phenomena from across the life sciences for which biologists seek mechanistic explanation. Each one is probabilistic. Although these phenomena are not meant to represent all of biology, they exhibit sufficient range to proffer the hefty extent to which probability permeates across the life sciences—and as such—the explanatory potential for the novel explanatory framework I will be offering in what is to come.

1.8.1 *Initiation of Electrical Activity in Post-synaptic Neurons (Synaptic Transmission)*

There are gaps between the neurons in our brains called *synapses* (see figure 1 below). Electrical signals must traverse these gaps in order to continue on their path through our nervous system. This process begins when a brief pulse of electricity called an *action potential* travels down a hollow tube in the neuron called an *axon*. This occurs because the axon is filled with (and surrounded by) an aqueous solution containing charged ions. At rest, the membrane

surrounding each neuron is polarized; its inner surface is negatively charged relative to its outer surface. Upon contact with the action potential, however, the charge of the axon's membrane rises enough to open specific gates in the membrane so as to allow positive ions (mainly Ca^{2+} and Na^{+}) into the cell causing the membrane to depolarize. Other positively charged ions (mainly K^{+}) flow out of the membrane to cause repolarization in its wake. What results is a wave of electricity flowing down the axon until it reaches the synapse, at which point a chemical called a *neurotransmitter* is released and moves across the synapse binding to specific proteins on the neighboring neuron.

This phenomenon coheres nicely with the mechanistic approach to explanation because it is composed of entities (e.g., axons, neurons, synapses, neurotransmitters, etc.) and activities (e.g., polarization, electricity flow, etc.) organized to produce an outcome (transmission of electrical activity across a synapse). Here we can think of the set-up conditions of the synaptic transmission mechanism as the presence of the requisite anatomical and chemical structures in a living organism; we can think of the start-up conditions as the arrival of an action potential at the start of an axon; and we can think of the termination conditions as the successful initiation of post-synaptic electrical activity.

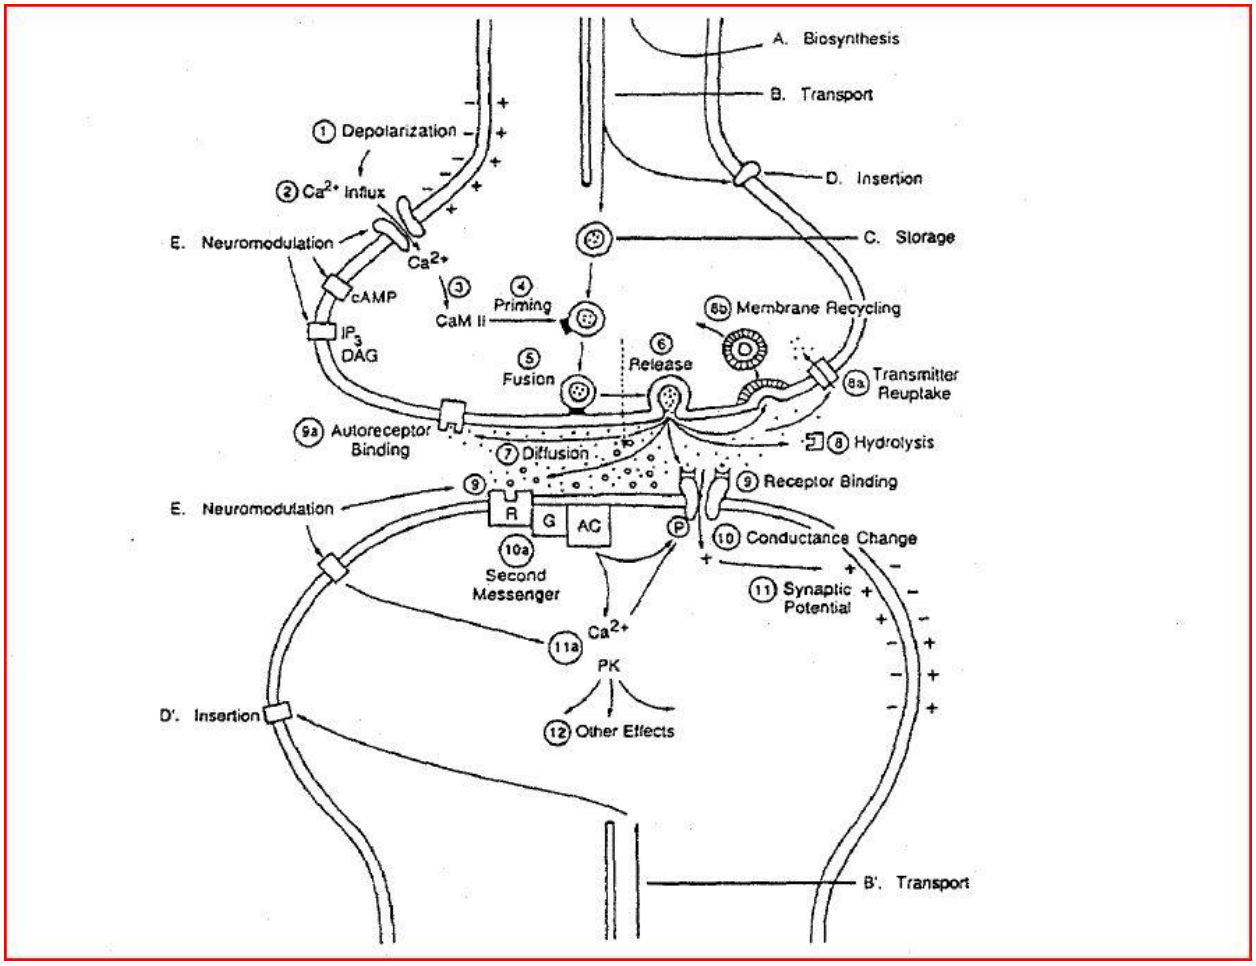

Figure 1. Biochemical mechanisms at chemical synapses. (From Machamer, Darden, and Craver 2000, p. 9).

This mechanism will be of particular interest to us in this dissertation because the process of synaptic transmission is not entirely dependable. In fact, according to Kandel et al.,

> The release of a quantum of transmitter is a random event. The fate of each quantum of transmitter in response to an action potential has only two possible outcomes—the quantum is or is not released. The event resembles a binomial or Bernoulli trial (similar to tossing a coin in the air to determine whether it comes up heads of tails). (Kandel et al. 2013, 270)

Kandel et al. estimate that the mean probability of synaptic transmission (from a single active zone), "varies widely among different presynaptic terminals, from less than 0.1 (that is, a 10% chance that a presynaptic action potential will trigger release of a vesicle) to greater than 0.9"

(ibid, 271). That is, the release of neurotransmitters can fail to result in the successful initiation of electrical activity in a particular postsynaptic neuron up to 90% of the time. Because synaptic transmission is (i) a process that fits within the mechanistic framework, and (ii) because scientists characterize it probabilistically, I will go on to argue that it is a prime target for analysis as a stochastic mechanism.

1.8.2 *Protein Synthesis*

In its most abstract schematization, protein synthesis takes place when DNA is transcribed into messenger RNA, which is translated into protein (see Figure 2).

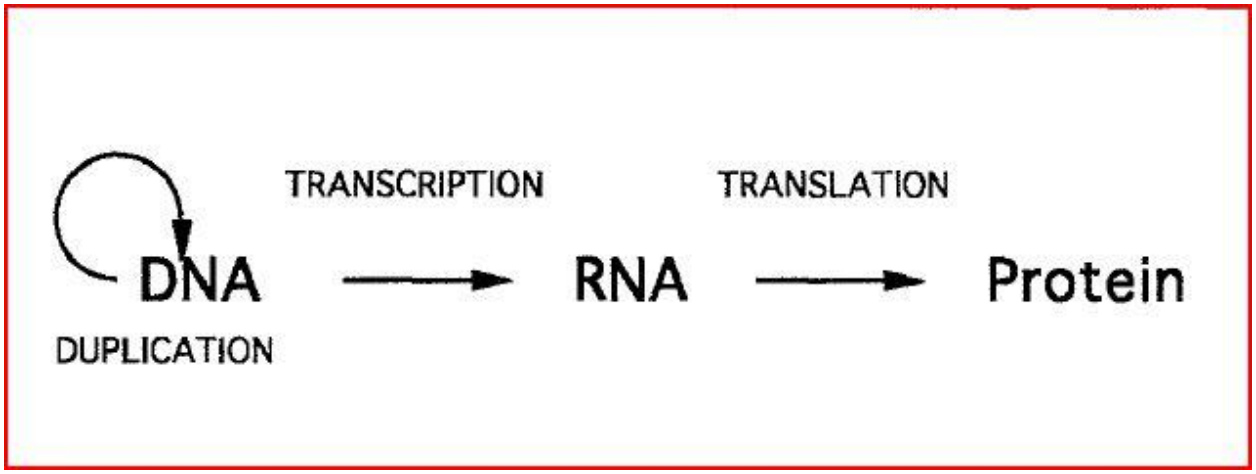

Figure 2. Watson's central dogma diagram (Redrawn, based on Watson 1965).

The process of protein synthesis also fits nicely within the mechanistic framework. It is composed of entities: DNA (deoxyribonucleic acid in the shape of a double helix), RNA polymerase (a core enzyme), and mRNA (ribonucleic acid, usually single stranded, which carries information to the protein synthesis machinery of the cell). And it has activities associated with it: *binding* of RNA polymerase to the DNA's promoter sequence, and *elongation* (when the RNA polymerase traverses the template DNA strand, using base pairing complementarity with the DNA template to create an RNA copy). Furthermore, protein synthesis has set-up conditions

(e.g., the presence of DNA in a living organism), start-up conditions (e.g., initiation of transcription), and termination conditions (e.g., the successful synthesis of a new protein).

Despite the universality of this process among living organisms, however, protein synthesis fails to operate exactly the same way in many circumstances. Different DNA sequences can produce different proteins. In bacteria, different RNA polymerases operate than in eukaryotes. Even among bacteria there is variation in how error-prone various different RNA polymerases are. Once again, this means that there is an important sense in which protein synthesis fails to work the same way each time it operates. As such, like synaptic transmission, it is best thought of as a probabilistic phenomenon (more on this in Chapter 2).

1.8.3 *Mutations During DNA Replication*

Deoxyribonucleic acid (DNA) comprises two polymers made up of units called a base and a backbone made of sugars and phosphate groups—all organized into the shape of a double helix. DNA replication begins when the parent molecule gets "unzipped" as the hydrogen bonds between the base pairs are broken. Once separated, the sequence of bases on each of the unzipped strands becomes a template for the insertion of a complementary set of bases. The new strands assemble in an order that complements the order of bases on the strand serving as the template. When the process is complete, two DNA molecules have been formed identical to each other and to the parent molecule. Here, again, we have entities: DNA, a purine base, individual molecules, etc. And we have activities: unzipping, separating, and inserting. We have start and set-up conditions: the presence of DNA in a living organism and the initiation of hydrogen bond separation. And we have termination conditions: the existence of a replicated strand of DNA.

However, just as with protein synthesis, there are several ways that the process of DNA replication can (and does) go wrong—resulting in mutation. One such way is when a base is

changed by the repositioning of a hydrogen atom, altering the hydrogen bonding pattern of that base resulting in incorrect base pairing during replication. Another way is when there is a loss of a purine base (A or G) to form an apurinic site (AP site). There can also be denaturation of the new strand from the template during replication, followed by renaturation in a different spot ("slipping"). This can lead to insertions or deletions.

With regard to the occurrence of the above sorts of replication errors, biologists use probabilities (a more detailed discussion of which will follow in Chapter 2) to describe them. So, yet again, we have a probabilistic phenomenon regarding which mechanistic explanation seems readily available.

1.8.4 *Natural Selection*

Nature tends to preserve those traits that afford their possessors the greater chance to survive and reproduce, and it tends to reject those that do not. The result is that species become increasingly matched to their respective environments; they become exquisitely adapted over time. In its most basic form, this is natural selection (NS).

Modern biologists (and philosophers of biology) give a somewhat more specific account of what is required for NS to occur. On one account, put forth by Skipper and Millstein (2005)[6], NS can be precisely formulated as follows:

I. Initial conditions
1. A population of *O*s exist.
2. *O*s vary according to forms of *T*, which are heritable.
3. *O*s are in environment *E* with critical factor *F*.

II. Interaction
1. *O*s in virtue of the varying forms of *T* interact differently with environment *E*.
2. Critical factor *F* affects that interaction.
3. This may lead to

[6] Heavily drawn from Skipper (1999), and Darden & Cain (1989)

III. Effects (1)
1. differential survival rates of *O*s across forms of *T* in *E*.
2. This may lead to

IV. Effects (2)
1. differential reproductive rates of *O*s across forms of *T* in *E*.
2. This may lead to

V. Effects (3)
1. differential representation in the population of *O*s across forms of *T* in *E*.

2. This may lead to

VI. Effects (4)
1. the predominance of *O*s with a certain form of *T* over other forms of *T* in *E*.
2. This may lead to

VII. Effects (5)
1. adaptation of the lineage with respect to *T* in *E*.

Filling instructions:
'*O*' is to be replaced by the name of some organism
'*T*' is to be replaced by some determinable organismic trait
'*E*' is to be replaced by the description of the environment of '*O*'
'*F*' is to be replaced by the description of a critical factor in '*E*'

(Fig. 1. From Skipper and Millstein 2005, p. 330)

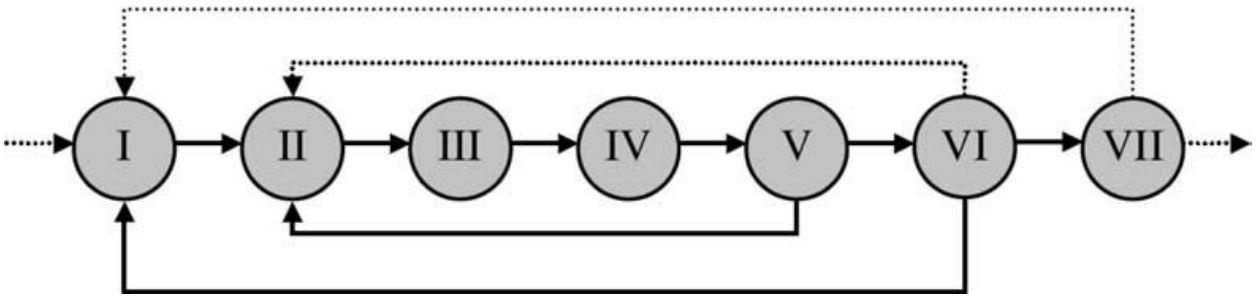

Figure 3. Natural selection schema (From Skipper and Millstein 2005, 331).

As we shall see in the coming chapters (esp. Chapter 3), natural selection is also a probabilistic process. This is, in part, because NS operates on genetic variation, and in evolutionary biology,

the evolutionary consequences of genetic mutation are conceptualized in terms of *the chance* (per unit of time) a gene has of changing from one state to another. But it is also in part due to the fact that natural selection is subject to many highly contingent environmental factors, factors which are too complex to predict.

As to whether NS should qualify as a mechanism, this is the subject of a heated debate in the literature: a debate I shall focus on in detail when we get to Chapter 3.

1.8.5 *Mendelian Segregation*

By breeding pea plants in his monastery garden, Gregor Mendel made invaluable strides towards understanding the mechanisms that rule heredity. When he crossed two pea plants with different traits—for example, a yellow pea producing plant with a green pea producing plant—he came to see a pattern in the results. In the first generation (which we now call the F_1 generation), all the peas were yellow. He, therefore, called the yellow trait the *dominant* one. But when he allowed the all-yellow peas to self fertilize in the subsequent generation (which we now call the F_2 generation), both yellow and green seeds appeared at a rate of near exactly 3:1. He, therefore, called the green trait the *recessive* one—since it had been hidden by the yellow ones in the previous generation.

Although Mendel himself lacked the modern resources for explaining this observed pattern, we now understand the mechanism underlying this phenomenon. Understanding it requires postulating the existence of paired material elements (later called "alleles")—one coming from the father and one coming from the mother—which segregate during meiosis. One of these alleles (capital 'A') is *dominant,* and the other (lowercase 'a') is *recessive*. In the pea plant example, these alleles combine in various ways such that some of the seeds have matching pairs of alleles (called homozygous) and some have mismatched pairs (called heterozygous).

This combination of alleles is called *genotype*, and its resulting character—yellowness or greenness—is called *phenotype.* To clearly visualize how alleles are distributed amongst descendents, we can use a Punnett Square:

	A	A
A	AA	Aa
a	aA	Aa

In the case of the green and yellow peas, the yellow character is dominant, so it is represented by the uppercase 'A', and the green character is recessive, so it is represented by the lowercase 'a'. The upper case 'A' always dominates over the lower case 'a', so the only time we get green peas is when there are two lowercase 'a's. In the initial generation, the yellow pea plant will give each of its progeny a dominant yellow allele, and the green pea plant will give each plant a recessive green allele. So all the pea plants in the F_1 generation will be yellow heterozygous. But in the F_2 generation, when the yellow heterozygous peas have interbred, the resulting progeny could have any of the combination of alleles (as illustrated by the Punnett Square)—showing a ratio of 3:1 yellow to green phenotypes. We now know that the mechanism works like this: 'A' and 'a' segregate in the formation of germ cells, which then combine randomly at fertilization to give offspring according to the formula: 1AA + 2Aa + 1aa, and this mechanism explains the 3:1 ratio that Mendel observed.

Importantly, however, the 3:1 ratio observed in the F_2 generation is not exact. Mendel's own results when breeding yellow and green peas in the F_2 generation were actually 6022 yellow: 2001 green (3.01:1). And when he bred round and wrinkled peas, he found 5474 round:

1851 wrinkled (2.96:1).[7] As such, Mendelian segregation is not fully deterministic, but rather must be considered *stochastic.* As with our previous examples, however, Mendelian segregation also seems to fit perfectly well into the mechanistic explanatory framework. It has entities: alleles, chromosomes, genes, etc. And it has activities: meiosis, interbreeding, segregation, etc. And it has intuitive start and termination conditions beginning with meiosis in the formation of germ cells and ending with the production of a new organism.

1.9 Anticipating the Central Task of this Dissertation

As we've just seen, many of the phenomena that biologists explain by appealing to mechanisms appear to behave probabilistically. But for something to be a mechanism, it is commonly thought to have to behave *regularly*. Depending on how we understand regularity, however, it seems difficult to understand how a regularly behaving mechanism could underlie and produce a probabilistic phenomenon. That is the central task I take on in what remains of this dissertation.

> In short, I aim to defend the following thesis:
>
> Dissertation Thesis: In this dissertation, I will argue that the best way to ground (at least some) of the probabilistic generalizations in the life sciences is by appeal to stochastic mechanisms.

Now let me set some goals I hope to achieve along the way.

1.10 Primary Goals of this Dissertation

In this dissertation, I hope to:

(I) successfully taxonomize all the ways a mechanism can be thought to behave regularly (and irregularly),

[7] Mendel, Gregor (1865)

(II) argue that regularity should not be seen to be a metaphysically demarcating feature of mechanisms—but that mechanisms with high degrees of regularity should be seen to hold an explanatorily privileged position,

(III) use this taxonomy of regularity to show how we can make sense of natural selection as an explanatorily privileged mechanism,

(IV) show that a version of a propensity theory is the best analysis of chance for underpinning an account of stochastic mechanism,

(V) explore the prospects for applying three important lessons from recent propensity interpretations of fitness and drift,

> (1) Propensities are aptly understood as having probability-weighted possibilia as their categorical base,
>
> (2) Propensities are quantifiable via a function of these probability-weighted possibilia,
>
> (3) If we can, we should avoid committing to a view which requires defending the causal efficacy of propensities (and dispositional properties in general),

(VI) use these lessons to specify a propensity interpretation of stochastic mechanism (PrISM),

(VII) explain how this novel PrISM analysis can handle the threat of metaphysical determinism.

1.11 What is to Come: Chapter Summaries

Chapter 2. Mechanisms and Regularity

> There is widespread disagreement in the mechanisms literature regarding how regularly a causal process needs to behave in order to qualify as a mechanism. In this Chapter, I explore a recent argument offered by Andersen (2012) in favor of placing a broadened regularity requirement on mechanisms. I argue, contra Andersen, that the reasons she gives for adopting this regularity requirement ultimately rest on a confusion between metaphysical, epistemological, and pragmatic considerations. Instead of requiring of all mechanisms that they behave regularly in order to *be* mechanisms, I'll suggest that

mechanisms come in degrees of regularity: some highly regular, some not very regular at all. In agreement with the spirit of Andersen's arguments, however, I'll concede that mechanisms with a high degree of regularity should hold a privileged position in scientific explanation.

Chapter 3. Is Natural Selection Regular Enough to be a Mechanism?

In this chapter, I apply several ideas developed in the previous Chapter to help sort out an active debate in the philosophy of biology: whether natural selection qualifies as an MDC mechanism. I argue that, by paying careful attention to some key distinctions drawn in the discussion of mechanistic regularity undertaken in Chapter 2 as well as some recent work on abstract mechanism schemas, natural selection can be seen to qualify as an MDC mechanism just fine—at least in the context of many legitimate explanatory contexts. More than a mere terminological dispute about what to call natural selection, I conclude by suggesting that this debate evinces a deeper point about the prospects for appealing to mechanisms to ground probabilistic generalizations in a contingent biological world.

Chapter 4. Stochastic Mechanisms and Theories of Chance

My goal in this Chapter is to further develop my characterization of stochastic mechanism by exploring how we should best understand the chanciness we attribute to them. To do this, I form a list of desiderata that any account of stochastic mechanism must meet. I then take the general characterization of mechanism offered by Machamer, Darden, and Craver (2000) and explore how it fits with several of the going philosophical accounts of chance: subjectivism, frequentism (both actual and hypothetical), Lewisian best-systems, and propensity. I argue that neither subjectivism, frequentism, nor a best-system-style account of chance will meet all of the proposed desiderata, but some version of propensity theory can. I conclude by showing the proposed account of propensity-backed stochastic mechanism has the added benefit of allowing us to escape many of the traditional objections to propensity theories of chance.

Chapter 5. A Propensity Interpretation of Stochastic Mechanism: Lessons from Fitness and Drift

The goal of this chapter is to introduce and motivate a propensity interpretation of stochastic mechanism (PrISM). I approach this, first, by examining some recent propensity interpretations of fitness and drift to see what lessons might be learned from them. I then lay out my own formal characterization of the propensity interpretation of stochastic mechanism and show how it enjoys an important advantage over propensity interpretations of fitness and drift. I conclude by hinting at how the PrISM might apply to one of the biological phenomena targeted for mechanistic explanation introduced in Chapter 1

Chapter 6. The Threat of Determinism: Synthesizing Emergent Chance and Multilevel Mechanisms

In this chapter, I explore whether—and if so, how—the propensity interpretation of stochastic mechanism (PrISM) developed thus far can handle the threat of metaphysical determinism. The answer I pursue is that it can. To show how, I first outline a few influential arguments for and against biological indeterminism and suggest that the only tenable philosophical conclusion available is an agnosticism as to whether there is genuine indeterminism at play in the biological world. As such, I pursue a disjunctivist strategy according to which what we say about the PrISM depends on which side of the biological determinism debate prevails. To accommodate the more difficult determinism disjunct, I draw on recent view from the philosophy of probability, emergent chance, according to which objective chance can exist at one level of description even if determinism holds at a different level. I then apply this idea to mechanism levels to show that genuine propensities may be able to exist at one mechanism level even if another mechanism level (within the same mechanism hierarchy) is characterized deterministically.

Chapter 7. Conclusion: Prospects for Future Research

In this concluding Chapter, I briefly summarize the main arguments offered in the dissertation; I revisit the central biological examples set forth in Chapter 1 in order to recapitulate some of the explanatory lessons afforded to them by my account; and I propose some avenues for future research uniquely afforded by this dissertation.

Chapter 2. Mechanisms and Regularity

Abstract: There is widespread disagreement in the mechanisms literature regarding how regularly a causal process needs to behave in order to qualify as a mechanism. In this Chapter, I explore a recent argument offered by Andersen (2012) in favor of placing a broadened regularity requirement on mechanisms. I argue, contra Andersen, that the reasons she gives for adopting this regularity requirement ultimately rest on a confusion between metaphysical, epistemological, and pragmatic considerations. Instead of requiring of all mechanisms that they behave regularly in order to *be* mechanisms, I'll suggest that mechanisms come in degrees of regularity: some highly regular, some not very regular at all. In agreement with the spirit of Andersen's arguments, however, I'll concede that mechanisms with a high degree of regularity should hold a privileged position in scientific explanation.

2.1 Introduction

The topic of this dissertation is to explore the prospects for appealing to mechanisms to ground probabilistic generalizations in the life sciences. As we saw in the previous Chapter, however, these prospects might seem dim at the outset because, throughout the history of philosophical thought, mechanisms have been associated with machine-like, deterministic behavior. But how could a deterministic machine produce results that are indeterministic, statistical, and/or probabilistic? As hinted in the introduction, my approach to answering this question will be to articulate and defend an account of mechanism that does not require mechanisms to behave like deterministic machines—an account of mechanism according to which (at least some) mechanisms operating in the natural world behave *stochastically*—an account of *stochastic mechanism for the life sciences*.

So far, however, nothing has been said about what these stochastic mechanisms are, what their defining characteristics might be, or how we should go about understanding them. By way of taking an initial step in this direction, in this Chapter I'll focus on one of the defining features of stochastic mechanisms. Whatever else they are, I suggest, stochastic mechanisms have the following characteristic: they fail to do exactly the same thing each time they operate. In other

words, stochastic mechanisms fail to operate entirely regularly. Indeed, this is the feature of stochastic mechanisms that distinguishes them from deterministic ones.

But why, you might wonder, should we begin by focusing on *this* aspect of stochastic mechanisms? Here is my reason. The failure of stochastic mechanisms to behave regularly seems to run up against a difficult problem. As we shall see, many philosophers have thought that a causal process *must* behave regularly in order to count as a mechanism at all. From the standpoint of these authors, it is a metaphysical prerequisite that mechanisms behave regularly—that they operate the same way (or very nearly the same way) each time they fire. If these authors are correct, the project of developing an account of stochastic mechanism seems it might be doomed from the start.

Given this seeming difficulty, the topic of this Chapter is to seek an answer to what I call *the regularity question*. The regularity question asks this: *to what extent is it a necessary feature of mechanisms that they work the same way (or near enough to the same way) each time they operate*? To explore this question, I'll look at the space of answers that have been given: those who argue that there are strong reasons for requiring that mechanisms behave highly regularly, and those who argue that mechanisms need not behave regularly at all. Because of its subtlety and rigor, I'll focus specifically on the recent defense of a broadened appeal to regularity offered by Holly Andersen (2012). Although Andersen's work has made invaluable strides towards laying the groundwork for addressing the regularity question, I'll disagree with many of her main conclusions. Specifically I will argue, contra Andersen, that the reasons she gives for adopting her broadened regularity requirement rest on a confusion between metaphysical, epistemological, and pragmatic considerations. Instead of requiring of all mechanisms that they behave regularly in order to *be* mechanisms, I'll suggest that mechanisms come in degrees of regularity: some

highly regular; some not very regular at all. In agreement with the spirit of Andersen's arguments, however, I'll suggest that mechanisms with a high degree of regularity should hold a privileged position in scientific explanation.

Here is my plan. In Section 2.2, I'll draw attention to a rift in the mechanisms literature regarding whether regularity should be included in philosophical characterizations of mechanism. In Section 2.3, I'll briefly outline Andersen's case for placing a broadened regularity requirement on mechanisms. In 2.4, I'll evaluate Andersen's case by: (i) highlighting and illustrating the positive attributes of her case, (ii) augmenting her taxonomy of mechanistic regularity with two considerations of my own, and (iii) offering some argument against the reasons to which she appeals to motivate her case for a regularity requirement. And In 2.5, I'll offer my own response to the regularity question.

2.2 A Rift on Regularity

Over the last two decades, a rift has emerged in the mechanisms literature. Some philosophers require that the component entities and activities of a mechanism behave in a regular fashion; some do not.

The most explicit appeal to regularity in the current literature appears in Machamer, Darden, and Craver's characterization of mechanism (henceforth, MDC). As we saw in the previous Chapter, MDC give the following characterization of mechanisms:

> MDC: Mechanisms are entities and activities organized such that they are productive of *regular* changes from start or set-up to finish or termination conditions (MDC 2000, 3. Italics added for emphasis).

For MDC, something is a mechanism just in case its entities engage in activities that produce *regular* changes from start to finish. They add, further, that "[M]echanisms are regular in that they work always or for the most part in the same way under the same conditions" (MDC 2000,

3). Similar appeals to regularity can be found in Darden 2008; Craver 2007; Craver and Darden 2013; Glennan 1996, 2002; and Bechtel and Abrahamson 2005. A careful analysis of what an appeal to regularity might mean in the context of mechanism is coming in the next section. But for now it will be helpful to gesture at what I take to be the main prima facie reason why we might want a characterization of mechanism with an appeal to regularity.

The main reason why it might be tempting to place a regularity requirement on processes we wish to deem as mechanisms is so that they can take over the roles that laws of nature once played in scientific explanation. That is, mechanisms, when understood to behave the same way (or near enough to the same way) each time they fire, can serve as the basis for making *scientific generalizations*, *predictions,* and *explanations*. Whereas, if no such regularity requirement is set forth, then they seem far less capable of taking over these roles. If mechanisms do not have to behave roughly the same way each time they operate, what basis do we have for generalizing about the phenomena that typically result from a certain kind of mechanism? What basis do we have for predicting the occurrence of a certain phenomenon that we believe will result from a certain kind of mechanism? And how can we explain general patterns of phenomena based on these mechanisms?

Despite the intuitive appeal of requiring mechanisms to behave regularly, there are some philosophers who argue against placing any kind of regularity requirement on mechanisms. James Bogen, (2005) concludes, "Mechanists need not include regularities and invariant generalizations in their account" (Bogen 2005, 399). Peter Machamer, influenced by Bogen, also drops his appeal to regularity. In a footnote (2004) Machamer writes, "I think 'regular' should be dropped from the definition. Jim Bogen argues forcefully that there might be mechanisms that operate only once in a while or even one that works only once" (Machamer 2004, 37). Stuart

Glennan, in a recent paper, offers another characterization of (certain type of) mechanism with no appeal to regularity. He calls them ‘ephemeral mechanisms’.

> I take an ephemeral mechanism to be a collection of interacting parts where:
>
> 1. the interactions between parts can be characterized by direct, invariant, change-relating generalizations
>
> 2. the configuration of parts may be the product of chance or exogenous factors
>
> 3. the configuration of parts is short-lived and non-stable, and is not an instance of a multiply-realized type. (Glennan 2010, 260)

For Glennan, ephemeral mechanisms are fleeting and short-lived; they only exist while the mechanism is operating and disappear once it is finished. This type of mechanism clearly makes no appeal to regularity at all—since it makes little sense to say of an ephemeral mechanism that it operates the same way each time it fires. By definition, it only ever fires once.

What reasons might these philosophers have for suggesting that the regularity constraint be dropped? The main reason is what Andersen calls *the argument from science.* In short, the science just doesn’t support a conception of mechanisms that requires them to behave the same way each time they operate. The fact is, opponents of the regularity requirement will point out, several (if not all) of the processes that scientists analyze mechanistically do *not seem to behave very regularly*. To support this claim, Bogen cites the transmission of electrical activity to postsynaptic neurons in the brain—a process that (as we saw in 1.8.1) fails up to 90% of time.

Much more detail is to follow on the synaptic transmission case in the next section. For now, let’s look at the general form of Andersen’s argument.

2.3 Andersen’s Case for a Broadened Appeal to Regularity

Now that we’ve seen that there is widespread disagreement in the mechanisms literature on whether mechanisms need to operate regularly, let us look at a particularly subtle and rigorous

case for placing a regularity requirement on mechanisms: that offered by Andersen (2012). I suggest Andersen's account is a good place to focus because, more than any of the other participants in this debate, (i) she makes explicit several plausible reasons why a regularity requirement seems attractive to philosophers interested in characterizing mechanisms, and (ii) she adeptly taxonomizes several of the many senses of regularity that are implicit in the literature. Although I'll ultimately disagree that the considerations Andersen offers constitute good reasons for placing a regularity requirement on mechanisms, I'll suggest that her taxonomy of mechanistic regularity lays an invaluable framework for addressing the regularity question—a framework I will borrow and add to in order to arrive at my own answer to the regularity question.

2.3.1 Reasons for Mechanistic Regularity

Andersen gives three main reasons why we should want to make some kind of regularity requirement on mechanisms. First, a regularity requirement makes it possible to individuate mechanisms in a non-arbitrary way. Second, a regularity requirement endows us with the means for drawing a distinction between mechanisms and causation in general. And third, a regularity requirement better facilitates the role of mechanisms for grounding generalizations in science.

Regarding the first reason, Andersen writes,

> When a mechanism operates with at least some kind of regularity, we have grounds by which to draw the mechanism boundary around the entities and activities that were actually involved in bringing about the phenomenon in question. (Ibid, 426)

And,

> The more often a mechanism occurs, and the more regularly it occurs, the clearer it becomes which entities and activities are part of the mechanism (i.e., what the mechanism is). (Ibid, 426)

The idea pursued by Andersen here is that, unless we require mechanisms to behave regularly (in some sense), we'll have a much harder time *individuating* mechanisms from their surroundings. If a mechanism has to do (roughly) the same thing each time it operates, on the other hand, it becomes much easier to figure out, by observing it over time, what its component entities and activities are—and thereby distinguish the mechanism from its surroundings.

Regarding the second reason, Andersen writes,

> One of the consequences of dropping the regularity requirement from the definition of a mechanism is that of collapsing the distinction between mechanisms and causation in general. Any chain of causes could be called a mechanism; we could add another causal interaction to a given chain, or take some away, and it would still be a mechanism. (Ibid, 428)

Here, Andersen suggests that placing a regularity requirement on mechanisms better facilitates a distinction between mechanisms and causation in general. Unless mechanisms are required to behave in a regular fashion, she contends, it seems much less clear what separates a mechanism from any kind of regular-old cause. To illustrate this point, consider the MDC characterization of mechanism with its appeal to regularity removed. Let's call it MDC-.

> MDC-: Mechanisms are entities and activities organized such that they are productive of changes from start or set-up to finish or termination conditions.

Now consider what sorts of causal process seem to meet this characterization. Andersen suggests the answer is: *any causal chain whatsoever.* But, for Andersen, this won't do. Whatever else mechanisms are, they have to be something more than run-of-the-mill causal chains.

Regarding the third reason, Andersen writes,

> The regularly recurring chains of entities and activities that constitute a mechanism make them more useful in explanations of other chains of causal-activity-connected entities because regularity grounds generalization from one instance to further instances. (Ibid, 428)

When mechanisms are conceived of as regular, Andersen points out, they also do a much better job of grounding the sorts of generalizations scientists appeal to mechanisms to make. If I know of a given mechanism (call it M_1) that it regularly produces a certain kind of outcome (call it O_1) I can then formulate a generalization of the following sort: mechanisms of the M_1-kind produce outcomes of the O_1-kind. But if mechanisms aren't required to behave regularly, no such generalization seems available.

Ultimately, I'll find these considerations inadequate for motivating the inclusion of a regularity constraint on processes we wish to deem as mechanisms. But my arguments against them will come in a later section. For now, let's look at the second part of Andersen's case.

2.3.2 Taxonomizing Mechanistic Regularly

As we've just seen, Andersen thinks there are powerful reasons for placing some kind of a regularity requirement on mechanisms. But this leaves her vulnerable to the difficult objection mentioned in 2.2. Namely, many of the processes that the mechanistic philosophy of science aims to explain mechanistically *do not seem to behave very regularly*. This objection has been articulated forcefully by Bogen (2005) who cites vesicle release when an action potential is present—a process that fails up to 90% of time—as a causal process that scientists analyze mechanistically, but one that cannot be considered to behave regularly. Since this is a paradigmatic target for mechanistic explanation, and it fails to behave very regularly at all, mechanists had better not require regularity of mechanisms.

In order to address this objection, Andersen develops a taxonomy of regularity for mechanistic explanation. She then makes a disjunctive argument for a broadened appeal to regularity. When regularity is properly taxonomized, Andersen argues, even Bogen's examples can be seen to qualify as regular—in at least one important sense of regularity.

On Andersen's taxonomy, there are three parameters along which a purported mechanism can exhibit regularity (or irregularity). These three parameters are:

- the *organizational location* of regularity in terms of stages within a mechanism;
- the *strength of connection* between component elements in a mechanism; and
- the *failure pattern* by which a mechanism could fail to operate always or for the most part while still exhibiting sufficient regularity to count as a mechanism. (Andersen 2012, 419)

The first parameter involves the organizational location of a mechanism's regularity. According to Andersen, there are four places in a mechanism in which they might behave regularly, or fail to behave regularly:

> R1) Regular occurrence of [a mechanism's] set-up conditions[8],
>
> R2) Regular triggering of the mechanism once its set-up conditions occur,
>
> R3) Regular operation of specific activities connecting the entities within a single mechanism,
>
> R4) Regular production of termination conditions once the mechanism has been triggered. (Ibid, 419, my use of 'R')

As a simple way of illustrating the forgoing distinctions between disparate organization locations of a mechanism's regularity, think for a moment about your household toaster. Ask yourself: how regularly does your toaster produce acceptably toasted pieces of toast? Suppose you estimate that your toaster only puts out an average of 2 pieces of toast per week. Now ask, what is the reason for this? One reason might be that you are on a low-carbohydrate diet, and simply do not allow yourself to indulge in toast on very many occasions. In this imagined scenario, your toaster may work perfectly wonderfully, in that it produces perfectly adequate toast every time you use it; you just don't happen to use it very much. This is a lack of regularity in the sense

[8] In her formulation of R1-R4, Andersen uses the term "start-up conditions" in R1 and R2. But, as I've defined "start-up conditions" (in 1.6), they necessarily result in triggering. As such, I have changed it to 'set-up' conditions—which is what I think she must have meant anyway.

described by Andersen's (R1). Your toaster fails to regularly produce yummy toast because its set-up conditions (e.g., your placing toast into it and depressing its lever) do not obtain very often. Suppose, instead, that you are an avid morning jogger who likes a bit of carbo-loading before you set off. You make toast nearly every morning. Your toaster, in this scenario, would have a much higher degree of (R1) regularity. Now suppose, staying with this (R1) regular scenario, that you put toast in your toaster and depress its lever every morning. But you only get acceptable toast on 4 out of 7 mornings. The other mornings, your bread either remains untoasted or it gets burnt beyond the point of being edible. What could account for this? It might be that your toaster fails to be (R2) regular: you put bread into the slot, you press the lever, but on some mornings your toaster fails to start toasting because its old and (for whatever reason) its lever doesn't always trigger the heating coils. Or perhaps it is because there is faulty wiring between the lever trigger and the heating coils such that, even when the lever is correctly depressed, an electrical short prevents the coils from beginning to heat. On this scenario, your toaster fails to be (R3) regular. Finally, it might be that, even when your toaster lever successfully triggers the heating coils to heat, your bread fails to pop-up at the right time resulting in charred toast on a few mornings. This would be a failure of (R4) regularity. As Andersen has it, when we speak of how regularly your toaster produces yummy toast, we might be referring to any one of these organizational locations.

The second parameter of Andersen's taxonomy or mechanistic regularity points to the fact that the term "regular" could also indicate a different *strength of connection* between the forgoing organizational stages of a mechanism. She cites four levels regarding the strength of connection between mechanism stages:

(a) Deterministic: cause cannot fail to act once conditions are appropriate

> (b) Reliable but not exceptionless: most of the time, the cause brings about the effect, but there are occasions on which it does not, and we may or may not be able to provide an explanation for the exception
>
> (c) Sporadic: the cause fails to act often enough that it cannot be considered merely an exception when this failure occurs
>
> (d) Infrequent: most of the time, the cause fails to bring about its effect, but once in a while it does. (Ibid, 420)

Returning to your toaster, it might be have a strength of connection of the sort described by level (a)—in which case it never fails to produce yummy toast when you put toast in it and press its lever. It might be (b)-level strength of connection, in which case it produces acceptable toast most of the time. Or, much to your dismay, you might have a toaster that has a strength of connection of (c) or (d), in which case your toaster either sporadically produces edible toast or rarely does.

Regarding the third and final parameter of Andersen's taxonomy of mechanistic regularity, Andersen suggests that there are at least two *failure patterns* by which a mechanism can fail at a given organizational location, but nevertheless exhibit some kind of regularity. If at least one of these failure patterns is discernible—even at an organizational location with low connection strength—she argues that the mechanism can still qualify as regular. These two failure patterns are:

> (i) Known statistical distribution of indeterminacy: the mechanism succeeds some consistent percentage of times, even though we may not be able to account for why it succeeds or fails when it does.
>
> (ii) Known interfering factors: when the mechanism fails at a given organization location, we can identify factors that interfered on this occasion, whether or not we know the precise quantitative impact of such factors on the mechanism function in general. (Ibid, 421)

Let's consider your toaster one last time. Suppose, like many of us, you know very little about the inner workings of your toaster. But you've astutely observed a very consistent pattern of its

success and failure: you get perfect toast on Monday, Wednesday, and Friday; but your toast is inedible on the other days. Despite the fact that it fails to produce acceptable toast more often than it succeeds, your toaster is (i)-regular. That is, it succeeds some consistent percentage of the time. Suppose, on the other hand, that on certain mornings your toaster fails and you know why. It's because you're trying to use too many kitchen appliances at once, and your kitchen fuse gets overloaded. On this scenario, your toaster is (ii)-regular despite failing on these mornings.

2.3.3 Countering the Argument from Science

Now that we've understood Andersen's taxonomy of mechanistic regularity, let's explore how it is meant to help support her case for regularity. Recall that Bogen's central example of an irregular process that scientists analyze mechanistically is the purported mechanism by which action potentials induce the release of neurotransmitters by vesicles in presynaptic neurons. Bogen argues that this process cannot be seen to be regular because this it fails up to 90% of the time. But scientists analyze this process mechanistically. So, he concludes, mechanisms should not be required to behave regularly.

Andersen responds to this charge, first, by arguing that there are really two distinct mechanisms operating here: one embedded in the other. While Bogen is right that the inner mechanism fails to qualify as regular (in any of the senses), Andersen argues that the embedding mechanism qualifies as regular enough. The inner mechanism she describes as the process by which the action potential induces the vesicles to move towards and fuse with the cell surface at the active zone of the presynaptic neuron. She calls this the *vesicle mechanism.* In agreement with Bogen, she acknowledges that "there is a relatively high failure rate—as many as 90% of vesicles fail to release their quanta of neurotransmitter when the triggering conditions are present" (Andersen 2012, 423). Applying her taxonomy to the vesicle mechanism, she writes.

> For vesicle release of neurotransmitter, then, the *organizational location* is the triggering of the mechanism given start-up conditions (R2); the *strength of connection* is sporadic—it is not an exception or unusual when a vesicle fails to dock and release transmitter (c). And the *failure pattern* is that of a known statistical distribution without known factors (i). Scientists are not entirely sure what factors block the mechanism running from action potential leading to release of neurotransmitter, but they have reliable means to calculate the probability of failure for a given vesicle, or a given neuron with many vesicles. In sum, for this mechanism, the taxonomy provides a label of (R2-c-i). (Ibid, 423)

She then asks us to consider the broader mechanism in which the forgoing vesicle mechanism is embedded: that of a presynaptic action potential triggering a postsynaptic depolarization. Call this the *depolarization mechanism.* According to Andersen, the vesicle mechanism is actually a stage in the depolarization mechanism. But when we consider the broader depolarization mechanism, the degree of regularity shoots up. Here is what she says.

> The previous mechanism now figures as a stage within this mechanism. In terms of the taxonomy, the *organizational location* in question is the release of neurotransmitters in the presynaptic cell given the triggering of the start-up conditions, which is now an intermediary stage leading from presynaptic cell firing to postsynaptic cell firing (R3). The *strength of connection* is now sufficient to count as reliable (b)… This parameter is based on the vesicles' likelihood of releasing sufficient neurotransmitter to trigger postsynaptic depolarization. Yet it is stronger than the strength of connection in the vesicle mechanism just discussed. This is because even though there is a low chance that a given quantum of neurotransmitter will release, there are multiple quanta that must all fail for the mechanism to fail. (Ibid, 424)

Despite the very low probability that any given vesicle mechanism will succeed in producing neurotransmitter, the overall depolarization mechanism in which the vesicle mechanism is embedded still has a relatively reliable strength of connection. The strength of connection goes up in the overall depolarization mechanism since there are many vesicles that would need to fail together for the overall mechanism to fail. Thus, she concludes:

> …there is an instance-by-instance indeterminacy: for any single action potential on a given occasion, we can't say whether any given vesicle will release neurotransmitter, and thus whether sufficient vesicles will release so as to trigger

> postsynaptic depolarization. However, in spite of this, there is a meta-regularity concerning the single-case indeterminacy that justifies calling the whole process a regularly occurring mechanism. (Ibid, 424)

In other words, despite the failures of the embedded vesicle mechanism to operate regularly, the broader depolarization mechanism qualifies just fine. Thus, when regularity is properly taxonomized, the argument from science against placing a regularity requirement on mechanisms fails.

2.4 Evaluating Andersen's Case

Now that we've outlined Andersen's case for a broadened regularity requirement, let's ask ourselves which aspects of it are convincing—and which are not. The answer that I'll pursue in 2.4.1 is that Andersen's taxonomy of mechanistic regularity takes several steps in the right direction towards laying the groundwork for answering the regularity question—and I'll show this by applying it to the cases of probabilistic phenomena I laid out in Chapter 1. In 2.4.2, however, I'll suggest there are a few features relevant to mechanistic regularity that Andersen either under-specifies or doesn't mention, and which I believe should be added to her taxonomy. And in 2.4.3, I'll conclude my evaluation of Andersen's case by offering some argument against the reasons she gives for motivating her disjunctive argument for a broadened appeal to regularity.

2.4.1 The Positive

I wholeheartedly agree with much of Andersen's approach. I take her taxonomy of regularity to be a helpful step toward laying a theoretical groundwork for answering the regularity question. To illustrate what I take to be the most helpful aspects of her taxonomy, I'll apply it to some of my own examples of probabilistic phenomena cited in Chapter 1.

With regard to (R1), we might ask: does this kind of regularity feature in any of the examples of probabilistic biological phenomena listed in the previous Chapter? And if so, what strength is it? I think it plausibly does exist in our examples, and I think it exists at varying degrees of strength. The set-up conditions for the protein synthesis mechanism (e.g., the presence of DNA, mRNA, and initiation of transcription) are met nearly universally in all living organisms; protein synthesis mechanisms have a very high degree of (R1) regularity. The same is true for synaptic transmission—the mechanism by which one nerve cell communicates with another. Kandel et al. estimate the human brain houses at least 10^{11} neurons—each one with a cell membrane that contains clouds of positive and negative ions spread over its inner and outer surfaces, and each one with an at-rest extracellular excess of positive charge and an intracellular excess of negative charge. (Kandel et al. 2013, 175). These neurons have all met the set-up conditions for synaptic transmission which starts up when an action potential travels down an axon causing depolarization in its wake. Thus, synaptic transmission has exceptionally high degree of (R1) regularity. There are, however, mechanisms we might be interested in whose start-up conditions occur quite rarely. We might consider, for example, the conditions for human evolution as an example of a mechanism with a very low degree of (R1)-type regularity. As Bogen (2005) points out, it (probably) only ever happened once that the conditions were precisely correct for humans to branch off the evolutionary tree from our primate ancestors.

There are also mechanisms that behave more or less regularly *once their set-up conditions have been met*. Mendelian inheritance appears to be an example of a mechanism with a high degree of (R2)-(R4) regularity. Once a pea plant's genes (call them A and a) segregate in the formation of germ cells, they will very nearly always combine at fertilization to give offspring according to the formula: 1AA + 2Aa + 1aa. On the other hand, there are some

biological mechanisms that have a much lower degree of (R2)-(R4) regularity. As we'll see, the mechanism for vesicle release when an action potential is present is an instance of a mechanism with quite a low strength of connection at the (R2) organizational location. Indeed, according to Kandel et al. (2013), it is estimated that the mean probability of synaptic transmission from a single active zone, "varies widely among different presynaptic terminals, from less than 0.1 (that is, a 10% chance that a presynaptic action potential will trigger release of a vesicle) to greater than 0.9" (Kandel et al., 271).

2.4.2 Augmenting Andersen's Taxonomy

As we've just seen, Andersen's taxonomy proves quite helpful for analyzing the many ways in which mechanisms can be thought to behave regularly (or irregularly). There are, however, a few ways in which I think her taxonomy could be augmented. Some of the features I'll propose adding to Andersen's taxonomy are, to varying degrees, implicit in her account but worthwhile to make explicit. Some are entirely absent from her account.

2.4.2.1 Mechanism Internal vs. External Sources of Irregularity

One way in which I think it would be helpful to add to Andersen's taxonomy is by pointing to the following fact: a mechanism can fail to be regular because of inhibitory sources either *internal* or *external* to the mechanism. And this, I suggest, makes a difference as to how regularly we should take a given mechanism to behave.

To illustrate this, think again about your toaster. Suppose that you and a colleague both put 12 slices of bread in your respective toasters during a week, but after depressing your toaster levers, you only end up with 4 adequate pieces of toast—whereas your colleague gets 12. Here, there is an important sense in which your toaster produces adequate toast less regularly than your

colleague's. But in this scenario, it is not a failure of regularity of type (R1) because the set-up conditions for both toasters have equally been met. Imagine, however, that the reason why your toaster only rendered 4 adequate pieces of toast this week is because you were attempting to use too many electrical appliances at once and this overloaded the electrical system and caused your kitchen's circuit board switch to flip off. In this case, there is a failure of regularity, but it is one caused by interfering factors *external to the mechanism*. That your kitchen's circuit board switch is in the on position is a feature of the background conditions of a working toaster mechanism, but not a feature internal to the toaster mechanism itself. Imagine, on the other hand, that the reason for your fewer number of adequate toast pieces is because you have an old toaster that overheats and shuts itself off on those mornings when you've attempted to use it too many times in a row. This would be a *mechanism-internal* source of irregularity because it's a problem within the toaster itself.

To further illustrate the importance of distinguishing between mechanism-internal and mechanism-external sources of irregularity, consider a few biological mechanisms of both of these sorts operating in the natural world. Synaptic transmission can be disrupted by a variety of factors outside of the mechanism itself. In Alzheimer's disease, patients have an abnormal aggregation of a microtubule-binding protein called 'tau'. These proteins form long, thin polymers that wind around one another to form what are called *neurofibrillary tangles* which accumulate in neuronal cell bodies, dendrites, and axons. These tangles, when present, interfere with axonal transport of electrical signals in the brain (Kandel et al. 2013, 78). This is an example of mechanistic irregularity due to inhibitory features (e.g., neurofibrillary tangles) *outside* the mechanism. However, synaptic transmission can also fail to occur due to problems *inside* the mechanism. Normal conduction of nerve signals in the brain can be disrupted, for

example, by defects in myelin proteins—proteins that insulate the very axons that carry electrical signals. This can result in serious disturbances in sensory and motor function (ibid, 91). Since axons are among the entities internal to the synaptic transmission mechanism, this should be seen as a mechanism-internal source of irregularity.

There are also examples of both mechanism-internal and mechanism-external sources of irregularity in the mechanism of DNA replication. DNA can fail to replicate fidelously because its purine base (internal to the mechanism) is changed by the repositioning of a hydrogen atom, altering the hydrogen bonding pattern of that base resulting in incorrect base pairing during replication. And since a particular purine base for a particular strand of DNA is internal to the DNA replication mechanism, this should be seen as a mechanism-internal source of irregularity. However, there may also be instances where DNA replication fails to occur perfectly due to outside inhibitory forces (e.g., epigenetic methylation).

Andersen implicitly acknowledges this distinction between mechanism-internal and mechanism-external sources of irregularity in her discussion of ecological succession.

> Ecologists may know the start-up conditions for ecological succession in a given ecosystem, involving disturbances to habitat or creation of new habitat. Yet not all start-up conditions lead to the termination conditions of climax communities, because there are a range of external factors governed by chance that could prevent the mechanism from working. (Andersen 2012, 422)

Implicit in this passage, Andersen acknowledges that mechanisms sometimes fail to produce regular outputs do to "external factors". But she doesn't say much about how this relates to whether we end up deeming a mechanism to be regular. For reasons that shall become clear in the next Chapter, it is worth making this feature of mechanistic regularity explicit.

In short, the point is as follows. Even the most regular mechanism can be made to fail given the right kind of external factors. But there is definitely a sense in which we still want to

say of a mechanism (e.g., a perfectly good toaster) that it is regular *even if it frequently gets prevented from working by sources external to it.* In other words, let's make sure we separate how regularly a mechanism works given that it's free from inhibitory influences, and how regularly a mechanism works given the presence of inhibitory features.

2.4.2.2 Abstract vs. Concrete Regularity

In addition to the three parameters offered by Andersen (organizational location, strength of connection between stages, and patterns of failure), I'll suggest that how regular a mechanism can be seen to behave also depends in interesting ways on (A) whether the mechanism in question is a *type* or a *token*, and if represented as a type (B) how abstractly it is schematized.

Andersen acknowledges that mechanisms come in types and tokens. But (again) she says nothing about how this relates to how regularly a mechanism can be seen to behave. Regarding the distinction between type and token, here is what she says,

> [T]he term "mechanism" may apply to either a type or a token. On one hand, the term can be used to pick out a single individual causal chain in the world. When a particular neuron fires on a given occasion, a mechanism led to that firing. On the other hand, the term is often used to indicate a type of causal chain, one that could recur on multiple instances: when a neuroscience textbook describes the mechanism for neuron firing, it does not describe a single instance, but rather a type of causal chain that presumably occurs on many occasions. In this way, mechanisms can explain both what happens on a single occasion, as well as what happens on all the occasions on which a neuron fires due to this mechanism. (Andersen 2012, 417)[9]

As Andersen rightly points out, scientists explain by appeal to both individual, actualized mechanisms (i.e., mechanism tokens) and general, idealized mechanisms (i.e., mechanism types). When appealing to the latter, a mechanism type often gets depicted as a *schema*—"a truncated abstract description" of a mechanism type in which entities are often depicted in boxes and their activities depicted as arrows (MDC, 15). These mechanism-type schemas vary in their *degree of*

[9] This feature of mechanisms has also been discussed in detail in Illari, P. M., & Williamson, J. 2010

abstraction. That is, mechanism-type schemas vary in how much detail they include. The more abstract a mechanism schema is, the more detail it leaves out. Here is an example. The protein synthesis mechanism can be schematized, on the one hand, like this:

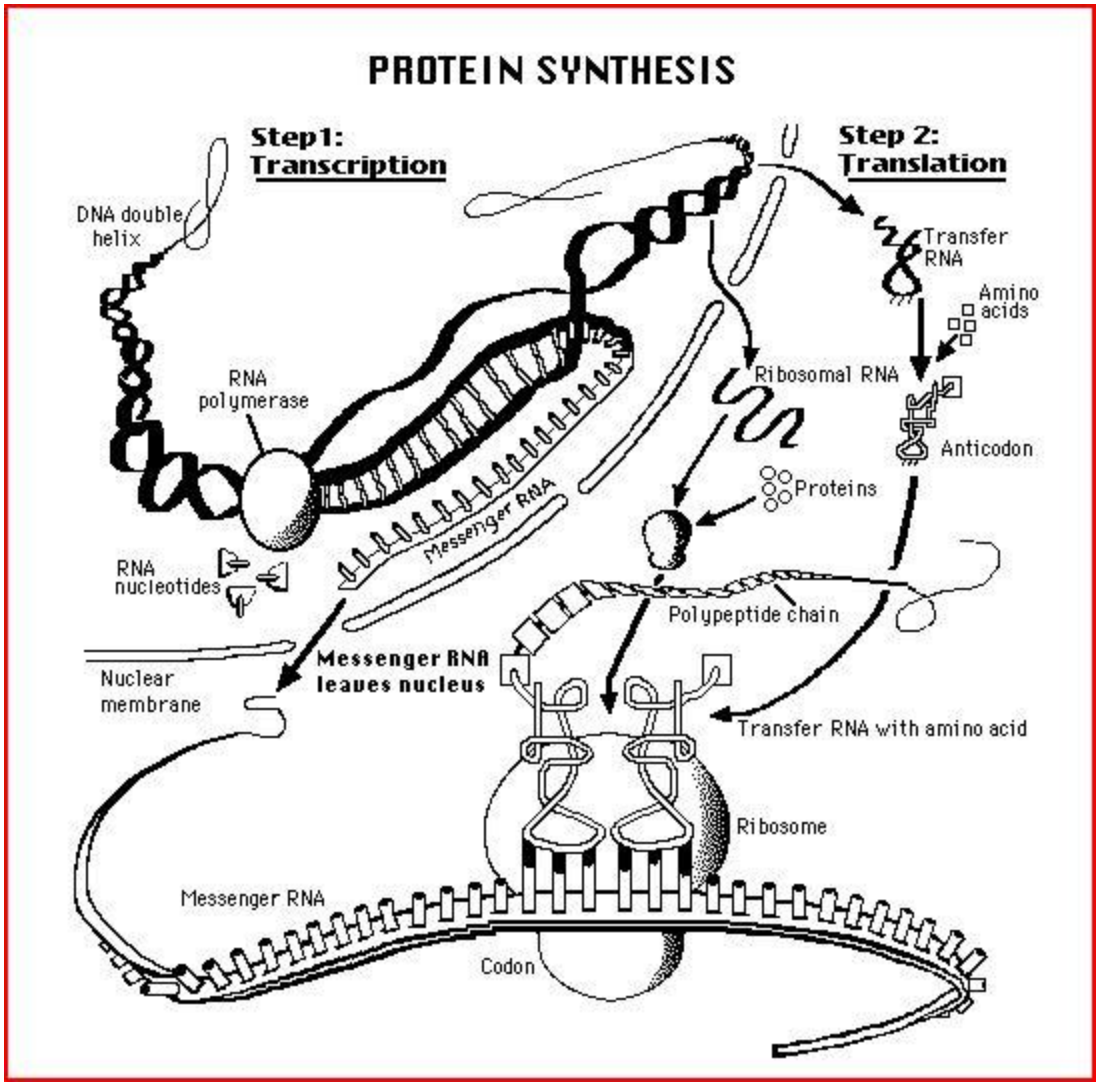

Figure 4. Concrete protein synthesis schema (From Genentech: http://www.gene.com).

Or, on the other hand, protein synthesis can be schematized like this:

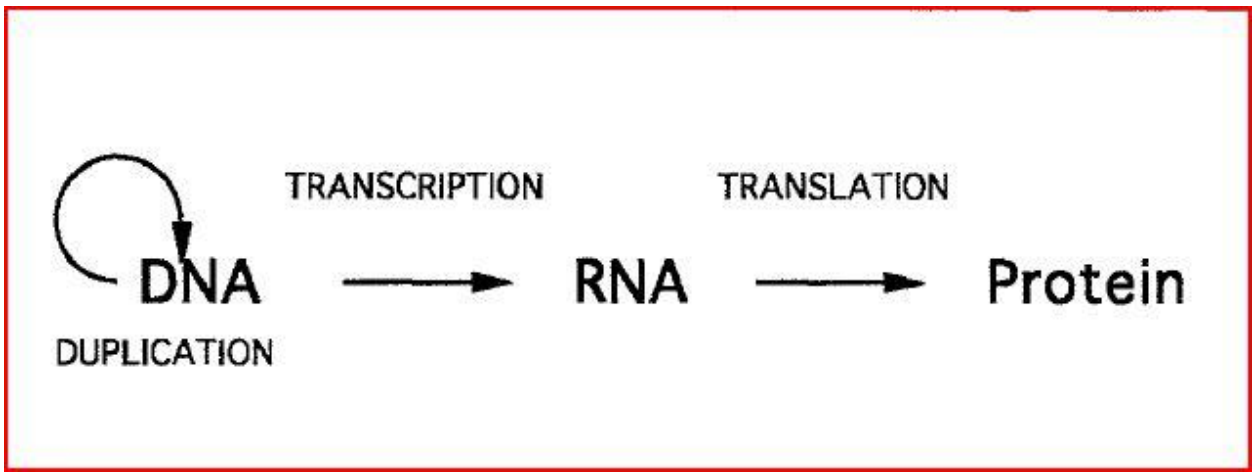

Figure 5. Watson's central dogma diagram (Redrawn, based on Watson 1965).

Both figures 4 and 5 are schematic representations of the mechanism-type of protein synthesis. Both are generalized pictorial representations rather than illustrations of a particular protein synthesis mechanism operating in the world. However, the first schema has a much lower degree of abstraction; it contains much more detail. The second schema, on the other hand, has a much higher degree of abstraction: it leaves a great deal of the detail out.

But how does any of this relate to regularity? The answer I suggest here is that there are interesting correlations between how abstractly a mechanism is schematized and how regularly it can be understood to behave. I'll describe two such correlations.

Given:

(a) the existence of a large number of working mechanism tokens with
(b) at least some features in common and
(c) a high degree of variance in the specific make-up of the features schematized,

CR1: the more abstractly a mechanism-type is schematized, the more of these mechanism tokens instantiate it—that is, the more of these mechanism tokens work in the way specified in the abstract mechanism schema.

CR2: the less abstractly a mechanism-type is schematized, the more it will be the case that the instances it covers behave in exactly that way.

To illustrate these correlations, let's think one more time about toasters. There are toaster tokens, like the one that actually sits on your kitchen counter. There is also a toaster type, a generalized idealization of a toaster's entities and activities. Like protein synthesis, toaster types can be represented with higher or lower degrees of abstraction. A toaster type with a low degree of abstraction would include many details (e.g., how many toast slots, how many toast settings, what kind of materials used, etc.). A toaster type with a high degree of abstraction would leave these details out—opting instead only to depict the bare-bones necessary features something needs to be a toaster. Regarding (CR1), the point I wish to emphasize here is that, (a) as long as there are a large number of working mechanisms, (b) these mechanisms have at least some features in common, and (c) there is a high degree of variance in the specific make-up of the schematized features, *the more abstractly a mechanism type is schematized, the more of these mechanism tokens it will cover*. In other words, (given the existence of a large number of working toaster tokens with some features in common and a high degree of variance in the specific make-ups of these schematized features) the more abstract a toaster type, the more toaster tokens will be subsumed under it. This, I suggest, is a new kind of regularity. Let's call it: *abstract regularity.* A highly abstract schematization of a toaster represents a kind of toaster regularly in this sense because (provided the schematization is abstract enough) *every single toaster that exists operates in the way specified by the schema*.

On the other hand, there is a sense in which less abstract mechanism schemas represent a different kind of regularity. As described in (CR2), the less abstract a mechanism schema is (i.e., the more detail it includes), the more it will be the case that the mechanisms instantiating it operate in exactly in that fashion. There just won't be many of them. Call this kind of regularity: *concrete regularity.*

2.4.3 Evaluating Andersen: The Negative

As articulated in the previous sections, I entirely agree with Andersen that there are many senses according to which a mechanism can be seen to behave regularly. Indeed, I've now added two other considerations to Andersen's helpful taxonomy[10]. Despite our agreement on the importance of taxonomizing mechanistic regularity, however, I disagree that a process needs to exhibit a high degree of any of these senses of regularity in order to *be* a mechanism. In short, I argue that the reasons to which Andersen appeals to motivate her case for a regularity requirement rest on confusion between metaphysics and epistemology/pragmatics.

Recall that Andersen appeals to three primary reasons in order to motivate her case for a broadened regularity requirement: (1) a regularity requirement makes it possible to individuate mechanisms in a non-arbitrary way, (2) a regularity requirement endows us with the means for drawing a distinction between mechanisms and causation in general, and (3) a regularity requirement better facilitates the role of mechanisms for grounding generalizations in science.

Regarding (1), I agree it is easier to individuate mechanisms that behave in a highly regular manner. But I take this to be an epistemological and pragmatic concern—not reason to include regularity as a metaphysical demarcating feature of mechanisms. To illustrate this point, think for a moment about classifying species. Some species have characteristics, whether they be morphological or behavioral, that make them relatively easy to identify and classify. It's pretty easy, for example, to tell if an organism is a western black-widow spider (*Latrodectus hesperus*). All members of this species are between 14 and 16 millimeters and have an hourglass-shaped red mark on their abdomen. Even someone who knows next to nothing about spiders would be able to look at a random sampling of spiders and pick out the western black widows. Now compare

[10] We'll go on to see, in the next Chapter, a concrete example in which these various new aspects of mechanistic regularity apply.

the western black widow to *Canus lupis*, the species that includes wolves, dingos, and dogs. More than any other species, there is an enormous amount of variability in the morophology and behavior across the members of this species. If a person had no background knowledge, it would be next to impossible for her to know that a Chihuahua, a Great Dane, and a Grey Wolf are all members of the same species.

Much like with mechanisms, it would be *easier* for biologists to individuate species if we required of them a high degree of regularity in their appearance and behavior. But this is an epistemological and pragmatic concern—one that deals with how easy it is to determine the boundaries around a species. It would be a mistake to conclude anything about the metaphysical demarcating features of species based on these epistemological and pragmatic concerns. Likewise with mechanisms. It may well be the case that imposing a regularity requirement on mechanisms would make them easier to individuate. But that, by itself, is a poor reason for imposing any metaphysical requirement on what it is to *be* a mechanism.

Regarding (2), is it really the case that we need mechanisms to be regular in order to be able to distinguish between mechanism and causation in general? Why might one think this? To show why, recall MDC- according to which "mechanisms are entities and activities organized such that they are productive of changes from start or set-up to finish or termination conditions." (MDC 2000 [minus 'regular']). Now consider the following example.

> *WWI*: on the 28th of June, 1914, Archduke Franz Ferdinand was shot and assassinated in Sarajevo. As a direct result, Austria-Hungary declared war against Serbia. This, in turn, directly caused Germany and Italy (countries allied with Austria-Hungary) to declare war on the United Kingdom, France and the Russian Empire (countries allied with Serbia). As a result, World War I began.[11]

This is a clear example of a singular causal chain; it only happened once. However, on the MDC characterization minus regularity, it apparently qualifies as a mechanism. In other words, we

[11] See DesAutels (2011) for more discussion of this example.

have entities (Franz Ferdinand and the forgoing countries) and activities (assassination and declarations of war) that are productive of change (the start of World War I). What we have, on such a characterization, is apparently a mechanism. But most of us would not consider this to be a mechanism, but rather, an instance of run-of-the-mill causation. So, it appears regularity is necessary for distinguishing causation in general from mechanisms.

By way of response to the above argument, I point to the following: it only succeeds if regularity proves to be *the only* means by which we can make the distinction between mechanisms and causation—but I don't think it is. To show why, consider again MDC-. At first glance, it seems Andersen is right that the MDC- characterization of mechanism is underequipped for making a distinction between mechanisms and run-of-the-mill causation. However, I disagree that this is the case. Even without its appeal to regularity, the MDC- characterization still makes an appeal to *organization.* It says, "Mechanisms are entities and activities *organized* so as to produce outcomes". Indeed, it is precisely this appeal to organization that I believe endows even MDC- with the capability to draw a distinction between mechanisms and causation in general. The reason why WWI isn't a mechanism is not the fact that it only ever happened once. Rather it's the fact that it lacks the right kind of organization. Regarding the organizational properties of mechanisms, the MDC view states that entities in a mechanism must be located, structured, oriented; and a mechanism's activities must have temporal order rate, and duration. The key feature of a mechanism's organization to focus on here is *structure.* Mechanisms have an organizational structure. Run-of-the-mill causation need not. Of course, much more would need to be said to make this story precise (see section 7.2.1 for some hints as to how this might go). However, that is a topic for another paper. For now, it suffices to have

shown that regularity need not be the *only* available resource for distinguishing between mechanisms and general causation. Organizational structure might serve this purpose just fine.

What about Andersen's reason (3) for placing a broadened regularity requirement on mechanisms? Is it true that mechanisms need to be regular in order for scientists to base generalizations on them? Analogously to my response to reason (1), I think this (again) rests on a conflation between metaphysics, epistemology, and pragmatics. Of course, it would make it easier for scientists to generalize based on highly regular mechanisms. But we shouldn't make requirements on the metaphysically demarcating features of mechanism solely based on what would make scientific explanations easier. We should look at the processes that successful science analyzes mechanistically, and figure out from there what can be said about their metaphysical characteristics.

2.5 My Answer to the Regularity Question

So far, I've drawn attention to a rift in the mechanisms literature regarding whether causal processes need to behave regularly in order to qualify as mechanisms. And because of its subtlety and rigor, I've outlined Andersen's case for a broadened appeal to regularity. By way of evaluating Andersen's case, (i) I've shown the applicability of Andersen's taxonomy of mechanistic regularity by showing how it applies to the cases of probabilistic phenomena I introduced in Chapter 1; (ii) I've augmented this taxonomy by drawing attention to a few additional features of mechanistic regularity; and (iii) I've argued, contra Andersen, that the reasons she gives for adopting this regularity requirement rest on a confusion between metaphysical, epistemological, and pragmatic considerations.

2.5.1 Regularity in Degrees

Here is my answer to the regularity question. Instead of requiring of all mechanisms that they behave regularly in order to *be* mechanisms, I suggest that mechanisms come in degrees of regularity: some highly regular, some not very regular at all.

There are two distinct but related reasons why I propose this answer to the regularity question. First, allowing regularity to come in varying degrees means that we can consider even highly irregular processes (e.g., vesicle mechanisms) count as mechanisms when the science tells us we should. Second, this approach gives the mechanistic framework more explanatory power.

To motivate the first reason, recall Andersen's response to the argument from science. Andersen's response to Bogen was to reframe the irregular vesicle mechanism in terms of the broader mechanism in which it is embedded. As we saw, Andersen argued that, when we turn our attention to the broader depolarization mechanism, its level of regularity becomes much higher—such that it becomes eligible for mechanism status. What I wish to point out, however, is that this strategy only works *at the cost of having to ignore the inner mechanism.*

Consider, however, what would happen if we accepted the following answer to the regularity question. *There is simply no threshold of regularity that should be required of a process to count as a mechanism*. Rather, mechanisms come in varying degrees of regularity on each of the dimensions outlined in the previous sections. The advantage of this approach, I suggest, is that we *can* analyze processes with low degrees of regularity in mechanistic terms just fine. And since that means that there are more mechanisms operating in the natural world, this gives the mechanistic philosophy of science more resources for explaining the natural world.

2.5.2 Privileging Regular Mechanisms

That said, there *is something compelling* about the spirit of Andersen's reasons for wanting mechanisms to have a regularity requirement. Namely, it is true that mechanisms do a better job of supporting scientific generalizations, predictions, and explanations if they operate regularly.

The crucial point I offer here is this: we don't need to require of all mechanisms that they operate regularly in order for them to fill these explanatory roles. Instead, let's admit that some mechanisms are better-suited for grounding generalizations/predictions/intervention strategies than others. The more regular they are, the better they do in this regard. Following this line of thought, I see no problem with bestowing a special privilege to the highly regular mechanisms viz. their superior explanatory role.

2.6 Conclusion

In this Chapter, I've explored the regularity question: to what extent is it a necessary feature of mechanisms that they behave regularly. To do this, I first drew attention to a rift in the mechanisms literature between those who require of candidate mechanisms that they behave regularly and those who make no such requirement. I focused specifically on Holly Andersen's (2012) case for a broadened appeal to regularity. By way of evaluating this case, (i) I highlighted the positive steps Andersen's taxonomy takes towards laying the groundwork for an answer to the regularity question; (ii) I added a few features to her taxonomy that were either under-specified or absent from Andersen's account; and (iii) I offered some negative argument against a few of her conclusions. The crux of my negative argument was to show that the considerations offered by Andersen to support her broadened appeal to regularity rest on a confusion between metaphysical, epistemological, and pragmatic considerations. Rather than requiring of a causal

process that it behave regularly (in any sense) in order to *be* a mechanism, we should allow that mechanisms come in varying degrees of regularity along a myriad of dimensions. This approach, I suggested, has the benefit of being able to analyze even the most irregular of processes as mechanistic when the science tells us we should. In agreement with the spirit of Andersen's arguments, however, I conceded that mechanisms with a high degree of regularity should hold a privileged position in scientific explanation.

Chapter 3. Is Natural Selection Regular Enough to be a Mechanism?

Abstract: In this chapter, I apply several ideas developed in the previous Chapter to help sort out an active debate in the philosophy of biology: whether natural selection qualifies as an MDC mechanism. I argue that, by paying careful attention to some key distinctions drawn in the discussion of mechanistic regularity undertaken in Chapter 2 as well as some recent work on abstract mechanism schemas, natural selection can be seen to qualify as an MDC mechanism just fine—at least in the context of many legitimate explanatory contexts. More than a mere terminological dispute about what to call natural selection, I conclude by suggesting that this debate evinces a deeper point about the prospects for appealing to mechanisms to ground probabilistic generalizations in a contingent biological world.

3.1 Introduction

In Chapter 2, I examined the question of how regular a process needs to behave in order to count as a mechanism. The answer I pursued was that regularity should not be conceived of as a metaphysical requirement for counting as a mechanism. Mechanistic regularity, I argued, is a matter of degree: some behave highly regularly, and some do not. However, there are good epistemological and pragmatic reasons for wanting to search for and describe mechanisms with some kind of regularity if mechanisms are to take over for laws of nature in their role as grounding generalizations, predictions, and explanations in science. The question I now undertake to explore is whether natural selection qualifies as regular enough to take on the role of one of these privileged mechanisms. The answer that I pursue is that it can.

To support this answer, I'll look at a recent argument from Skipper and Millstein (2005) in which they suggest that natural selection fails to operate regularly enough to qualify as an MDC mechanism. I then appeal to the expanded taxonomy of mechanistic regularity offered in the previous Chapter to clarify precisely what kind of regularity natural selection has.

Here is my plan. I begin in Sections 3.2 and 3.3 by characterizing NS and offering prima facie reasons for why it seems to fit within the framework of a mechanistic philosophy of science. I then, in Section 3.4, review one significant argument against NS as an MDC

mechanism: that it fails to meet the regularity requirement set forth in the MDC characterization of mechanism. In Section 3.5, I review a response offered to this argument by Barros (2008) and, in Section 3.6, offer two reasons for Barros's response falls short of answering Skipper and Millstein's regularity critique. In Section 3.7, I draw attention to a couple of important resources from Chapter 2 as well as some current literature on abstract mechanisms to dismantle the regularity critique, showing that (at least in certain legitimate explanatory contexts) NS can be seen to behave regularly enough to qualify as a MDC mechanism. I conclude in Section 3.8 by suggesting that, rather than being a mere terminological dispute about what to call NS, this debate evinces a deeper point about grounding probabilistic generalizations in a contingent biological world.

3.2 Clarifying the Analysandum: What is Natural Selection?

"Nothing in biology makes sense except in the light of evolution." (Dobzhansky 1973)

"I am fully convinced that species are not immutable… Furthermore, I am convinced that Natural Selection has been the main but not exclusive means of modification." (Darwin 1859, 6)

"This preservation of favourable variations and the rejection of injurious variations, I call Natural Selection." (Darwin 1859, 81)

In the late 1830s, after his voyage on the H.M.S. Beagle, Darwin slowly came to the realization that, contrary to the prevailing view of his time, species are not immutable: species were not created in definitive packages to remain static and unchanging. His own experience as a breeder of animals showed that, by choosing to mate individuals with certain characteristics with each other, he could affect change on the form of future generations. He, therefore, knew that species *could* change over time. He further argued, by analogy, that nature affects its own kind of change to species over time. Nature tends to preserve those traits that afford their possessors the greater chance to survive and reproduce, and it tends to reject those that do not. The result is that species

become increasingly matched to their respective environments; they become exquisitely adapted over time. In its most basic form, this is natural selection (NS).

For NS to occur, Darwin argued that three conditions must obtain. (1) There must be variation among members of a population (e.g., some wolf pups are born with faster running abilities than others); (2) These variations must be heritable (e.g., those faster wolf pups will tend to have faster offspring); and (3) There must be, what he termed, a 'struggle for existence' (e.g., there must be more wolf pups in a given generation than can survive)[12]. Given, these three preconditions, it follows that those individuals with advantageous variations will tend to survive and reproduce while those with deleterious ones will tend to die off.

Modern biologists and philosophers of biology give a somewhat more complicated account of what is required for NS to occur. On one account, put forth by Skipper and Millstein (2005)[13], NS can be precisely formulated as follows:

I. Initial conditions
1. A population of *O*s exist.
2. *O*s vary according to forms of *T*, which are heritable.
3. *O*s are in environment *E* with critical factor *F*.

II. Interaction
1. *O*s in virtue of the varying forms of *T* interact differently with environment *E*.
2. Critical factor *F* affects that interaction.
3. This may lead to

III. Effects (1)
1. differential survival rates of *O*s across forms of *T* in *E*.
2. This may lead to

IV. Effects (2)
1. differential reproductive rates of *O*s across forms of *T* in *E*.
2. This may lead to

[12] There is debate as to whether (3) is really required for NS. Since nothing I go on to say in what remains of the Chapter turns on this debate, I shall go on under the presumption that SFE is required.
[13] It should be noted that Skipper and Millstein's characterization of NS draws heavily from Darden and Cain (1989) and Skipper (1999).

V. Effects (3)
1. differential representation in the population of *O*s across forms of *T* in *E*.

2. This may lead to

VI. Effects (4)
1. the predominance of *O*s with a certain form of *T* over other forms of *T* in *E*.
2. This may lead to

VII. Effects (5)
1. adaptation of the lineage with respect to *T* in *E*.

Filling instructions:
'*O*' is to be replaced by the name of some organism
'*T*' is to be replaced by some determinable organismic trait
'*E*' is to be replaced by the description of the environment of '*O*'
'*F*' is to be replaced by the description of a critical factor in '*E*'

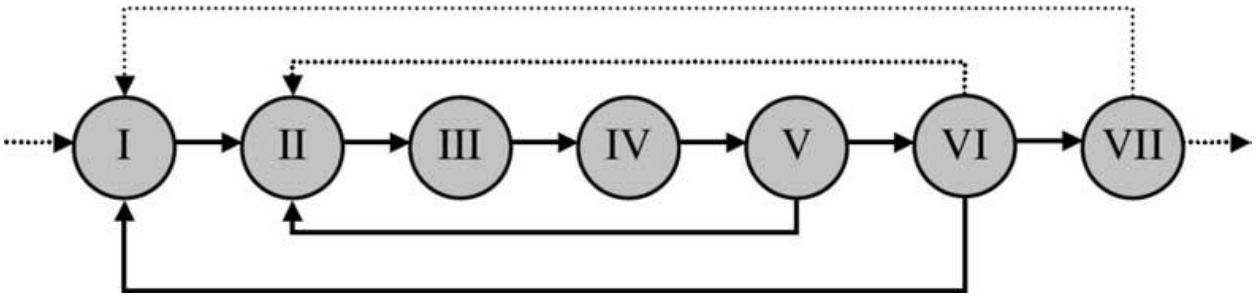

Figure 6. Natural selection schema (From Skipper and Millstein 2005, 331).

To illustrate their formulation of NS, Skipper and Millstein refer to Darwin's famous Galapagos finches.

> There exists a population of finches, Os. The finches vary according to their beak length T, the form of a trait that is heritable. And the finches are in a rocky environment E where there is a critical factor, F—variously shaped edible seeds varying in availability. Notice that the critical factor F in conjunction with differences in T is what sets up the selective interaction. Certain beak lengths enable some of the finches to obtain seeds that other finches, lacking the appropriate beak length, are unable to obtain. This constraint on environmental resources that leads to a struggle for existence and a check on the population is the interaction step of the schematic. Given the initial conditions and the causal

> interaction between the environment and the finches, we expect various downstream effects. (Ibid, 331)

As helpful as the above schematization of NS is, it leaves open the answer to a question that has long interested philosophers of biology is: what *kind of a thing* is NS? How should we represent *what sort of process it is*?

Indeed, Darwin himself seems to have been unsure about the answer to this question. In *the Origin of Species*, He referred to NS in a myriad of ways: an "action" (1859/1964, 90,108, 129, 133, 211), a "doctrine" (5, 95), a "means" (6, 246), a "power" (43, 109, 205, 238, 454), a "theory" (237, 245, 281, 320, 325, 338, 345, 460, 462, 472, 474, 478), a "principle" (80, 95, 116, 127, 188, 206, 239, 475), and a "process" (93, 104, 109, 179, 203, 235, 280, 350). In part due to Darwin's own apparent indecision in the *Origin,* contemporary philosophers of biology have been hard at work arguing for a more precise understanding of what kind of a process NS actually is. Some have argued that NS is best understood as a *force* (Sober 1984); some that NS is a *purely statistical trend* manifesting in natural histories (Matthen and Ariew 2002, 2009); some that NS results from *causal processes operating at the individual level* (Glennan 2009); some that NS is a *causal process, but one that necessarily operates at the population level* (Millstein 2013), and some that it is a multi-staged *mechanism characterizable on both the individual and population level* (Barros 2008). Each of these positions has outspoken critics. But, here I'll focus on one particular critical position: a recent argument against NS as an MDC mechanism (Skipper and Millstein 2005).

In what follows, I hope to show that there are some good reasons for wanting to understand NS as an MDC-type mechanism, and by appealing to some resources afforded by the previous Chapter, NS can be seen to escape at least one of the serious problems set forth against it counting as such.

3.3 Prima Facie Reasons for NS as a Mechanism

It would be nice if NS could be characterized as a mechanism. For reasons already articulated in this dissertation, mechanistic explanation (at least in the life sciences) affords several important advantages over laws-based deductive-nomological accounts of explanation. To review, mechanisms, unlike laws, physically exist in the world, so are more easily investigable by empirical science. Furthermore, generalizations based on mechanisms need not be exceptionless, necessary, or contain universals as is commonly required of laws of nature. Mechanistic explanation also matches intuitions, originally highlighted by Salmon (1984) that giving a scientific explanation must involve laying bare the causal structure of the world: locating a puzzling phenomenon in its causal nexus. And, as a matter of coherence with current practice, biologists *actually do* engage in searching for and describing mechanisms.

Prima facie, NS seems amenable to a mechanistic analysis. There are several reasons for this. For one thing, central to the notion of mechanism is the idea that mechanisms, among other things, are *set up for something*; mechanisms are *productive of some regularly observed behavior or phenomenon*. This feature of mechanisms is highlighted in many of Stuart Glennan's early works. In his 1996 article, "Mechanism and the Nature of Causation", Glennan points out the "one cannot even identify a mechanism without saying what it is that the mechanism does" (Glennan 1996, 52). Put another way, whatever else a mechanism is, it at least needs to have a function; it needs to be set up to do something. Here, it seems NS fares quite well. Quite clearly, NS is a system *for* something: it is that which brings about adaptation[14].

[14] In making this point, I don't mean to attribute any problematic teleology to NS. It isn't that NS is directed towards any specific goal (e.g., to approach perfection or the like). Rather, I mean only to draw attention to the fact that NS has a *function* (in the Cummins sense of function). It is that which brings about adaptation.

Another central feature of mechanisms is that they support reductionist explanation. That is, one of the reasons why mechanisms are explanatorily useful is that they can be decomposed to their component parts and operations, and by doing so, crucial insight into the *why* the phenomenon in question regularly occurs can be achieved. This feature of mechanistic explanation gets emphasized in Bechtel and Richardson's seminal 1993 book, *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research.* In their discussion of (what they call) 'complex localization', Bechtel and Richardson write, "Complex localization requires a decomposition of systemic tasks into subtasks, localizing each of these in a distinct component. Showing how systemic functions are, or at least could be, a consequence of these subtasks is an important element in a fully mechanistic explanation" (Bechtel and Richardson 1993, 125). At first glance, there is a clear sense in which NS fits into this reductionistic explanatory framework: it seems decomposable into its component parts as well as the tasks these parts perform. NS, we might think, is composed of entities (e,g,, populations of organisms with varying traits and a critical environmental factor) and activities (e.g., interaction of organisms with the environment and differential reproduction). In this respect, NS fits in quite well with the reductionistic approach to scientific explanation afforded by the mechanistic approach.

3.4 The Regularity Argument against NS as a Mechanism

Despite the ease with which various aspects of NS appear to cohere with the mechanistic approach, there has been a spate of authors arguing, for varying reasons, that there are significant problems with understanding NS as a mechanism. In this section, I review one especially forceful argument put forward by Skipper and Millstein: that NS fails to qualify as an MDC mechanism because it doesn't operate regularly enough.

In their 2005 paper, Skipper and Millstein write, "With respect to regularity on MDC's account, recall that '[m]echanisms are regular in that they work always or for the most part in the same way under the same conditions.' However, we see natural selection as probabilistic" (Skipper and Millstein 2005, 342). To demonstrate the force of this objection, they ask us to return to Darwin's finches. We are to suppose there is a collection of one hundred equal-sized populations of finches with the same distribution of beak lengths, all located in the same environment. We are to suppose, further, that these finches engage in their usual survival and reproduction activities. Skipper and Millstein ask us then to imagine examining the distributions of beak length in the subsequent generation. They write,

> In this thought experiment, we would not expect the same distributions of beak length in each of the one hundred populations. In some of the populations longer beaks may prevail, in others, shorter beaks may prevail, and in some populations the distribution may be roughly equal… *The bottom line is that natural selection is not regular in the way that MDC require (presumably because natural selection is not regular in the way the mechanisms that MDC discuss are, such as DNA transcription and protein synthesis).* (Skipper and Millstein 2005, 342-343, italics added)

This thought experiment is supposed to reinforce the reader's intuition there is there is something inherently probabilistic about NS. And due to its probabilistic nature, NS should not be seen to behave regularly enough to be an MDC mechanism. In a collection of 100 populations of finches with identically distributed beak lengths—all of which engage in their normal survival and reproduction activities—we simply should not expect the evolutionary outcomes of each of these populations to be identical after several generations. Why?—presumably because there are just too many factors involved in deciding the eventual evolutionary outcome that might go slightly differently in each respective evolutionary iteration. The eventual distribution of beak lengths depends, after all, on spontaneous, chance mutations at the molecular level as well as all kinds of contingent environmental factors. Indeed, it would be something of a miracle if all of these 100

populations ended up with identical beak-length distributions after only a few generations. But, as Skipper and Millstein point out, if NS is to be an MDC mechanism, it needs to behave regularly (or nearly regularly). And if NS were to behave regularly, then *we should expect* there to be identical (or near enough to identical) distributions of beak lengths after a few generations. So much the worse, we are to suppose, for NS as an MDC mechanism.

In argument form, the regularity critique looks like this.

P1. MDC requires that mechanisms behave regularly (i.e., they 'work always or for the most part in the same way under the same conditions').

P2. But NS is inherently probabilistic.

P3. Given (2), NS cannot behave regularly.

C1. Therefore, NS cannot be a mechanism.

Let us call this argument the *Skipper and Millstein (S&M) Regularity Critique.*

3.5 Barros's Reply to Skipper and Millstein

Before developing my own response to this charge, it is worth taking a look at another recent attempt at circumventing the S&M Regularity Critique, one that I suggest ultimately falls short. In his paper "Natural Selection as a Mechanism" Barros (2008) attempts to respond to the S&M Regularity Critique by offering a novel way of characterizing NS: one that avoids the force of their critique. As Barros understands it, NS is (what he terms) a 'two-level multistaged stochastic mechanism'.

To understand what this means and why Barros thinks that understanding NS this way affords an escape from the S&M Regularity Critique, let's look at a few of Barros's definitions. He defines *stochastic mechanisms* as "those whose outcome can be predicted in advance in terms

that are probabilistic" (311). He further distinguishes two types of stochastic mechanism: *biased* and *unbiased*. He writes,

> …stochastic mechanisms should be divided into two categories. First, *unbiased stochastic mechanisms* are those that operate with outcome probabilities of 50% or less. The fair-coin flipping mechanism is an example. Another example is the mechanism that releases neurotransmitters and initiates electrical activity in postsynaptic neurons described by Bogen (2005). Looking forward at any particular operation of a neurotransmitter mechanism, it is at best possible to say that the probability of the initiation of electrical activity is a percentage less than 50%... *biased stochastic mechanisms* are those that operate with probabilities of greater than 50% and whose operations can be characterized by terms like 'are likely to' or 'probably will'. Natural selection, discussed further below, is an example of a biased stochastic mechanism.

In addition to distinguishing between biased and unbiased stochastic mechanisms, Barros suggests, further, that the only way to fully characterize NS is by referencing two levels: *the individual level* and *the population level*. To illustrate his claim, he appeals to the case of predatory crabs (*C. maenas*) selecting for shell shape in a population of intertidal snails (*L. obtusata*). In the population of intertidal snails, there are high-spired shells and low-spired shells, but crab predation is more successful on the high-spired shells because they are more easily crushed by the crab pincers. Selective pressure thereby leads to the gradual prevalence of low-spired rather than high-spired shells. This is clearly an instance of NS in action. For Barros, however, this instance of NS can only be fully encapsulated by appeal to two levels[15]: the individual and the population. Figure 7 depicts the former, and Figure 8 depicts the latter.

[15] Barros does not do much to specify exactly what sense of 'levels' he means.

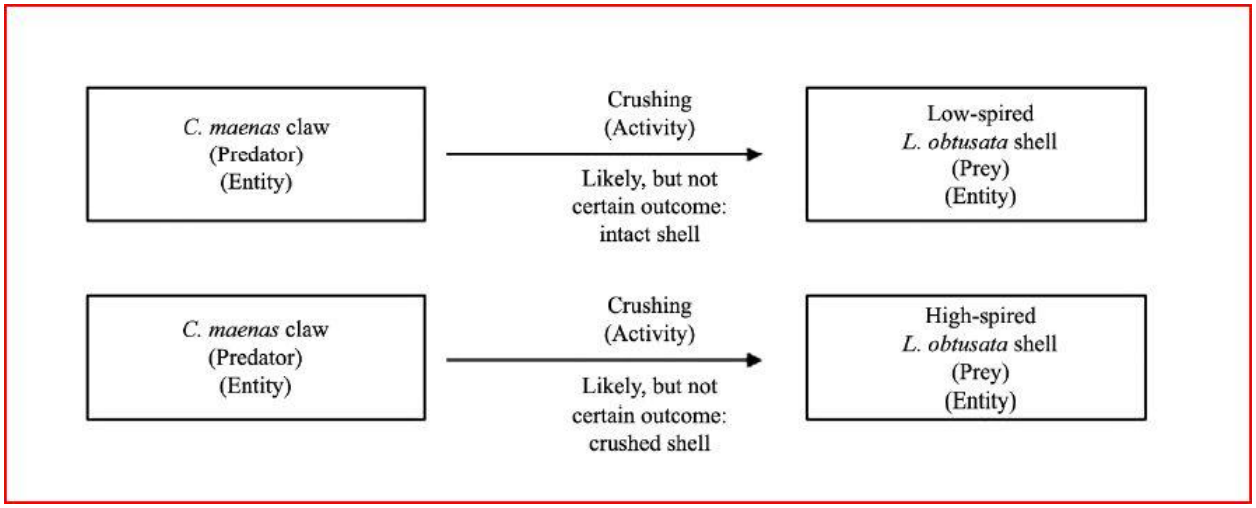

Figure 7. Predatory crabs (*C. maenas*) selecting for shell shape in a population of intertidal snails (*L. obtusata*): *individual level* (From Barros 2008, 315).

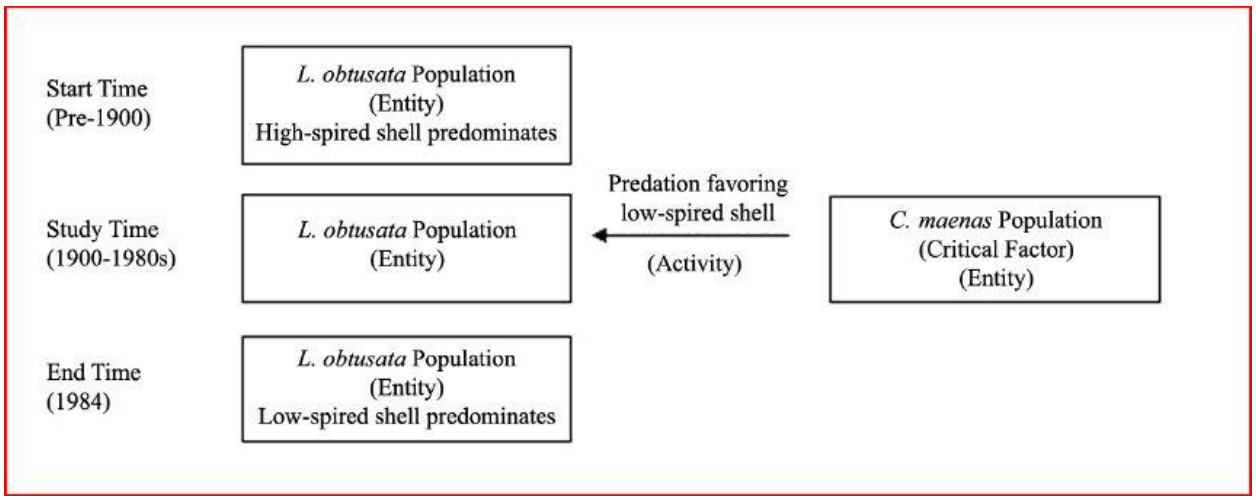

Figure 8. Predatory crabs (*C. maenas*) selecting for shell shape in a population of intertidal snails (*L. obtusata*): *population level* (From Barros 2008, 316).

Barros's central argument is that NS must be represented at both levels because neither one alone can capture it. Only at the level of individual crabs and snails are there differences in shell-crushing abilities. However, to fully capture NS, one must make reference to the population-level effects over time: namely, the predomination of low-spired shell types. No one level alone captures NS. He thus defines NS as,

> …a two level, multistage stochastic mechanism that explains the phenomenon of adaptation. It is two-level because the phenomenon of adaptation cannot be fully explained using either individual or population level mechanisms alone. (Barros 2008, 318)

Barros argues that this new conception of NS allows an escape from the S&M Regularity Critique. Barros rests his case on his characterization of NS as a *biased stochastic mechanism*. He writes, "[it] occurs when a particular trait gives an individual organism a selective advantage over other individuals who do not possess the trait. In marked contrast to the example of drift, it is possible to look forward from the starting conditions and predict that it is likely that the trait's prevalence in the population will increase because of the selective bias in favor of the trait" (Barros, 319). The point here seems to be that NS, unlike drift, is predictable. And since NS is predictable, it must be likely to occur. And if it's likely, then it must have greater than a 50% probability of occurring. So, NS is a biased stochastic mechanism. And if NS is a biased stochastic mechanism, then it meets the requirement set for by MDC that mechanisms operate regularly, as in it produces the anticipated result always or *for the most part*.

3.6 A Couple of Problems for Barros's Reply to Skipper and Millstein

As we've just seen, Barros's general point appears to be that NS, characterized as a biased stochastic mechanism, endows it with enough regularity to pass muster as an MDC-type mechanism. By virtue of being a biased stochastic mechanism, on Barros's account, we are to understand NS to have over a 50% probability of success in a given instance of selection. As such, we are licensed to speak of it as if it 'probably will' or 'is likely to' occur.

As will soon become clear, I am sympathetic to this general strategy. In fact, in the next Section, I will follow Barros in suggesting that *just because a mechanism is stochastic doesn't necessarily mean it behaves in a problematically irregular fashion*. I will also follow Barros in

suggesting that there are certain advantages afforded by characterizing NS on both *the concrete, individual level* and *the abstract, population level*. That said, some of the details of Barros's account are worrisome. For one thing, Barros's definition of stochastic mechanisms as "those whose outcome can be predicted in advance in terms that are probabilistic" (311), appears to identify stochasticity with a lack of predictability, and that makes his notion of stochasticity look entirely epistemic. This runs counter to an argument I'll offer in the next chapter (4.3.1) against a subjective understanding of stochastic mechanism. But for now, I'll say this. It may be true that the outcomes of stochastic processes are not predictable. But, my view is that this unpredictability is not the defining characteristic of a stochastic mechanism; it is merely a byproduct of its mind-independent stochasticity—a feature of mechanisms that I will define (in Chapter 5) as a function of probability-weighted ways a mechanism might fire.

Let's put aside for a moment our definitional quibbles regarding stochasticity. Recall that Barros rests his response to Skipper and Millstein on his claim that NS is a biased stochastic mechanism: one that succeeds more often than it fails. However, Barros troublingly doesn't offer any evidence for this biased nature of NS—which, it seems quite clearly to me—is an *empirical claim*. Is it really the case that NS succeeds over 50% of the time? What would this even mean? All Barros says by way of defense of this claim is that NS is biased because "[it] occurs when a particular trait gives an individual organism a selective advantage over other individuals who do not possess the trait. In marked contrast to the example of drift, it is possible to look forward from the starting conditions and predict that it is likely that the trait's prevalence in the population will increase because of the selective bias in favor of the trait" (Barros 2008, 317). Is it really possible to look forward and predict at a degree of certainty over 50% that a particular population will adapt to a particular key environmental factor? Are there examples in the

literature supporting this claim? It strikes me that, for Barros to substantiate his claim that NS is a biased stochastic mechanism, more empirical evidence is needed.

There is another, even more troubling, aspect of Barros's account worth pointing out. Barros seems to suggest that biased stochastic mechanisms (mechanisms that succeed over 50% of the time) meet the regularity requirement set forth by the mechanists, while unbiased ones (mechanisms that succeed less than 50% of the time) do not. And, recall, that Barros's idea of success has to do entirely with whether we can predict its outcome. However, it seems to me that there is a confusion here. Consider a fair roulette wheel with 99 wedges, one third of which are red, one third of which are black, and one third of which are green. Now ask yourself what the probability of landing on green is after a fair spin of the wheel. The answer is .3 repeating, no? So, on Barros's account, the roulette wheel is an unbiased stochastic mechanism when it comes to the result of landing on green (or either of the other colors). And because it's unbiased, we cannot say that any particular outcome is likely to occur. Here's the worry. Although the probability of landing on green is less than 50%, wouldn't we be wrong to think that this is a mechanism incapable of supporting prediction, that this is an *irregular process*? In this scenario, it seems to me, we should be able to make a perfectly good prediction: that the wheel will land on green 1/3 of the time. Put another way, it seems arbitrary to draw the cutoff at 50% certainty such that mechanisms with results above it are regular enough, and those below it are not. More on this in the next section.

It may very well be that Barros could smooth out his account to accommodate my confusion. But, as it stands, more needs to be said to eliminate the threat of the S&M Regularity Critique.

3.7 Dismantling the Regularity Critique

In this section, I recapitulate a couple of distinctions drawn in Chapters 1 and 2; I highlight some recent work on abstract mechanistic explanation; and I then set about employing these theoretical resources to tease out a novel solution to the S&M Regularity Critique.

3.7.1 Regularity Isn't a Metaphysically Demarcating Feature of Mechanisms

The first thing to say by way of response to the S&M Regularity Critique is that, contra (P1), mechanisms shouldn't be required to behave regularly to begin with. As I argued in Chapter 2, considerations supporting this requirement rest on a confusion between metaphysical, epistemological, and pragmatic considerations. Rather than requiring of a causal process that it behave regularly in order to *be* a mechanism, we should acknowledge that mechanisms come in degrees of regularity: some highly regular, some not very regular at all.

However, as I also suggested in Chapter 2, the pragmatic and epistemological considerations cited by Andersen (and others) in support of placing a regularity requirement on our characterization of mechanisms are very real. Mechanisms that *are* regular do a much better job of grounding predictions, generalizations and explanations in science. And since these are the roles for which scientists aim to search for and describe mechanisms, we should really be interested in which are these explanatorily privileged mechanisms. In other words, even if regularity isn't a metaphysically demarcating feature of mechanisms, we should still be interested in determining which *are* these regular mechanisms that should be given a privileged role in our scientific explanations.

Furthermore, there are good reasons to suppose that NS is exactly the sort of candidate mechanism that scientists rely on to fill these privileged explanatory roles. Do biologists appeal

to NS to ground predictions? Do they appeal to NS to support generalizations? The answer is undoubtedly, yes.

Given these considerations, we might well simply re-construe the S&M regularity critique in the following way:

> P1*. For a mechanism to hold an explanatorily privileged position, it needs to behave regularly (i.e., it needs to 'work always or for the most part in the same way under the same conditions').
>
> P2*. But NS is inherently probabilistic.
>
> P3*. Given (2), NS cannot behave regularly.
>
> C1*. Therefore, NS cannot be an explanatorily privileged mechanism.

Let's call this the *S&M+ Regularity Critique* for its added reference to explanatory privilege.

3.7.2 Eliminating External Irregularity

So now let's ask whether the S&M+ Regularity Critique is sound. I argue that it is not. The first strategy I undertake to begin to dismantle the S&M+ Regularity Critique has to do with the notion of *regularity* itself. My argument will be two-pronged: first I claim that Skipper and Millstein have not appreciated a few important distinctions regarding regularity in the context of mechanisms. When these distinctions are made salient, I argue, it becomes clear that *only certain forms of irregularity should be seen to inhibit a process from counting as regular enough to be a (privileged) mechanism.* The second prong of my argument will be to show that, *to the extent that NS fails to behave regularly, it only does so in a fashion that should not be seen to threaten its status as a (privileged) MDC mechanism.*

Regarding the first prong, I will again rely on Andersen's taxonomy of regularity originally outlined in Chapter 2.3. Andersen's taxonomy of mechanistic regularity, recall,

separates out four organizational locations where a mechanism can behave (or fail to behave) regularly:

> R1) Regular occurrence of [a mechanism's] set-up conditions,
>
> R2) Regular triggering of the mechanism once set-up conditions occur,
>
> R3) Regular operation of specific activities connecting the entities within a single mechanism,
>
> R4) Regular production of termination conditions once the mechanism has been triggered. (Andersen 2012, 419, my use of 'R')

It also cites four levels regarding the strength of connection between these mechanism stages:

> (a) Deterministic: cause cannot fail to act once conditions are appropriate
>
> (b) Reliable but not exceptionless: most of the time, the cause brings about the effect, but there are occasions on which it does not, and we may or may not be able to provide an explanation for the exception
>
> (c) Sporadic: the cause fails to act often enough that it cannot be considered merely an exception when this failure occurs
>
> (d) Infrequent: most of the time, the cause fails to bring about its effect, but once in a while it does. (Ibid, 420)

In 2.3, I gave some reasons why each of these senses of regularity is important for understanding probabilistic biological phenomena by appeal to underlying stochastic mechanisms. However, I now wish to suggest that some of these senses of regularity are more important for mechanisms to instantiate than others—at least from the standpoint of how well they can ground generalization, predictions, and explanations in science.

First, a bit of review is in order. Recall that the MDC Regularity Constraint requires mechanisms to operate "always or for the most part in the same way". Recall further that, as I argued in 2.4-2.5, while regularity should not be considered a metaphysically demarcating feature of mechanisms, we should privilege regular mechanisms for the role that they can play in scientific explanation.

What I argue now is that the failure of a mechanism to behave (R1) regularly should not be seen to matter as to whether a mechanism counts as regular enough to qualify as an explanatorily privileged MDC mechanism. To begin to illustrate why, think again about your household toaster. I might have a toaster that produces perfect toast every time I depress its lever and just so happen not to depress its lever much at all. This toaster would have a very low degree of (R1) regularity. Here is my central point. Even though this toaster fails to get regularly used, *when used*, it operates perfectly regularly. And as such, *it is repeatable even if it has not (as a matter of fact) been repeated*. A failure for a mechanism to get triggered very often should not be seen, by itself, to undermine its ability to count as an explanatorily privileged mechanism. The importance of this point for undermining the S&M+ regularity critique will become clear in a moment. But for now, let us turn our attention to another potentially helpful resource from Chapter 2.

Another resource from Chapter 2 that I suspect might help to undermine the S&M+ regularity critique is the distinction we drew in 2.4.2.1 between *mechanism-external* and *mechanism-internal* sources of irregularity. What I wish to suggest now is that mechanism-external sources of irregularity should not be seen to inhibit a process from qualifying as a privileged MDC mechanism. Let's think about toasters one more time. If my toaster fails to regularly produce adequate toast because the wiring in my kitchen is bad, this would be a mechanism-external source of irregularity. This is because the reason it fails to achieve its output conditions is not due to anything *within* the mechanism; it is due, rather, to inhibitory conditions in its surrounding. Let us ask, once again, whether a failure of regularity of this sort should be seen to undermine the toaster's status as a privileged mechanism. Here is the general point. Even the most regular mechanism can be made to fail to produce its outcome if it is inhibited

externally. Synaptic transmission, protein synthesis, DNA replication, all considered regular biological mechanisms, can be made to fail to occur by the right kind of inhibitory interference (e.g., if a test subject is killed). What matters for the mechanistic explanatory framework is that mechanisms function with some degree of regularity—*when free from inhibitory influence*. The only point I wish to emphasize is that sensitivity to external sources of irregularity should not, by itself, be seen to threaten whether a mechanism can function in scientific intervention and prediction strategies. And if this is so, then external irregularity should not, by itself, constitute a reason for rejecting a process from counting as an explanatorily privileged MDC mechanism.

If successful, the above arguments show that not all irregularity should be seen to pose a problem for whether a process counts as regular enough to be an explanatorily privileged MDC mechanism. If a mechanism simply has a low frequency of getting triggered, but operates perfectly regularly once it is triggered, it should still count as regular enough. If a mechanism fails to achieve its termination conditions regularly due entirely to external inhibiting factors, it needn't be seen to fail to meet the MDC regularity requirement.

Having established these points, let us return our attention to the S&M+ Regularity Critique. The second prong of my argument is to show that the sort of irregularity attributed to NS by Skipper and Millstein is exactly the sort of irregularly I've just described as entirely unthreatening to its status as a privileged MDC mechanism. To show why, recall (once again) Skipper and Millstein's Galapagos finch example. They suggested that, given 100 equal-sized populations of finches in the same environmental conditions performing the same survival and reproductive activities, we should not expect the average beak lengths to be the same across the imagined populations after a few generations. On the basis of this thought-experiment, they concluded, "The bottom line is that natural selection is not regular in the way that MDC require

(presumably because natural selection is not regular in the way the mechanisms that MDC discuss are, such as DNA transcription and protein synthesis)" (Skipper and Millstein 2005, 343). I am now in a position to counter this claim. Here is my argument.

> P4. Because Skipper and Millstein's thought experiment presupposes a starting condition of 100 *identical populations* of finches engaging in *the very same survival/reproductive activities*, the reason why we wouldn't expect to see identical beak distribution lengths across these populations after a few generations must be because of non-critical environmental contingencies.
>
> P5. This constitutes a mechanism-external source of irregularity.
>
> P6. However, as I've already argued, irregularity due entirely to mechanism-external inhibiting factors should not be seen to preclude a mechanism from qualifying as a privileged MDC mechanism.
>
> C2. So, Skipper and Millstein's finch thought experiment should not be seen to threaten an understanding of NS as being regular enough to qualify as an explanatorily privileged MDC mechanism.

Some of these premises need more motivation. Regarding (P4), recall that Skipper and Millstein's finch example stipulates that we are to imagine 100 *identical populations of finches* all engaging in *the very same survival and reproductive activities*. If this is the case, I suggest, the only reason we wouldn't see the same evolutionary outcomes after a few generations must be differences in contingent, non-critical environmental factors between imaginative iterations of these finches' evolutionary histories. Why?—because everything else is to remain constant by stipulation. If all of the survival and reproductive activities are to remain entirely fixed across iterated versions of this though-experiment, then we should expect these finches to be doing the exact same foraging for food, the exact same escaping of predators, and the exact same choosing of mates. So if all of these remain entirely constant, what else could account for a difference in beak-length distribution in these iterated examples? Perhaps, in some of these imagined evolutionary scenarios, a portion of finches get killed in a tropical storm; or perhaps some of the

finches get disturbed by human encroachment. And these are the imagined reasons for the resulting non-identical beak-lengths. Indeed, if each of the 100 finch populations truly is identical in genetic make-up at the imagined start point, and each of these populations engages in exactly the same activities with the same critical environmental factors (e.g., seeds of the same sizes located in the same places), what else could explain the variation in beak-length at the end point besides such non-critical environmental contingencies?

Suppose this is right, and the only factors that could explain the variation in beak-length at the imagined end point are contingent non-critical environmental factors. What makes these mechanism-external sources of irregularity—as (P5) suggests they must be? To answer this question requires saying something about how we might draw the boundaries around a particular token natural selection mechanism. And, admittedly, this isn't easy to do. It's not like natural selection comes in an obvious package like a toaster. So how *do* we draw the boundaries around a token NS mechanism? Although it won't be possible to give a complete story regarding mechanism individuation here, I can point to at least some constraints on how mechanism individuation might proceed. Most obviously, a given NS mechanism is constituted by a population of organisms with variation in fitness-relative traits. In our example, these are the finches with varying beak-lengths. But this can't be the end of the story for individuating this token NS mechanism. For NS to occur, there must also be some kind of critical environmental factor. In the case of the finches, this is usually presented as seeds with differing husks and seeds which are located in places such that certain of the finches' beaks are better suited to forage for them, and certain of the finches' beaks are worse suited for foraging for them. Outside of this critical environmental factor, the rest of the environment, I submit, should not be taken to be constitutive of the token NS mechanism. Why?—because every mechanism token (no matter

what kind of mechanism it is) exists in some kind of environment. Including the whole of the finches' environment within the boundaries of the mechanism would rob us of any ability to conceive of this mechanism as having any boundary at all. So, on this line of reasoning, with the exception of the critical environmental factor, the rest of the finches' environment should not be seen to be constitutive of the token NS mechanism.

But if what I'm suggesting is so, then the source of the irregularity exhibited in this thought experiment must be entirely due to non-critical environmental perturbations to the purported mechanism of NS. And like the faulty wiring in my kitchen, these non-critical environmental perturbations must be considered *mechanism-external inhibitory factors* (i.e., P-6). And if this source of irregularity is mechanism-external, it need not constitute a failure to meet the MDC Regularity Constraint. So, on this line of thought, Skipper and Millstein's thought experiment need not preclude an understanding of NS as regular enough to be an MDC mechanism.

Consider the following objection. Skipper and Millstein might agree with me that mechanism-external sources of irregularity should not be seen as problematic for meeting the MDC Regularity Constraint. But they might deny that the irregularity described in their thought experiment must be seen as having a mechanism-external source. They might, for example, suggest that the reason why we would not expect to see identical beak-lengths after a few generations is not necessarily due to non-critical environmental factors. It may be that the reason we shouldn't expect to see identical beak lengths in the finches is because, for NS to take place at all, there must be *variation* in a population. And one important source of variation is random mutation during DNA replication. But if these variations are random, then we should not expect the same mutations to occur across all 100 populations of finches. And if mutations occur in a

disparate fashion across these populations, then (of course) we would not expect the eventual beak-length distribution to be the same after a few generations. Furthermore, since varied members of populations are entities *inside* the purported NS mechanism, then this should be considered a mechanism-internal source of irregularity: one that most certainly does threaten NS's ability to be seen as regular enough to be an MDC mechanism.

My first response to this potential objection is to (again) point to Skipper and Millstein's stipulation that these 100 populations start out *identical*. Presuming metaphysical determinism, whatever mutations occur within these populations should be the same *unless* some difference in their environments brings about disparate mutations across these populations. And if it environmental differences that bring about these mutational differences, then it is still a mechanism-external source of irregularity. Perhaps, however, Skipper and Millstein argue that we need to take seriously the possibility that (at least some) genetic mutation occurs in a genuinely *indeterministic* manner (say, of the sort argued for by Weber 2005). On a view like Weber's, the very same precipitating conditions might result in different mutational outcomes due to some real (perhaps quantum-level) chanciness. Since, as we'll see, I will go on to argue (in Chapter 6) that the only intellectually responsible position regarding the possibility of such genuine metaphysical indeterminism is a disjunctivist/agnosticism regarding the metaphysical indeterminism question, I'll briefly consider a response to the indeterministic disjunct to this objection.

3.7.3 Regularity *Once Mutations Occur* Might be the Best We Can Do

Let me briefly rehearse the dialectic so far. Skipper and Millstein use their finch thought experiment to motivate a claim that NS does not meet the MDC Regularity Constraint. First I countered this claim by arguing that, given Skipper and Millstein's stipulation that the finch

populations (both in form and activity) are identical, it must be environmental differences that account for irregular beak-length outcomes. But environmental contingencies, I suggested, are mechanism-external sources of irregularity and so should not be taken to threaten NS's status as an MDC mechanism. Not so, Skipper and Millstein might object. What if the reason for the irregular beak-lengths is actually that there is some genuine metaphysical indeterminism operating at the molecular level causing mutations during DNA replication? If this were so, then the source of irregularity would be mechanism-internal, and thereby a real threat to NS's ability to meet the MDC regularity constraint. Let us call this the *indeterministic disjunct* of the S&M Regularity Critique.

Responding to this objection requires that I make something of a concession. Presuming genuine indeterministic mutation during DNA replication, it must be admitted that some of the finches on whose variation NS operates will mutate at different times and in different ways across the imagined populations of finches. This is why we shouldn't expect to see identical beak-length distributions after a few generations. However, I maintain that there is a sense of regularity that remains. Here is my central claim.

> CC: Even if we concede that sources of variation on which NS operates are irregular, *what NS does with them may not be*.

To motivate this claim, think back to Barros's crab and snail example. Recall that the predatory crabs in this example have an easier time crushing their snail prey when these snails have higher spires on their shells. In this regard, we should expect the snails with lower-spired shells to be more successful in avoiding crab predation—and consequently—for snails with lower spires to be selected. Now imagine, a la Skipper and Millstein, iterating this selection history 100 times under the presumption that mutations in snail-shell-spire-height occur differently in each iteration. Just as Skipper and Millstien argue regarding the finches, we should not expect the

eventual shell height distribution to be identical across all imagined iterations. However, and this is the key, suppose we take each population of differentially mutated snails as *separate inputs to the NS mechanism.* In other words, suppose we take variable populations as the starting point for the NS mechanism to operate. The point I want to make is this: *once the mutations have occurred, what NS does with them should be considered quite regular.* E.g., NS will favor those members of the snail population who have the lower-spired shells. Regardless of how regularly these shells get distributed throughout the population, once they are, NS will carry out its task of preserving those individuals who are better suited to avoid crab predation. And in this respect, NS should still be seen to operate always (or for the most part) in the same fashion.

Skipper and Millstein might accept that NS operates regularly once environmental factors and mutation distributions are held steady. But they might simply reply that this is an empirically untenable position. From the standpoint of science, they might suggest, we will never be in such a position to control for these complexities. The result being that, as a matter a scientific practice, NS cannot be regular enough to play the role mechanists need it to: e.g., to support prediction and intervention strategies. In response to this, I have one more strategy.

3.7.4 NS as an Abstract Mechanism Type

My third strategy for dismantling the S&M Regularity Critique has to do with two other important features of how we represent a given mechanism from Chapter 2: whether we represent a given mechanism as a *token* versus a *type,* and what degree of *abstraction* we use to represent a given mechanism.

Regarding the difference between token and type mechanisms, recall that the former is an actualized instantiation of a given mechanism in the world, and the latter is a idealized

description of the general way a given mechanism operates. Recall again the way Andersen puts it,

> [T]he term "mechanism" may apply to either a type or a token. On one hand, the term can be used to pick out a single individual causal chain in the world. When a particular neuron fires on a given occasion, a mechanism led to that firing. On the other hand, the term is often used to indicate a type of causal chain, one that could recur on multiple instances: when a neuroscience textbook describes the mechanism for neuron firing, it does not describe a single instance, but rather a type of causal chain that presumably occurs on many occasions. In this way, mechanisms can explain both what happens on a single occasion, as well as what happens on all the occasions on which a neuron fires due to this mechanism. (Andersen 2012, 417)

As applied to NS, on the one hand, we can give a characterization of the mechanism of NS as a specific, actualized instance of NS at work (like Barros's crab/snail example or Darwin's finches). And, on the other hand, we can describe NS as a mechanism type (like the schema given by Skipper and Millstein [figure 6 in Section 3.2]).

Before showing why the token/type distinction is important, there is one more idea, the explication of which will help tremendously in addressing the remaining issues for NS as a mechanism: the role of abstraction in mechanistic explanation. In a recent article entitled, "Abstraction and the Organization of Mechanisms" (2013), Bechtel and Levy make the case that certain mechanistic explanations are successful by virtue of the fact that the mechanism(s) referred to in these explanations are allowed to remain abstract. So what, exactly is meant by 'abstract' in this context? Bechtel and Levy write, "Abstraction is, of course, a matter of degree. Indeed to say of a description that it is abstract is to imply, or at least suggest, that a more concrete description is possible" (Bechtel and Levy 2013, 242). They go on to argue that, in the case of at least some of the phenomena targeted for mechanistic explanation, "It is always possible and, we argue, often desirable to overlook the more concrete aspects of a system and represent its organization abstractly as a set of interconnections among its elements." (Ibid, 255).

I take their main idea to be this: when attempting to explain highly dynamic mechanistic systems with many interconnected parts, it is often beneficial to leave out some of the details. The reason is, as they put it, "The concrete relations that are pertinent to organization differ across different systems". But, and this is key, sometimes detail-poor explanatory models enable us to "track those features of the system that make a difference to the behavior being explained" (Ibid, 256).

Bechtel and Levy, in this article, point to a peculiar feature of scientific explanation, one that has been recognized before. Think for a moment about an example famously articulated by Hillary Putnam (1975) in which we are trying to explain why a one-inch square peg won't fit into a circular hole with a one-inch diameter. We might explain this by appeal to a detail-rich account of why the specific micro-physical properties of *this particular peg* impede its ability to pass through the molecular make-up of *this particular hole*. But the more explanatorily satisfying answer here, Putnam points out, might actually be an abstract one having to do with the generic properties of squareness and roundness of the peg and hole as such. In giving such an explanation we can track the features of the system (e.g., the squareness of the peg and roundness of the hole) that make a difference to the behavior we're interested in explaining. And those are not necessarily the micro-physical properties. The moral of the story, here, is meant to be that the best explanation isn't always to the most detailed; sometimes explanations are good because they are abstract.

So how do the concepts of token vs. type mechanisms and degrees of abstraction help us to address the S&M Regularity Critique? My answer lies in the following argument.

> P7. The degree of regularity of a particular mechanism is sensitive to two aspects of how the mechanism is represented: (1) whether it is characterized as a *type* rather than a *token*, and if characterized as a type, (2) its *degree of abstraction*.

> P8. When conceived as a type with a relatively high degree of abstraction, NS should be understood as behaving quite regularly—at least in one of the senses that is required for it to function well for grounding certain generalizations, predictions, and explanations.
>
> P9. There are legitimate explanatory contexts (a la Putnam's pegs) in which NS should be characterized as a mechanism type with a high degree of abstraction.
>
> C3. Given (P7)-(P9), there are legitimate explanatory contexts in which NS qualifies as regular enough to meet the MDC requirement.
>
> C4. So for at least for these explanatory contexts, NS escapes the S&M+ Regularity Critique.

Premises (P7)–(P9) need more motivation to show that (C3) and (C4) follow. Starting with (P7), recall from section 2.4.2.2 of the previous Chapter that the idea is this. How regular a mechanism is seen to behave depends (1) on whether it is represented as a type (rather than a token) and (2) the degree of abstraction of its representation. And as we saw, there are two interesting correlations between abstract mechanism schemas and regularity:

> Given:
>
> (a) the existence of a large number of working mechanism tokens with
> (b) at least some features in common and
> (c) a high degree of variance in the specific make-up of the features schematized,
>
> CR1: the more abstractly a mechanism-type is schematized, the more of these mechanism tokens instantiate it—that is, the more of these mechanism tokens work in the way specified in the abstract mechanism schema.
>
> CR2: the less abstractly a mechanism-type is schematized, the more it will be the case that the instances it covers behave in exactly that way.

To illustrate this point, think about a vending machine. There are actual vending machines in the world; these are vending machine tokens. There is also a vending machine type: an idealized representation of a machine that dispenses items selected by its user when adequate payment is inserted. At the token level, there are a great many differences in how vending machines operate. Some have coils that push candy out; some have rectangular compartments that collapse when an

item is selected; some have digital readout screens, and others don't; some take dollar bills, and others don't; etc. At this token level of description, vending machines do not operate always (or even for the most part) in the same way. They, therefore, do not seem to meet the MDC Regularity Constraint. However, when we characterize the vending machine *type* (rather than token), there is a significant amount of regularity in the way vending machines operate. They are all machines that display items for sale, take money for such items, and when payment is made, they all dispense purchased items in a collection area accessible by the patron. Furthermore, the higher degree of abstraction used to represent the vending machine type (provided the above qualifications [a]-[c] are met), the more regularity it exhibits. There are abstract candy vending machine types; there are abstract soft-drink or coffee vending machine types—all of which share regularities in their constitution and behavior. However, there is also the most abstract vending machine type that captures each of these.

Applying this idea to NS is straight-forward. Like the aforementioned vending machine tokens, actualized instances of NS (e.g., Barros's crabs and snails) will be constituted by specific entities (e.g., *C. maenas* and *L. obtusata*) engaging in specific activities (e.g., snail-shell-crushing) in specific environmental surroundings (e.g., rocky coast-lines). NS operating on Darwin's finches will be constituted by different entities, activities, and surroundings. These concrete NS descriptions may only qualify as regular in the sense described by CR2. And that kind of regularity does do particularly well at grounding generalizations of the sort required from privileged MDC mechanisms. However, when idealized as an abstract mechanism type, NS can be depicted as a general kind of process (e.g., figure 6 in 3.2). And when represented with this high degree of abstraction, it overcomes the individual differences between its actualized token instances. Represented like this, NS does operate the same way every time.

Furthermore, I suggest that there are legitimate explanatory contexts in which NS as an abstract mechanism type does real explanatory work—work that mere NS tokens cannot do. As with Putnam's pegs, there are instances where we aren't interested in the specifics of why a particular population has ended up with the morphological make-up that it does. But rather, we might be interested in why, in general, low-spired snail shells prevail in tidal pool regions across the globe. Here we would not want to appeal to a specific population of snails/crabs. We would instead want to depict the NS type as it applies to snail predation across a variety of specific tidal pool regions. We may even have legitimate questions about what instances of selection share in common at its *most* abstract level—in which case an answer would be to appeal to the highest degree of abstraction (e.g., figure 6 in 3.2).

In this section, I have argued that the amount of regularity a mechanism is seen to have depends on (1) whether it is characterized as a type (rather than a token) and (2) the degree of abstraction used to represent a given mechanism. I have also argued that, with regard to the NS mechanism type, its regularity increases with the degree of abstraction used to represent it. I suggested further that (just as with Putnam's pegs) there are legitimate explanatory contexts in which appeals to abstract NS mechanism types give answers that actualized NS tokens cannot give. All of this goes to show that (at least in certain explanatory contexts) NS should be seen to meet the MDC Regularity Constraint just fine. And if that is the case, then the S&M Regularity Critique seems to lose much of its purchase against NS as an MDC mechanism.

3.8 More than a Mere Terminological Debate

In the end, does it really matter what we call natural selection? Is there anything of philosophical import that turns on whether we term it a "mechanism", or a "causal-process" or a "force" or a "hamburger"? Perhaps, one might think, the only thing that matters is that *it happens*—and that

(whatever we call it) it has been responsible for the exquisite adaptations we observe all around us in the natural world.

Like my imagined interlocutor, I am entirely against indulging in shallow terminological disputes for their own sakes. But, in this case, I firmly believe that there is a deeper issue belying this debate—and the exercise I've just undertaken to make arguments on one side serves to illustrate this point in a helpful manner.

Following Beatty (1995), one might take the biological world to be riddled with contingency: too complex to formulate laws governing it. One might object to a mechanistic approach to explaining the living world for similar reasons: when it comes to biological phenomena, we don't have something that looks like clockwork. We have a teeming, seething, mess. Sensitive to this exception-ridden, highly complex nature of the living world, proponents of the mechanistic approach, as we've seen, either soften their regularity requirement (e.g., MDC's now oft-quoted "mechanisms are regular in that they work always or for the most part in the same way") or they get rid of it altogether (e.g., Bogen, [late] Machamer, Glennan). What I take the arguments in this Chapter to have shown is that proponents of MDC mechanisms can do better than this; they can be more nuanced in how they respond to problems regarding irregularity; and indeed careful attention to this dissertation can provide crucial resources for understanding how appealing to mechanisms can explain probabilistic phenomena in the natural world.

So rather than a mere attempt to vindicate the use of the word 'mechanism' to describe NS (a project which alone matters very little), I take these arguments to demonstrate a general strategy for thinking about appealing to mechanisms to explain stochastic biological phenomena

in the contingent natural world. The central features of this strategy might be summarized as follows:

- Be aware of the nature and source of the irregularities that riddle your causal process of interest. Not all irregularity is equally damning to a process's mechanistic status.
- Be aware that stochasticity takes different forms, some of which do better than others at supporting prediction and intervention strategies.
- Understand that, whether a causal process behaves regularly enough to count as a mechanism, depends crucially on the degree of abstraction used in its representation.

3.9 Conclusion

In this Chapter, I have explored whether the account of stochastic mechanism developed thus far in this dissertation can shed any light on a live debate in the mechanisms literature: whether natural selection operates regularly enough to be an explanatorily privileged MDC mechanism. I have argued that it is. To do so, I first offered prima facie reasons for why we might want to think about NS in mechanistic terms. I then summarized Skipper and Millstein's regularity critique against NS as an MDC mechanism. Next, I looked at a recent response to Skipper and Milstein—one which I ultimately found unsuccessful. I then appealed to some recent work on abstract mechanistic explanation as well as a few distinctions drawn in earlier Chapters in order to alleviate the regularity critique. I concluded that, more than a mere terminological dispute, this debate evinces a deeper point about the prospects for appealing to mechanisms to ground probabilistic generalizations in a contingent biological world.

Chapter 4. Stochastic Mechanisms and Theories of Chance

Abstract: My goal in this Chapter is to further develop my characterization of stochastic mechanism by exploring how we should best understand the chanciness we attribute to them. To do this, I form a list of desiderata that any account of stochastic mechanism must meet. I then take the general characterization of mechanism offered by Machamer, Darden, and Craver (2000) and explore how it fits with several of the going philosophical accounts of chance: subjectivism, frequentism (both actual and hypothetical), Lewisian best-systems, and propensity. I argue that neither subjectivism, frequentism, nor a best-system-style account of chance will meet all of the proposed desiderata, but some version of propensity theory can. I conclude by showing the proposed account of propensity-backed stochastic mechanism has the added benefit of allowing us to escape many of the traditional objections to propensity theories of chance.

4.1 Introduction

As we saw in the previous Chapters, there are significant questions having to do with the notion of *regularity* employed in extant characterizations of mechanism. In Chapter 2, I argued that regularity should not be a requirement for metaphysically demarcating mechanisms from non-mechanisms. However, I suggested that mechanisms with a high degree of regularity should hold a privileged position in scientific explanation. I then employed this framework, in Chapter 3, to help sort out an active debate in the philosophy of biology: whether natural selection is regular enough to be one of these privileged MDC mechanisms. I concluded that, when we are sensitive to the sources of irregularity that should be seen to threaten a process from counting as a privileged mechanism, we can see that natural selection isn't highly irregular in any of these ways. As such, there are no problems with understanding natural selection as an explanatorily privileged MDC mechanism.

In this Chapter, I focus on a different but related question about stochastic mechanisms. Specifically, I ask how we should further understand mechanisms that are irregular in a certain way. Namely, I ask how we should further understand mechanisms whose set-up conditions obtain but whose termination conditions may not (i.e., stochastic mechanisms with a less that deterministic strength of connection at the R2-R4 organizational locations). My goal in this

Chapter is to sort out which, if any, of the going conceptual understandings of chance is best suited to pair with stochastic mechanisms.

Here is the plan. In Section 4.2, I first outline and explain several desiderata which I believe any adequate account of stochastic mechanism should meet. In Section 4.3, I then go about showing how subjectivist, frequentist, and best-system-style analyses of stochastic mechanism fail to meet one (or several) of these desiderata. In 4.4, I argue that a propensity-backed account of stochastic mechanism is the most promising candidate for meeting these desiderata; it also has the added benefit, I suggest, of pairing nicely with the varied senses of mechanistic regularity emphasized in Chapters 2 and 3. In 4.5, I put forth a few serious objections to a propensity understanding of chance and show (in Sections 4.6 and 4.7) how making explanatory reference to the underlying mechanisms themselves allows an escape from (at least some of) the force of these objections.

4.2 Desiderata for an Adequate Account of Stochastic Mechanism

Before addressing the question of what sort of chance we want to pair with an understanding of stochastic mechanism, it is necessary to lay out a few general desiderata for an adequate account of stochastic mechanism. The list that follows by no means exhausts the features that might be desirable from a characterization of stochastic mechanism. Rather, the following desiderata are best thought of as a short list of *constraints* on any adequate account of stochastic mechanism.

I begin by articulating and motivating three desiderata for an adequate analysis of stochastic mechanism.

> *COHERENCE*: any adequate account of stochastic mechanism must cohere with the general practice of biologists using mechanisms to explain natural phenomena.

> *CAUSAL EXPLANATION*: any adequate account of stochastic mechanism must allow for descriptions of underlying mechanisms to feature in causal explanations of regularities seen in nature.
>
> *ONE-OFF*: any adequate account of stochastic mechanism must allow (at least some) mechanisms whose initial conditions are only ever met once to have definable, non-extremal chances of firing.

By way of briefly motivating this list of desiderata, I'll say a few words about each.

COHERENCE states that whatever else our account of stochastic mechanism is, it must fit with the way biologists actually appeal to mechanisms to explain puzzling phenomena. I take this to be uncontroversial. Indeed, I take it that one of the central purposes of developing this account is to supply some theoretical and conceptual foundations to a concept of mechanism already widely in use in contemporary, empirically successful life science.

The next two desiderata are more controversial. Beginning with *CAUSAL EXPLANATION*, I follow Wesley Salmon (1989) and James Woodward (2003) who both argue forcefully that giving a scientific explanation of a phenomenon requires doing more than subsuming it under a covering law—as the once-received deductive-nomological view of explanation required. To give a scientific explanation, one must lay bare the inner causal workings of nature. To answer *why* something happens, in the context of science, requires showing what caused it. Mechanistic explanation is a particularly strong form of causal explanation. When one gives a mechanistic explanation of a phenomenon, one does more than just describe its underlying cause; one describes the causal structure—both the entities and activities—that gives rise to its outputs. Furthermore, (as we've seen) one of the primary advantages of a mechanistic philosophy of science is that it provides a theoretical basis for life scientists to explain the uniformity we see in the natural world without necessarily having to appeal to laws of nature. Without going too far astray into the hotly debated issue of whether

there are any laws of biology, recall that there are many who doubt the existence of exceptionless and metaphysically necessary laws governing the natural world (Cartwright 1983, Beatty 1995). Ceterus paribus laws are just as fraught with controversy (Fodor 1991, Earman and Roberts 2002). But even if that were not the case, the fact remains that life scientists *actually do* search for mechanisms (as opposed to laws) to ground their explanation of regularities observable in nature. So if we are to have any hope of achieving a working conception of stochastic mechanism that coheres with scientific practice, such an account better allow us to appeal to these mechanisms to causally explain observed regularities.

Furthermore, as stated in *ONE-OFF*, I suggest that our account of stochastic mechanism needs to accommodate an ability to assign non-extremal chances to (at least some) mechanisms whose initial conditions for firing are only ever met once. A good example of this is comes from Bogen (2005), and it is the mechanism for human evolution. In all likelihood, Bogen points out, the initial conditions for the mechanism of human evolution were only ever met once—and never will be again. However, I now suggest, most of us would say of that mechanism that there was some non-extremal chance it failing. Suppose, for example, a giant meteor hit at precisely the right time to destroy all of our closest primate ancestors. This would have been unlikely, but possible. Given this intuition, we had better allow (at least some) one-off mechanisms to have non-extremal chances.

So now having gained some understanding of the above constraints on an adequate account of stochastic mechanism, we can get on with the work of seeing how various accounts of chance fare with respect to each one of them.

4.3 Subjectivism, Frequentism, and Best Systems Analyses of Chance

This section comprises arguments against various philosophical theories of chance as ways of underpinning an account of stochastic mechanism. I argue by elimination that a propensity-style approach to stochasticity is the only adequate means of laying a foundation for our account of stochastic mechanism.

What is a philosophical theory of chance? For the purposes at hand, theories of chance should be thought of as systematic attempts to give the truth conditions for chance statements. When you flip a fair coin, you assign the chance of it landing heads: 1/2. When you role a fair six-sided die, you assign the chance of it displaying a 6: 1/6. What makes these statements true? This is the question philosophical theories of chance attempt to answer.

4.3.1 Against Subjectivism

The first account of chance that might be considered as a candidate to underpin an understanding of stochastic mechanism is a subjectivist one. On a subjectivist account, there are no objective chances: only credences. When we say of a given outcome that it has a certain chance of occurring, we ought to mean nothing more than that we should have a certain degree of belief in that outcome.[16] Chance, on this type of account, gets replaced by credence or rational confidence level that some event will occur. What would an account of stochastic mechanism look like if we understood stochasticity in this manner? It might go something like this: when we say that the mechanism responsible for the release of electrical activity in postsynaptic neurons has a 10% chance of firing at any given time, we are not ascribing any kind of chanciness to the synaptic mechanism itself. Rather we mean only to assert that we ought to have a rational degree

[16] A classical example of this is Bruno de Finetti's (1937) account of subjective probability.

of belief of .1 that this mechanism will fire on any given instance when its start-up conditions obtain.

My view is that applying this type of subjectivism about chance to our understanding of mechanism cannot give us an adequate analysis of stochastic mechanism. Here is my argument[17]:

P1. Life scientists give mechanistic explanations of objective facts.

P2. Some of these mechanistic explanations of objective facts are probabilistic.

P3. The probabilities in probabilistic explanations of objective facts must be objective.

P4. So (on pain of violating COHERENCE) the probabilities in mechanistic explanations of objective facts must be objective

C1. Therefore, we have good reason to reject a subjectivist account of stochastic mechanism.

Premise (P1) is uncontroversial: As we saw in Chapter 1, scientists appeal to mechanisms to explain facts about the natural world. Proteins come to exist from DNA molecules because of the mechanism of protein synthesis. Alleles segregate in the formation of germ cells because of the mechanism of Mendelian segregation. Electrical signals cross synapses in the brain because of the mechanism of synaptic transmission. Barring radical idealism or scientific anti-realism, the facts explained by these mechanisms are taken to be objective.

As we also saw in Chapter 1, some of the mechanistic explanations of these objective facts are probabilistic (P2). Although Mendelian segregation occurs at rate of nearly 3:1 in the F_2 generation, it does not do so perfectly. Protein syntheses and DNA replication mechanisms, although highly successful, are error prone to various degrees. And vesicle release of neurotransmitter upon the presence of an action potential fails up to 90% of the time.

[17] This is a modified version of an argument defended by Lyon (2011) against a subjectivist understanding of the probabilities in classical statistical mechanics.

But why think, as (P3) states, that the probabilities in probabilistic explanations of objective facts must be objective? To help understand why, consider the alternative. We might think that, rather than objective probabilities, the probabilities in probabilistic explanations of objective facts might simply be a measure of our ignorance. As Lyon (2011) puts it, "…an explanation involving probability is not automatically a probabilistic explanation—it could be a probability of explanation" (Lyon 2011, 423). In other words, it may be not be that these probabilistic mechanistic explanations are themselves probabilistic, but rather it may be that the probabilities in these explanations are merely a measure of how strongly we should believe that the candidate explanation is the correct one. Following Lyon's strategy, however, I don't think this can be right. In other words, I don't think that the probabilities in mechanistic explanations merely reflect our ignorance in the way that would be appropriate for understanding them as probabilities of explanation. A detailed defense of this premise is beyond the purview of this Chapter. However, convincing arguments have been given to this effect. An especially relevant (and convincing) one can be found in Millstein (2003b), in which she argues that the probabilities in evolutionary theory cannot be mere measures of our ignorance. Rather than measuring the factors in evolutionary processes of which we are ignorant, she argues, many of the probabilities in evolutionary theory represent causal factors about which we have knowledge—but knowledge we *choose to ignore.* She writes, "[T]his 'ignorance interpretation overlooks the fact that we are aware of more causal factors than are included in the transition probability equation; for example, we know things about the predator and the color of the butterflies. Thus, we chose to ignore these causal factors, rather than being ignorant of them" (Millstein 2003b, 1321). Sober (2010) is another good example of someone who convincingly argues that the probabilities in evolutionary biology are objective.

Of course, these examples aren't enough to show that all of the probabilities in probabilistic mechanistic explanations are objective. But, I suggest, it's enough to provide some theoretical basis for accepting (P3). If Millstein and Sober are correct, then at least the probabilities in evolutionary theory are objective. And since evolutionary theory is one of the primary arenas for mechanistic explanation, this is significant support of (P3). And if we have significant support for the premise that the probabilities in probabilistic explanations of objective facts *are* objective, then it follows that the probabilities in mechanistic explanations of objective facts are objective. And if these probabilities are objective then this rules out a subjectivist understanding of stochastic mechanism.

Note, however, that this argument does nothing to undermine subjectivist understandings of probability in all contexts (e.g., Bayesianism). These subjectivist accounts certainly have plenty of uses. But, as I've suggested here, my only point is that they don't do well cohering with the way scientists actually appeal to mechanisms to explain the objective world.

4.3.2 Against Frequentism

On an actual frequentist (AF) view of chance, the chance of a given event occurring in a finite reference class is just the frequency of actual occurrences of that event relative to that reference class.[18] How would an account of stochastic mechanism look on such an analysis? An obvious application of AF to MDC mechanism would be this: the chance of a given mechanism firing is the frequency of that mechanism successfully firing relative to instances over the history of the world where that mechanism's initial conditions actually obtained.

I argue that this view of stochasticity does not cohere with what we want from an account of stochastic mechanism. My argument is this:

[18] See Venn 1876 for the origin of this view.

P5. On an AF view of stochastic mechanism, the chance of a given mechanism producing its outcome is the actual relative frequency of the occurrence of the outcome given the instantiation of its initial conditions throughout history.

P6. Given (P5), however, the chance of a one-off mechanism must be extremal, because, by definition, its initial conditions are only ever met once.

P7. But *ONE-OFF* require that (at least some) one-off mechanisms be allowed to have definable, non-extremal chances.

C2. Therefore, an AF view of mechanism is unacceptable.

(P5) is uncontroversial; it's just the combination of MDC mechanism with an AF view of chance. Premises (P6) and (P7) need to be justified.

Recall that ONE-OFF states that any adequate account of stochastic mechanism must allow that (at least some) mechanisms whose initial conditions are only ever met once should have definable, non-extremal chances of firing. The mechanism for human evolution, I suggested, is a plausible example of such a one-off mechanism: its initial conditions were only ever met once, but we nevertheless would have assigned it a definable, non-extremal chance of failing to succeed. But since the initial conditions for this one-off mechanism were only ever met once (and plausibly only ever will be met once), on an AF understanding, the mechanism for human evolution must be given an extremal chance of firing. But this violates *ONE-OFF*. Indeed, on an AF understanding of stochastic mechanism, every one-off mechanism whose initial conditions are only ever met once *must* be given an extremal probability of firing (1 if the mechanism successfully achieves its termination conditions, and 0 if it doesn't). But this is exactly what *ONE-OFF* precludes from being the case. Thus, it appears that an AF understanding of chance cannot give us what we want from an account of stochastic mechanism.

How about a hypothetical frequentist theory of chance? Does it fare any better with respect to our desiderata for an adequate account of stochastic mechanism? On an HF view of

chance, the chance of a given event occurring is the limiting relative frequency of that event occurring relative to a hypothetical, infinite (or very large) series of trials of that event.[19] The result of combining this type of theory of chance with our understanding of mechanism would be this: the chance that a given mechanism will fire (given that its start/set-up conditions obtain) is just the frequency of the mechanism achieving its expected outcome over a hypothetical, infinite (or very large) series of trials. Notice that this kind of account clearly avoids the problems of an AF understanding of stochastic mechanism: it has no problem rendering non-extremal chances to one-off mechanisms.

Unfortunately, I see different problems with this understanding of stochastic mechanism. Here is my argument.

> P8. On an HF analysis of stochastic mechanism, the stochasticity of a given mechanism is the limiting relative frequency of it achieving its outcome given the instantiation of its initial conditions over a hypothetical, infinite (or very numerous) series of (non-actual) trials of that event.
>
> P9. Give (P8), the chance of a given stochastic mechanism firing is grounded on a counterfactual.
>
> P10. Life scientists, however, appeal to the chanciness of underlying mechanisms to causally explain actual output frequencies.
>
> P11. But if the chanciness of a mechanism is grounded on a counterfactual, it's difficult to see how it can causally explain output frequencies of actual mechanisms.
>
> P12. So, given (P8)-(P11), an HF account of stochastic mechanism fails to meet *COHERENCE* and *CAUSAL EXPLANATION*.
>
> C3. Therefore, a (HF) view of stochastic mechanism is not viable.

Premise (P8), once again, is just the result of combining our understanding of mechanism with a HF theory of chance. As before, however, premises (P9)-(P12) need more defense.

[19] Some classic examples of hypothetical frequentists include Reichenbach 1949 and von Mises 1957.

Suppose a molecular biologist observes that the mechanism of DNA replication in a particular population of fruit flies is significantly error prone. She notices, let's say, that the DNA of flies in a given generation is only 95% identical to those in the previous generation. After observing several generations with similar results, she thereby generalizes that the mechanism of DNA replication for these fruit flies has a 5% error rate. On a HF understanding of stochastic mechanism, this is by virtue of the following true counterfactual: if the sequence of generations continued indefinitely, then the relative frequency of errors in DNA replication would limit to 5%.

Here is the problem. The scientist in this example set out to explain *actual* output frequencies of a stochastic mechanism. That is, she set out to say why we see the frequency of DNA replication errors that we do in an actual population of fruit flies. On a mechanistic approach to explanation, the answer is that the mechanism for DNA replication fails 5% of the time. But an HF understanding of stochasticity grounds this chance on a counterfactual: namely, the non-actual world where some infinitely large (or very large) number of trials took place. Here is the vexing question for the HF account: how can stochasticity grounded on counterfactual, non-actual world causally explain anything observed in the actual world? It strikes me that it cannot.[20] And if it cannot, then an HF understanding neither coheres with the practice of life scientists appealing to mechanisms to explain output frequencies, nor can it meet our CAUSAL EXPLANATION desideratum. As such, we cannot accept an HF understanding of stochastic mechanism.

Suppose, however, that the proponent of HF were to respond as follows. There are plenty of perfectly good causal explanations that appeal to counterfactuals. Indeed, both Lewis (1973B) and Woodward (2003) offer accounts of causal explanation in which counterfactuals feature

[20] For a particularly forceful articulation of the relevance problem for counterfactual explanation, see Salmon 1988.

prominently. For Lewis, to say of some even *E* that it is causally dependent on *C* is just to say that if *C* had not occurred, then *E* would not have occurred. That is, causal dependence *just is* a counterfactual notion. Similarly, Woodward offers an account of causation according to which what it means to say that some event *E* was caused by *C* is that, *were* we to have intervened on *C* in the right way, *E* would not have occurred. Like Lewis, Woodward clearly thinks that causation is (in some way) to be understood in by appeal to counterfactuals. But if these two authors are correct, perhaps there is nothing wrong with an HF account of stochastic mechanism according to which the chance of a mechanism firing is grounded on counterfactuals. Counterfactuals already feature in our causal explanations.

I cannot here present anything close to a detailed case against counterfactual analyses of causation. That said, there are many well-known objections to them—objections I find convincing enough to raise serious doubts about whether they constitute grounds for rejecting my argument. It's far from clear, for example, whether Lewis's counterfactual analysis can deal with causal preemption cases.[21] But even if this were not the case, there are other reasons why we might disagree that the notion of causation necessary involves an appeal to counterfactuals. To illustrate this, consider a few of Woodward's own remarks in the opening pages of his (2003) book, *Making Things Happen A Theory of Causal Explanation.* He says, "The account that I present is not reductive…" (Woodward 2003, 20). He adds that his account is set up to "…test or elucidate the content of particular causal and explanatory claims" (Ibid, 22). And "…the theory should enable us to make sense of widely accepted procedures for testing causal and explanatory claims" (Ibid, 24). If we look carefully at these claims, we can see that Woodward's account is not meant to tell us *what causation is.* It's explicitly nonreductive. Rather, on Woodward's own

[21] Lewis's own solution to preemption cases (1973b) is to appeal to a notion of 'causal chain' which is itself in want of analysis.

admission, his counterfactual analysis is meant to provide a theory for *testing causal claims.* One could absolutely agree with Woodward that *testing* causal claims involves seeing what would have happened if the purported causal event had not occurred (or occurred differently)—but nonetheless disagree that causation, itself, necessarily has anything to do with counterfactuals. Indeed, this is precisely the point made by Anscombe when she advocates her analysis of causation as a brute fact. She writes (1971), "If *A* comes from *B,* this does not imply that every *A*-like thing comes from some *B*-like thing or set-up or that every *B*-like thing or set-up has an *A*-like thing coming from it; or that given *B, A* had to come from it… Any of these may be true, but if any is, that will be an additional fact" (Anscombe 1971, 388). For Anscombe, causation at its core consists simply and brutally as a "derivativeness". As she says, "effects derive from, arise out of, come of, their causes". No further analysis of causation is needed or possible. If Anscombe is correct, and I suspect she is, then causation need not be understood counterfactually. And if causation need not be understood counterfactually, then counterfactuals need not figure into causal explanations. And if counterfactuals need not figure into causal explanations, then the relevance problem for HF accounts of stochastic mechanism may well stand.

Given these considerations, I maintain that—if possible—we should avoid an analysis of stochastic mechanism according to which their chanciness is understood on an HF theory of chance. HF chances are grounded in counterfactuals. But it is far from obvious how counterfactually grounded chances can play any causal-explanatory role in the explaining the actual world. And since scientists appeal to mechanisms to explain the actual world, we have good reason for rejecting an HF understanding of them.

4.3.3 Against Best-System Analysis (BSA)

Another candidate theory of chance, first put forth by David (1980, 1986), is called the Best-System Analysis (BSA).[22] According to BSA, the chance of any given outcome occurring is whatever the best systematization of the Humean mosaic of particular facts tells us it is. What makes a particular systematization of the Humean mosaic of particular facts better than all the others? On a BSA view, the best system is the one which achieves the most balance between the theoretical virtues of simplicity, strength (informativeness), and fit to the data—where the three are thought to tradeoff in some fashion. On a BSA understanding of chance, therefore, the stochasticity of a mechanism should also simply be whatever the best system tells us it is.

Prima facie, there are many challenges for a BSA theory of chance. How are we to understand how these three virtues trade-off? By what measure are we supposed to determine which is the system that achieves the most balance between these virtues? And what, precisely, is meant by 'simplicity'? Is it the number of entities postulated in a given system? Is it the number of variables required to formulate the axioms of a given system? Is it the number of predicates used to describe a given system? Doesn't it matter what language we use to describe the system? Is there any hope of achieving a canonical language where all of its predicates correspond perfectly to natural kinds (as Lewis thought we could)? Put these questions aside for the moment, and assume that a coherent version of the BSA is achievable.[23] I argue that, even still, the BSA theory of chance is not amenable to our notion of stochastic mechanism.

My argument is this:

P13. On a BSA account of stochastic mechanism, the chance of a given mechanism firing

[22] Lewis first articulated a best-system analysis of laws (1973a) and later extended it to apply to chance (1986).
[23] Hoefer (1997) and Cohen and Callender (2009) have made considerable efforts to save the BSA account of lawhood which might be extended to apply to the BSA account of chance. That said, I still believe (for reasons outside the purview of this paper) that they have fallen short of offering and articulating a coherent BSA analysis.

> is whatever our best systematization of the Humean mosaic of facts (the most balanced between simplicity, strength, and fit to the data) tells us it is.
>
> P14. However, given (P13), the chances we ascribe to mechanisms arise from merely systematizing particular facts and thereby cannot causally explain these particular facts (as *CAUSAL EXPLANATION* requires).
>
> C4. Therefore, a (BSA) account of stochastic mechanism fails to give us what we want from an account of stochastic mechanism.

Here, again, the middle premise (P14) needs support.

Recall again that *CAUSAL EXPLANATION* requires that, any adequate account of stochastic mechanism must allow for descriptions of underlying mechanisms to feature in causal explanations of regularities seen in nature. It strikes me, however, that a BSA account renders difficult to imagine how this is supposed to take place. Recall that the best system is the one that systematizes all of the local facts in the most balanced fashion, and the chances are whatever the system says they are. But which are the sorts of local facts that would inform the attribution of chances in the best system? It seems to me, the relevant facts must (at least much of the time) be the frequencies of particular kinds of events. If the Humean mosaic of particular facts includes the fact that roughly half of coins tossed have landed heads, then the system with the most simplicity, strength, and fit should ascribe a chance of .5 to a coin landing heads. Like an actual frequentist view, best-system chances depend on what the frequencies happen to have been. But, if the chances depend on the actual frequencies, then it becomes unclear what explanatory work the mechanisms with these chances can do by way of explaining those frequencies. As Abrams puts it, "Best System probabilities sometimes depend on whatever the frequencies happen to be, without requiring that these frequencies have any causal explanation at all" (Abrams 2012, 3). We want stochastic mechanisms to causally explain regular frequencies observable in nature. On a BSA view, however, the stochasticity we attribute to a mechanism already depends on the

known frequencies. Put another way, the best system *systematizes* the local facts. *It doesn't explain them*—at least not in the way a mechanist requires.[24]

A proponent of the BSA might respond in the following way. Yes, BSA chances depend on whatever the local matters of fact happen to have been. And some of these facts will be the very mechanism output frequencies that scientists aim to explain by appeal to chancy mechanisms. However, the BSA proponent might point out that, on Lewis's original view, causal facts are also part of the Humean mosaic. That is, causal dependencies are counterfactual dependencies, and counterfactual dependencies are grounded in BSA laws, which also arise from the best systematization of local facts. So, BSA chances *are* causal—at least in the sense that they are ultimately grounded in causal facts.

By way of response to this, I'll agree that BSA chances (understood in the above Lewisian terms) may ultimately be *grounded* in causal facts. But, it still isn't clear to me that this renders them capable of *causally explaining* any of these facts. To see why, think about what it means to give a causal explanation. Explanations are answers to 'why' questions. And causal explanations are answers to 'why' questions that proceed by identifying the cause of the phenomenon in question. But are stochastic mechanisms with BSA chances capable of doing this? BSA chances *depend* on local matters of fact: some of which are causal. But all of these causal matters of fact are already known. That's what allows them to be systematized. How can the chances that supervene on these facts we already knew add anything explanatory regarding these facts? I can't see an easy way.

Despite my inability to see how BSA chances can play any causal-explanatory role in explaining mechanism output frequencies, it should be noted that some proponents of the BSA

[24] It may be that a BSA analysis of stochastic mechanism would allow for other types of explanation (unificationism perhaps). But what I suggest here is that life scientists seek the sort of explanation where describing the underlying causal structure of an observed fact is what does the explaining.

approach have attempted to make strides in this direction (cf. especially Helen Beebee 2000). However, even if Beebee is correct that there is some sense to be made of the explanatory role of BSA chances, I maintain that there are sufficient other problems with the BSA approach to make it an unattractive candidate for underpinning our account of stochastic mechanism. The most forceful reason is this. The BSA places an incredibly important role on theoretical virtues (simplicity, strength, and fit). However, many have noticed that, in the realm of biology, there's no reason to expect the facts should conform (or even be constrained by) these theoretical virtues (cf. especially Sober 1990). Why should we expect the world to conform to the best way we can think of to systematize facts? Why should it matter one bit whether a systematization of facts is simple? Sure, this makes it easier for us to understand. More than anything else, however, parsimony is an *aesthetic* virtue (see Craver 2007 for this point). If the life-sciences tell us nothing else, they show that the natural world is far more complicated than we could ever imagine.

4.4 A Propensity-Backed Account of Stochastic Mechanism

What I hope to have shown in the forgoing sections is that none of the theories of chance heretofore considered (subjectivism, frequentism [both actual and hypothetical], or a Lewisian best-system analysis of chance) is capable of cohering with what we want from an account of stochastic mechanism. There are, however, other theories of chance left to be explored, namely, traditional propensity theories. My aim for this section is to show that a version of propensity theory is the best theory of chance to pair with stochastic mechanism; at the very least, it doesn't fail to meet the desiderata laid out in 4.2. It also has the added benefit, I suggest, of pairing nicely with the many senses of irregularity outlined in Chapter 2.

On a propensity theory, chance is a dispositional property or tendency of a given type of physical situation to produce certain outcomes over others.[25] Following Gillies (2000), we can distinguish between *long-run* and *single-case* accounts of propensity. On a long-run account, a given event has a chance *c* if and only if the kind of experimental setup which can generate this type of event possesses a dispositional property to generate this kind of event with a corresponding characteristic relative frequency in the long run of trials of the setup. On a single case account, a given event has chance *c* if and only if the actual token experimental setup which can generate this event possesses a dispositional property (tendency) with that degree of strength. Because of the *ONE-OFF* desideratum set forth in 4.2, however, I'll adhere to the single-case account.

So what might a single-case propensity-backed account of mechanism look like? It might go something like this:

> Propensity-Backed Stochastic Mechanism: a stochastic mechanism has chance *c* if and only if the actual token mechanism setup which can generate this event possesses a dispositional property (tendency) to produce that outcome with that degree of strength.

What I hope to show now is that this is the *only* type of account (of those considered) which doesn't fail any of the desiderata set forth in 4.2. Rather than an all-out articulation or defense of a propensity-backed account of stochastic mechanism, however, my only hope here is to show (i) that a propensity theory (in its most general form) passes muster with regard to our set of desiderata; and (ii) that it pairs nicely with the kinds of regularity distinguished in the previous Chapter.

[25] The origins of this type of account can be traced back to Peirce (1910) and Popper (1957).

4.4.1 Does a Propensity-Backed Account of Stochastic Mechanism Fail our Desiderata?

Starting with *COHERENCE*, we might ask the following: Does a propensity-style account of stochastic mechanism fail to cohere with the general practice of scientists searching for mechanisms to explain puzzling phenomena in the natural world? The answer, it seems to me, is no. When molecular biologists search for the mechanism for genetic mutation or protein synthesis, what they might well be looking for is a structure in the world that itself has chancy properties. When evolutionary biologists speak of the chance of natural selection endowing adaptive characteristics to a given population, they plausibly take this chanciness to be a feature of the mechanism of natural selection itself; the same goes for the release of electrical activity in post-synaptic neurons. This process fails up to 90% of the time, a neuroscientist might suggest, because the mechanism itself has chancy properties. I certainly do not claim to have access to what scientists actually mean when they use the term. Rather, I am content to suggest here that a propensity-backed account of stochastic mechanism is *capable* of cohering (without any glaring inconsistency) with what scientists actually do when they search for mechanisms to describe puzzling phenomena.

Furthermore, a propensity account of stochastic mechanism, by virtue of the fact that it locates the chanciness of a mechanism in the world, does not run into the problems associated with an AF or HF account. A propensity-backed stochastic mechanism does not define the stochasticity of a given mechanism as a relative frequency of outputs (whether actual or hypothetical)—so has no problem dealing with one-off mechanisms. And under a propensity account, their stochasticity is an objective feature of the actual world. They are, therefore, perfectly well-suited for grounding causal explanations without appeal to counterfactuals.

Finally, a propensity account of stochastic mechanism doesn't require that the chances we ascribe to mechanisms arise out of the best systematization of local facts. And this means that it does not suffer from the problem we found with a BSA understanding of stochastic mechanism. That is, propensity-backed stochastic mechanisms don't have the same trouble explaining output frequencies as a BSA understanding of stochastic mechanism seemed to have, because these propensities aren't constrained by these very frequencies.

4.4.2 Propensity-backed Mechanisms and Andersen's Taxonomy of Regularity Revisited

In Chapter 2, recall, we explored Andersen's taxonomy of mechanistic regularity—adding to it along the way. It's worth pausing for a moment and pointing to the fact that a propensity understanding of stochastic mechanism has some prima facie advantages with regard to cohering with some elements of that taxonomy.

Andersen's taxonomy of mechanistic regularity, recall, separates out four organizational locations where a mechanism can behave (or fail to behave) regularly:

> R1) Regular occurrence of [a mechanism's] set-up conditions,
>
> R2) Regular triggering of the mechanism once set-up conditions occur,
>
> R3) Regular operation of specific activities connecting the entities within a single mechanism,
>
> R4) Regular production of termination conditions once the mechanism has been triggered. (Andersen 2012, 419, my use of 'R', and my use of 'set-up' rather than 'start-up')

It also cites four levels regarding the strength of connection between these mechanism stages:

> (a) Deterministic: cause cannot fail to act once conditions are appropriate
>
> (b) Reliable but not exceptionless: most of the time, the cause brings about the effect, but there are occasions on which it does not, and we may or may not be able to provide an explanation for the exception

> (c) Sporadic: the cause fails to act often enough that it cannot be considered merely an exception when this failure occurs
>
> (d) Infrequent: most of the time, the cause fails to bring about its effect, but once in a while it does. (Ibid, 420)

In support of this taxonomy, I suggested (in 2.4.1) that each of these senses of regularity applies in some way to the probabilistic phenomena targeted for mechanistic explanation (synaptic transmission, protein synthesis, DNA replication, Mendelian segregation, and natural selection). For that reason, it would count in favor of a propensity understanding of stochastic mechanism if it could shed some further light on this aspect of Andersen's taxonomy of regularity. My suggestion is that it can.

The main reason why I think a propensity-backed understanding of stochastic mechanism coheres especially nicely with Andersen's taxonomy is this: *propensities are dispositional properties of physical situations. These dispositional properties capture/are sensitive to/partially depend on the many ways a physical process might go wrong.*

Let me illustrate this with a simple example. The glass picture window in my living room is fragile—which is to say that it has a dispositional property to break relatively easily when exposed to forceful perturbations. Put another way, it has a propensity to break easily. The very notion that my window has such a propensity contains within it a lot of background knowledge about the many ways that it could break. Bricks heaved, baseballs hit, birds flying, hurricane winds: these are all ways my picture window could break. To say of my picture window that it has a fragile disposition, to say that it has the propensity to break easily, is to affirm (at least some) implicit knowledge of the many ways it might be shattered.

The same is true of mechanisms. When I say of a synaptic transmission mechanism that it has the propensity to fail more often than it succeeds, I am affirming (at least some) implicit

knowledge that there are many ways it could go wrong. That propensity, when correctly ascribed, is capable of capturing these varied possibilities. Note that several of the other candidate understandings of chance don't seem to capture such information, at least not nearly as easily. On an AF understanding of chance, as we've seen, the chance of a mechanism firing is just the actual relative frequency of its successful firings. No more further explanation of why this frequency has obtained is available. The BSA interpretation's insistence on supervening on actual facts makes it more difficult to incorporate any knowledge of modal facts regarding how the world might have been. The only other interpretation considered thus far that could match this explanatory advantage is HF. But, as we've seen, HF seems to have other problems.

Propensities carry within them information about the many myriad of ways the world *might* go. This notion will be made much more precise in the next Chapter. But for now, it suffices to say that this at least provides a prima facie advantage of offering explanatory resources for accommodating the many ways a mechanism can behave (and fail to behave) regularly. And since, in the previous Chapters, this proved to be a helpful resource for addressing the regularity question, it seems this virtue is quite worth noticing.

4.4.3 The Upshot for a Propensity-backed Account of Stochastic Mechanism.

The upshot, therefore, is that a propensity theory is the most promising philosophical understanding of chance (of the ones heretofore considered) to give us what we want from an account of stochastic mechanism. Of course, much more will need to be said about the details of this propensity account—work that I shall take on in the next Chapters. For now, let's consider some immediate objections to propensity views of chance and see what can be said in reply.

4.5 Three Objections and Replies

The goal of this section is to consider a few serious objections to propensity theories of chance and show why they are not devastating to the particular propensity-backed account of stochastic mechanism on which we have been zeroing in thus far.[26] To do this, I draw from an important line of thought put forward in both Marshall Abrams's recent account of "Mechanistic Probability" (2012) and Michael Strevens's "Probabilities out of Determinism" (2011).

4.5.1 Single-case Propensities Cannot Explain Frequencies

Recall that the propensity theory we have tentatively chosen to back our account of stochastic mechanism is a view according to which a stochastic mechanism has chance *c* if and only if the actual token mechanism setup which can generate this event possesses a dispositional property (tendency) to produce that outcome with that degree of strength. Recall, further, that part of the goal of a propensity-backed account of stochastic mechanism is to allow these underlying mechanisms to causally explain regularities observed in nature. In effect, then, what *CAUSAL EXPLANATION* requires is that stochastic mechanisms have to be able to explain the regular frequencies of a given event observed to occur in nature. Many have questioned, however, whether propensities connect up with frequencies at all.[27] The worry is this: if chance is just a propensity of an experimental set-up to produce a particular result on a specific occasion, then there does not seem to be any conceptual connection to the stable, long-run frequencies we observe in nature. It seems, then, that (at least a single-case) propensity-backed view of stochastic mechanism cannot meet a key one of our desiderata.

---

[26] There are several other objections on which I could have here focused. Due to constraints on space, however, I've chosen those I take to be the most serious. See Eagle (2004) for others.

[27] C.f. objection 14 "The Horizontal/Vertical Problem" and Objection 17 "The Method of Pure Postulation" in Eagle 2004.

4.5.2 Propensities are Unobservable and Un-underminable, and as Such, are not Amenable to Empirical Science

Because, as we have just seen, (at least single-case) propensities have no obvious relation to frequencies, it seems impossible to undermine any purported propensity by observing any (apparently) anomalous frequency. Suppose, for instance, that a scientists conjectures that a stochastic mechanism has a propensity of .7 of producing a given outcome. But upon running several tests, the frequency of successful firings is much lower. Since there is no connection between propensities and frequencies, there appears to be no reason to give up on the conjectured propensity. Furthermore, propensities (like other dispositional properties) cannot be directly seen with the naked eye, and so cannot be an object of investigation for empirical science. This means that a propensity view of stochastic mechanism fails another one of our own desiderata (*COHERENCE*).

4.5.3 Stipulating what Propensities Do is not Enough to Show that They Exist

Propensities are able to give us what we want from a theory of chance, in large part, because proponents have stipulated that they are precisely the sorts of properties to play the roles we desire from an account of chance. However, as Hajek points out, "To be sure, one can *stipulate* that they do so, perhaps using that stipulation as part of the implicit definition of propensities. Still, it remains to be shown that there really are such things — stipulating what a witch is does not suffice to show that witches exist" (Hajek 2011). Put another way, one might object that propensity accounts seem to give an empty account of chance, à la Molière's 'dormative virtue'. [28]

[28] Sober makes this point. (2000, 64).

4.6 Drawing on a Recent Line of Thought from Abrams and Strevens

Before explicitly addressing these important objections to propensity accounts of chance, I briefly draw from a line of thought recently formulated by Marshall Abrams (2012) and Michael Strevens (2011)—which I will, in the following section, make use of to address the forgoing objections.

In their respective papers, "Mechanistic Probabilities" (2012) and "Probabilities out of Determinism" (2011), Abrams and Strevens both argue for a general interpretation of probability according to which global determinism need not be a threat to objective probabilities. My current project of seeking a working characterization of stochastic mechanism differs largely in scope and details from both of these authors. Specifically, the project I am here undertaking is different from theirs in one important respect: I am not seeking a *general* interpretation of probability, but rather a more limited understanding of stochastic mechanism that coheres with scientific practice. That said, there is one move made by both of the above authors which will prove quite helpful in addressing the above objections to propensity theory. This move, I suggest, is quite independent of the parts of their respective theories about which we differ, and can be quite easily separated from the rest of their theories. The main idea is this: *stable output frequencies of a mechanism can be adequately explained entirely by reference to facts about the mechanism itself and the kinds of initial conditions feeding into it.*

In his 2012 paper "Mechanistic Probabilities", Abrams is interested in the probabilities that we assign to outputs resulting from certain kind of causal devices. Some devises, Abrams suggests, have a causal structure such that it matters very little what pattern of inputs the device is given in repeated trials. The pattern of outputs remains roughly the same. Take, for example, a

standard (fair) roulette wheel with equally-sized wedges alternating between red and black. He writes,

> …if the ratio of the size of each red wedge to that of its neighboring black wedge is the same all around the wheel, then over time such a device will generally produce about the same frequencies of red and black outcomes, no matter whether a croupier tends to give faster or slower spins of the wheel. (Ibid, 349)

Why is this? He answers,

> The wheel of fortune divides a croupier's spins into small regions [which Abrams calls "bubbles"] within which the proportion of velocities leading to red and black are approximately the same as in any other such region As a result, as long as the density curve of a croupier's spins within each bubble is never very steep, the ratio between numbers of spins leading to red and leading to black within each bubble will be roughly the same. The overall ratio between numbers of red and black across all spins will then be close to the same value. In order for frequencies to depart from this value, a croupier would have to consistently spin at angular velocities narrowly clustered around a single value, or produce spins according to a precisely periodic distribution. (Abrams 2012, 350)

Abrams illustrates this point with the following picture:

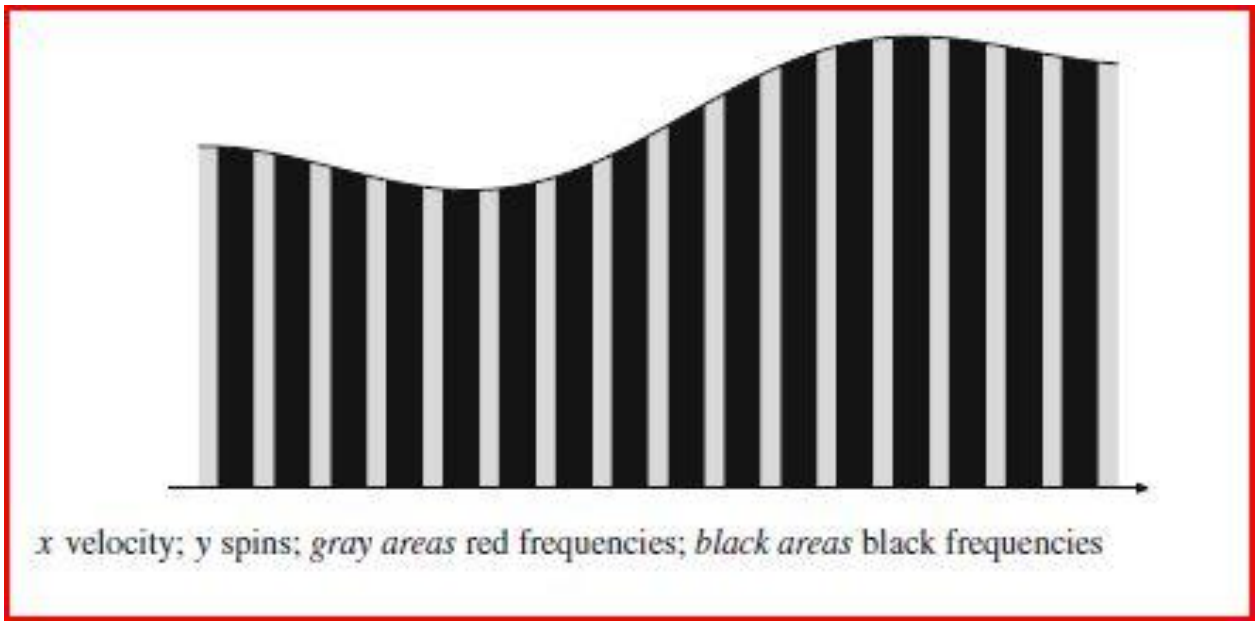

Figure 9. Roulette wheel output frequency distribution (from Abrams 2012, 350).

He then goes on to describe the general features that a device (like the roulette wheel) has to have in order to have this peculiar characteristic. He calls it a causal map device (figure 10).

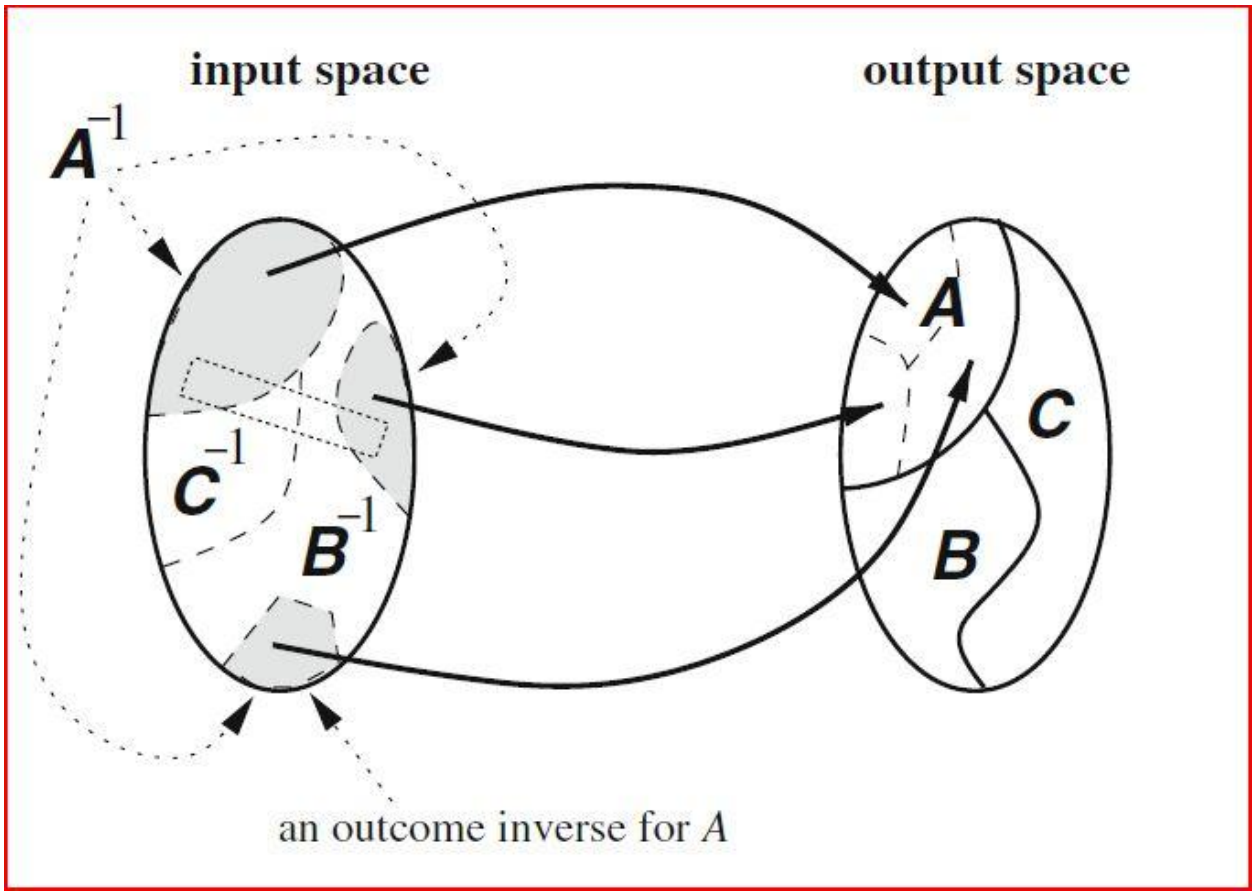

Figure 10. Causal map device (From Abrams 2012, 350).

Without getting too bogged down in the details of Abrams (very complicated) account, the short story about what he is doing is as follows. He is trying to construct an account of a certain kind of probability that exists in certain situations. In his words, he seeks to: "describe a realist, ontologically objective interpretation of probability, 'far-flung frequency (FFF) mechanistic probability'. FFF mechanistic probability is defined in terms of facts about the causal structure of devices and certain sets of frequencies in the actual world" (Ibid, 340).

Why focus on Abram's FFF account? What is important about this account, for our purposes, is this: stable output frequencies of a particular mechanism get explained by mapping and describing the underlying mechanism and delimiting the kinds of inputs it requires. No more explanatory work is needed.

This point is echoed by Strevens (2011) who also seeks to come up with a general determinism-friendly interpretation of probability and so, like Abrams, differs in scope from my own project. However, Strevens makes a similar move when explaining stable frequencies. He begins by asking us to imagine a coin-toss. He writes,

> In the case of the coin toss… robustness [or output frequencies] is explained by certain dynamic facts about the physiology of the coin-tossing together with some statistical information about actual pre-toss states. The same is true more generally. (Strevens 2011, 31)

The relevant point here is the same: we can explain stable output frequencies of a given mechanism (in this case, a human coin-tosser) entirely with reference to facts about that mechanism (in this case, facts about human physiology) and statistical information about the kinds of preconditions feeding into the mechanism (e.g., background conditions of the toss).

Having pointed to a couple of features of Abrams's and Streven's recent accounts that I hope will prove helpful for addressing the objections to propensity theories of chance, let me pause to make explicit a few additional ways in which I'm departing from both of these authors. First, unlike my account, Abrams's account of mechanistic probability is a version of actual frequentism. He writes, "FFF mechanistic probability is a kind of actual frequency theory in that it depends on actual frequencies in many natural collections" (Abrams 2012, p. 370). Second, Abrams explicitly states that his account of mechanistic probabilities is not meant to apply to single trials. He says, "there is no sense in which mechanistic probabilities should be seen as

probabilities of outcomes of single trials" (Ibid, 370). Strevens's account also differs from my own in that he explicitly affirms its dependence on statistical information about inputs.

Given these disparate features of their accounts from my own, an objection might arise. Recall that one of the reasons I gave for adopting a propensity-backed understanding of stochastic mechanism was that such an account needn't depend on the very frequencies that the mechanism was meant to explain. However, on both Abrams's and Strevens's accounts, there is an explicit dependence on frequencies. And there is no mention of propensity. By way of response to this, I suggest that the roles that the frequencies play in both of these author's accounts can be taken on by the propensity—thereby eliminating any reliance on frequencies. As emphasized in 4.4.2, propensities (by their very nature) carry with them modal information about what would happen to their possessors given certain kinds of interactions with the world. The fact that my picture window is fragile means it has a propensity to break relatively easily when perturbed in various ways. The assignment of a propensity to break easily already carries information about what kinds of inputs to it would result in what kinds of outputs. The same goes for mechanisms. Assigning a propensity to a synaptic transmission mechanism to fail more often than it succeeds already carries with it information about what kinds of inputs lead to successful initiation of electrical activity in the postsynaptic neurons. My suggestion, therefore, is this. The extent to which Abrams and Strevens both rely on statistical information about input and output frequencies can be taken over by the propensity we ascribe to the mechanism. And if this information can be taken over by the propensity, then there is no need to depend on these frequencies in a way that makes stochastic mechanisms explanatorily inefficacious.

But even if the roles that frequencies play for Abrams and Strevens can be taken over by propensities, why should we want them to be? Doesn't this just add a further, more mysterious

element to their already effective theories? The answer I propose is that it is worthwhile to relegate the explanatory work from the frequencies to the propensity because of the advantages I've already attributed to propensity theories in 4.3. Namely, a propensity-backed understanding of stochastic mechanism passes muster with regard to the desiderata laid out in 4.2. Whereas, frequency accounts do not.

4.7 Objections Revisited and Addressed

My view is that the three serious objections briefly presented in 4.5 are not devastating to the propensity-backed account of stochastic mechanism I've heretofore been attempting to motivate. To show this, I'll take each one in turn and show what responses I see as available to a proponent of a propensity account of stochastic mechanism. The main thrust of my responses to each of the above objections is the following: while each of these challenges may pose a significant threat to propensity interpretations of chance in general, conjoining propensities with mechanisms provides (at least some resources) for a novel way out of each objection. In other words, *propensities as properties of mechanisms are explanatorily stronger than either is alone*. Let me say why.

Recall that the first charge against propensity accounts of chance is that they are incapable of explaining frequencies. This would be a major problem for my account of propensity-backed stochastic mechanism because my own *CAUSAL EXPLANATION* desideratum requires that mechanisms be able to explain regularities (output frequencies) observable in nature. To address this serious issue, I suggest we follow the line of thought just attributed to Abrams and Strevens and focus on features of the mechanisms themselves. In short, what I contend is this: while (at least single-case) propensities alone may be incapable of explaining frequencies, facts about the mechanisms themselves—seen as instantiating these

propensities—gives us additional resources to undertake this explanatory work. Here, an example from genetics is instructive. Gregor Mendel, a monk and schoolteacher living in isolation in Brno, experimented with peas in his monastery garden. In doing so, he made invaluable strides towards understanding the mechanisms that rule heredity. This discovery came about between 1856 and 1863 as he traced the inheritance patterns of certain traits in pea plants. When he crossed two pea plants with different traits (for example, a yellow pea producing plant with a green pea producing plant), he came to see a pattern in the results. He noticed strong tendency for a 3:1 ratio of yellow to green characters in (what we now call) the F_2 generation. What Mendel initially observed was a *frequency* of outcomes such that it would have been natural for him to conclude that the chance of a hybrid pea plant to produce yellow peas in the F_2 generation was .75. On a propensity view of chance, this means that some features or properties of the physical set-up conditions give rise to a tendency for these frequencies to result. However, as the objection we have considered points out, postulating this propensity alone does not seem to *explain* the frequencies Mendel observed.

What I wish to point out, on the other hand, is that conceiving of these propensities as instantiated by the underlying mechanism of Mendelian segregation allows for a fully adequate explanation of the resulting frequencies. In this case, what is needed is a description of alleles—one coming from the father and one coming from the mother—which segregate during meiosis. In the pea plant example, these alleles combine in various ways such that some of the seeds are homozygous and some are heterozygous. To clearly visualize how alleles are distributed amongst descendents, we can use a Punnett Square:

	A	A
A	AA	Aa
a	aA	Aa

Recall from Chapter 1, that in the case of the green and yellow peas, the yellow character is dominant, so it is represented by the uppercase 'A', and the green character is recessive, so it is represented by the lowercase 'a'. The upper case 'A' always dominates over the lower case 'a', so the only time we get green peas is when there are two lowercase 'a's. In the initial generation, the yellow pea plant will give each of its progeny a dominant yellow allele, and the green pea plant will give each plant a recessive green allele. So all the pea plants in the F_1 generation will be yellow heterozygous. But in the F_2 generation, when the yellow heterozygous peas have interbred, the resulting progeny could have any of the combination of alleles (as illustrated by the Punnett Square)—showing a ratio of 3:1 yellow to green phenotypes. In short, the mechanism works like this: 'A' and 'a' segregate in the formation of germ cells, which then combine randomly at fertilization to give offspring according to the formula: 1AA + 2Aa + 1aa, and this mechanism explains the 3:1 ratio that Mendel observed.

This example shows that attributing the 3:1 ratio of yellow to green peas in the F_2 generation to a propensity alone may not seem to explain the frequency of yellow and green peas. But when these propensities are seen as instantiated by the mechanism of Mendelian segregation of alleles, the frequency gets fully explained by explaining features of the mechanism itself. Just as Abrams explains the consistent outputs of a roulette wheel by appeal to the features of the wheel itself, we can explain the output frequencies of pea plant crossings by

appeal to the mechanism of Mendelian segregation itself. And in this way, propensities together with mechanisms *can* explain frequencies.

Recall the second objection: that given the dubious connection between propensity and frequency, there is reason to doubt whether any conjectured propensity could be undermined by the observation of a surprising frequency. And furthermore, propensities are not directly observable and as a result not amenable to empirical science. My response here is similar. As suggested by Abrahms and Strevens, there needn't be any explanatory gap between propensity and stable output frequencies. Facts about the mechanism (and its initial conditions) bridge that gap. And if this bridge is sound, then anomalous output frequencies can give us reason to revise our initial propensity assignment—either that or reason to revise our background beliefs about the mechanism. Furthermore, although propensities cannot be observed directly, the mechanisms that instantiate them surely can. And if facts about these mechanisms are what ground the propensities they instantiate, then science can do enough by observing these mechanisms.

Regarding the third objection that we've stipulated (in some ad hoc fashion) the roles that propensities play without ever having shown that they actually exist. To fully address this objection would require diving head-first into the realm of the ontology of dispositional properties. I cannot do that here. So I'll say two things: (1) I do not see any reason to hold propensities to be any more mysterious than other every-day dispositional properties. How do you show that something is fragile? You do different things to it, and see how easily it breaks. How do you check if something is poisonous? You poke tiny bits of poison into test subjects and see what happens. Likewise, we can check to see if a mechanism has a certain propensity by testing it. Initiate its start-up conditions and see how often you get the output you expect. If the frequency stabilizes, you can reasonably conclude that it has the corresponding propensity. No

deep mystery, at least not any special one. And (2), I'll go in Chapter 5 to give a concrete suggestion as to one way we might describe what a mechanistic propensity actually *is*. I'll suggest that it can be metaphysically characterized in terms of probability-weighted possibilia. If this suggestion is plausible, it will show that, indeed, more can be said about propensities than simply what role they play.

4.8 Conclusion

In this Chapter, I have attempted to take steps towards a working characterization of stochastic mechanism: one that coheres with the way life scientists actually seek out mechanisms to explain the natural world. To do this, I have taken the MDC characterization of mechanism and held it up to the main types of philosophical theories of chance (subjectivism, frequentism, Lewisian best-systems, and propensity). If the forgoing arguments are sound, then a propensity theory is the most promising type of theory of chance to give us what we want from an account of stochastic mechanism. To further motivate this propensity-backed account of stochastic mechanism, I have taken three serious objections to propensity theories of chance and shown that, following a line of thought from Abrams and Strevens, they can be avoided.

Chapter 5. A Propensity Interpretation of Stochastic Mechanism: Lessons from Fitness and Drift

Abstract: The goal of this chapter is to introduce and motivate a propensity interpretation of stochastic mechanism (PrISM). I approach this, first, by examining some recent propensity interpretations of fitness and drift to see what lessons might be learned from them. I then lay out my own formal characterization of the propensity interpretation of stochastic mechanism and show how it enjoys an important advantage over propensity interpretations of fitness and drift. I conclude by hinting at how the PrISM might apply to one of the biological phenomena targeted for mechanistic explanation introduced in Chapter 1.

5.1 Introduction

The thesis pursued thus far in this dissertation is that it makes good sense to think of the truth of (at least some of) the probabilistic generalizations made in the life sciences as grounded on biological mechanisms in the world, biological mechanisms that underlie and produce these observable phenomena, biological mechanisms which are themselves (in some sense) chancy: *stochastic mechanisms*. To arrive at this answer, we have seen in Chapter 2 that mechanistic regularity comes in various senses, and while regularity should not be considered a metaphysically demarcating feature of mechanisms, mechanisms with more regularity should hold a privileged position in our scientific explanations. We have also seen in Chapter 3 that appreciating these varied senses of regularity helps clarify whether natural selection should count as an explanatorily privileged MDC mechanism. In Chapter 4, we saw that there seem to be significant problems with underpinning an account of stochastic mechanism with other philosophical analyses of chance besides a propensity interpretation. But, as of yet, not much has been said about what, precisely, such a propensity interpretation of stochastic mechanism would amount to. I address that here.

In what follows, I draw a few important lessons from recent propensity interpretations of fitness and drift (PIF and PID) in order to present a novel propensity interpretation of stochastic mechanism (PrISM) according to which stochastic mechanisms are thought to have probabilistic

propensities to produce certain outcomes over others. This understanding of stochastic mechanism, once fully fleshed-out, will provide the benefits of (1) allowing the stochasticity of a particular mechanism to be an *objective property* in the world, a property investigable by science, (2) a way of *quantifying* the stochasticity of a particular mechanism, and (3) a way to *avoid committing to the problematic causal role of propensities* (and dispositional properties in general).

Here is my plan. In 5.2, I discuss biological fitness and drift, paying close attention to some of the key features of Grant Ramsey's recent propensity interpretations of both. In 5.3, I motivate and defend three lessons from these recent accounts: two positive and one negative. In 5.4, I offer some considerations to bolster the negative lesson. In 5.5, I follow these lessons to lay out my own formal characterization of the PrISM—addressing potential objections along the way. And, in 5.6, I demonstrate how the PrISM might work by suggesting how it applies to the phenomena of initiation of electrical activity in postsynaptic neurons.

5.2 Motivation from Fitness and Drift

A natural place to start looking for resources for developing a propensity interpretation of stochastic mechanism is by looking at other concepts in biology that have been given propensity interpretations and see what lessons we might learn from them. Two biological concepts that have received recent propensity interpretations are fitness and drift.

5.2.1 Fitness

Biological fitness is a probabilistic notion. Intuitively, it seems that there are many ways an organism's life might turn out depending on its particular genome and how it interacts with its

environment[29]: some of these possibilities resulting in many progeny; some not. Beginning with Brandon (1978) and Mills and Beatty (1979), the propensity interpretation of fitness (PIF) has been defended by several philosophers of biology over the past few decades[30]. In its most general form, the PIF holds that an organism's fitness is its probabilistic propensity to produce offspring. But why think of fitness this way? One of the primary motivations for conceiving of fitness in this manner is to avoid, what has been called, the "tautology problem"[31]. If fitness is not conceived of as a probabilistic propensity, and is instead defined in terms of an organism's actual number of offspring (also known as 'realized fitness'), then fitness *cannot explain* these actual outcomes in any way that isn't tautologous and thereby vacuous. In other words, if an organism's fitness is defined by how many progeny it actually has, no appeal to an organism's fitness level can be made to explain its particular reproductive outcome; the two are by definition the same. Just as realized fitness cannot explain an organism's actual reproductive outcome, neither can it serve as the basis for *predicting* a living organism's reproductive outcomes. On a realized fitness view, an organism's fitness level can only be determined after it has finished reproducing; so there can be no way to base predictions about reproductive outcomes on an individual organism's fitness level. Similarly, if fitness is defined as the actual number of progeny that an organism produces (and not as its propensity to produce offspring), no adequate distinction can be made between the *property of being fit* and the *outcome resulting from being fit*; the two are by definition one and the same. Furthermore, unless fitness is distinguished from actual reproductive outcomes, we cannot think of fitness playing a *causal* role in how many progeny an organism has. We cannot say, for example, that an organism had many progeny *because* it was

[29] Following Sober (2010) and Ramsey (2012), I take it that this need not constitute a denial of metaphysical determinism. More on this is in Chapter 6.
[30] Cf., Brandon and Carson (1996) and Beatty and Finsen (1989)
[31] Mills and Beatty (1979), and Pence and Ramsey (2013)

very fit—at least not where 'because' is understood causally. The PIF was thereby introduced as a dispositional property of organisms—one that is ostensibly capable of (A) explaining the actual number of offspring an individual organism produces, (B) grounding predictions regarding the number of progeny an organism produces, (C) grounding a distinction between fitness as a property versus the outcomes that result from an organism's fitness, and (D) underpinning an understanding of fitness as causal.

Given these sorts of considerations, Mills and Beatty (1978) and Brandon (1978) offer probabilistic propensity definitions of fitness of the following sort[32]:

> *PIF: x* is fitter than *y* in [environment] *E* = *x* has a probabilistic propensity >.5 to leave more offspring than *y.*

Despite its intuitive appeal, however, many have noticed significant problems with this definition. The most serious one, articulated forcefully by Rosenberg and Bouchard (2008), is that it is false. They write, "…there are many circumstances in which the organism of greater fitness has the propensity to leave fewer immediate offspring than the organism of lower fitness; as when for example, the larger number of a bird's chicks all die owing to the equal division of a quantity of food which would have kept a smaller number viable" (Rosenberg and Bouchard 2008). Put another way, it simply is not the case that an organism with a higher propensity to leave more immediate offspring will end up with the higher number of viable offspring in the end. Environmental contingencies can get in the way. In response to this problem, attempts were made to advance more abstract schematizations of this definition—or to hedge it with various *ceteris paribus* clauses—but other problems seem to crop up (cf., Sober 2001; Walsh, Lewens, and Ariew 2002; Matthen and Ariew 2002; Ariew and Lewontin 2004).

[32] Taken from Rosenberg and Bouchard (2008)

Grant Ramsey (2006) offers a novel way of characterizing the PIF, one that does not appear to suffer the problems plaguing the original PIF approach. He calls his characterization "Block Fitness". He writes, "Fitness, I will argue, is best conceived as a function of the probability distribution of all the possible numbers of offspring the individual might produce" (Ramsey 2006, 487-488). In his 2012 paper, Ramsey gives this helpful description.

> Consider an organism *O* with genome *G* in environment *E*. Assuming that *O*'s fitness is non-zero, there are a number of distinct ways that such an *O* with *G* can interact with its environment. It might be eaten by a predator early in life and die without leaving behind any progeny, or it might live a long life and leave behind a large number of progeny. Let's designate each of these possible ways *O* could live its life in *E* (henceforth *O*'s *possible lives*) with *L*. Thus *O* has a large set of possible lives, $L_1, L_2, \ldots, L_n$. Each of these possible lives will have a probability associated with it. The understanding of fitness as a propensity, then, can be explicated in terms of the properties of this set of possible lives (with their associated probabilities). Holding *E* constant, a change from one *G* to a different genome *G*! will change the properties of the L_i (i.e., different genes can lead to differences in fitness)… The fitness of *O* consists in the properties of *O*'s set of possible lives (with their associated probabilities). Fitness is thus quantified via a function on *O*'s probability-weighted possible lives. (Ramsey 2012, 6)

As seen here, Ramsey characterizes an organism's fitness as a probabilistic propensity. However, this propensity does not merely take features of an organism's actual life as its categorical base[33], but instead is a function of all of an organism's probability-weighted possible lives. More on the specifics of how this is meant to work is coming in subsequent sections. But, for now, it's worth pointing to a couple of the benefits this approach is meant to afford its adherents. Since Ramsey conceives of fitness as consisting in properties of the whole set of an organism's possible lives rather than the actual number of offspring it has, he can still maintain the benefits of the original PIF (A-D listed above). Because, on Ramsey's account, an organism's fitness is a propensity (albeit one that gets explicated in a new way), it can still *explain* the actual reproductive outcomes of an organism in a way that isn't obviously tautologous. It can still serve as the basis

[33] *Aka*: metaphysical or supervenience base

of *predicting* what the actual reproductive outcomes of its possessor will be. It can still ground a distinction between fitness as a property vs. an outcome. And (at least on Ramsey's account), it can still serve as the basis for speaking of fitness as *causally efficacious*. In addition to maintaining these benefits of the original PIF, Ramsey's account isn't subject to the same objection leveled by Rosenberg and Bouchard. Rather than characterizing one organism as fitter than another merely based on its having a higher propensity to leave more immediate offspring, Ramsey's notion of block fitness requires that an organism's fitness be a function of all of the possible ways its whole life might go. On this view, information regarding how many progeny (e.g., baby birds) can get adequately fed until reaching maturity gets included in the L_i.—thereby eliminating the kind of counterinstances described by Rosenberg and Bouchard in which having more progeny might actually result (in the end) in lower fitness.

Ramsey's PIF contains a few features that I suggest provide an excellent template for the propensity interpretation of stochastic mechanism. But before making these explicit, let's look at another biological concept regarding which it is appealing to give a propensity interpretation: drift.

5.2.2 Drift

Biologists think of drift as that which happens when a population changes over time—where these changes are *not* the result of selection. Consider, for example, a population of snails half of which are pink, and half are which are yellow[34]. Imagine further that the yellow snails are twice as fit as the pink ones because of the greater resistance to the sun that their coloring affords them. Scientists observing these snails expect, therefore, that the population of yellow to pink snails should increase from 1/2 of the population to 2/3 in the subsequent generation. However,

[34] Example owes to Roberta Millstein (1996, S15)

after observing these snails for one generation, they find that the population of yellow snails actually decreases from 1/2 to 2/5 of the snail population. Because this change in the population is not due to selection (selection would have increased the relative proportion of yellow snails), they attribute this unexpected result to drift. Despite the seeming straight-forward nature of this common understanding of drift as non-selective change in a population, however, more precise philosophical analyses of drift vary greatly. There are those who argue that drift cannot be distinguished from selection and so cannot be conceived of as a distinct evolutionary process at all (e.g., Matthen and Ariew 2002). There are those who argue that drift is a distinct evolutionary process from selection, but it can only be defined in terms of actual evolutionary outcomes (e.g., Brandon 2005). And there are some (e.g., Millstein 2002) who argue that drift is distinct from selection by virtue of the kind of process it is (indiscriminate vs. discriminate sampling of a population). The details of these debates are beyond the purview of this paper. My goal, as with fitness, is only to see whether there are any features of a propensity interpretation of drift (PID) that might be useful in fleshing out my propensity-backed understanding of stochastic mechanism.

Again, Ramsey's recent account may be instructive. Ramsey (2012) argues that, rather than a population-level process of indiscriminate sampling (a la Millstein), drift should be characterized as an individual-level probabilistic propensity which he calls 'driftability'. His characterization of driftability, as we'll see, shares many of the same features as his PIF. He writes, "Driftability, then, can be identified [equated] with intra-organismic heterogeneity [differences] in the L_i [set of an individual's possible lives] and can be quantified by a function on this heterogeneity" (Ramsey 2012, 7). As with fitness, an organism's driftability is a propensity—understood as consisting in properties of the set of possible ways that organism's

life might go. Each of these possible lives has a probability associated with it. And its overall propensity for drift is quantifiable by a function on these differences. Again, we'll say more about what this means in later sections. But for now, it's worth reemphasizing the purported benefits of conceiving of drift in this way. On Ramsey's view, we should conceive of drift in this way (once again) because it allows us to divorce the *outcomes* of evolutionary drift from the *process* of drift; it makes driftability an *objective* property of individual organisms (as opposed to a measure of our ignorance). And, as with fitness, this interpretation enables us to speak *causally* about drift (e.g., "There are more white rabbits than grey ones in this population *because* of drift.").

Rather than offering a critical examination of Ramsey's version of the PID, however, I aim instead to focus on a couple of its key features as starting points for my own propensity interpretation of stochastic mechanism.

5.3 Three Lessons from Fitness and Drift

Before adducing lessons from Ramsey's PIF/PID, let's pause and take stock of where we are in my overall project. Recall that I have taken it as my goal in this dissertation to develop a way of grounding probabilistic generalizations in the life sciences in the context of mechanistic explanation. To do this, I have focused on the (now widely accepted) MDC characterization of mechanism:

> MDC: Mechanisms are entities and activities organized such that they are productive of regular changes from start or set-up to finish or termination conditions (MDC 2000, 3)

As we saw in the last Chapter, combining the MDC characterization of mechanism with other major interpretations of chance (subjectivism, frequentism, and best system analyses) violated one (or more) of the desiderata we set forth for an adequate account of stochastic mechanism. As

such, I settled on a propensity-backed account of stochastic mechanism which I characterized quite broadly as this.

> Propensity-Backed Stochastic Mechanism: a stochastic mechanism has chance *c* if and only if the actual token mechanism setup which can generate this event possesses a dispositional property (tendency) to produce that outcome with that degree of strength.

At the end of Chapter 4, however, there were many important questions about this propensity-backed understanding of stochastic mechanisms that were left unanswered. Metaphysically speaking, what *is* a mechanism's propensity? What does a mechanism's propensity have as its categorical base? Can a mechanism's propensity to produce a certain outcome be quantified? If so, how? And what, if any, is the causal role played by a mechanism's propensity?

Now, I believe we are in the position to take some important steps toward answering these questions by drawing three important lessons from Ramsey's PID/PIF. Two of these lessons are positive; one is negative. That is, the first two lessons are ways in which I suggest Ramsey's PIF/PID provides potentially *beneficial* resources for understanding propensity-backed stochastic mechanisms, and the third is a way in which I think our account of stochastic mechanism should *depart* from Ramsey's approach.

3 Lessons from Fitness and Drift:

> (L1) Propensities are aptly understood as having probability-weighted possibilia as their categorical base.
>
> (L2) Propensities are quantifiable via a function of these probability-weighted possibilia.
>
> (L3) If we can, we should avoid committing to a view which requires defending the causal efficacy of propensities (and dispositional properties in general).

An analysis of how (and whether) these lessons can be applied to the propensity interpretation of stochastic mechanism will be made explicit in the later sections of this Chapter. But, for now, I can make a few general remarks by way of motivating the first two lessons. One reason why I

take lesson (L1) to be potentially beneficial for an account of stochastic mechanism has to do with an analogy to the tautology problem for realized fitness. Recall that if an organism's fitness is simply defined as the actual number of offspring it produces, then an organism's fitness can't explain these reproductive outcomes in any way besides one that is tautological. It strikes me that there is a corresponding worry for stochastic mechanisms. If we take a view of stochastic mechanism according to which its chance of firing is merely defined as its relative frequency of successful firings given the instantiation of its initial conditions, then there can be no appeal to a mechanism's stochasticity as a means to explain these actual outcome frequencies. From the point of view of biology, however, the point is (at least sometimes) to say *why* we see certain probabilistic output frequencies from biological mechanisms. If the only answer available is that the mechanism is stochastic—where this is simply defined in terms of its frequency of outputs—then this explanation is vapid and empty. Another reason why (L1) seems beneficial when applied to an account of stochastic mechanism is that it endows stochastic mechanisms with *predictive power* they wouldn't have if their stochasticity were understood merely in terms of its relative frequency of actual outputs. Just as with the PIF/PID, if a mechanism's stochasticity were defined solely in terms of its actual output frequencies, then it would make little sense to base predictions on this stochasticity; the outputs are already known. Furthermore, In Ramsey's terms, Lesson (L1) provides a way of differentiating between *the stochastic mechanistic process* and *the product of the stochastic mechanistic process*. That is, a propensity-backed understanding of stochastic mechanisms allows us to divorce its nature as a probabilistic process from the results of that process. If we understood the stochasticity of a mechanism solely by virtue of its actual output frequencies, it seems no such conceptual divide is possible. And lastly, following lesson (L1) means that the propensities instantiated by stochastic mechanisms *need not*

be deeply, metaphysically mysterious. We aren't just saying what the propensity does without saying what it is. Lesson (L1) endows us with (at least the potential for) some further metaphysical understanding of what the propensity of a given mechanism is, on what it is instantiated, and where it is to be found. The propensity of a given mechanism *is* a dispositional property with possibilia as its categorical base. Of course, much more needs to be said about how this lesson should get applied. However, the forgoing considerations, I hope, provide some prima facie reasons for taking seriously the potential advantages of following Ramsey's approach in the way specified in (L1).

Regarding lesson (L2), Ramsey's PIF/PID gives an important hint as to what might be needed to precisely calculate the value of a mechanism's propensity to produce a given outcome. Following Ramsey's approach means that we can conceive of the value of a given mechanism to produce a certain output as quantifiable via a function of the mechanism's probability-weighted possibilia (much more on this in 5.5.1). Without diving into the details here, however, it's (once again) worthwhile to suggest that following (L2) would amount to the following prima facie advantage: it would mean that a definite value for a given mechanism's propensity to produce a certain output can, in principle, be attained. This seems an especially worthwhile feature if scientists are to appeal to stochastic mechanisms to make precise predictions.

I'll say more about the benefits from applying lessons (L1) and (L2)—as well as their potential shortcomings—when it comes time to explicitly formulate the PrISM and apply it to a specific example (5.5-5.7). But first, I need to say something about one important way in which I intend to depart from Ramsey's approach: lesson (L3).

5.4 Lesson Three: Are Propensities Causally Efficacious?

In this section, I distance myself (in part) from Ramsey's PIF/PID. On Ramsey's view, in order to speak coherently, we must allow for fitness and drift, conceived of as propensities, to be causally efficacious. In what follows, however, I offer some reasons why I disagree that dispositional properties should be conceived of as, themselves, causally efficacious. I then admit that this appears to lead to an inconsistency with an argument made in Chapter 3. I resolve this inconsistency by appeal to a distinction made by Jackson and Pettit (1990) between causal efficacy and causal relevance.

5.4.1 When a Wine Glass Breaks in the Sink

There seem to be good reasons for defenders of propensity interpretations of fitness and drift to want these propensities to be causally efficacious. If these propensities are conceived of as causally efficacious, we can coherently speak of reproductive and evolutionary outcomes as having been caused by fitness and drift(ability) respectively. We can say that this snail had more progeny *because* he was more fit; and that the driftability of the individual snails in a given population *caused* its particular evolutionary outcome.

While there is no disputing the appeal of being able to coherently make such utterances, it comes at a cost. Namely, defenders of the causal role played by these propensities have to explain how propensities (and dispositional properties in general) can cause anything. Ramsey realizes that this might be difficult. Regarding drift, he says,

> There are of course long-standing debates in metaphysics over the nature of dispositional properties, their relationship to their categorical bases, and whether (and how) dispositions have causal efficacy. By characterizing driftability as a dispositional property, the cogency of driftability and the justification for the claim that driftability causes evolution therefore rests on particular views of dispositional properties. I will try to remain as neutral as possible about these debates and point out that all that my view needs is for dispositions to be causally

> efficacious with respect to their manifestations. Thus, I need it to be true that glasses can break because they are fragile, where 'because' is understood causally… What is required is the claim that dispositions can at times (correctly) be said to cause their manifestations. (Ramsey 2012, 10)

As indicated here, Ramsey's account requires that dispositions be causally efficacious with respect to their manifestations in order to garner the benefit of being able to speak causally about fitness and drift. As he says, he needs it to be the case that glasses can break *because* they are fragile. To illustrate why this (apparently modest) claim might not be so easy to defend, think for a moment about a cheap wine glass. Suppose that when it comes time to do the washing-up after an evening of Dionysian indulgence, you accidentally knock a wine glass over in the sink, and it cracks to pieces. What was the cause of this? More specifically, *what caused the glass to break*? Putting aside the herculean task of untangling the literature on philosophical analyses of 'cause', let's focus instead on what is required for a good causal inference. Here I follow Cartwright (1983) whose view is that we make our best causal inferences "…where our general view of the world makes us insist that a known phenomenon has a cause; where the cause we cite is the kind of thing that could bring about the effect and there is an appropriate process connecting the cause and the effect…" (Cartwright 1983, 4). Let's apply Cartwright's criteria to the case of the cheap wine glass breaking in the sink. I suspect there are few who would argue that the phenomenon of the wine glass breaking lacked any cause at all. With regard to her second and third criteria, we might ask ourselves: what kind of thing could bring about the breaking of a cheap wine glass? What sort of process would we deem appropriate to have brought this about? As he states above, Ramsey's view is committed to the fact that *the fragility* of the glass caused it to break. But is fragility (conceived of as a dispositional property) the kind of thing that could have brought that about? Put another way, are dispositional properties like fragility causally efficacious? My inclination is that they are not. To show why, consider the following contrastive query. Which

makes more sense: (1) it was the *fragility* of the inexpensive stemware that caused it to break, or (2) it was *the force of impact* on the stainless-steel sink together with *the particular molecular structure of the glass* that caused it to shatter? If your intuitions match mine, (2) is much the more reasonable answer. The fragility of the glass didn't cause the break. Indeed fragility *doesn't do anything*. In Cartwright's terms, fragility isn't the kind of thing that brings about effects. The glass breaking was a causal result of it forcefully impacting against the rigid surface of the sink.

Of course, Ramsey is correct that we want to say that the cheap wine glass broke *because* it is fragile. However, we should notice that there is an important ambiguity in the word 'because'. Sometimes it is used causally; and sometimes it is used *explanatorily*. I submit that, in the case of the wine glass, fragility comes in when we *explain* why the glass broke. Why did your tipping-over the glass cause it to shatter?—on account of the fact that it is cheap and fragile. The glass's fragility provides an explanatory reason why it broke. When we say that it broke because it is fragile, we mean to explain the occurrence of it breaking rather than to elucidate the specific causal history of its breaking. To use an example from biology, consider life-expectancy. Life-expectancy, like fitness and driftability, is a dispositional property. But, following Sober (1984), it hardly seems like an organism's life-expectancy causes it to die or continue living. Rather, life-expectancy is a value we assign an organism based on the sorts of causes that might result in its death and how likely we think it is that any of these causes might come about. We wouldn't say that an organism's life expectancy caused it to die; however, life expectancy can help explain why an organism dies when it does.

On the basis of these considerations, I conclude that (at the very least) it is difficult to explain how propensities are themselves causally efficacious[35]. As such, if we can avoid an

[35] This worry is amplified by the so-called 'Humphrey's Paradox'. (Cf. Fetzer 1981)—wherein propensities, if conceived of as causal, cannot make sense of inverse conditional probabilities.

account of stochastic mechanism that commits to the causal efficacy of its propensities, we should. Hence, lesson (L3).

5.4.2 Objection: can we still meet *CAUSAL EXPLANATION*?

Suppose we accept that propensities (and dispositional properties in general) do not seem to be causally efficacious. In doing so, we may have opened ourselves to a difficult objection regarding something argued for in Chapter 4. Recall that one of the key desiderata we employed for sorting out which interpretation of chance to adopt for our account of stochastic mechanism was:

> *CAUSAL EXPLANATION*: any adequate account of stochastic mechanism must allow for descriptions of underlying mechanisms to feature in causal explanations of regularities seen in nature. (3.2)

Indeed this desideratum played a key role in dismissing several of the alternative interpretations of chance (e.g., frequentism and BSA). However, by arguing as I have above that propensities (and dispositional properties in general) should not be seen as causally efficacious, it appears we might have undercut our own ability for propensity-backed stochastic mechanisms to causally explain.

5.4.3 Response: Distinguishing Causal Efficacy from Causal Relevance

In their 1990 paper "Program Explanation: A General Perspective", Jackson and Pettit make a distinction between *causal efficacy* and *causal relevance,* and correspondingly, a distinction between two kinds of causal explanation: *process explanation* and *program explanation.* Here, I will argue that, when properly understood, these distinctions show how propensities can meet the proposed *CAUSAL EXPLANATION* desideratum offered in Chapter 4 despite not being causally efficacious. Specifically, I argue that, even though propensities are not

themselves causally efficacious, they are nevertheless *causally relevant*. And by virtue of this causal relevance, they are entirely apt for featuring in causal explanation—albeit of the sort described by Jackson and Pettit as 'program explanation' rather than 'process explanation'. To motivate this strategy, I'll need to lay out some of the formal features of Jackson and Pettit's argument.

Jackson and Pettit begin by making the following uncontroversial assumption about causal explanation: "A causal explanation of something must direct us to a causally relevant property as opposed to a causally irrelevant property of the factor it identifies as explanatory: a property relevant to the causal production of the effect explained" (Jackson and Pettit 1990, 108). One way in which properties are causally relevant, Jackson and Pettit suggest, is by being *causally efficacious*. They define a *causally efficacious property* with regard to an effect as "a property in virtue of whose instantiation, at least in part, the effect occurs; the instance of the property helps to produce the effect and does so because it is an instance of that property" (ibid, 108). A property *F* fails to be causally efficacious of an effect *e*, on the other hand, if it meets all of the following conditions:

> (i) there is a distinct property G such that F is efficacious in the production of e only if G is efficacious in its production;
>
> (ii) the F-instance does not help to produce the G-instance in the sense in which the G-instance, if G is efficacious, helps to produce e; they are not sequential causal factors;
>
> (iii) the F-instance does not combine with the G-instance, directly or via further effects, to help in the same sense to produce e (nor of course, vice versa): they are not coordinate causal factors. (Ibid, 108)

Like me, Jackson and Pettit do not take fragility to be a causally efficacious property. This is because, as they see it, fragility meets all three of the conditions above. They write:

> The property of fragility was efficacious in producing the breaking only if the molecular structural property was efficacious: hence (i). But the fragility did not help to produce the molecular structure in the way in which the structure, if it was efficacious, helped to produce the breaking. There was no time-lag between the exercise of the efficacy, if it was efficacious, by the disposition and the exercise of the efficacy, if it was efficacious, by the structure. Hence (ii). Nor did the fragility combine with the structure, in the manner of a coordinate factor, to help in the same sense to produce e. Full information about the structure, the trigger and the relevant laws would enable one to predict e; fragility would not need to be taken into account as a coordinate factor. Hence (iii). (Ibid, 109)

I take the forgoing argument to be further demonstration of the thesis I offered in 4.4.1: dispositions like fragility are not causally efficacious. But the help that Jackson and Pettit can offer my account does not stop there. Even more important than the above argument against the causal efficacy of fragility, Jackson and Pettit argue that *causal efficacy is not the only kind of causal relevance.* Indeed, there is another kind, and it just so happens that dispositional properties (like fragility) fit into it quite nicely.

To illustrate this second kind of causal relevance, Jackson and Pettit appeal to the notion of a computer program. They write,

> A useful metaphor for describing the role of the [causally relevant but non-causally efficacious] property is to say that its realization *programs* for the appearance of the productive property and, under a certain description, for the event produced. The analogy is with a computer program which ensures that certain things will happen - things satisfying certain descriptions - though all the work of producing those things goes on at a lower, mechanical level. (Ibid, 114, italics added)

The realization of an abstract, higher-order dispositional property (like fragility), on Jackson and Pettit's view, *programs* for the appearance of causally efficacious properties at the level of the stuff doing the causing. While it's the physical bits and pieces of machinery inside my computer that do the work of causally producing the letters that are now appearing on my screen as I'm typing, there are many bits of programming code that constrain how this physical causation can occur. Fragility works the same way. Although the fragility of a glass doesn't physically cause it

to break, its realization ensures that many different kinds of physical interventions would cause it to break. Just as the programming in my computer is causally relevant to the effect of words appearing on my screen, so is fragility causally relevant to the effect of a cheap wine glass breaking in my sink. This shows that the property of being fragile can be seen to be *causally relevant without being causally efficacious*. It also shows that explanations appealing to first-order, concrete causal properties are not the only kinds of causal explanations we can give about the world. In addition to these first-order causal explanations. which Jackson and Pettit call *process explanations*, there are also explanations that appeal to these higher-order, abstract, properties. These are called *program explanations.*

These distinctions from Jackson and Pettit, I suggest, are exactly what is needed to undermine the objection considered in 4.4.2. Here is the precise point. *While propensities are not causally efficacious, they are nevertheless causally relevant. And causal relevance is all that is needed to meet CAUSAL EXPLANATION.* Put another way, the propensity of a given stochastic mechanism is causally relevant to that mechanism's output in exactly the same way that fragility is causally relevant to the event of a cheap wine glass breaking in my sink. This is because the realization of a propensity *programs for the realization of lower-order efficacious properties and, in these circumstances, for the occurrence of the event in question*.[36]

5.4.4 Objection: But Aren't BSA Chances Causally Relevant?

Suppose that Jackson and Pettit's distinction between causal relevance and causal efficacy (as well as the corresponding distinction between program explanation and process explanation) does the work I've argued that it does. That is, suppose that propensities can be seen to meet *CAUSAL EXPLANATION* in the sense of being causally relevant even if they are not

[36] This conclusion has the added benefit of cohering nicely with two recent papers defending the causal relevance of dispositions: Mckitrick (2005) Vicente (2002).

causally efficacious. A new objection seems to arise. Namely this: aren't BSA chances causally relevant in just the same way that propensities are? And if so, wouldn't this negate the argument I offered (in 4.3.3) against a BSA understanding of stochastic mechanism?

Before I offer my response to this objection, let's ask why it might seem that BSA chances are causally relevant in the same way that propensities are. Recall that, on the BSA interpretation, the chance of any given outcome occurring is whatever the best systematization of the Humean mosaic of particular facts tells us it is. BSA chances might seem to be causally relevant in the following sense. Just as the word processing program I'm currently using constrains the kinds of causally efficacious interactions I can have when typing these words, so too does BSA chance amount to a constraint on the space of possible causal events that can take place in the world. When the BSA, for example, tells us that there is a 1/6 chance of a six-sided fair die landing on six when I roll it, what it is doing (in effect) is giving us some information regarding what kinds of constraints there are on the ways that I can be causally efficacious in rolling a six with a fair die. E.g., I shouldn't expect to be able to roll a six ten times in a row. If this is correct, then it seems BSA chances are causally efficacious in just the same way that propensities are. And if this is correct, then it seems we no longer have any theoretical basis for dismissing a BSA interpretation of stochastic mechanism on the grounds that it fails to meet *CAUSAL EXPLANATION*.

Despite its apparent force, I argue that this objection rests on a mistake. Specifically, I suggest that, on the BSA, the facts constrain the chances; not the other way around. So BSA chances aren't causally relevant in the way that propensities are.

To see why, consider again the example of my word processing program. On Jackson and Pettit's view, what makes this program causally relevant is the fact that "[it] ensures that certain

things will happen - things satisfying certain descriptions - though all the work of producing those things goes on at a lower, mechanical level" (Ibid, 114). Now ask yourself, do BSA chances *ensure* that things will happen? Put another way, do BSA chances *place constraints* on the way that causal events can occur in the world? My intuition is that the answer to both of these questions is no. Rather, it seems to me that (by their very definition), BSA chances *are constrained by the causal facts—not the other way around.* Indeed, the central point of the BSA account of chance is that the chances supervene on the Humean mosaic of particular matters of fact. Given this central feature of the BSA account, I argue, it must be that those facts constrain the chances; it doesn't work the other way. And if this is so, BSA chances are not causally relevant in the way that my word processing program is. My word processing program, given that it is realized on my computer, *makes it* such that certain ways of poking my keys will produce the appearance of certain symbols on my screen (and not others). But BSA chances don't *make* anything be the case in the natural world. As such, I take the objection offered in this section not to threaten the arguments I gave in 4.3.3 after all.

5.5 A Propensity Interpretation of Stochastic Mechanism

Having now motivated (L3) and defended it against a couple of possible objections, I now propose to examine the prospects for applying the two positive lessons, (L1) and (L2), that we adduced from Ramsey's PIF/PID in order to begin to formulate my own propensity interpretation of stochastic mechanism (PrISM).

Before addressing (L1) and (L2) specifically, however, let's see what the basic form of the PrISM is going to be. Following the template from Ramsey, we can describe the PrISM in the following manner.

> Consider a mechanism *M* operating in an environment *E*. There are a number of factors (both internal to *M* and from *E*) that influence whether the mechanism successfully fires. The set-up conditions for a particular *M* might or might not obtain. The particular entities and activities might get interfered with by the *E* once the mechanism is triggered. And an *M*'s termination conditions may or may not occur even after triggering. Consequently, there are a number of possible ways the mechanism can act. Let's designate all the possible ways the mechanism can act with *F* (as in 'possible firing'). Each of these ways the mechanism can go (F_1, F_2, F_3,… F_n) will have a probability associated with it. The propensity of a given stochastic mechanism can be understood as metaphysically based on the properties of the entire set of *Fs*. (Call this set F_i.). More specifically, the propensity of a given stochastic mechanism is a function over its probability weighted *F*s.

On the basis of this description, I characterize the PrISM in its most general form as this.

> *PrISM*: the propensity (*Pr*) of a given stochastic mechanism to fire can be identified with heterogeneity [differences] in the *F*i [set of a mechanism's possible firings] and can be quantified by a function on this heterogeneity.

So far, this is just the direct application of Ramsey's approach to the PIF/PID to notion of stochastic mechanism. What needs to happen next is to more carefully examine whether Ramsey's approach *works* when applied in this way to our notion of stochastic mechanism. To do this, I'll look at the prospects of applying both (L1) and (L2) individually.

5.5.1 On the Prospects for Applying Lesson 1

Recall that the first feature of Ramsey's PIF/PID that seemed it might be beneficial to apply to an account of stochastic mechanism was this.

> (L1) Propensities are aptly understood as having probability-weighted possibilia as their categorical base.

Recall, further, that the reasons motivating the application of (L1) to stochastic mechanisms were directly analogous to Ramsey's own reasons for understanding fitness and drift(ability) in this manner. Namely, (L1) means we have some resources for offering an analysis of propensities such that they aren't entirely mysterious. If propensities can be understood as having probability-weighted possibilia as their categorical base, then we can have some idea (metaphysically

speaking) of what they are. And this would, at the very least, offer the some response to critics who argue that propensity theorists merely say what propensities do without saying what they are.

If we follow (L1), what can we say about what propensities are? At the very least, we can say what their categorical base is. Just as fitness (as a propensity) can be explicated in terms of the properties of the set of an organism's heterogeneous possible lives, so too can the stochasticity of a mechanism (as a propensity) be explicated in terms of the properties of the set of a mechanism's heterogeneous possible firings. The propensity of a particular vesicle mechanism to successfully fire only 10% of the time can be given further analysis. It can be explicated in terms of the properties of the set of the possible ways this mechanism could operate under various conditions. More on the details of this will come in 5.7. But before we get to the details of how these propensities get calculated, it will be helpful to pause and consider a few objections to applying (L1) to our account of stochastic mechanism.

5.5.1.1 Objection: Didn't We Argue (contra HF) that Non-actual States Can't be Explanatory?

The first serious objection that needs to get considered is this. Didn't we argue in 4.3.2 that the problem with a hypothetical frequentist interpretation of stochastic mechanism is that it ultimately grounds the stochasticity of a mechanism on counterfactuals? And wasn't our reason for not wanting to do this that it doesn't make sense to causally explain actual output frequencies of mechanisms by reference to counterfactuals? However, isn't that precisely what is going on here when we apply (L1)? In other words, aren't we ultimately appealing to non-actual states (possible mechanism firings) as a metaphysical analysis of the very propensities we're supposed to be using to causally explain actual output frequencies?

By way of response to this objection, I want us to think carefully about *what* is doing the explaining. As I argued earlier in this Chapter, propensities themselves are not causally efficacious. But, following Jackson and Pettit, they can be seen to be causally relevant. That is, just like my word processing program, propensities constrain the kinds of causal interactions its possessor can accomplish. My response to the above objection regarding HF is the same one I gave to the BSA objection. The counterfactuals appealed to on an HF understanding of chance are not causally relevant in this way. Why? – because (once again) the counterfactual long-run frequencies appealed to by HF don't *make* anything be the case in the actual world. To see why, consider again my chance of rolling a six with a six-sided fair die. On an HF account, my chance is 1/6 because, on a counterfactual infinite (or very large) series of trials of me rolling that die, the relative frequency of instances of the die landing on six will eventually draw ever closer to reaching a limit of 1/6. But, just as we did with the BSA objection, ask yourself whether this HF counterfactual is causally relevant in the same way my word-processing program is. I submit that, once again, it is not. My word processing program *makes* it the case that certain symbols appear on my screen when I type. The counterfactual infinite (or very large) series of trials in which I role a fair six-sided die *doesn't make* it the case that I will role a six roughly $1/6^{th}$ of the time here in the actual world.

Now ask yourself whether a propensity fairs any better in this regard. I think it does. To see why, recall the example of my picture window. It has a dispositional property of being fragile. That is, it has a propensity to break relatively easily when struck by things like baseballs, bricks, flying birds, and hurricane-force winds. Does the property of being fragile in this way *make* it the case that it will react by breaking when causally interacted with by baseballs, bricks and the like? It seems to me plausibly so. It is *by virtue* of instantiating the property of fragility,

that my picture window is susceptible to breaking in all of these possible ways. Just as Jackson and Pettit suggest, being fragile *programs* for this to be the case—just as Microsoft Word *programs* for it to be the case that my font switches to italics when I press control 'i'.

Granted, much more needs to be said in order to fully specify why propensities are causally relevant and hypothetical frequencies aren't: more than I can say here. But I can say one last thing that helps motivate this claim. Propensities, by definition, are objective properties in the actual world. Just like computer programs are objectively realized on my computer. Even though they carry in them (they have as part of their content) information about modal possibilities, propensities do exist as part of the furniture of the actual universe. The hypothetical frequency of my infinite roles of a six-sided fair die, on the other hand, exists nowhere in this universe. And perhaps this is part of the reason why hypothetical frequencies seem less equipped for featuring in causal explanations of the actual world than do propensities.

5.5.1.2 Objection: the Ramsey Approach Leads to a Vicious Regress

Even if I have succeeded in showing that propensities meet *CAUSAL EXPLANATION* even if they carry modal information, there remains another serious objection to applying (L1) to an account of stochastic mechanism. And I fear it is an even harder one to deal with.

Recall that (L1) states that propensities are aptly understood as having *probability-weighted* possibilia as their categorical base. It seems any follower of (L1) owes some kind of story about what these probabilities are, where they come from, and how we get them. The problem is, as we'll see, it's unclear what (if anything) can be said in answer to these questions without running into some kind of trouble.

Consider first the following kind of answer. I don't care where you get the probabilities weighting these possible ways a mechanism could fire. Get them wherever you want. I'm not

trying to offer a general interpretation of how to understand all probabilities—in all instances where they occur. The important thing is that you do the best you can to assign probabilities to these possible mechanism firings given whatever evidence you have. And once they get assigned, (however they get assigned) we can calculate the propensity of the mechanism to achieve various output conditions via a function of these probability weights. If this process leads to the identification of a propensity that varies widely from the results we go on to observe when testing the mechanism in question, then we can always go back and adjust our initial probability weight assignments.

The problem with this approach, however, is that it seems to undermine the very advantage that (L1) was supposed to bestow. Namely, if we say nothing about what these probabilities are, then the mysterious aspect of propensities that we were trying to mitigate against (by offering a further analysis in terms of probability-weighted possibilia) simply gets moved back one step to the probabilities we assign to the possible mechanism firings on which the propensity is categorically based. In other words, rather than making propensities less mysterious, (L1) merely relocates the mystery one step below. And this seems like a serious problem.

Perhaps, then, if we are to maintain the benefit of applying (L1), we *do* owe some story about what these underlying probabilities are. Sadly, telling this story may prove difficult. The reason is that it seems we may, by the very same arguments offered in 4.3, end up having to say that these underlying probabilities have to (themselves) be propensities. But then, those propensities, if we are to understand what *they* are, will also have to analyzed in terms of probability-weighted possibilia. And those underlying probabilities will also have to be analyzed as propensities. And on, and on. In short, it seems we have a vicious regress on our hands.

What (if anything) can be done to avoid this regress? One option would be to explore an alternative route for understanding these propensities—one that does not follow (L1).

5.5.1.3 An Alternate Route: Following Abrams

Suppose the forgoing arguments succeed in showing that the prospects for applying (L1) to our analysis of stochastic mechanisms are quite dim. Suppose we now find ourselves convinced that (L1) either pushes the mystery of propensity back a step or it results in a pernicious explanatory regress. Does that put the proverbial final nail in the coffin for the PrISM? Not necessarily. There's another way to proceed.

The other way is this. Rather than following Ramsey's approach of grounding propensities on the heterogeneity of the underlying probability-weighted possible ways a mechanism might fire, we might take an approach inspired by Abrams (2012). The Abrams-inspired approach does just the opposite of what the Ramsey approach does. Rather than grounding an understanding of a mechanism's propensities in terms of the heterogeneity of their underlying probability-weighted possibilia, we might understand the heterogeneity of possible ways a mechanism might fire in terms of the very mechanisms themselves. That is, we might ground our understanding of the stochasticity of mechanisms by appeal to features of the mechanisms themselves. On this way of looking at it, structural features of the mechanism itself specify the propensities it has to operate in various ways. Recall that, as we saw in 4.6, Abrams explains the smooth probability distribution we see in the outputs of a fair roulette wheel (despite highly variable inputs) by appeal to facts about the mechanism itself. This smooth output distribution can be understood, according to Abrams, by paying attention to facts like the equal-size of the wedges, and the ways in which the input space partitions map to the output possibility space.

Just as Abrams appeals to facts about the roulette wheel to explain the probability we assign to its outputs, we might explain a mechanism's propensity to fail some percentage of the time by appeal to facts about that mechanism's structure. And in doing so, we don't need to appeal to further (and equally mysterious) underlying probabilities.

Even if it succeeds at avoiding the regress attributed to the Ramsey approach, however, the Abrams-inspired approach comes at its own costs. First, the Abrams-inspired approach seems to leave us back where we started in terms of the brute, unanalyzability of propensities. Yes, we can learn about the structural features of the mechanisms that instantiate these propensities. But we can't say much at all about what the propensities are. And second, as we saw in 4.6, Abrams explicitly affirms a reliance on frequencies in a way that we've argued against doing.

So where is there left to go from here?

5.5.1.4 Gesturing at an Argument from Balance

In the spirit of intellectual honesty, it may be that we need to admit that neither the Ramsey-inspired approach nor the Abrams-inspired approach are without serious drawbacks. And as such, it may be that all we can do is ask which are the theoretical costs we're more comfortable absorbing.

To sum up, applying the Ramsey-inspired (L1) has the following real advantages and real costs:

> Ramsey-inspired (L1) Payoff: it offers a resource for further analyzing the metaphysics of propensities—thereby making them less mysterious
>
> Ramsey-inspired (L1) Costs: on pain of merely pushing the mystery back a step, we owe a story about what the probabilities weighting the Fs are. And depending how this story goes, there is a serious risk of explanatory regress.

If we follow the Abrams-style approach instead, there are also real advantages and disadvantages:

Abrams-inspired Payoff: avoids the risk of explanatory regress

Abrams-inspired Costs: leaves propensities mysterious, and relies heavily on frequencies. Given these difficult trade-offs, the question now becomes, on which approach is the price right? I cannot give an answer to this question that takes anything like the form of a deductive argument. In the end—as is always the case with such things—what matters is how much we end up *caring* about the garnering the respective benefits vs. how much we care about avoiding the respective costs. Having said that, my view is that the traction that the Ramsey approach has gotten for understanding other concepts in biology in terms of propensities may (itself) lend additional motivation for trying to follow his lead. If only for the sake of unification with his other promising accounts, it seems that (if possible) it would be a good idea to see how far we can take the lessons adduced from Ramsey's well-received PIF/PID accounts.

5.5.2 On the Prospects of Applying Lesson 2

Suppose we decide to stick with the Ramsey approach after all. There still remains the difficult task of exploring the prospect for applying (L2). Recall that (L2) states: propensities are quantifiable *via a function* of these probability-weighted possibilia. Regardless of whether there is an adequate story to be told about the probabilities we use to weight these possibilia, we still need to figure out, with regard to stochastic mechanisms, what this function might possibly be.

In the space I have remaining in this (already long) Chapter, I won't be able to tackle this project in a high degree of detail. However, I can at least say *something* about what a candidate function might look like as well as a few words on its viability in the context of science.

So what function might be required in order for us to be able to calculate the propensity of a mechanism to achieve a particular output? Suppose for example, we wanted to calculate the overall propensity of a given synaptic transmission mechanism reach its termination conditions (e.g., to result in post-synaptic electrical activity). In its most general form, we can place the following constraint on such an equation (where T signifies achieving termination; $\mathscr{f}$ signifies some function; and p signifies some set of probability-weighted possibilities).

$$Pr(T) = \mathscr{f}(p)$$

That is, whatever else we say about the prospects of applying (L2), the basic form of any equation calculating the propensity for a given stochastic mechanism to reach termination conditions has to take the forgoing general form. The equation has to tell us that the propensity of achieving termination conditions equals a function of its probability-weighted possibilities.

Now, the question becomes: what is this function? The only suggestion I can think to offer is also the most obvious. It's the standard way of expressing the average probability of some event given a bunch of known conditional probabilities of that event given various background conditions—each with its own know probability: *the law of total probability*. Informally, the law of total probability says that, given the occurrence of some event A, with known conditional probabilities given a range of possibilities (B_n)—each with its own known probability itself—the total probability of A equals the sum of the probability of A given B_1 times the independent probability of B_1 plus the probability of A given B_2 times the independent probability of B_2 plus the probability of A given B_3 times the independent probability of B_3, and so on until all the Bs are included in the sum.

Formally, the equation looks like this: Pr(A) =

$$\sum_{i=1}^{n} P(A|Bi)P(Bi)$$

This is easy to apply to the PrISM: the total propensity of a given stochastic mechanism to achieve termination conditions [Pr(T)] =

$$\sum_{i=1}^{n} P(T|Fi)P(Fi)$$

Now that we've seen one obvious candidate for applying (L2) to the PrISM, let's ask briefly whether it is a good one.

One consideration that might count in favor of this candidate function is that it was the very mathematical model first introduced by Brandon (1978) and Mills and Beatty (1979) in the very first formulations of the PIF. If it worked for an organism's overall fitness—conceived of as a propensity—then it might well work for understanding the overall propensity of a stochastic mechanism to reach termination. The problem, however, is that the above function *didn't* end up working very well as a way of understanding the PIF. Without going into the details, the above way of calculating an organism's overall fitness ran into several counterexamples (cf. Pence and Ramsey 2013 for specifics on these counterexamples). In short, the above way of calculating an organism's overall fitness can't handle *variance* (in possible offspring over an organism's lifespan, of the timing of selection, and the timing of reproduction).

One positive feature of the prospects of using this candidate function for the purposes of applying (L2) to the PrISM is that stochastic mechanisms aren't subject to the same problems with variance that the fitness seems to be. The overall propensity for a mechanism to reach termination doesn't necessary vary over the time the mechanism exists, and it isn't sensitive to the timing of selection and reproduction the way that fitness is. And if this is so, our candidate

function may not need to take ever increasingly complex form—as it seems to have done in the case of the PIF.

That said, adopting the above candidate function as a means of applying (L2) to the PrISM does suffer at least one major shortcoming. It doesn't really offer anything *new* to the PrISM. Statisticians and mathematicians have known and used the law of total probability for the better part of the 20th century. If this is the manner in which we choose to calculate the overall propensity of a stochastic mechanism to achieve termination, what is the rest of the PrISM really adding that is useful to science?

In response to this apparent shortcoming, I can only say this. Even if it is the case that this feature of my PrISM analysis is not (itself) novel, that doesn't preclude its *application* from being novel and interesting. In other words, even if statisticians and mathematicians have availed themselves of the law of total probability for decades, it could still be worthwhile to have found a new and interesting place to apply it. And if the mechanistic approach to the philosophy of life sciences succeeds in supplying a framework to replace the laws-based deductive-nomological explanatory approach, as I hope it will, this novel application could add to its success by supplying a method for scientists to calculate overall propensities of stochastic mechanisms for supporting probabilistic generalizations in the life sciences.

5.5.3 Following lesson 3: An important advantage of the PrISM

The final lesson we have yet to explore the prospects for applying to the PrISM is (L3). As we saw in 5.4, there are some good reasons for avoiding a commitment to the causal efficacy of propensities (and dispositional properties in general). What I wish to point out here is that departing from Ramsey's approach in this way paves the way for the PrISM to enjoy a significant advantage over Ramsey's PIF/PID. My argument is this.

P1. Ramsey's PIF/PID requires propensities to be causal.

P2. As seen in 4.4, the causal efficacy of propensities is (at best) difficult to defend.

P3. But the PrISM need not endorse the causal efficacy of propensities *because the mechanisms themselves can be seen to do the causal work.*

P4. Thus, the PrISM can retain the general advantages of a propensity interpretation (e.g., objectivity, causal explanation, scientific investigability, and grounding of predictions/interventions) *without* the metaphysical cost of having to defend the causal efficacy of dispositional properties.

C1. Given (iv), the PrISM therefore holds an important advantage[37] over Ramsey's PIF/PID.

We have already seen that Ramsey explicitly commits to premise (P1). The truth of premise (P2) depends on the arguments given in 4.4. Premises (P3) and (P4) need more support to show that (C1) follows. In an effort to provide such support, let's look at an example of the PrISM in action.

5.6 Example: Synaptic Transmission

To illustrate the PrISM in action—as well as its advantage over Ramsey's PIF/PID—let's see how it might apply to an example from neuroscience.

Initiation of Electrical Activity in Post-synaptic Neurons

There are gaps between the neurons in our brains called *synapses*. Electrical signals must traverse these gaps in order to continue on their path through our nervous system. Very roughly, the chemical process goes like this: a brief pulse of electricity (called an 'action potential') travels down a hollow tube in the neuron called an *axon*. This occurs because the axon is filled with (and surrounded by) an aqueous solution containing charged ions. At rest, the membrane surrounding each neuron is polarized; its inner surface is negatively charged relative to its outer surface. Upon contact with the action potential, however, the charge of the axon's membrane rises enough to open specific gates in the membrane so as to allow positive ions (mainly Ca^{2+} and Na^{+}) into the cell causing the

[37] When I say that the PrISM holds an important advantage over the Ramsey's PIF/PID, I do not mean to suggest that it can do a better job of explaining fitness and drift than his can. I do not intend to construe the PRISM as a proper alternative to Ramsey's interpretations. How much work the PrISM can take over from the PIF/PID depends on to what extent the phenomena that these conceptual frameworks are intended to explain are mechanistic.

> membrane to depolarize. Other positively charged ions (mainly K^+) flow out of the membrane to cause repolarization in its wake. What results is a wave of electricity flowing down the axon until it reaches the synapse, at which point a chemical (called a *neurotransmitter*) is released and moves across the synapse binding to specific proteins on the neighboring neuron. However, this process is not very dependable. In fact, according to some estimates, the release of neurotransmitters can fail to result in the successful initiation of electrical activity in postsynaptic neurons up to 90% of the time.[38]

This example constitutes a probabilistic phenomenon of interest to life scientists. And importantly for us, neuroscientists explain this phenomenon mechanistically. The successful release of electrical activity in postsynaptic neurons fails up to 90% of the time because of features of the mechanism of postsynaptic actuation. Because this phenomenon is both probabilistic and explained by appeal to a mechanism, the mechanism underlying and productive of this phenomenon is clearly a stochastic mechanism. This example is, therefore, ripe for analysis by the PrISM.

Here is how it might go. Call the mechanism for transmission of electrical activity across a particular synapse in my brain *M*. Now consider a point in time (call it *t*) when the membrane surrounding each neuron is at rest and polarized. At *t*, there are many ways that *M* might fire, depending on the environmental factors *E* particular to the chemical make-up of my brain at *t*. Call the set of these possible firings F_i. Each of these Fs (F_1, F_2, $F_3 \ldots F_n$) has a probability weight assigned to it. For example, we might say that the probability that depolarization occurs but no neurotransmitter signal is released is .6; the probability that polarization occurs and the neurotransmitter is released is .1; the probability that insufficient release of positive ions leads to no polarization at all is .3; etc. We might also know the probabilities of this mechanism reaching termination conditions given each of these scenarios. If we do know all of this, according to the

[38]For detailed mechanistic analysis of this phenomenon, see the original MDC (2000) paper and Craver's book *Explaining the Brain* (2007). For discussion of the science behind successful vs. unsuccessful instances of neurotransmitter release, see Kandel et al. (2000). For detailed discussion of how this case relates to MDC's regularity requirement, see Bogen (2005) and Andersen (2012).

PrISM, we can calculate of the overall propensity of this *M* to lead to successful termination conditions (e.g, the release of electrical activity across the synapse) as a function of these probability-weighted possibilities such that the output of this function is the propensity value *Pr* that the stochastic mechanism will successfully terminate.

Where the PrISM departs from Ramsey's PIF is on the issue of causality. As we've seen, Ramsey's view is committed to saying that the propensity (whether it's fitness or driftability) causes the reproductive or evolutionary outcome. However, if we follow lesson 3 (as argued for in 5.4.), then it seems difficult (at best) to see how this could go. At this point, a real advantage of the PrISM becomes salient. *If we adopt the PrISM, we can relegate the causal work to the token mechanism rather than the propensity*. It is not that the disposition to fail at a rate of .9 caused the mechanism not to fire in a given instance. Rather, this particular actuation mechanism, the particular entities and activities out of which it is composed, does the causing. Of course, we do want to say that we observe the output frequency of this mechanism *because* it has a propensity of .9 of failing to fire. But, as I suggested in 4.4, this need not mean that we must commit ourselves to the causal efficacy of the mechanism's propensity. This propensity, rather, should be seen to explain why the actual mechanism does the causal work when it does.

Applying the PrISM to this case makes salient several virtues of this kind of approach. When neuroscientists say that successful transmission of electrical activity across a given synapse fails up to 90% of the time, we now have a helpful conceptual framework for explaining what makes this true. It is true because the mechanism underlying and productive of this phenomenon is a stochastic one with a propensity of .9 to fail. Whether we decide to follow the Ramsey-inspire approach or the Abrams-inspired approach, there is more that we can say about this propensity than we couldn't without the PrISM. Plus the analysis offered in this Chapter

means that the chanciness of this mechanism exists in the world, instantiated by a real, existing, mechanism. It can be investigated by science. And once discovered, it can be the basis for explaining and predicting actual outcomes.

5.7 Conclusion

In this chapter, I have drawn on a few recent propensity interpretations of fitness and drift in an effort to adduce some important lessons for my own propensity interpretation of stochastic mechanism. If successful, the PrISM may prove a novel way of grounding probabilistic generalizations in the life sciences: one which allows for the stochasticity of a mechanism to an objective feature of the world, one which endows descriptions of stochastic mechanism with explanatory and predictive power, and one which need not commit to the causal efficacy of dispositional properties.

Chapter 6. The Threat of Determinism: Synthesizing Emergent Chance and Multilevel Mechanisms

Abstract: In this chapter, I explore whether—and if so, how—the propensity interpretation of stochastic mechanism (PrISM) developed thus far can handle the threat of metaphysical determinism. The answer I pursue is that it can. To show how, I first outline a few influential arguments for and against biological indeterminism and suggest that the only tenable philosophical conclusion available is an agnosticism as to whether there is genuine indeterminism at play in the biological world. As such, I pursue a disjunctivist strategy according to which what we say about the PrISM depends on which side of the biological determinism debate prevails. To accommodate the more difficult determinism disjunct, I draw on recent view from the philosophy of probability, emergent chance, according to which objective chance can exist at one level of description even if determinism holds at a different level. I then apply this idea to mechanism levels to show that genuine propensities may be able to exist at one mechanism level even if another mechanism level (within the same mechanism hierarchy) is characterized deterministically.

6.1 Introduction

~ Given for one instant an intelligence which could comprehend all the forces by which nature is animated and the respective situation of the beings who compose it—an intelligence sufficiently vast to submit these data to analysis—it would embrace in the same formula the movements of the greatest bodies of the universe and those of the lightest atom; for it, nothing would be uncertain and the future, and the past, would be present to its eyes". (Laplace 1814, 4)

~ There is no chance without chance. If our world is deterministic there are no chances in it, save chances of zero and one. Likewise, if our world somehow contains deterministic enclaves, there are no chances in those enclaves". (Lewis 1986, 120)

The suggestion put forth thus far in this dissertation, recall, is that we can ground the truth of (at least some) probabilistic generalizations in the life sciences by appeal to stochastic mechanisms in the world: mechanisms that have probabilistic propensities to result in certain outputs over others. In the previous Chapter, we explored the prospects for understanding these propensities as (1) categorically based on, and (2) quantifiable via a function of probability-weighted possible ways a particular mechanism might fire. The question now arises: what is meant by possible ways a mechanism *might* fire? For example, what are we committed to when we say of a particular synaptic transmission mechanism in the brain that it *might* result in an initiation of electrical activity across a particular synapse—or when we say of a particular DNA replication

mechanism in a living organism that it *might* result in the mutation of a particular strand of DNA during replication—or when we say of a particular Mendelian segregation mechanism in a hybrid pea plant that it *might* result in the growth of yellow peas, or it *might* result in the growth of green peas?

One natural way of interpreting these statements is this: once the start up conditions of the mechanism underlying these phenomena obtain, there are multiple ways it might turn out—some more probable, and some less. On an interpretation like this, these utterances appear to commit us to a 'garden of forking paths'[39] stemming out from a given mechanism such that whichever mechanism output actually obtains is just one of many results that might have obtained. On a view like this, there appear to be *genuine alternate possibilities* for mechanism end-states. On a view like this, it seems *the very same mechanism history can lead to different futures.* On a view like this, there is *genuine biological indeterminism*.

If there is genuine indeterminism at play in the natural world, and it extends up to biological mechanisms, then the propensity interpretation of stochastic mechanism (PrISM) laid out in the previous chapters seems to be a readily available way to characterize them. The propensities argued for and defined in the previous chapter, we might easily suggest, *just are* the dispositional tendencies for these mechanisms to behave in the indeterministic manner just described.

However, as we shall see, there seems to be plenty of good reason to doubt whether there is genuine indeterminism at play in the biological world. Indeed there seems to be some convincing reasons to think that many, maybe all, of the stochastic biological phenomena we seek to explain mechanistically are either themselves fully deterministic or the result of underlying processes that are ultimately deterministic. From the point of view of this

[39] To borrow Borges's landmark phrase

dissertation, this is a cause for justified concern. The reason is as simple as it is worrisome. Supposing global determinism holds, then it would seem that any philosophical theory that makes reference to propensity fails on the outset. Why?—Because propensities seem straightforwardly antithetical to determinism. A traditional understanding of propensity, recall, is that a propensity is a dispositional property of certain states of affairs such that the same generating conditions (e.g., flip of a fair coin) sometimes end up going one way (e.g., heads) and other times in other ways (e.g., tails). Global determinism is the very denial that this ever occurs. On a deterministic worldview, the entire history, past and present, of the world supervenes on the world's laws of nature such that there is only one way that any given event can turn out. In the eyes of the infamous Laplacian demon, there simply are no chances to be found. Therefore, on a deterministic worldview, the unfortunate upshot appears to be the following. *If global determinism holds, the probabilities we assign to mechanism outputs are nothing more than measures of our ignorance about the complex ways mechanism inputs lead to outputs.* In other words, on a deterministic worldview, there is no sense to be made of stochastic mechanisms other than a subjectivist one. Since I have already argued (in 4.3.1) that a subjectivist understanding of stochastic mechanism is untenable for the life sciences, it appears that determinism is a real threat to the PrISM proposed thus far.

My goal in this chapter is to do my best to dismantle this threat of determinism. Here's how I plan to go about it. In section 6.2, I briefly articulate some of the arguments for and against indeterminism operating at the level of biology. After demonstrating that any ability to arbitrate between these opposing positions rests on empirical matters uninvestigable in this dissertation, I adopt, in 6.3, an agnosticism about metaphysical determinism. In 6.4, I suggest that the most reasonable strategy for moving forward, given our inability to arbitrate the determinism issue, is

to maintain a kind of disjunctivism—whereby what we say about the PrISM splits in two directions depending on whether we assume indeterminism or determinism in the biological world. Because it is relatively easy to situate the PrISM on the biological indeterminism disjunct, I spend the remainder of the chapter showing how it fares on the biological determinism disjunct. In 6.5, I briefly outline List and Pivato's recent view of emergent chance. And in 6.6, I lay out Carl Craver's analysis of multilevel mechanism—pausing to point out, in 6.7, a few important features regarding which the two approaches are complementary. In 6.8, I then attempt to synthesize the two approaches using the example of spatial memory in rats as a plausible instance where chance emerges at one mechanism level even if others are deterministic. In 6.9, I anticipate a serious objection to this approach: that this determinism friendly version of propensity is no propensity at all. I then respond, in 6.10, with an argument for why, even on the biological determinism disjunct, we can reasonably locate this view in the propensity camp.

6.2 For and Against Biological Indeterminism

There is an open debate as to whether there is genuine metaphysical indeterminism operating in the biological world. Some argue that there are good reasons for thinking that there is (Brandon and Carson 1996, Stamos 2001, Glymour 2001, Sober 2010). Some argue that there is not (Rosenberg 1994, Graves et al. 1999, Millstein 2003[a and b], Weber 2005).

Proponents of genuine indeterminism in biology give at least two general kinds of arguments to support their position[40]. Let's call the first type: *Argument from Quantum Indeterminism (QI),* and the second type: *Argument from Autonomy*. Brandon and Carson (1996) provide a good example of the first type of argument. Citing population genetics as their example, the general form of their argument is this.

[40] This distinction owes to Millstein (2003a)

Argument from QI:

> P1. There are scenarios where a single mutational event affects the end state of an entire population of organisms.
>
> P2. Single mutational events occur at the molecular level, and as such, are subject to the effects of quantum indeterminism.
>
> C. Therefore, the end state of a population of organisms is the result of indeterministic processes. (Brandon and Carson 1996)[41]

Prima facie, this argument appears unproblematic. Biological populations, as with everything in the biological world, are the product of evolution by natural selection. And mutation is the grist for the mill of natural selection. If the biological populations are affected by mutations that occur as the result of indeterminism at the quantum level, then it seems to follow the end state of the biological population is the result (at least in part) of indeterministic processes.

Some, however, have noticed problems with this line of reasoning. For one, there isn't much consensus among physicists as to *whether there is indeterminism* operating at the quantum level. Whether there is quantum indeterminism depends crucially on how one interprets Heisenberg uncertainty—a topic that is infamously divisive in study of the foundations of physics[42]. But even if we allow an interpretation of quantum mechanics according to which there is random chance operating at the molecular level, some have pointed out that there is, at best, a tenuous relationship between the chances we assign at the quantum level and the chances we assign at the biological level—where these chances often fail to match up in a way that would suggest quantum chances bubble up to the biology level[43]. It would seem that proponents of the argument from QI, at the very least, owe an explanation of the disparate chance values between the quantum and macromolecular levels.

---

[41] see also Stamos (2001).
[42] See Beller (1999) for a nice survey of the issues surrounding Heisenberg uncertainty.
[43] Cf. Rosenberg (2001) on this point.

Motivated in part by reasons like these, some philosophers take the approach of arguing for biological indeterminism in a way that does not rely on quantum indeterminism. Bruce Glymour (2001) provides a good example of the second type of argument for biological indeterminism. Focusing specifically on what he calls 'random search strategies', he argues roughly in the following manner.

> Argument from Autonomy:
>
> P4. It is selectively advantageous for predators to generate search paths for prey that maximize the amount of the patch searched while minimizing energy expenditure.
>
> P5. Doing this requires that the search path be generated by a mechanism characterized by a random variable *R*.
>
> P6. *R* is essentially indeterministic.
>
> C2. Therefore, we have good reason to believe that there is indeterminism operating at the biological level. (Glymour 2001, 525 – 526)

Notice that the above argument makes no reference whatever to quantum level indeterminism. As Glymour has it, genuine indeterminism can be demonstrated to be a requirement for the optimal foraging patterns exhibited by predators searching for their prey. Without delving into the details of Glymour's account, it is worth mentioning that there are some controversial assumptions made. For one, we might question the role that optimality assumptions should play in biological models. Whether we are entitled to believe that evolution by natural selection settles on optimality is a topic of heated debate among philosophers of biology[44]. Furthermore, even if it can be demonstrated that an optimal search path is best modeled with a random variable, it is a further step to assume that these search strategies are actually adopted in nature. And even if we can make the move from optimality to actual biological phenomena, questions remain about whether the apparent randomness of foraging behavior is the result of underlying

[44]Cf. Beatty (1980), Richardson (2003), Potochnik (2009), Bolduc and Cezilly (2012)

processes that are ultimately deterministic. It would be a mistake to assume that, just because a phenomenon is best modeled using a random variable, that the physical processes which bring it about exhibit genuine metaphysical indeterminism.

Putting these concerns aside, let me re-emphasize that my goal in this chapter is not to advocate or critique either of these types of arguments for biological indeterminism. Rather, it is to offer these arguments as examples of why we might think there is genuine indeterminism operating at the biological level. If either of the above arguments is sound, or indeed any other argument for biological indeterminism, then the PrISM seems to be an easy way to characterize the mechanisms that generate this indeterminism. Whether the indeterminism results from QM or independently, the propensity instantiated by a given stochastic mechanism can simply be characterized as its disposition to behave indeterministically. (More on this is in Section 6.4.)

What about biological determinism? What reason, if any, do we have for thinking that the processes operating at the biological level are, in fact, fully deterministic? One sort of argument for biological determinism is the following (Rosenberg 1994, Horan 1994).

> P7. Quantum indeterminism, if it exists, is only at play at the microphysical level, and only in systems that are sufficiently isolated.
>
> P8. Biological systems are macroscopic systems that strongly interact with their environments.
>
> C3. Therefore, their behavior is only subject to deterministic physical laws and should be considered entirely deterministic.

In other words, indeterministic quantum effects, if there are any, can only exist at (i) a size-level which is sufficiently microscopic so as to subvert the physical laws governing macroscopic objects and (ii) in at a sufficient level of isolation from other systems. These effects disappear, however, as we move upwards from the level of atoms and chemical bonds to systems the size of a living cell or above—as these objects fail to meet the requirements set forth in both (i) and (ii).

Therefore, we should understand biological systems as fully deterministic. As Weber puts it "If biological systems behave stochastically—which they certainly do—this stochasticity is not of the *objective* kind known from QM [quantum mechanics], for example, as in radioactive decay. Instead, biological stochasticity is only *apparent*; it reflects our inability to predict the behavior of complex systems" (Weber 2005, 664).

Returning to the topic of mechanism, Richardson makes the following related charge regarding our, by now familiar, *initiation of electrical activity in post-synaptic neurons* example:

> Return for a moment to the 90% rate for the nerve transduction failure. This is descriptive. It captures what I call the phenomena. The *phenomena* seem to be clearly stochastic. The other indeterminacy would concern what I'd think of as the mechanisms. So there are diffusion rates for the diffusing molecules, which are essentially features characterized by statistical mechanics. The two are of course connected. The indeterminacy of the synaptic transmission rate depends, after all, on the diffusion rate. But the diffusion rate is one settled by molecular dynamics. It's important to be careful here. The systems here are thought of as deterministic, and not stochastic. Diffusion is typically thought of in Newtonian terms, and so deterministic. The *models* are stochastic. The *phenomena* are stochastic too, but that's because they abstract from the factors that would fix them. Put another way, the aggregate patterns are stochastic, but every case is deterministic." (Robert Richardson [personal correspondence])[45]

As Richardson rightly points out, questions about determinism can be directed towards at least three disparate explananda: *models*, *phenomena* (used in the Bogen/Woodward [1988] sense), and *single cases*. And we must take care to keep these separate. It may well be that the best way to *model* a complicated, albeit ultimately deterministic, phenomenon is stochastically. And it may equally be that a given phenomenon, conceived of in the regularistic Bogen/Woodward sense, are best thought of as chancy—but this chanciness only arises as a feature of abstracting from the deterministic factors that fix single-cases. The fact that we might best model something stochastically, or characterize phenomena as stochastic, need not have any bearing on whether each individual case is deterministic. In fact, some of our best science (e.g., molecular dynamics)

[45] This position is further supported by Weber (2005).

seems to suggest that, in the case of synaptic transmission, these single-case instances of synaptic transmission are best thought of as deterministic. If this is indeed the case, then stochasticity may only appear as a feature of the way we *model* synaptic transmission as a *generalized phenomenon*—and if so, the actual case-by-case mechanisms may still be best thought of as fully deterministic.

The upshot of this discussion, I believe, is that the debate about biological indeterminism is at best complicated and at worst intractable. But where does that leave us? For the purposes of developing the PrISM for grounding probabilistic generalizations in biology, where should we stand on the issue of biological indeterminism vs. biological determinism? Below, I propose an answer that won't settle the matter, but one that I believe to be prudent in the face of uncertainty.

6.3 Agnosticism about Biological Determinism/Indeterminism

In the previous section, I highlighted several claims that feature in various arguments regarding the role determinism plays at the biological level.

(QI) There is genuine indeterminism operating at the quantum level.

> (QI Manifestation) Under the right circumstances (e.g., when mutations affect populations), quantum indeterminism can manifest at the biological level.
>
> (Autonomous I) There is indeterministic behavior operating at the biological level (e.g., in optimal foraging patterns) independently of whether there is QI.
>
> (Single-case D) Biological stochasticity (at least in the case of synaptic transmission) is only a feature of our models and of the way we aggregate single-cases into a given phenomenon. Single-case outcomes are best understood to be deterministic.

All of these claims, as we've seen, have been marshaled in arguments for or against biological indeterminism. However, it must be noted that each of these claims has an empirical flavor. Whether we have justification for affirming their truth depends on the science; a priori, armchair reasoning won't (and shouldn't) settle the matter. That said, the science simply isn't in. With

regard to QI, as von Plato puts it, "The various attitudes towards this microphysical indeterminism are reflected in the interpretations of the Heisenberg uncertainty relations. They have been given a purely ontological, as well as an epistemological interpretation, to mention the most extreme positions" (von Plato 1982, 53). Thus, I submit that the existence QI is not a secure foundation for any argument regarding biological indeterminism.

Even less understood are which, if any, are the avenues on which quantum indeterminism, if it did turn out to exist, could travel up to the macro level. And even if we could settle this matter, there remains the difficulty of squaring up the different chance values we ascribe between levels. As to the issue of autonomous biological indeterminism, it may well be that foraging patterns are best modeled stochastically using a random variable. But as we've just seen, the fact that a phenomenon is best modeled using a random variable needn't mean that the phenomenon is indeterministic. And even if the phenomenon is indeterministic, that needn't mean that every single-case foraging pattern is fully indeterministic. All of these issues rest on the science.

Because I would not deign to settle these issues, I suggest that the only tenable philosophical position (following Millstein 2000) is to adopt an *agnosticism* about the above claims regarding biological determinism and see where to go from there.

6.4 Disjunctivism

Supposing agnosticism about biological determinism is the only prudent philosophical position given the lack of scientific consensus, how do we proceed? Recall that, as stated in 5.1, the PrISM seems to handle biological indeterminism quite easily, but seems not to square up with determinism very well at all. Given this observation, I submit that what we say about the PrISM depends crucially on which way the biological indeterminism debate ends up going. In other

words, given our inability to arbitrate the biological determinism debate, any answer to the threat of determinism should take a disjunctive form. *Either* indeterminism exists at the biological level, or it doesn't. The diagram below illustrates my disjunctivist strategy.

Disjunctivism:

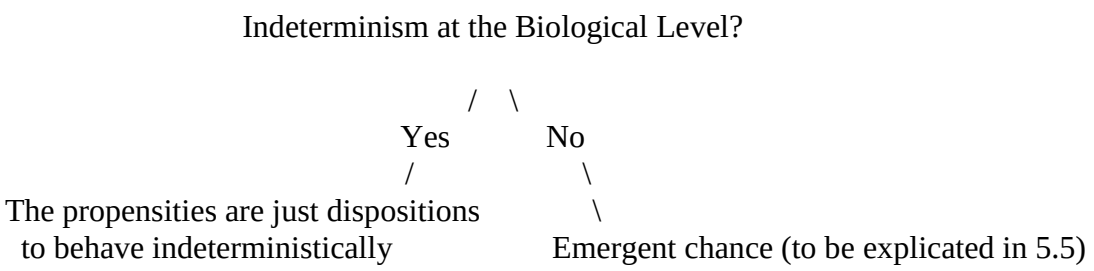

As the picture suggests, I take it that the PrISM is relatively easily characterized if there is indeed genuine indeterminism operating at the level of biological mechanisms: the propensities described in the previous chapter *just are* the dispositions of stochastic mechanisms to behave indeterministically. Suppose, for example, the argument from QI is sound, and biological populations take the form that they do because of mutation events taking place at the molecular level: mutation events that are due (at least in part) to indeterministic quantum events. It strikes me that we can say, without any difficulty, that a particular DNA replication mechanism has a propensity to fail to replicate fidelously. And this propensity just is the disposition for this mechanism to behave indeterminsitically.[46]

Something similar can be said of the argument from autonomy. If it is true that optimal foraging patterns require a random variable, then we can say of the mechanism productive of these foraging patterns that it has a propensity to produce one foraging pattern over another. And the propensity of this mechanism just is its tendency to produce random paths.

[46] It is worth noting here that—if this genuine indeterminism exists—it may only do so on one level of explanation. In which case, we may still need emergent chance even on the indeterminism disjunct.

Suppose, on the other hand, that the arguments against biological indeterminism ring true. Suppose, for example, that apparent random behavior at the quantum level is an artifact of the processes by which we measure quanta. Or suppose that, if there is quantum level randomness, this indeterminacy fails to manifest at the biological level (e.g., at the level of chromosomal assortment, synaptic transmission, DNA replication, etc.). If either is the case, then the stochasticity we observe in the biological world seems to be either (i) merely a feature of the models we use to represent its complexity or (ii) a byproduct of our aggregation of deterministic single-cases under a given regularistic phenomenon. Call this scenario the *difficult determinism disjunct.* Since this disjunct is a perfectly consistent worldview, we'll need a way of making sense of the PrISM that squares up with it. As we shall see, I am cautiously optimistic that one is available. The answer, I go on to suggest, lies in an analysis of emergent chance according to which objective chanciness can exist at one level of description even if it doesn't at another. So let's see how this might work.

6.5 Emergent Chance

Having affirmed an agnosticism about whether determinism holds together with a disjunctivist strategy to proceed, we can now begin to take on the difficult deterministic disjunct. The central question is this: is there any reasonable construal of propensity such that it can coexist with determinism? Following a recent view in the philosophy of probability, I argue that there is a sense in which it can. Specifically, I advocate a strategy formulated in a 2013 paper by List and Pivato (henceforth, L&P) which they call 'emergent chance' and apply it to the multilevel understanding of mechanism advocated by Craver (2007). The goal is to show that genuine indeterministic propensities can be seen to exist at certain mechanism levels even if others are

fully determined. This approach, if successful, will open the door for there to be propensities in existence at the biological level—even if other levels operate in a fully deterministic fashion.

Because L&P's view is quite complex, I'll have to gloss over many of the details of their account and focus on their central ideas. In its most simplified form, L&P can be understood to argue as follows:

> L&P's Central Argument:
>
> P9. All it takes for a probability function (or family of probability functions) to count as objective chance is for it (them) to relate in the right way to various other pertinent concepts surrounding chance (e.g., credence, possibility, the future, intrinsicness, lawfulness, and causation).
>
> P10. Given a certain formal framework together with a few uncontroversial assumptions, it can be shown that these criteria can be met at one level of description even if it cannot at another.
>
> C4. Therefore, objective chance can exist at one level even if it does not at another.

To motivate this argument, I'll say something about both premise (P9) and (P10).

Beginning with (P9), L&P adopt a modified version of Shaffer's (2007) list of criteria for objective chance. Summarized roughly, they are:

> *Chance-credence*: If you have information about the objective chance of an event, you should set your credence level to match that information.
>
> *Chance-possibility*: If you assign a non-zero chance to an event, it must be possible for that event to occur.
>
> *Chance-future*: to say (at some time *t*) that some event has a non-extremal objective chance of occurring requires that the event take place in the future.
>
> *Chance-intrinsicness*: If you assign an objective chance to an event occurring after a certain history, then you must assign the same chance to any intrinsic duplicate of such an event with such a history.
>
> *Chance-causation*: If some event appears causally relevant to another event, then the first event must happen before the other.

> *Chance-lawfulness*: The laws operating at a given level must be seen to determine the chances at that level.

Following Shaffer, L&P consider these criteria to be platitudes about objective chance. Each of them constitutes an important relationship that we would (pre-theoretically) expect objective chance to hold with respect to other important concepts surrounding chance. According to Shaffer and L&P, any probability function that meets these criteria qualifies as what they call *non-degenerative, objective chance* (3). In other words, a probability function that meets these criteria should not be seen to degenerate into merely subjective, epistemic chance. By way of contrast, consider an example where the probabilities we assign fail to meet all of these conditions. Suppose that you role a pair of fair dice, but your friend covers them with her hand before you see them. Suppose you then are asked to assign the chance that the value of the dice totals up to 7. Since you know that there are 36 possible rolls of the two dice—all of which are equally probable. And there are 6 ways to roll 7. You know that the probability of rolling 7 is 6/36 or 0.1667. But this cannot be an objective chance. Why?—because it fails the chance-future desideratum. You've assigned a non-extremal chance to an event that has already happened: you've already rolled the dice. So what you are evaluating is not the objective chance that the value of the dice comes up 7, but the degree of confidence that you have that this *was* the result. In this case, the chance you assign to a value of 7 degenerates into merely epistemic chance.

Returning to L&P's central argument, let's look at what kind of definitions and assumptions they make to get emergent chance off the ground. First we need to get straight on a bit of the technical terminology they employ—some of which should be familiar from 1.7. Among the main technical notions we need from L&P are *determinism* and *indeterminism*. But to get adequate formal definitions of these, we'll need to borrow their definitions of 'time', 'state', and 'history' (L&P 2013). According to L&P, we can define *time* as a set of points

(called *T* points) ordered linearly; and we can define the *state* of a given system at each time as given by an element of some *state space S*. A *history*, on L&P's framework, is a temporal path through some state space, (histories can be seen to play the role of possible worlds—the collections of which they represent with Ω). Furthermore, for any history *h* and any time *t*, L&P designate *ht* to denote the *truncated history* up to time *t* (defined as the restriction of the function *h* to all points in time up to *t* in the relevant linear order). Together, these notions allow for straightforward formal definitions of deterministic and indeterministic histories:

> A history *h* is *deterministic* if, for every time *t*, its truncation *ht* has only one possible continuation in the set of possible histories, Ω, where a *possible continuation* of *ht* is a history *h'* such that *h't* = *ht*. (Ibid, 3)

Conversely,

> A history *h* is *indeterministic* if, for some time *t*, its truncation *ht* has more than one possible continuation in Ω. Thus indeterministic histories allow *branching*, while deterministic histories do not. (Ibid, 3)

In addition to their definitions of deterministic and indeterministic histories, we also need an understanding of what L&P's take *levels* to be. As L&P have it, and this is crucial, levels are not global size-levels[47]. Rather, their use of 'level' refers to a level of *description* whereby what counts as a given level depends on which state space has been delimited. They write,

> The state space *S* could be, for example, the set of all possible microphysical states, and Ω the set of all possible microphysical histories. Often, however, we wish to employ *higher-level descriptions*, for example by describing the state of water as liquid or frozen, rather than as a complex configuration of individual molecules, or by describing a tossed coin as landing heads or tails, rather than as following a particular finely specified physical trajectory. (Ibid, 11)

As this quotation demonstrates, L&P understand levels not to be entirely mind-independent substrates in the world, but as contingent on the what sort of description one aims to give. In this

[47] traditionally implied, for example, in the work of Wimsatt (1976)

sense, what makes 'on the same level' as another thing is whether they both exist in the state space delimited by the one doing the describing.

The last bit of theoretical scaffolding needed from L&P is two of their key assumptions:

(i) Higher-level states and histories *supervene* on lower-level states and histories.

(ii) Higher-level states are often *multiply realizable* by lower-level states.

Regarding (i) supervenience is a kind of dependence. To use an easy example, we might say that the top speed of a vehicle supervenes, all things equal, on its engine. In saying this, we mean that no change can be made to the top speed of a vehicle without, all things equal, having made some change in its engine. Regarding (ii) for something to be multiply-realizable at one level, there must be more than one lower-level configuration that gives rise to it. For example, the higher level description of water as frozen admits of several lower level realizers; there are a myriad of ways that hydrogen and oxygen molecules can be precisely arranged to realize frozen water. Like L&P, I take these assumptions to be uncontroversial.

With these technical definitions and assumptions in place, we can now introduce the central feature of L&P's account of emergent chance: *course-graining* of disparate levels of description. L&P state,

> The relationship between the different levels can be formally captured by the idea of *coarse-graining*: each higher-level state corresponds to an equivalence class of lower level states, consisting of all its possible lower-level realizations. Call a partition of the state space *S* into such equivalence classes a *coarse-graining* of *S*, and let *S** denote the set of all equivalence classes [my use of star]... let σ denote the function that maps each lower level state *s* in *S* onto its higher level state s* onto S*. (Ibid, 11)

L&P apply the same idea to histories:

> For each lower-level history *h* in Ω, the corresponding higher-level history *h* is the function from *T* into *S* such that, for each *t* in *T*, $h(t) = \sigma(h(t))$. The set of all possible higher-level histories is the projection of Ω under σ, formally $\Omega = \sigma(\Omega)$. (Ibid, 12)

This notion of course-graining is the crucial final piece of the puzzle to allow for emergent chance. The idea is this: For any course graining function σ and sufficiently large set of possible histories Ω, there can be indeterministic histories in Ω* *even when all histories in Ω are deterministic*. To illustrate how, L&P provide an example of course-graining from the real numbers to the integers (figure 11). Part (a) depicts a simple system at the lower level of description (Ω). Time is plotted on the horizontal axis (T={1,2,3,...}), and the state of the system on the vertical one. Here the state space S is the set of all real numbers. The figure displays five deterministic histories, from time 13 t = 1 to time t = 25. Part (b) shows the same system at a higher level of description (Ω*), obtained by coarse-graining the state space S into S*. Specifically, S* is the set of all integers. The coarse-graining function σ maps each real number s in S to the closest integer s* in S* (with the usual rounding convention). In this coarse-grained description, there is now one indeterministic history, supervenient on the 5 lower-level deterministic ones. In particular, they all coincide up to time t = 9 before diverging from one another.

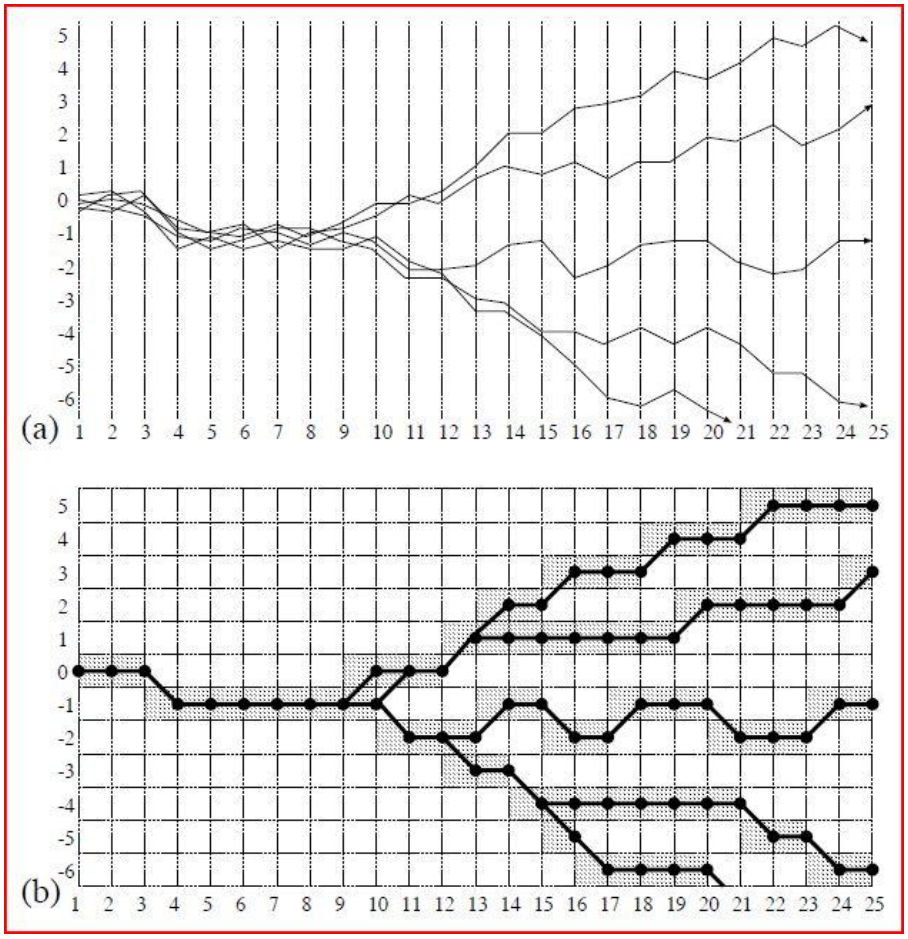

Figure 11. Course graining (From List and Pivato 2013, 13).

As figure 11 shows, the result of course-graining (a) into (b) is that there are genuine forking paths in (b) while there are not in (a). While each of the 5 histories plotted in (a) is deterministic, there is a single history in (b) resulting in several different outcomes.

Importantly, this transformation is not simply the result of ignoring the fine-grained information in (a). It is rather that, when we use the language at the (b) level of description, this model meets all the criteria for non-degenerative, objective chance. As L&P put it, "when evaluating chance and (in)determinism at a higher level of description, *only higher-level language is available*" (Ibid, 14). When we limit ourselves to one level of description, we *must* only allow ourselves the language available in that state-space. And when we do so, chance

literally emerges. We know this because, at the (b) level of description, all of the platitudes for non-degenerative, objective chance are met.

6.6 Multilevel Mechanisms

Supposing L&P's account of emergent chance is sound, why does it matter for the project of dismantling the threat of determinism for the propensity interpretation of stochastic mechanism? In answer to this question, it will help to draw on the work of Carl Craver—who has produced the most thorough analysis of multilevel mechanisms in the mechanisms literature. Using his analysis as well as one of his central examples, we show a direct application of emergent chance to multilevel mechanisms. And in doing so, we will have shown a plausible instance where genuine, non-degenerate chance may be able to exist at one mechanism level even if determinism holds at another. And if the level at which chance emerges can be shown (plausibly) to be the biological level, then there may be room for propensities at play in stochastic mechanisms after all.

Carl Craver (2007) introduces an account of multilevel mechanisms as a way of capturing "the central explanatory sense in which explanations in neuroscience (and elsewhere in the special sciences) span multiple levels" (Craver 2007, 163). Because explanations in our most successful special sciences appeal to multiple levels, and many of these explanations are mechanistic, there seems to be a need for a philosophical analysis of multilevel mechanisms. However, as Craver correctly points out, talk of 'levels' in science is often ambiguous between many senses of the term. As such, Craver carefully disambiguates different understandings of levels, offering arguments against most of them, and settling on his own preferred understanding. For example, Craver argues that mechanism levels should not be understood merely as levels of different sized objects (e.g., societies, organisms, cells, molecules, atoms). Additionally he

claims, “levels of mechanisms are not monolithic divisions in the structure of the world” (Ibid, 190).

Craver’s central idea is that mechanism levels are *levels of behaving components*. As he puts it, “Lower-level components are made up into higher-level components by organizing them spatially, temporally, and actively into something greater than a mere sum of the parts” (Ibid, 189). Regarding organization, Craver says “organization is the interlevel relation between a mechanism as a whole and its components” (Ibid, 189). Figure 12 represents a Craver-style multilevel mechanism. At the top level is an entity *S* engaging in an activity *ψ*. At the second level, there are four *X*’s *θ*-ing organized in *S*’s *ψ*-ing, And at the bottom level, four *P*’s are *ρ*-ing such that they are organized in *X*3’s *θ*-ing. On Craver’s view, mechanism levels are local and only definable within a compositional hierarchy. Examples of multilevel mechanisms operating in the biological world are the mechanisms in the spatial memory system, the circulatory system, the osmoregulatory system, and the visual system (Ibid, 191).

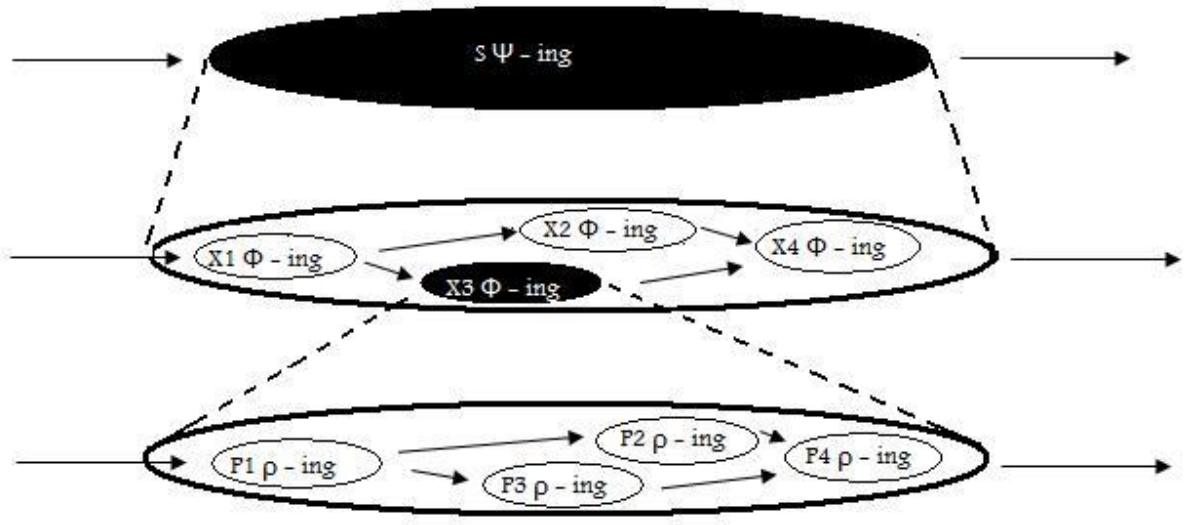

Figure 12. Nested multilevel mechanism (Reproduced from Craver 2007, 189).

So to sum up, according to Craver, levels of mechanisms are levels of composition, but this composition relation is not to be understood as merely spatial or material; instead, levels of mechanisms relate to one another by virtue of the relationship of composition. As Craver has it, mechanism levels are distinguishable by virtue of a compositional relationship. Different mechanism levels need not correspond to different size levels or spatial locations.

6.7 Complementarity of L&P's Emergent Chance and Craver's Multilevel Mechanisms: Toward a Solution to the Difficult Determinism Disjunct

Before moving to a concrete example in which both L&P's account of emergent chance and Craver's analysis of multilevel mechanism apply, it's worth pausing to highlight a few features of these two accounts regarding which they seem to complement each other quite nicely.

A. Neither view espouses global, monolithic levels.

B. Both views assume or imply the acceptance of supervenience and multiple realizability—so are amenable to emergence

Let me say something about each. With regard to (A), neither L&P nor Craver accepts a view of levels according to which they are understood monolithically. That is, on neither of these views does levelhood merely correspond to size or objecthood. Rather, both views accept a degree of relativism regarding levelhood. For L&P, the content of a level is relative to the state-space delimited by the inquiry and can only be described by the language endemic to that state-space. On Craver's view, the content of a level is relative to the organization of the mechanism responsible for the phenomenon in question.

With regard to (B), L&P's view explicitly assumes supervenience and multiple-realizability between levels—such that objective chance can emerge at a higher level—even when it isn't present at a lower one. However, it is important to notice that Craver implicitly endorses emergence when he claims of multilevel mechanisms that "Lower-level components

are made up into higher-level components by organizing them spatially, temporally, and actively *into something greater than a mere sum of the parts*" (italics added for emphasis). As this quotation demonstrates, Craver endorses a view of mechanism levels according to which the active entities at one level constitute a higher level that is 'greater than the sum of its parts'. Emergence, as it's commonly defined, *just is this relation*. So on Craver's view, higher mechanisms emerge from lower ones.

Having just highlighted these important complementary features of Craver's multilevel mechanism and L&P's account of emergent chance, there is an important qualification to make.

> C. L&P's account of emergent chance does not require levels to stand in a compositional relation.

As evidenced in the previous section, Craver endorses a view of multilevel mechanism according to which mechanism levels, by definition, stand in a compositional relation with one another. L&P, on the other hand, make no such requirement of disparate levels of description. That said, however, this need not prevent the successful union of these views. After all, there is a clear sense in which different levels of composition (a la Craver) can still be understood as different levels of description (a la L&P). Recall that, for L&P, what it takes to specify a distinct level is only to delimit the state space encompassing the entities at that level. Craver's levels of composition approach merely places a restriction on what sorts of state spaces we can delimit. Namely, the state spaces delimited must correspond to what composes state spaces at a higher level. As such, there is no conflict between the two approaches—at least when applied to mechanism levels.

6.8 Synthesizing Emergent Chance and Multilevel Mechanism: Spatial Memory in Mice and Rats

I cannot undertake a detailed defense of the two views outlined above. So here is a modest claim: *If* L&P succeed at formally systematizing emergent chance, and *if* Craver offers a plausible analysis of multilevel mechanisms, *and* the two positions complement each other, *then we have the resources to show that it may be possible for genuine propensities to exist at the certain mechanism levels even if determinism holds at others. And in doing so, we will have offered a route for the PrISM to escape the threat of determinism.* Given the conditional nature of this claim, I'll proceed under the assumption that both L&P's account of emergent chance and Craver's analysis of multilevel mechanism are successful—and that they cohere with one another without contradiction. Having done so, we can proceed to explore a concrete example wherein the synthesis of emergent chance and multilevel mechanism can begin to show its merits.

6.8.1 A Multilevel Mechanistic Analysis of Spatial Memory in Mice and Rats

Consider the spatial memory of mice and rats. To investigate this phenomenon, researchers have placed these animals in a variety of testing environments (e.g., radial arm mazes, sunburst mazes, and three-table problems). In one particularly revealing test, the Morris water maze[48], test subjects are placed in a circular pool containing opaque fluid covering an occluded platform. The mice and rats are then trained to use various environmental cues (e.g., colored shapes) to locate the platform and escape the water. During the test, various parameters are recorded (e.g., time spent in each quadrant of the pool, time taken to reach the platform, and total distance traveled). Successful escape from the water appears to reinforce a desire to quickly find the platform, because on subsequent trials, with the platform in the same position, subjects demonstrate an ability to locate the platform increasingly quickly. The thought is that this

[48] Developed by Richard Morris (1981)

improvement in performance occurs as a result of *learning and memory* for the location of the hidden platform relative to the visual cues. After a sufficient number of trials, a high-performing rat is able to swim directly from any release point to the platform.

This phenomenon lends itself nicely to multilevel mechanistic explanation. In other words, the mechanism productive of the ability for mice and rats to learn and remember the location of the occluded platform admits of a multilevel analysis. As Craver describes in detail (ibid, 167), part of the explanation of mice and rats' spatial memory learning is at the *gross anatomical level*. More specifically, researchers now think that the hippocampus, a region in the medial temporal lobe, is required for spatial memory learning. This is evidenced by the fact that subjects with lesions to the hippocampus show significant deficits in maze learning. Some researchers even suggest that we have good reason to think the hippocampus, in itself, functions as a kind of spatial map.

But how is this apparent spatial-map-functionality of the hippocampus explained? According to Craver, the answer seems to lie at the next mechanism level down: which he calls the *cellular-electrophysical level*. At the cellular-electrophysical level, spatial memory appears connected to a process called "Long Term Potentiation" (LTP). LTP is understood as a long-lasting enhancement in signal transmission across neurons resulting from synchronous stimulation. LTP has for several decades been thought to play a role in memory. More specifically, neuroscientists have speculated that the brain *stores* memory by altering the strength of connections between neurons that are simultaneously active—and that much of this takes place in the hippocampus.[4950]

[49] LTP occurs at synapses that terminate on dendritic spines (a small membranous protrusion from a neuron's dendrite that typically receives input from a single synapse of an axon) and uses the transmitter glutamate. Researchers have discovered that several of the major pathways within the hippocampus fit this description.
[50] This idea was formalized by Donald Hebb (1949)

But the mutlilevel mechanism responsible for spatial memory learning in rats doesn't bottom out at the electro-physical level with LTP—for there are even lower-level mechanisms which give rise to LTP in the hippocampus. Craver calls this mechanism level, the *molecular level* (169). He writes,

> When the presynaptic neuron is active, it releases glutamate (GLU). This glutamate binds to N-methyl-D-aspartate (NMDA) receptors on the postsynaptic cell. The NMDA receptors change their conformation, exposing a pore through the cell membrane. If the postsynaptic cell is inactive, the channel remains blocked by large Mg2+ ions. But if the postsynaptic cell is depolarized, these Mg2+ ions float out of the channel, allowing Ca2+ to diffuse into the cell. The rising intracellular Ca2+ concentration sets in motion a long chain of biochemical activities" (Craver 2007, 169).

Researchers know that NMDA receptors play an essential role in LTP taking place in the hippocampus because drugs that interfere with NMDA receptors block LTP and have major effects on some types of memory, especially spatial memory.

Figure 13 (from Craver 2007, 166) illustrates the multiple levels in the mechanism of learning and memory in rats. As the figure shows, the phenomenon of rats learning to navigate through mazes is explained by a mechanism with multiple levels: the spatial map functionality of the hippocampus at the gross anatomical level, synapses inducing LTP at the electro-physical level, and NMDA receptor binding and bending at the molecular level. And these levels are organized into a hierarchy of part-whole relations to one another such that each of the lower-level entities and activities are components of the higher-level mechanism level.

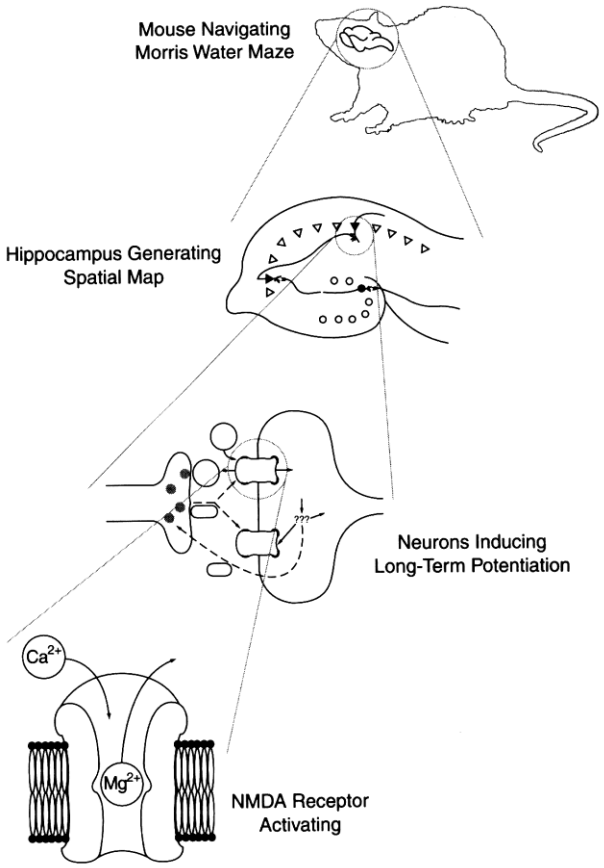

Figure 6.4 Levels in the hierarchical organization of the mechanism of spatial memory.

Figure 13. Multilevel mechanism of spatial learning in rats (From Craver 2007, 166).

6.8.2 Spatial Memory and Emergent Chance

With a concrete example of a multilevel mechanism operating in the biological world now in hand, let's see whether L&P's account of emergent chance can shed any light on it. What I suggest is this: the afore-described example of the multilevel mechanism of spatial learning in rats is a plausible instance where *genuine (non-degenerative) chance may exist at one mechanism level even if determinism holds at different level*. In the case of rats' spatial learning, this could take place in two ways:

> A) The molecular level is fully deterministic, but chance emerges at the electro-physical level.

> B) Both the molecular and electro-physical levels are fully deterministic, but chance emerges at the gross anatomical level.

Important disclaimer: my aim here is not to show that either of these possibilities *actually is the case*. Much more modestly, rather, I offer these as instances where the requirements for emergent chance *may* plausibly be met between mechanism levels in the natural world. To motivate this modest claim, I'll subject the example of spatial learning in rats to the criteria required for L&P-style emergent chance and demonstrate why I think it might be able to meet them.

As articulated in Section 5.5, there are several characteristics that must obtain between levels of description in order for L&P-style emergent chance to exist. First, and most obviously, the phenomenon in question must admit of multiple levels of description—each with its own defined state-space. Second, higher levels must stand in a supervenience relation to lower levels such that higher levels are multiply-realizable by their supervenience bases; i.e., each higher-level state must correspond to an equivalence class of lower level states, consisting of all its possible lower-level realizations. Third, there must be an available coarse-graining function that maps each lower level state onto its higher level state. I'll say something about each.

To begin with, as summarized above, I take Craver to have demonstrated quite clearly the sense in which the case of spatial memory in rats admits of a multilevel description. And just as L&P require, each mechanism level delimits a distinct state-space with a corresponding language of description. The gross-anatomical level defines as its state-space a specific region of the brain: the hippocampus; and adopts descriptive terms like 'spatial maps' and 'memory storage'. The electro-physical level defines as its state space particular neurons inside the hippocampus, and takes as its descriptive terms: 'axons', 'synapses', 'long term potentiation', and 'signal transmission'. And the molecular level defines as its state space particular molecules organized

in and around these neurons; it takes as its descriptive terms: ‘glutamate’ (GLU), ‘N-methyl-D-aspartate’ (NMDA), ‘Mg^{2+} ions’ and ‘Ca^{2+}’.

Furthermore, with regard to the multiple realizability/supervenience requirements, it seems uncontroversial to espouse a supervenience relation between these mechanism levels. In other words, no change in the memory storage of particular rat’s hippocampus will occur without some change at the electro-physical level. The same goes between the electro-physical level and molecular levels. Additionally, and only slightly more controversially, there is a clear sense in which higher mechanism levels seem plausibly to be multiply realizable by lower ones—at least if what this means is that *the same phenomenon* can be brought about by (and explained by) underlying mechanisms that either have different components or components acting in different ways. Craver seems to admit this when he writes, “There are no doubt epistemic difficulties of determining when two mechanistic contexts are equal, but there is no conceptual difficulty seeing how the same type of phenomenon could be explained by different components in different contexts” (Craver 2007, 160). The phenomenon of rats’ memory storage, for example, can occur via different specific instances of LTP. Likewise, LTP can be instantiated by different specific molecular activities. All that is necessary for the spatial memory example to meet L&P’s supervenience/multiple realizability requirement is that each higher level state (e.g., memory storage in the hippocampus, or LTP across neurons) corresponds to an equivalence class of lower level states (e.g., the set of possible electro-physical or molecular realizers of these higher level states). It strikes me that such a regimentation is achievable.

The only thing left, then, is to explore whether there might be a course-graining function that obtains between these levels. In other words, we must now ask whether there is a function that maps each lower level state *s* (e.g., molecular level or electro-physical level state) in *S* (set of

possible lower-level states) onto its higher level state s* (e.g., gross anatomical level state) in S* (the set of possible higher level-states). My view is that there may indeed be such a coarse-graining function available. However, I lack both the space and technical prowess to produce one here. That said, the rough idea is that we should be able to come up with a function that will map each of the molecular states in the set of possible states that make up the electro-physical supervenience base onto a particular electro-physical state—just as L&P have a function from the real numbers to the integers. And in doing so, we will have shown that there may be multiple deterministic histories at the molecular level even if there is one indeterministic history at the electro-physical level. The same strategy applies between the electro-physical level and the gross anatomical level.

6. 9 Objection: But are these Non-degenerate Chances 'Propensities'?

Suppose we accept everything argued for thus far: that (i) L&P's account of emergent chance is a plausible construal of how non-degenerative chance can exist at one level of description even if it does not at another level; (ii) Craver offers a plausible analysis of multilevel mechanisms, (iii) these two accounts seem complementary in some important respects; and (iv) given these three forgoing conclusion, it may well be the case that genuine, objective, non-degenerative chance can exist at one mechanism level even if another mechanism level is fully deterministic. Suppose we even accept that the example of spatial memory in rats provides a good application of the union of these views. What reason do we have to designate these emergent, biological-level chances as *propensities*? After all, none of this theoretical scaffolding does us any good unless it defends the *propensity* interpretation of stochastic mechanism against the threat of determinism. Without any such independent argument for understanding these emergent chances as propensities, we may not have rescued the PrISM at all.

6.10 Reply

I offer such an independent argument here. My strategy is this. I adopt a list of criteria offered by von Plato for something to count as a propensity interpretation, and show that the emergent chance view of multilevel mechanisms meets these criteria just fine.

von Plato's writes,

> According to the propensity interpretation of probability, at least the following is required of probabilities as objective propensities:
>
> (1) a propensity is a physical property of the experimental arrangement;
>
> (2) it is a dispositional property, having an underlying structural basis;
>
> (3) and it concerns individual events." (von Plato 1982, 62 [numbering added for organizational purposes])

Following this list, we shall now investigate whether the synthesis of Craver's analysis of multilevel mechanism and L&P's view of emergent chance meets these qualifications.

Regarding (1), the chance that emerges at higher mechanism levels is a physical property of the experimental arrangement. Namely, the propensity is a property instantiated by a given rat's spatial memory mechanism as it interacts with the water maze. Regarding (2), the propensities exhibited by these rats have as their underlying structural basis their respective spatial memory mechanisms. Regarding (3), there are propensities for stochastic decision-making at each individual maze-navigation task; so there are propensities of individual events.

As such, von Plato's requirements for a propensity interpretation are fully satisfied. And we, therefore, have good reason to construe the chances that emerge at higher mechanism levels as propensities.

6.11 Conclusion

In this chapter, I have surveyed a few of the general ways philosophers argue for and against biological indeterminism. I have concluded that the only philosophically prudent position regarding this debate is to adopt an agnosticism about the central claims made therein. As such, I have advocated a disjunctivist approach whereby what we say about the PrISM splits in two directions depending on which side of the biological indeterminism debate prevails. I suggested that, if the arguments for biological indeterminism are sound, the propensities postulated by the PrISM just are the dispositions of stochastic mechanisms to behave indeterministically (whether this indeterminism is the result of quantum-level indeterminacy or it can be seen to exist independently). On the other hand, if determinism holds, I argued that a recent view of emergent chance together with an understanding of multilevel mechanisms may provide a way of dismantling the threat of determinism. Provided that the explanadum in question meets the criteria for multilevel mechanistic analysis, and provided that a course-graining function obtains between these levels, then it may well be that genuine, non-degenerate chance can exist at one mechanism level even if others are fully deterministic. The example of spatial memory in rats, I suggested, may be a good example where these conditions are met. I concluded by offering an independent argument for why these emergent chances should be understood as propensities. The result of these arguments, I hope to have shown, is that determinism may not be such a threat to the PrISM after all.

Chapter 7. Conclusion

Abstract: In this concluding Chapter, I first briefly summarize the main arguments offered in the dissertation. I then undertake two remaining tasks: (1) to suggest a few future avenues of research in the philosophy of science opened by this dissertation and (2) to suggest a few potential applications of my account to further areas of empirical study.

7.1 Dissertation Summary

In this dissertation, my goal has been to develop an account of stochastic mechanism for grounding probabilistic generalizations in the life sciences. To do this, in Chapter 1, I first offered some historical context for the reemergence of mechanistic philosophy in the life sciences. I then offered examples of probabilistic phenomena in biology for which an account of stochastic mechanism seems explanatorily necessary and useful: synaptic transmission in the brain, protein synthesis, DNA replication, and evolution by natural selection. To begin to understand whether these probabilistic phenomena can be explained mechanistically, in Chapter 2, I examined the extent to which processes should be required to behave regularly in order to count as mechanisms. To do this, I examined and augmented a taxonomy of mechanistic regularity offered by Andersen (2012). But I argued, contra Andersen, that regularity (of any sort) should not be considered a metaphysically demarcating feature of mechanisms. However, mechanisms with a high degree of regularity should hold a privileged role in scientific explanations. In Chapter 3, I then set about using the taxonomy of mechanistic regularity developed in Chapter 2 to help sort out an active debate in the philosophy of biology: is natural selection regular enough to count as a (privileged) MDC mechanism? My answer was yes, and from this debate I drew some deep consequences for the prospects for grounding probabilistic biological generalizations in a mechanistic explanatory framework. In Chapter 4, I then set about exploring how we should understand the chance we ascribe to (a certain sort of) stochastic mechanisms: those whose outcomes are irregular even when their stet-up and triggering

conditions obtain. To do this, I formed a list of desiderata that any account of stochastic mechanism must meet. I then took the general characterization of mechanism offered by Machamer, Darden, and Craver (2000) and explored how it fit with several of the going philosophical accounts of chance: subjectivism, frequentism (both actual and hypothetical), Lewisian best-systems, and propensity. I argued that neither subjectivism, frequentism, nor a best-system-style account of chance can meet all of the proposed desiderata, but some version of propensity theory can. By pairing an account of stochastic mechanism with a propensity interpretation of chance, I suggested, some of the shortcomings of each position could be avoided. Borrowing from recent propensity accounts of biological fitness and drift, in Chapter 5, I then went on to explore the prospects of developing a propensity interpretation of stochastic mechanism (PrISM) according to which propensities are (i) metaphysically analyzable and operationally quantifiable via a function of probability-weighted ways a mechanism might fire and (ii) not causally efficacious but nonetheless explanatorily useful. By appealing to recent analyses of deterministic and emergent chance, in Chapter 6, I argued further that this analysis need not be vulnerable to the threat of metaphysical determinism.

In what remains of this concluding Chapter, I set forth two further tasks: (1) to suggest a few future avenues of research in the philosophy of science opened by this dissertation and (2) to suggest a few applications of the analysis of stochastic mechanism heretofore offered in this dissertation to further areas of empirical study. Far from an exhaustive list of future projects opened by my account, I take the following potential areas of applicability to merely be a small sample of the myriad of ways the work done in this dissertation might prove to be fruitful in the future.

7.2 Avenues for Future Philosophy of Science Research

In this section, I propose a few potential avenues for future research in the philosophy of science opened by this dissertation. I focus on three: (i) additional work on the metaphysics of mechanism, (ii) an exploration of the epistemology of propensity, and (iii) potential implications for debates on statisticalism vs. causalism with regard to natural selection. I'll say something about each.

7.2.1 More on the Metaphysics of Mechanism

Recall that, in section 2.4.3 of Chapter 2, I briefly discussed whether Andersen is correct that mechanists need to appeal to regularity as a means of distinguishing between one-off mechanisms and mere single causal chains. I suggested that mechanists do not need to rely on regularity for this reason, but may instead be able to appeal to mechanisms' organizational structure as a means of underpinning a distinction between mechanism and causation in general. But I admitted that much more would need to be said to make this precise. One avenue of future research opened by this dissertation would be to attempt to augment MDC's list of organizational properties in order to better establish a metaphysical difference between processes we wish to deem as mechanisms and those we would only want to understand as run-of-the-mill causation. Here is a hint of how it might go.

The MDC view, recall, states that entities in a mechanism must be (1) located (2) structured, (3) oriented; and a mechanism's activities must have (4) temporal order (5) rate, (6) duration. In future research, I propose to add to this list. At first pass, here are some further organizational properties of instantiated by mechanisms but not by mere single causal chains: (7) connectedness, (8) existence absent triggering, and (9) temporal persistence. As a precautionary note, it will be important to stress that *not all* mechanisms instantiate each of these properties. As

such, my guess is that (7) – (9) will best be thought of as the sort of *family resemblance* characteristics offered for other tenuous biological concepts (e.g., species or gene). What is important for the purpose of distinguishing single causal chain from one-off mechanism is that one-off mechanisms generally instantiate (at least) some of these properties; and single causal chains generally do not. Of course, much more will need to be said to make this precise. For now, it will suffice to briefly indicate what I take each of these proposed organizational properties to be. But I'll leave the rest for future research.

> 7. One of the properties that mechanisms sometimes have, that mere single causal chains do not, is (what I'm calling) *connectedness.* A one-off mechanism comprises several entities *connected* to one another. The entities in WWI only bump into each other.
>
> 8. A second property of a one-off mechanism that mere single causal chains do not have is (what I'm calling) *existence absent triggering*. One-off mechanisms exist as an organized system even if they never gets triggered. The events that brought about WWI did not exist as an organized structure until the assassination of the Archduke.
>
> 9. Similarly, we might describe a third property of one-off mechanism that is absent from a mere single causal chain: (what I'm calling) *temporal persistence.* One-off mechanisms persists as a system of organized entities for the period of time after its construction until its detonation. The events leading up to WWI only existed at the instants that they occurred.

MDC also say of mechanisms that they produce regular outcomes in ways that are "non-accidental and support counterfactuals" (Ibid, 7-8). They write,

> For example, if this single base in DNA were changed and the protein synthesis mechanism operated as usual, then the protein produced would have an active site that bonds more tightly. This counterfactual justifies talking about mechanisms and their activities with some sort of necessity (Ibid, 8).

As indicated by this quotation, another potential means for distinguishing the metaphysical properties of mechanisms from those held by run-of-the-mill causation might have to do with counterfactual support and natural necessity. Thus, another future project might be to more carefully specify what (if any) special relationship to counterfactuals held by mechanisms. And if

a special relationship could be found, what (if any) conclusions might be drawn about mechanisms' relationship to natural necessity.

7.2.2 An Exploration of the Epistemology of Propensities

In Chapters 4 and 5 of this dissertation, I laid out several considerations in favor of a propensity interpretation of stochastic mechanism. I anticipate, however, that many philosophers of probability will have objections to this. One of the main reasons why philosophers of probability have not taken kindly to propensity accounts of objective chance is that they are epistemically untenable. Carl Hoefer (2007), for example, writes:

> …propensities are epistemologically hopeless (i.e., one can only claim that statistics are a reliable guide to propensities *via* arguments that are all, in the end, ineffective—usually, circular). (Hoefer 2007, 551)

Here, Hoefer suggests that proponents of propensity accounts of objective chance run into a serious epistemological problem. They can't give an adequate account of how propensities come to be known. If the propensity proponent attempts to argue that statistics reveal propensities, they run the risk of circularity. After all, we might imagine Hoefer wondering, how could one expect statistics to be a reliable guide to propensities unless one already presumed propensities to generate the observed statistics? In other words, the only way to establish that statistics are any kind of reliable guide to propensities is to already have knowledge that propensities are what brought about these statistics in the first place.

Although I cannot here offer such an argument, my suspicion is that the PrISM might offer some potential resources for answering to these kinds of epistemic concerns about propensities. In short, my suspicion is this: *stochastic mechanisms could provide an epistemic entry point between the propensities and the statistics—thereby breaking the vicious circle*.

As I argued in Chapter 4, propensities together with stochastic mechanisms seem to afford special advantages when paired with one another, advantages absent from either alone. One of the most important of these advantages is that mechanisms are actual, existing structures in the world. This allows mechanisms to be the objects of empirical study by scientists. Indeed, it is through experimentation on these actual, existing mechanisms that scientists might get some independent epistemological traction on what lies between propensities and statistics. On my view, propensities are dispositional properties instantiated by stochastic mechanisms. This allows there to be antecedent science for determining a given stochastic mechanism's propensity to successfully fire, science that *goes beyond* merely keeping track of output frequency statistics of a given mechanism. The mechanism can be intervened on to determine details about *why* the output frequency statistics are what they are. If my account offers hints at a methodology for determining why mechanisms have the output frequencies they have, then postulating propensities for a given mechanism is not just a matter of observing its output frequency statistics.

7.2.3 Implications for Debates between Statisticalists vs. Causalists about NS

In Chapter 3, recall that I argued that natural selection can survive the regularity critique leveled at it by Skipper and Millstein. As such, natural selection seems perfectly amenable to an analysis via the account of stochastic mechanism developed in this dissertation. If this is so, there is one current debate in the philosophy biology that a PrISM analysis of natural selection might help to sort out. Namely, is natural selection a causal process or a non-causal, purely statistical phenomenon? (cf., Millstein 2006; Forber and Reisman 2007; Abrams 2007; Matthen and Ariew 2002, 2009; Glennan 2009; Lange 2013).

Here, again, my suspicion is that an understanding of natural selection as a propensity-backed stochastic mechanism would supply some potentially unificatory answers to several of the concerns motivating both sides of this debate. On the one hand, token instances of natural selection (e.g., Barros's crab example) are understood as fully causal. Individual entities (e.g., crabs and snails) engage in causally productive activities (e.g., predation and reproduction), with the result that more fit organisms tend to survive (e.g., snails with lower-spired shells will come to dominate). On the other hand, however, typed schematizations of natural selection are not causal. They are abstract representations which generalize over specific token instances of selection. Furthermore, as I argued in Chapter 4, the chance we ascribe to natural selection should be conceived of as propensities which are themselves causally inefficacious, but which are nevertheless causally relevant. If fully fleshed out, my suspicion is that many of the worries on both sides of the casualist/statisticalist debate would fall away.

7.3 Potential Applications of the PrISM to Further Areas of Empirical Study

In addition to the forgoing topics in the philosophy of science, this dissertation has potential applications to several further areas of empirical study. Here, I mention four such possible avenues: (i) human neuroscience, (ii) foundations of physics, (iii) game/decision theory, and (iv) medicine.

7.3.1 Human Neuroscience

Recently, neuroscientists have become increasingly interested in studying various stochastic processes taking place in the human brain. It strikes me as highly plausible that these stochastic processes might fit nicely into the framework I've offered in this dissertation.

One example of a stochastic process recently studied by neuroscientists interested in the human brain is something called *binocular rivalry*. Researchers have found that, when humans are presented with stimuli dichoptically, their perception alternates between the two stimuli in a stochastic manner. Over the last four decades, there has been rigorous debate (cf., Blake et al. 1971; Tong 2001; Urte 2011) over how to explain this phenomenon. However, there seems to be a growing consensus that binocular rivalry is the result of a dynamic competition between multiple levels of the visual system's hierarchy (Roumani and Moutoussis 2012). According to Roumani and Moutousis, "The role of low- and high-level adaptation mechanisms in controlling these perceptual alternations has been a key issue in the rivalry literature. Both types of adaptation are dispersed throughout the visual system and have an equally influential, or even causal, role in determining perception" (Roumani and Moutoussis 2012). This explanation coheres with the account offered in this dissertation because (i) it involves postulating a mechanistic basis for this stochastic phenomenon, and (ii) this mechanistic basis spans multiple levels of description. The considerations laid out in Chapters 2-4 of this dissertation, I suspect, might help precisify the manner in which these stochastic brain phenomena (like binocular rivalry) have a mechanistic basis. And the work I've done in Chapter 6 in which I propose a synthesis between levels of mechanisms and emergent chance might do well to buttress the inter-leveled nature of this phenomenon.

7.3.2 Foundations of Physics

The foundations of physics is another area of empirical study to which this dissertation might have potential applicability, although my lack of expertise in this area certainly necessitates that I would need collaborators to accomplish it.

Microphysicists have long debated the extent to which quantum mechanics is fundamentally stochastic (cf., for example, Aron 1979, 1986; Prugovečki 1984, Butterfield 2007; and Khrennikov 2011). Although the account offered in this dissertation has been developed, in large part, to provide metaphysical and explanatory resources for biology, it would be worth exploring the extent to which it might also provide similar resources to physics. The possibility of pursuing such a project is made more plausible by recent work by Mauricio Suárez, a philosopher of physics who has recently published several highly-regarded papers in which he defends a role for propensity theory in quantum physics (Suárez 2007, 2013, and forthcoming).

7.3.3 Game/Decision Theory

Beyond human neuroscience and foundations of physics, this dissertation might have potential applicability to several topics in game theory and decision theory.

In a 1997 paper entitled "Stochastic Evolution of Rationality", Falmagne and Doignon investigate "possible mechanisms explaining how preference relations are created and how they evolve over time." They go on to postulate

> …a preference relation which is initially empty and becomes increasingly intricate under the influence of a random environment delivering discrete tokens of information concerning the alternatives. The framework is that of a class of real-time stochastic processes having interlinked Markov and Poisson components. (Falmagne and Doignon 1997, 107).

Once again, the work I've done in this dissertation seems applicable here because (i) mechanisms play a significant explanatory role for these researchers, and (ii) there are stochastic processes at play in the way humans make decisions.

In an experiment on the stochastic nature of human decision-making, Sopher and Narramore (2000), explore the relative merits of stochastic vs. deterministic models of human choice. This seems to be another excellent opportunity for my account to weigh in. Specifically,

my appeal to a determinism-friendly understanding of stochastic mechanism might, once again, provide an interesting middle ground between these stochastic vs. deterministic understandings of human decision-making. Similar applicability might arise from current debates surrounding the nature of revealed-preferences in stochastic environments (cf., Avahami et al. 2005 and Heufer 2011).

7.3.4 Medicine

In a couple of recent papers on cancer cells, Jean-Paul Capp argues that stochastic gene expression is "the driving force of cancer" (Capp 2011, 2012). Prominent philosopher of biology, Joel Velasco, also has a recent paper on whether the probabilities associated with stochastic gene expression should be understood objectively or subjectively (Velasco 2012). He suggests that

> …while our best models of the phenomena in question are stochastic models, this fact should not lead us to automatically assume that the processes are inherently stochastic. After distinguishing between models and reality, I give a brief introduction to the philosophical problem of the interpretation of probability statements. I argue that the objective vs. subjective distinction is a false dichotomy and is an unhelpful distinction in this case (Velasco 2012, 5).

This statement by Velasco shows many commonalities with the conclusions I have drawn throughout this dissertation. In particular, Velasco recognizes that stochastic models need not indicate stochasticity in the phenomenon being modeled—a key feature of my account. As such, another interesting potential application of my work on stochastic mechanisms would be to bring it to bear on some of these issues in stochastic gene-expression as it relates to cancer research.

7.4 Conclusion

In this concluding Chapter, I have briefly summarized the main conclusions of this dissertation. I then offered three potential avenues for future research in the philosophy of science opened by

my account: more on the metaphysics of mechanism, an exploration of the epistemology of propensity, and potential implications for debates between causalists and statisticalists about natural selection. I also pointed to four further areas of empirical research to which my account has potential applicability: human neuroscience, foundations of physics, game/decision theory, and medicine. Far from an exhaustive list of future projects opened up by this dissertation, I take these merely to be a small sample of the myriad of ways this dissertation might be fruitful.

Bibliography

Abrams, M. 2007. "How Do Natural Selection and Random Drift Interact?" *Philosophy of Science* 74 (5): 666-679.

Abrams, M. 2012. "Mechanistic Probability," *Synthese* 187 (2): 343–375.

Anscombe, E. 1971. "Causality and Determination: An Inaugural Lecture," In J. Kim, D. Korman, and E. Sosa (eds.) *Metaphysics: An Anthology* (2nd ed.) Malden, MA: Wiley-Blackwell Press, 386-397.

Andersen, H. 2012. "The Case for Regularity in Mechanistic Causal Explanation," *Synthese* 189 (3): 415–432.

Ariew, A., Lewontin, R. C. 2004. "The Confusions of Fitness," *British Journal for the Philosophy of Science* 55 (2): 347–363.

Aron, J. C. 1979. "A Stochastic Basis for Microphysics," *Foundations of Physics* 9 (3-4): 163-191.

Aron, J. C. 1986. "A Stochastic Approach to the Hadron Spectrum," *Foundations of Physics* 16 (10): 1021-1060.

Armstrong, D. 1983. *What Is a Law of Nature?* Cambridge: Cambridge University Press.

Avrahami, J., Güth, W., and Kareev, Y. 2005. "Games of Competition in a Stochastic Environment," *Theory and Decision* 59 (4): 255-294.

Barros, B. 2008. "Natural Selection as a Mechanism," *Philosophy of Science* 75 (3): 306-322.

Beatty, J. 1980. "Optimal-Design Models and the Strategy of Model Building in Evolutionary Biology," *Philosophy of Science* 47 (4): 532-561.

Beatty, J. 1995. "The Evolutionary Contingency Thesis," In G. Wolters and J. Lennox (eds.) *Concepts, Theories, and Rationality in the Biological Sciences*, Second Pittsburgh-Konstanz Colloquium in the Philosophy of Science. Pittsburgh, PA: University of Pittsburgh Press, 45-81.

Beatty, J. H.; Finsen, S. 1989. "Rethinking the Propensity Interpretation of Fitness: A Peek Inside Pandora's Box," In M. Ruse (ed.), *What The Philosophy of Biology is: Essays for David Hull*, Ann Arbor, MI: Kluwer Academic, 17–30.

Bechtel, W. 2006. *Discovering Cell Mechanisms: The Creation of Modern Cell Biology*. Cambridge: Cambridge University Press.

Bechtel, W. and Abrahamsen, A. 2005. "Explanation: A Mechanist Alternative", in C. F. Craver and L. Darden (eds.), Special Issue: "Mechanisms in Biology," *Studies in History and Philosophy of Biological and Biomedical Sciences* 36: 421-441.

Bechtel, W. and Levy, A. 2013. "Abstraction and the Organization of Mechanisms," *Philosophy of Science* 80 (2): 241-261.

Bechtel, W. and Richardson, R. C. 1993. *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*, Princeton, NJ: Princeton University Press.

Bechtel, W. and Richardson, R. C. 2010, *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*, (2nd. ed.) Cambridge, MA: MIT Press.

Beebee, H. 2000. "The Non-governing Conception of Laws of Nature," *Philosophy and Phenomenological Research* 61 (2): 571-594.

Beller, M. 1999. *Quantum Dialogue*, Chicago: University of Chicago Press.

Blake, R. R., Fox, R., and McIntyre, C. 1971. "Stochastic Properties of Stabilized-image Binocular Rivalry Alternations," *Journal of Experimental Psychology* 88 (3): 327.

Bogen, J. and Woodward, J. 1988. "Saving the Phenomena," *The Philosophical Review* 97 (3): 305-352.

Bogen, J. 2005. "Regularities and Causality; Generalizations and Causal Explanations," In Special Issue: "Mechanisms in Biology," In C. Craver and L. Darden (eds.) *Studies in History and Philosophy of Biological and Biomedical Sciences* 36 (2): 397-420.

Bolduc, J-S; and Cezilly, F. 2012. "Optimality Modeling in the Real World," *Biology and Philosophy* 27 (6): 851-869.

Boyle, R. 2000. *The Works of Robert Boyle*, Hunter, M., and Davis, E. B. (eds.), 14 vols., London: Pickering and Chatto.

Braithwaite, R. 1953. *Scientific Explanation*, Cambridge: Cambridge University Press.

Brandon, R. N. 1978. "Adaptation and Evolutionary Theory," *Studies in History and Philosophy of Science* 9 (3): 181–206.

Brandon, R. N. 2005. "The Difference between Selection and Drift: A reply to Millstein," *Biology and Philosophy* 20 (1): 153–170.

Brandon, R. N., Carson, S. 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.

Bromberger, S. 1966. “Why Questions,” in R. G. Colodney (ed.) *Mind and Cosmos*, Pittsburgh, PA: University of Pittsburgh Press.

Butterfield, J. 2007. “Stochastic Einstein Locality Revisited,” *British Journal for the Philosophy of Science* 58 (4): 805-867.

Capp, J-P. 2011. “Stochastic Gene Expression is the Driving Force of Cancer,” *Bioessays* 33 (10): 781-782.

Capp, J-P. 2012. “Stochastic Gene Expression Stabilization as a New Therapeutic Strategy for Cancer,” *Bioessays* 34 (3): 170-173.

Cartwright, N. 1983. *How the Laws of Physics Lie*, Oxford: Oxford University Press.

Cartwright, N. 2009. Presidential Address: American Philosophical Association Pacific Division. Vancouver, APA. 82.

Cohen, J. and Callender, C. 2009. “A Better Best System Account of Lawhood,” *Philosophical Studies* 145 (1): 1-34.

Cottingham, J., Stoothoff, R., Murdoch, D., Kenny, A. 1985, 1984, and 1991. *The Philosophical Writings of Descartes, Volumers, 1, 2, and 3*, Cambridge University Press.

Craver, C. F. 2007. *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*, New York: Oxford University Press.

Craver, C. F. and Darden, L. 2013. *In Search of Mechanisms: Discoveries across the Life Sciences.* Chicago, IL: University of Chicago Press.

Darden, L. 2008. "Thinking Again About Mechanisms," *Philosophy of Science* 75 (5): 958-969.

Darden, L. and Cain, J. 1989. “Selection Type Theories,” *Philosophy of Science* 56 (1): 106-129.

Darwin, C. 1859. *On the Origin of Species*. (Facsimile of the First Ed). Cambridge, MA: Harvard University Press.

de Finetti, B. 1937. “La Prévision: Ses Lois Logiques, Ses Sources Subjectives”, *Annales de l'Institut Henri Poincaré*, 7: 1–68; translated as “Foresight. Its Logical Laws, Its Subjective Sources,” in *Studies in Subjective Probability*, H. E. Kyburg, Jr. and H. E. Smokler (eds.), Robert E. Krieger Publishing Company, 1980.

DesAutels, L. 2011. “Against Regular and Irregular Characterizations of Mechanisms,” *Philosophy of Science* 78 (5): 914-925.

Dobzhansky, T. 1973. "Nothing in Biology Makes Sense Except in the Light of Evolution," *American Biology Teacher* vol. 35 (March 1973) reprinted in *Evolution versus Creationism*,

Eagle, A. 2004. “Twenty-one Arguments against Propensity Analyses of Probability,” *Erkenntnis* 60 (3): 371-416.

Earman, J., Roberts, J. 2002. “Ceteris Paribus, Lost,” *Erkenntnis* 57 (3): 281-301.

Elgin, M. 2006. “There May be Strict Empirical Laws in Biology, After All.” *Biology and Philosophy* 21 (1): 119-134.

Falmagne, J.-C., Doignon, J.-P. 1997. “Stochastic Evolution of Rationality,” *Theory and Decision* 43 (2): 107-138.

Fetzer, J. H. 1981. “Probability and Explanation,” *Synthese* 48 (3): 371-408.

Fodor, J. 1991. “You can Fool Some of the People all of the Time, Everything Else Being Equal; Hedged Laws and Psychological Explanations,” *Mind* 100 (397): 19-34.

Forber, P., and Reisman, K. 2007. “Can There be Stochastic Evolutionary Causes?” *Philosophy of Science* 74 (5): 616-627.

Garber, D. 2013. “Remarks on the Pre-history of the Mechanical Philosophy,” In D. Garber and S. Roux (eds.) *The Mechanization of Natural Philosophy*, Boston, MA: Springer: 3-27.

Gardiner, P. 1959. *The Nature of Historical Explanation*, Oxford: Oxford University Press.

Gillies, D. 2000. “Varieties of Propensity,” *British Journal for the Philosophy of Science* 51 (4): 807-835.

Glennan, S. 1996. "Mechanisms and the Nature of Causation," *Erkenntnis* 44 (1): 49-71.

Glennan, S. 2002. "Rethinking Mechanistic Explanation," *Philosophy of Science* 69 (3): S342-S353.

Glennan, S. 2009. “Productivity, Relevance, and Natural Selection,” *Biology and Philosophy* 24 (3): 325-339.

Glennan, S. 2010. “Ephemeral Mechanisms and Historical Explanation,” *Erkenntnis* 72 (2): 251-266.

Glymour, B. 2001. “Selection, Indeterminism, and Evolutionary Theory,” *Philosophy of Science* 68 (4): 140-157.

Goodman, N. 1947. “The Problem of Counterfactual Conditionals,” *Journal of Philosophy* 44

(5): 113–128.

Graves, L., Horan, B., and Rosenberg, A. 1999. "Is Indeterminism the Source of the Statistical Character of Evolutionary Theory?" *Philosophy of Science* 66 (1): 140 – 157.

Harvey, W. 1628. *Exercitatio Anatomica de Motu Cordis et Sanguinis in Animalibus*. Frankfurt. (Translated 1889) as *On the Motion of the Heart and Blood in Animals.* London: George Bell and Sons.

Hajek, A. 2009. "Fifteen Arguments against Hypothetical Frequentism," *Erkenntnis* 70 (2): 211-235.

Hajek, A. 2011. "Interpretations of Probability," *The Stanford Encyclopedia of Philosophy* (Winter 2014 Ed.), E. N. Zalta (ed.), URL = <http://plato.stanford.edu/entries/probability-interpret/

Hausman, D. 1992. *The Inexact and Separate Science of Economics*, Cambridge University Press, Cambridge.

Hebb, D. O. 1949. *The Organization of Behavior.* New York: Wiley.

Hempel, C. 1942. "The Function of General Laws in History," *Journal of Philosophy* 39 (2): 35–48; reprinted in Hempel 1965, 231–244.

Hempel, C. 1965. *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*, New York: Free Press.

Heufer, J. 2011. "Stochastic Revealed Preference and Rationalizability," *Theory and Decision* 71 (4): 575-592.

Hoefer, C. 1997. "On Lewis's Objective Chance: 'Humean Supervenience Debugged'," *Mind* 106 (422): 322 – 334.

Hoefer, C. 2007. "The Third Way on Objective Probability: A Sceptic's Guide to Objective Chance," *Mind* 116 (463): 549-596.

Horan, B. L. 1994. "The Statistical Character of Evolutionary Theory," *Philosophy of Science* 61 (1): 76–95.

Hume, D. A., 1739. *A Treatise of Human Nature*, in L.A. Selby-Bigge (ed.), London: Oxford University Press, 1888, (reprinted 1960).

Illari, P. M., and Williamson, J. 2010. "Mechanisms are Real and Local," In P. M. Illari, F. Russo, and J.Williamson (eds.) *Causality in the sciences*. Oxford: Oxford University Press: 818-844.

Jackson, F., and Pettit, P. 1990. "Program Explanation: A General Perspective," *Analysis* 50 (2): 107-117.

Jacob, F. and Monod, J. 1961. "Genetic Regulatory Mechanisms in the Synthesis of Proteins," *Journal of Molecular Biology* 3: 318-356.

Kandel, E. R., Schwartz, J. H., Jessell, T. M., Siegelbaum, S. A., Hudspeth, A. J. 2013. *Principles of Neural Science*, (5th ed.) McGraw-Hill Inc.

Khrennikov, A. 2011. "Prequantum Classical Statistical Field Theory: Schrödinger Dynamics of Entangled Systems as a Classical Stochastic Process," *Foundations of Physics*. 41 (3): 317-329.

Kondrashov, A. S. and Crow, J. F. 1993. "A Molecular Approach to Estimating Human Deleterious Mutation Rate," *Human Mutation* 2: 229-234.

Laplace, P. 1814. *Essai Philosophique sur les Probabilités* forming the introduction to his *Théorie Analytique des Probabilités*, Paris: V Courcier; repr. F.W. Truscott and F.L. Emory (trans.), *A Philosophical Essay on Probabilities*, New York: Dover, 1951.

Lange, M. 2000. *Natural Laws in Scientific Practice*, Oxford University Press, New York.

Lange, M. 2002. "Who's Afraid of Ceteris Paribus Laws: or: How I Learned to Stop Worrying and Love Them," *Erkenntnis* 57 (3): 407-423.

Lange, M. 2013. "Really Statistical Explanations and Genetic Drift," *Philosophy of Science* 80 (2): 169-188.

Lewis, D. 1973a. *Counterfactuals*. Boston, Harvard University Press.

Lewis, D. 1973b. "Causation," *Journal of Philosophy* 70 (17): 556-567.

Lewis, D. 1980. "A Subjectivist's Guide to Objective Chance," In R. C. Jeffrey (ed.) *Studies in Inductive Logic and Probability*. University of California Press. 83--132.

Lewis, D. 1986. "A Subjectivists Guide to Chance," *Philosophical Papers* 83–132.

List, C. and Pivato, M. 2013. "Emergent Chance," Forthcoming

Lyon, A. 2011. "Determinisitc Probability: Neither Chance nor Credence," *Synthese* 182 (3): 413-432.

Machamer, P., Darden, L., and Craver, C. 2000. "Thinking About Mechanisms," *Philosophy of Science* 67 (1): 1-25.

Machamer, P. 2004. "Activities and Causation: The Metaphysics and Epistemology of Mechanisms," *International Studies in the Philosophy of Science* 18 (1): 27-39.

Matthen, M., and Ariew, A. 2002. "Two Ways of Thinking about Fitness and Natural Selection," *The Journal of Philosophy* 99 (2): 55–83.

Matthen, M., and Ariew, A. 2009. "Selection and Causation," *Philosophy of Science* 76 (2): 201-224.

Mckitrick, J. 2005. "Are Dispositions Causally Relevant?" *Synthese* 144 (3): 357 - 371.

Mendel, G. 1865. "Experiments on Plant Hybrids," in C. Stern, and E. Sherwood (eds.) (1966), *The Origin of Genetics, A Mendel Source Book,* San Francisco: W. H. Freeman, 1-48.

Mills, S. K., and Beatty, J. H. 1979. "The Propensity Interpretation of Fitness," *Philosophy of Science* 46 (2): 263–286.

Millstein, R. L. 1996. "Random Drift and the Omniscient Viewpoint," *Philosophy of Science* 63 (3): S10-S18.

Millstein, R. L. 2000, "Chance and Macroevolution," *Philosophy of Science* 67 (4): 603-624

Millstein, R. L. 2002. "Are Random Drift and Natural Selection Conceptually Distinct?" *Biology and Philosophy* 17 (1): 33–53.

Millstein, R. L. 2003a, "How Not to Argue for the Indeterminism of Evolution: A Look at Two Recent Attempts to Settle the Issue," In A. Huttemann (ed.) *Determinism in Physics and Biology*. Mentis: 91-107.

Millstein, R. L. 2003a, "Interpretations of Probability in Evolutionary Theory," *Philosophy of Science* 70 (5): 1317-1328.

Millstein, R. L. 2006. "Natural Selection as a Population-level Causal Process," *The British Journal for the Philosophy of Science* 57 (4): 627–653.

Mitchell, S. 1997. "Pragmatic Laws," *Philosophy of Science* 64 (4): 468-479.

Mitchell, S. 2000. "Dimensions of Scientific Law," *Philosophy of Science* 67 (2), pp. 242-265.

Mitchell, S. 2002. "Biological Contingency and Laws," *Erkenntnis* 57 (3): 329-350.

Mitchell, S. 2003. *Biological Complexity and Integrative Pluralism.* Cambridge University Press, New York.

Morris, R. G. M. 1981. "Spatial Localization Does Not Require the Presence of Local Cues," *Learning and Motivation* 2 (2): 239–260.

Nagel, E. 1961. *The Structure of Science: Problems in the Logic of Scientific Explanation*, New

York: Harcourt, Brace and World.

Pence, C., and Ramsey, G. 2013. "A New Foundation for the Propensity Interpretation of Fitness," *British Journal for the Philosophy of Science*. 64 (4): 851-881.

Peirce, C. S. 1910. "Note on the Doctrine of Chances," in C. Harteshorne and P. Weiss (eds.) *Collected Papers of Charles S. Peirce*, Volume II, Cambridge, MA. Harvard University Press: 408.

Pietroski, P., and Rey, G. 1995. "When Other Things Aren't Equal: Saving Ceteris Paribus Laws from Vacuity," *British Journal for the Philosophy of Science*. 46 (1): 81-110.

Popper, K. R. 1957. "The Propensity Interpretation of the Calculus of Probability and the Quantum Theory," in S. Körner (ed.), *The Colston Papers* 9: 65–70.

Popper, K. R. 1959. *The Logic of Scientific Discovery*, translation of *Logik der Forschung*, London: Hutchinson.

Potochnik, A. 2009. "Optimality Modeling in a Suboptimal World," *Biology and Philosophy* 24 (2): 183-197.

Prugovečki, E. 1984. "The Stochastic Quantum Mechanics Approach to the Unification of Relativity and Quantum Theory," *Foundations of Physics*. 14 (12): 1147-1162.

Putnam, H. 1975. "Philosophy and Our Mental Life," In *Mind, Language, and Reality.* Cambridge: Cambridge University Press, 291-303

Ramsey, G. 2006. "Block Fitness," *Studies in History and Philosophy of Biological and Biomedical Sciences* 37 (3): 484-498.

Ramsey, G. 2012. "Driftability," *Synthese* 190 (17): 3909-3928.

Reichenbach, H. 1949. *The Theory of Probability*, Berkeley: University of California Press.

Richardson, R. C. 2003. "Adaptationism, Adaptation, and Optimality," *Biology and Philosophy* 18 (5): 695-713.

Roumani, D.. and Moutoussis, K. 2012. "Binocular Rivalry Alternations and their Relation to Visual Adaptation," *Frontiers in Human Neuroscience* 6.

Rosenberg, A. 1994, *Instrumental Biology or the Disunity of Science*, Chicago: University of Chicago Press.

Rosenberg, A. 2001, "Discussion Note: Indeterminism, Probability, and Randomness in Evolutionary Theory," *Philosophy of Science* 68 (4): 536–544.

Rosenberg, A., and Bouchard, F. 2008. "Fitness," *The Stanford Encyclopedia of Philosophy* (Winter 2012 ed,), Edward N. Zalta (ed.), URL = http://plato.stanford.edu/archives/win2012/entries/davidson/

Salmon, N. 1988. "Review of *On the Plurality of Worlds*," *Philosophical Review* 97 (2): 237-244.

Salmon, W. 1984. *Scientific Explanation and the Causal Structure of the World*, Princeton: Princeton University Press.

Salmon, W. 1989, *Four Decades of Scientific Explanation*, Minneapolis: University of Minnesota Press.

Schiffer, S. 1991. "Ceteris Paribus Laws," *Mind* 100 (397): 1-17.

Sharratt, M. 1994. *Galileo: Decisive Innovator*. Cambridge: Cambridge University Press.

Skipper, R. A. Jr. 1999. "Selection and the Extent of Explanatory Unification," *Philosophy of Science* 66 (3): S196-S209.

Skipper, R. A., Jr. and R. L. Millstein. 2005. "Thinking about Evolutionary Mechanisms: Natural Selection," In Special Issue: "Mechanisms in Biology," in C. F. Craver and L. Darden (eds.) *Studies in History and Philosophy of Biological and Biomedical Sciences* 36: 327-347.

Smart, J. J. C. 1963. *Philosophy and Scientific Realism.* Routledge and Kegan Paul, London.

Sober, E. 1984. *The Nature of Selection*, London: MIT Press.

Sober, E. 1990. "Explanation in Biology: Let's Razor Ockham's Razor," *Royal Institute of Philosophy Supplement* 27: 73-93.

Sober, E. 1997. "Two Outbreaks of Lawlessness in Recent Philosophy of Biology," *Philosophy of Science* 64 (3), pp. S458-S467.

Sober, E., 2000. *Philosophy of Biology*, (2nd ed.) Boulder, CO: Westview Press.

Sober, E. 2001. "The Two Faces of Fitness,". In E. Sober (ed.), *Conceptual Issues in Evolutionary Biology*, (2nd ed.) London: MIT Press. 25–38.

Sober, E. 2010. "Evolutionary Theory and the Reality of Macro-probabilities," In E. Ellery & J. H. Fetzered (eds.) *The Place of Probability in Science*, New York: Springer. 133–161.

Sopher, B., and Narramore, C., 2000. "Stochastic Choice and Consistency in Decision Making Under Risk: An Experimental Study," *Theory and Decision*. 48 (4):323-349.

Stamos, D. 2001. "Quantum Indeterminism and Evolutionary Biology," *Philosophy of Science* 68 (2): 164–184.

Strevens, M. 2011. "Probability out of Determinism," In *Probabilities in Physics*, C. Beisbart and S. Hartmann (eds.), Oxford University Press: 339-364.

Suárez, M. 2007. "Quantum Propensities," *Studies in History and Philosophy of Science*. Part B: 38 (2): 418-438.

Suárez, M. 2013. "Propensities and Pragmatism," *Journal of Philosophy*. 110 (2): 61-02.

Suarez, M. (ed.) (forthcoming). *Probabilities, Causes and Propensities in Physics*. Springer.

Tong, F. 2001. "Competing Theories of Binocular Rivalry: A Possible Resolution," *Brain and Mind* 2 (1): 55-83.

Urte, R. 2011. "A difference in the Brain's Ability to Automatically Encode Serially Deterministic and Stochastic Regularities," *Frontiers in Human Neuroscience* 5.

Velasco, J. D. 2012. "Objective and Subjective Probability in Gene Expression," *Progress in Biophysics and Molecular Biology*. 110: 5-10.

Venn, J., 1876, *The Logic of Chance*, (2nd ed.), London: Macmillan; reprinted, New York: Chelsea Publishing Co., 1962.

Vicente, A. 2002. "How Dispositions Can be Causally Relevant," *Erkenntnis* 56 (3): 329-344.

von Mises R. 1957. *Probability, Statistics and Truth*, revised English ed., New York: Macmillan.

von Plato, J. 1982. "Probability and Determinism," *Philosophy of Science* 49 (1): 51-66.

Walsh, D., Lewens, T., and Ariew, A. 2002. "The Trials of Natural Selection and Drift," *Philosophy of Science* 69 (3): 452-473.

Watson, J. 1965. *Molecular Biology of the Gene.* New York: W. A. Benjamin.

Weber, M. 2001. "Determinism, Realism, and Probability in Evolutionary Theory," *Philosophy of Science*. 68 (3): S213-S222.

Weber, M, 2005. "Indeterminism in Neurobiology," *Philosophy of Science* 72 (5): 663-674.

Westfall, R. 1971. *The Construction of Modern Science: Mechanisms and Mechanics*, New York, NY: Cambridge University Press

Whewell, W. 1837. *History of the Inductive Sciences from the Earliest to the Present Times, Vol*

2. London: J. W. Parker.

Wimsatt, W. C. 1976. "Reductionism, Levels of Organization, and the Mind-Body Problem," In G. Globus, G. Maxwell, and I. Savodnik, (eds.) *Brain and Consciousness: Scientific and Philosophic Strategies*, New York: Plenum Press: 199-267.

Woodward, J. 2000, "Explanation and Invariance in the Special Sciences," *British Journal for the Philosophy of Science*, 51 (2): 197-254.

Woodward, J. 2001. "Law and Explanation in Biology: Invariance is the Kind of Stability that Matters," *British Journal for the Philosophy of Science*. 68 (1): 1-20.

Woodward, J. 2002. "There is No Such Thing as a Ceteris Paribus Law," *Erkenntnis* 57 (3): 303-328.

Woodward, J. 2003. *Making Things Happen: a Theory of Causal Explanation*. New York: Oxford University Press.

Woodward, J. 2009. "Scientific Explanation," *The Stanford Encyclopedia of Philosophy* (Winter 2014 Ed.), E. N. Zalta (ed.), URL = <http://plato.stanford.edu/entries/scientific-explanation/>