Toward a Propensity Interpretation of Stochastic Mechanism for the Life Sciences

Lane DesAutels

§1. Introduction

The life sciences are rife with probabilistic generalizations[1]. Mendel (1865) discovered that the chance of a hybrid green and yellow pea plant to produce yellow peas in 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 (Kandel et al. 2013, 271). 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 (Sober 2010). 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 what follows, I suggest that it makes good sense to think of the truth of (at least some of) the probabilistic generalizations made in the life sciences as metaphysically grounded in biological mechanisms in the world. These biological mechanisms underlie and produce the observable phenomena, and these biological mechanisms are themselves—in some sense—chancy. I call them *stochastic mechanisms*[2]. But how should we understand such stochastic mechanisms?

To begin to answer this question, I formulate two desiderata that any adequate account of stochastic mechanism should 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 meets the proposed desiderata, but some version of propensity theory can. Because I will not be able to consider every possible interpretation of chance, the ensuing arguments should not be taken to be strict, deductive, arguments for a propensity-backed

[1] Following Sober (2010), I don't take 'probabilistic' here to be incompatible with determinism. Rather, I mean it to encapsulate generalizations that are both fundamentally probabilistic (i.e., those that emerge out of genuine indeterministic processes) and those that are statistical (i.e., those that emerge out of deterministic processes).
[2] This term finds its origins with Jeffrey's use of 'stochastic process' (1969)—and later gets briefly mentioned in Salmon (1989). Stochastic mechanisms do not, however, receive any detailed discussion until Glennan (1992) and (1997)—though he offers no explicit analysis of which philosophical theory of chance to understand them with.

understanding of stochastic mechanism. Rather, I will take these arguments to suggest that a propensity-style approach to stochasticity, when compared to several other leading contenders, enjoys a few crucial advantages as a means of further understanding stochastic mechanisms and their explanatory uses in the life sciences. Having motivated a propensity-style understanding of stochastic mechanism, I then go on to draw a few important lessons from a recent propensity interpretation of biological fitness (PIF). From these lessons, I proceed 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 potential way of *quantifying* the stochasticity of a particular mechanism, and (3) a way of maintaining the causal relevance of propensities *without a problematic commitment to their causal efficacy*.

Here is my plan. In §2, I outline and explain two desiderata which I believe any adequate account of stochastic mechanism should meet and then go about showing how subjectivist, frequentist, and best-system-style analyses of stochastic mechanism fail to meet one (or both) of these desiderata. In §3, I argue that a propensity interpretation of chance can meet these desiderata. In §4, I discuss biological fitness, paying close attention to some of the key features of Grant Ramsey's recent propensity interpretation of fitness. In §5, I adduce three lessons from this recent propensity account of fitness: one about the causal role of propensities, one about their metaphysical base, and one about the way in which they can be quantified. And, in §6, I demonstrate how the PrISM might work by suggesting how it applies to the phenomena of initiation of electrical activity in postsynaptic neurons.

§2. Mechanism and Theories of Chance

For many centuries, the heartbeat was deeply mysterious. It wasn't until William Harvey (1628) described its role in the mechanism for blood circulation in animals that it was fully explained. 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.[3] (MDC 2000, 3)

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 vexing once it is acknowledged that many of the very phenomena which MDC was designed to explain behave probabilistically: synaptic transmission (a process that fails more often than it succeeds), protein synthesis and DNA replication (both of which are significantly error-prone), and natural selection (a process highly sensitive to environmental contingencies and which operates on spontaneous genetic mutations). If new mechanists are to appeal to mechanisms to explain these phenomena, it would help to have an account of stochastic mechanism[4] on hand.

2.1 Desiderata for an account of stochastic mechanism

As a way of taking steps toward such an account, I begin by briefly articulating and motivating a couple of central desiderata for an adequate analysis of stochastic mechanism:

> *DESCRIPTIVE ADEQUACY*: an account of stochastic mechanism must cohere with the general practice of biologists using mechanisms to explain natural phenomena.
>
> *CAUSAL EXPLANATION*: an account of stochastic mechanism must allow for descriptions of underlying mechanisms to feature in causal explanations of regularities seen in nature.

---

3 I cite MDC here because it is the most widely known. But other philosophical characterizations of mechanism have been offered by Glennan (1996) as well as Bechtel & Abrahamson (2005).

4 As one of my reviewers notes, a complete account of stochastic mechanism should also specify the locus of a particular mechanism's stochasticity: where among a mechanism's entities and activities the stochastic element emerges. I cannot here undertake this project. But recent work by Andersen (2012) helpfully taxonomizes the various ways in which mechanisms fail to behave regularly.

By way of briefly motivating these desiderata, I'll say a few words about each. *DESCRIPTIVE ADEQUACY* 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 that (philosophers have convincingly argued) applies to the mechanisms widely discussed, postulated, and studied in contemporary, empirically successful life science.

Regarding *CAUSAL EXPLANATION*, I follow Wesley Salmon (1984, 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, 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, it suffices to say 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 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.

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 both of them.[5]

2.2 Contra Subjectivism, Frequentism, and Best Systems Analyses of Chance

This section comprises arguments against various philosophical theories of chance[6] as ways of underpinning an account of stochastic mechanism. I cannot here consider every possible interpretation of chance; so what follows should not be taken to be a strict, deductive argument for a propensity-backed understanding of stochastic mechanism. Rather, I take the ensuing arguments to suggest that a propensity-style approach to stochasticity, when compared to several other leading contenders, enjoys several prima facie advantages as a means of further understanding stochastic mechanisms and their explanatory uses in the life sciences.

2.2.1 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.[7] 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

---

[5] There is a sense in which the argument strategy for this section mirrors the first chapter of Brandon's (1990) book, *Adaptation and Environment.* However, rather than natural selection, it is applied to mechanisms.

[6] Following Schaffer (2007), I define *Objective Chance* as*:* an understanding of probabilities that meets a set of commonly accepted platitudes regarding its relationship to several related concepts 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. *Probability*, on the other hand, I understand as merely a measure on a likelihood that an event will occur: where this likelihood may be understood either as a subjective degree of belief or as an objective chance.

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

has a 10% chance of firing[8] 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 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[9]:

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 DESCRIPTIVE ADEQUACY) 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: 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.

Further, it seems uncontroversial that 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.

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

[8] From the standpoint of the mechanisms literature, 'firing' can be defined as a general way of saying the mechanism has begun operation.

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

objective facts might simply be 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 would be appropriate for understanding them as probabilities of explanation. A detailed defense of this premise is beyond the purview of this paper. 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.

2.2.2 Frequentism

On a frequentist view of chance, the chance of a given event occurring is just the frequency of occurrences of that event relative to a relevant reference class. The initial formulation of this account of chance, given by Venn (1876), was an *actual frequentism* according to which 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. The problems with actual frequentism are many and well-known[10]. So I won't consider it as a viable candidate for bolstering an account of stochastic mechanism. Unlike actual frequentism, however, hypothetical frequentism (HF) still holds a fair amount of intuitive appeal. 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.[11] 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.

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

> P5. 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.
>
> P6. Give (P5), the chance of a given stochastic mechanism firing is grounded on a counterfactual.
>
> P7. Life scientists, however, appeal to the chanciness of underlying mechanisms to causally explain actual output frequencies.
>
> P8. 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.

[10] Cf, especially Hajek's famous "Fifteen Arguments Against Finite Frequentism" (1997)
[11] Some classic examples of hypothetical frequentists include Reichenbach 1949 and von Mises 1957.

P9. So, given (P5)-(P8), an HF account of stochastic mechanism fails to meet *DESCRIPTIVE ADEQUACY* and *CAUSAL EXPLANATION*.

C2. Therefore, a (HF) view of stochastic mechanism is not viable.

Premise (P5) is just the result of combining our understanding of mechanism with a HF theory of chance. Premises (P6)-(P9) need more defense.

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.[12] 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

[12] For a particularly forceful articulation of the relevance problem for counterfactual explanation, see Salmon (1989).

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.[13] 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…" (20). He adds that his account is set up to "…test or elucidate the content of particular causal and explanatory claims" (22). And "..the theory should enable us to make sense of widely accepted procedures for testing causal and explanatory claims" (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 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*-

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

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" (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.

2.2.3 Best-system Analysis (BSA)

Another candidate theory of chance, first put forth by David Lewis (1980), is called the Best-System Analysis (BSA).[14] 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

[14] Lewis first articulated a best-system analysis of laws (1973a) and later extended it to apply to chance (1980).

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.[15] I argue that, even still, the BSA theory of chance is not amenable to our notion of stochastic mechanism.

My argument is this:

> P10. On a BSA account of stochastic mechanism, the chance of a given mechanism firing 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.
>
> P11. However, given (P10), the chances we ascribe to mechanisms arise from merely systematizing particular facts and thereby cannot causally explain these particular facts (as *CAUSAL EXPLANATION* requires).
>
> C3. 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 (P11) 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 fair coins tossed have landed heads, then the system with the most simplicity, strength, and fit should ascribe a chance of .5 to a fair 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

[15] 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 BSA analysis which would comfortably cohere with what we want from an account of stochastic mechanism..

all" (Abrams 2012, 345). 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.[16]

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.

I take the above arguments to have shown that neither subjectivism, frequentism, nor a best-systems analysis of chance can give us what we want from an account of stochastic mechanism. Here, it is important to note, however, that I do not take these arguments to constitute a general refutation of these interpretations of chance. Nothing I say in this section precludes the fact that these analyses of chance may play useful roles in other areas of scientific discourse. I only wish to have shown that they do not fit readily with what we need from an account of stochastic mechanism for explanatory use in the life sciences.

[16] 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.

§3. Propensity and Mechanism

What I hope to have shown in the forgoing sections is that none of the theories of chance heretofore considered (subjectivism, frequentism, 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 section 2.1.

On a propensity theory, chance is a dispositional property or tendency of a given type of physical situation to produce certain outcomes over others.[17] So what might a propensity-backed account of mechanism look like? At first pass, 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.

Before exploring what this view might amount to in any detail, what I hope to show first is that this is the *only* type of account of those considered which meets (or at least doesn't fail) any of the desiderata set forth in 2.1.

Starting with *DESCRIPTIVE ADEQUACY*, 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[18] itself; the same goes for the release of electrical activity in post-synaptic neurons. This process fails up to 90% of the time, a

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

[18] It is worth noting that there has been some dispute as to whether natural selection should count as an MDC mechanism. Skipper & Millstein(2005) argue that it should not. Barros (2008) argues that it should. I need not take a stand on this debate here, however, because my only point is that a propensity interpretation of stochastic mechanism is capable of cohering with the actual practice of scientists searching for and describing mechanisms (whatever the scientists take these mechanisms to be).

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 'mechanism'. 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 HF account. A propensity-backed stochastic mechanism does not define the stochasticity of a given mechanism as a relative frequency of outputs. 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.

On the basis of having passed muster with regard to the desiderata set forth in 2.1, I thereby conclude that there is some good prima facie reason to accept the following propensity interpretation of stochastic mechanism (PrISM):

> PrISM: the propensity *Pr* of a stochastic mechanism to produce a given output is a dispositional property of that mechanism (given the instantiation of its start-up conditions) to produce that outcome.

Many questions now arise. *What* are these propensities? *Where* are they? What is their *metaphysical base*? How do we *discover* them? How do we *quantify* them? We're told that propensities are dispositional properties, but how should we understand *dispositional property* in this context?[19]

[19] See Eagle (2004) for a detailed summary of some of the main objections to propensity accounts. Since my aim is only to defend a local version of propensity theory (apt for achieving a better understanding of stochastic mechansims), I don't take it as necessary to fend off all of them.

Full-fledged answers to these questions are out of reach in the space remaining in this paper. But to begin to find answers to some of these questions, let us look at another concept in biology that has been given a well-worked-out propensity interpretation: fitness.

§4. 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[20]: 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[21]. 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? To get hints at what motivates this view, we can contrast it to another candidate understanding of fitness: *realized fitness*. On a realized fitness view, an organism's fitness is defined in terms of its actual number of offspring. One of the primary motivations for eschewing a realized fitness view in favor of a propensity view is to avoid, what has been called, the "tautology problem"[22]. If fitness is not conceived of as a probabilistic propensity, and is instead defined in terms of an organism's actual number of offspring, then fitness *cannot explain* these actual outcomes in any way that isn't tautologous and thereby vacuous. 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

[20] Following Sober (2010) and Ramsey (2012), I take it that this need not constitute a denial of metaphysical determinism.
[21] Cf., Brandon and Carson (1996) and Beatty and Finsen (1989)
[22] Mills and Beatty (1979), and Pence and Ramsey (2013)

because it was 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 (as well as many others), Mills and Beatty (1978) and Brandon (1978) offer probabilistic propensity definitions of fitness of the following sort[23]:

> *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).

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

[23] Taken from Rosenberg and Bouchard (2008)

> 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[24], 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). 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 may provide a working template for the propensity interpretation of stochastic mechanism.

[24] *Aka*: metaphysical or supervenience base

§5. Lessons from Fitness

At the end of §3, recall, there were many important questions about PrISM that were left unanswered. Metaphysically speaking, what is a mechanism's *Pr*? What does a mechanism's *Pr* have as its metaphysical base? Can a mechanism's *Pr* be quantified? If so, how? What, if any, is the causal role played by a mechanism's *Pr*?

5.1 Lessons from Ramsey's Block Fitness

There are three features of Ramsey's account that seem particularly helpful in further developing the PrISM offered at the end of §3.

> 3 Lessons from Ramsey's PIF:
>
> (1) In order for them to explain outcomes, propensities should be understood as playing some kind of causal role.
>
> (2) Propensities are aptly understood as having probability-weighted possibilia as their categorical base.
>
> (3) Propensities are quantifiable via a function of these probability-weighted possibilia.

An analysis of how (and whether) these lessons can be applied to the propensity interpretation of stochastic mechanism will be the subject of the remaining sections of §5. But before delving into the details, let us see what a basic Ramsey-style propensity interpretation of stochastic mechanism would look like.

5.2 A Ramsey-Style Propensity Interpretation of Stochastic Mechansim

If we take these lessons to heart, one obvious strategy we might undertake would be to merely extend a Ramsey understanding of propensity to our notion of stochastic mechanism. It might go something like this: let's call it the Ramsey-style propensity interpretation of stochastic mechanism (R-S PrISM). Following the template from Ramsey, we can describe the R-S 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 start-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 (F1, F2, F3,… F_n) will have a probability[25] 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 dispositional property that manifests in a probability distribution, the various values of which can be quantified by a function on the heterogeneity in the F_i.

On the basis of this general description, I characterize the R-S PrISM as this.

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

Just as Ramsey's propensity interpretations of fitness is endowed with objectivity and predictive/explanatory power, the PrISM allows descriptions of stochastic mechanisms to (A) explain the actual number of successful firings of a given mechanism, (B) ground predictions regarding the number of future successful firings to expect of a given mechanism, (C) ground a distinction between stochasticity as a property of a mechanism versus the actual outcomes that result from a mechanism firing, and (D) allow for a causal-explanatory role for propensities in explaining mechanistic outputs.

Now having seen what a basic Ramsey-style interpretation of stochastic mechanism might look like, let us ask if it is any good. The answer, as we'll see, is complicated.

5.3 Lesson One: On the Causal Role of Propensities

On Ramsey's view, in order to speak coherently, we must allow for fitness, conceived of as propensity, 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 §2. I resolve

[25] Here, it is important to note that I am not committed to any particular view about how these probabilities should be interpreted or where/how we get them. By advocating a propensity interpretation of stochastic mechanism, I do not, thereby, mean to endorse a general, one-size-fits-all propensity view of chance/probability. Indeed I can be a pluralist about metaphysical interpretations of chance/probability because all I'm doing is arguing that propensity is a useful notion *in certain explanatory contexts*: namely, ones where we seek mechanistic explanations for probabilistic phenomena in the life sciences. It may well be that other interpretations of chance/probability are useful in other contexts.

this inconsistency by appeal to a distinction made by Jackson and Pettit (1990) between causal efficacy and causal relevance.

5.3.1 When a Wine Glass Breaks in the Sink

There seem to be good reasons for defenders of propensity interpretations of fitness to want these propensities to be causally efficacious. If these propensities are conceived of as causally efficacious, we can coherently speak of an organism's reproductive outcome as having been caused by its fitness. We can say that this snail had more progeny *because* he was more fit.

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. 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… 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. 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 begin 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 what I've said so far amounts only to intuition-pumping and would be question-begging against any defender of the causal efficacy of dispositional properties. That said, there are good deductive arguments against the causal efficacy of dispositional properties. One such argument comes from Jackson and Petit (1990). According to Jackson and Pettit, a *causally efficacious property* with regard to an effect is "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 that dispositions like fragility are not causally efficacious.

5.3.2 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 §2. 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*: an 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.

5.3.3 Response: Distinguishing Causal Efficacy from Causal Relevance

In order to respond to this objection, we can look again to Jackson and Pettit (1990). On their view, we can 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 §2 despite not being causally efficacious. Specifically, I argue that, even though propensities are not themselves causally efficacious, they are nevertheless *causally relevant*.

To illustrate their notion of causal relevance absent causal efficacy, Jackson and Pettit appeal to the example 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 5.5.2. 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*.[26] Indeed, it is on the basis of these considerations that I believe we can follow the first lesson from Ramsey's PIF—albeit in a qualified manner: we *can* assign a causal role to the propensities instantiated by stochastic mechansims. It's precisely the role described by Jackson and Petit as causal relevance absent causal efficacy.

5.4 On the Prospects for Applying Lesson 2

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

---

[26] An 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 2.2.3) 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 2.2.3 after all.

> (L2) Propensities are aptly understood as having probability-weighted possibilia as their categorical base.

The reasons motivating the application of (L2) to stochastic mechanisms are directly analogous to Ramsey's own reasons for understanding fitness in this manner. Namely, (L2) 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 (L2), 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 §6. But first we must consider a potential objection to applying (L2) to the PrISM.

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

Didn't we argue in 2.2.2 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 (L2)? 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 the previous section, 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 essentially this: the counterfactuals appealed to by HF are not causally relevant in the same way the propensities are. Why? – because the counterfactual long-run frequencies appealed to by HF don't *make* anything be the case in the actual world. To see why, consider 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 ask yourself whether this HF counterfactual is causally relevant in the same way my word-processing program is. I submit that 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, consider 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.

5.4.2 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 (L2) to an account of stochastic mechanism. And I fear it is an even harder one to deal with.

Recall that (L2) states that propensities are aptly understood as having *probability-weighted* possibilia as their categorical base. It seems any follower of (L2) 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 (L2) 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, (L2) 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 (L2), 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 §2, 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 (L2).

5.4.3 An Alternate Route: Following Abrams

Suppose the forgoing arguments succeed in showing that the prospects for applying (L2) to our analysis of stochastic mechanisms are quite dim. Suppose we now find ourselves convinced that (L2) 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.

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 devices, 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. (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 illustrates this point with the following picture:

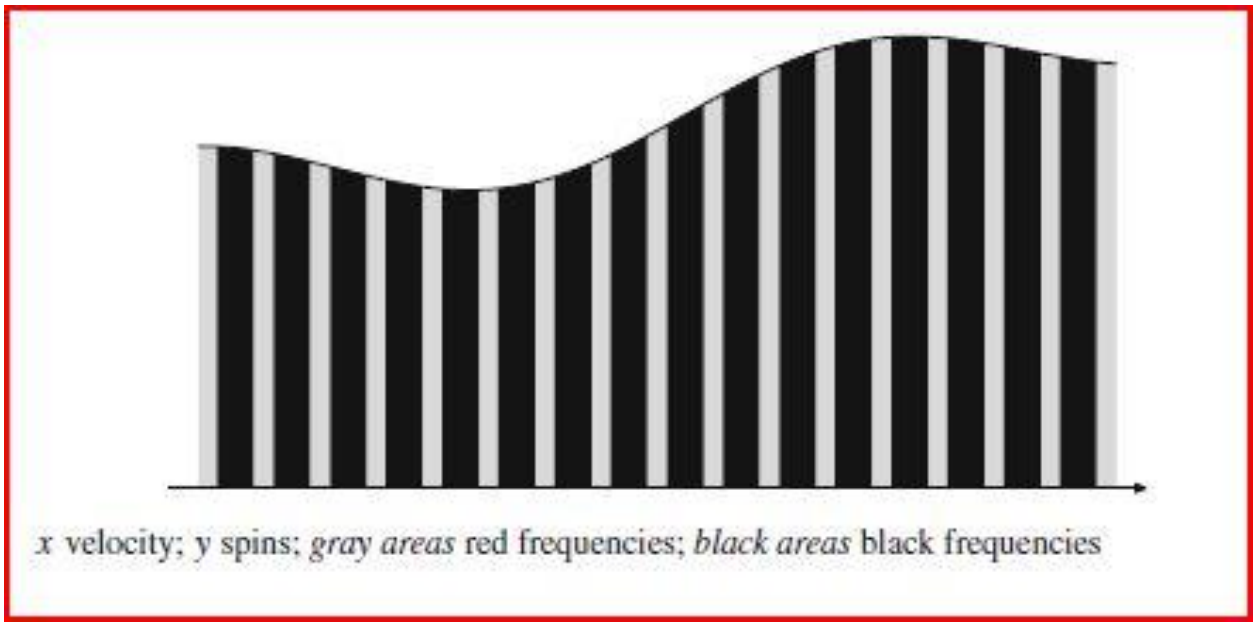

Figure 1. Roulette wheel output frequency distribution (from Abrams 2012, 350; reproduced by permission)

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 2).

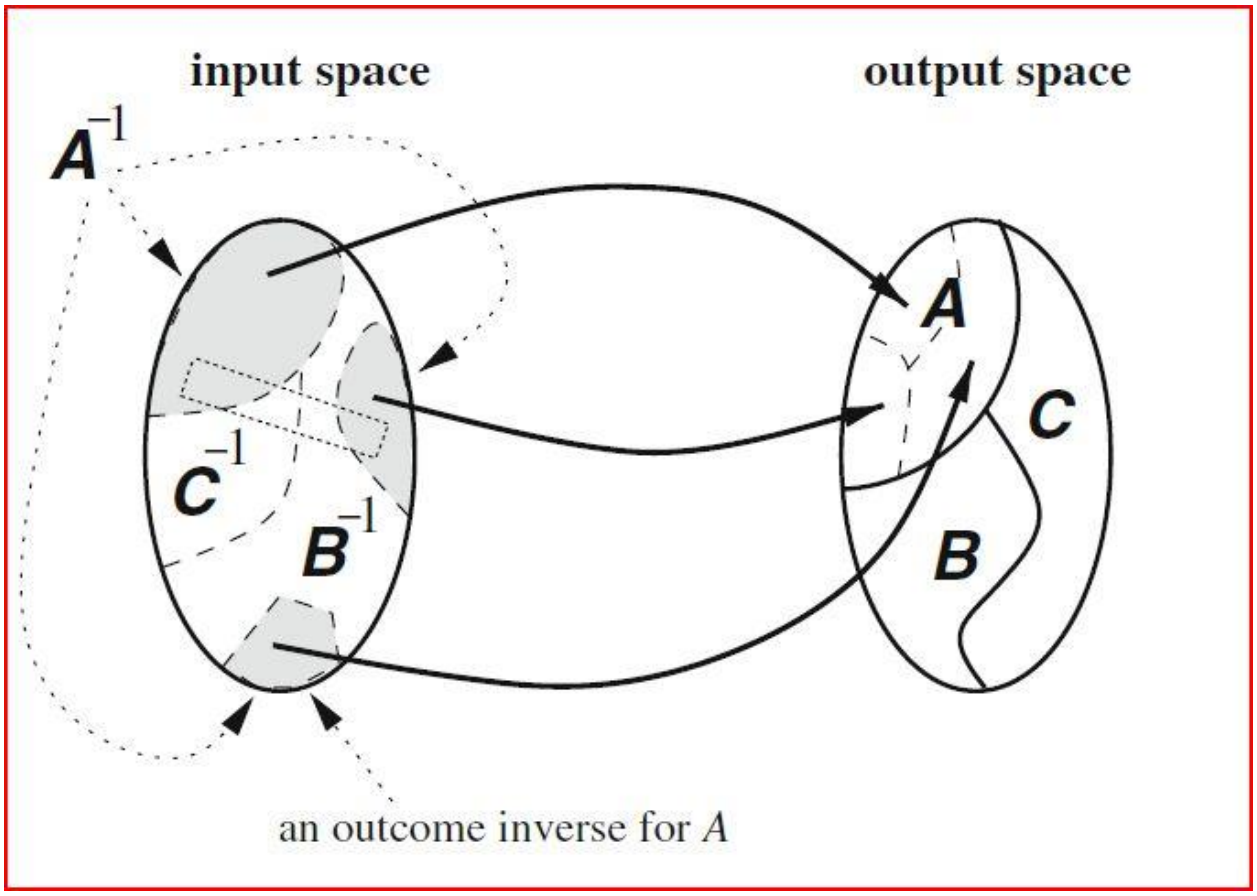

Figure 2. Causal map device (From Abrams 2012, 350; reproduced by permission).

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. 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. Let us call this the *Abrams-inspired propensity interpretation of stochastic mechanism* (A-I PrISM).

Whether or not we choose to follow Abrams in this regard, it is worth pointing to a general feature of his account shows an important advantage enjoyed by the PrISM over other applications of propensity (e.g., the PIF). On the PrISM, it matters less if sense can be made of the metaphysical nature of propensities. This is because much of the causal-explanatory work can be relegated to the *actual token mechanisms themselves*. On the PIF accounts, it is much less clear what specific processes underlie the propensities they postulate. In short, propensities together with mechanisms are symbiotic. Propensities help by offering an objective, causally relevant, and quantifiable understanding of a given mechanism's stochasticity. And mechanisms help propensities by providing ontic explanatory resources for explaining why the mechanism has the propensity it has.

Even if it succeeds at avoiding the regress attributed to the Ramsey approach, however, the Abrams-inspired approach comes at its own costs. Most importantly, 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.

So where is there left to go from here?

5.4.4 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 (L2) has the following real advantages and real costs:

> Ramsey-inspired (L2) Payoff: it offers a resource for further analyzing the metaphysics of propensities—thereby making them less mysterious
>
> Ramsey-inspired (L2) 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.

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 the concept of fitness 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 account, 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 account.

5.5 On the Prospects of Applying Lesson Three

Suppose we decide to stick with the Ramsey approach after all. There still remains the difficult task of exploring the prospect for applying (L3). Recall that (L3) 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, 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; f signifies some function; and p signifies some set of probability-weighted possibilities).

$$Pr(T) = f(p)$$

That is, whatever else we say about the prospects of applying (L3), 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 known 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 (L3) 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 (L3) to the PrISM is that stochastic mechanisms may not be 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 (L3) 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.

§6 Example: Synaptic Transmission

Now that we've explored the prospects for applying L1)-(L3) to our propensity interpretation of stochastic mechanism, let's see how the PrISM 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 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.[27]

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

[27] 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).

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 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.

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-inspired approach or the Abrams-inspired approach, there is more that we can say about this propensity than we could without the PrISM. Plus the analysis offered in this paper 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.

§7 Conclusion

In this paper, I have offered several arguments in favor of understanding the chanciness we ascribe to stochastic mechanisms as propensities. And I have drawn on a recent propensity interpretation of fitness 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 makes more explicit the causal role of propensities, one which allows for the stochasticity of a mechanism to an objective feature of the world, and one which endows descriptions of stochastic mechanism with explanatory and predictive power.

References:

Abrams, M. (2012). Mechanistic probability. *Synthese*, 187(2), 343–375.

Andersen, H. (2012). The case for regularity in mechanistic causal explanation. *Synthese*, 189(3), 415–432.

Anscombe, E. (1971). Causality and determination: an inaugural lecture. In J. Kim, D. Korman, & E. Sosa (eds.) *Metaphysics: An Anthology* (2nd ed.) Malden, MA: Wiley-Blackwell Press, 386-397.

Ariew, A., & Lewontin, R. C. (2004). The confusions of fitness. *British Journal for the Philosophy of Science*, 55(2), 347–363.

Barros, B. (2008). Natural selection as a mechanism. *Philosophy of Science*, 75(3), 306-322.

Beatty, J. (1995). The evolutionary contingency thesis. In G. Wolters & 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. & Abrahamsen, A. (2005). Explanation: a mechanist alternative. in C. F. Craver & L. Darden (eds.), Special Issue: "Mechanisms in Biology," *Studies in History and Philosophy of Biological and Biomedical Sciences* 36(2), 421-441.

Bogen, J. (2005). Regularities and causality; generalizations and causal explanations. In Special Issue: "Mechanisms in Biology," in C. F. Craver & L. Darden (eds.), Special Issue: "Mechanisms in Biology," *Studies in History and Philosophy of Biological and Biomedical Sciences* 36(2), 397-420.

Brandon, R. N. (1978). Adaptation and evolutionary theory. *Studies in History and Philosophy of Science*, 9(3), 181–206.

Brandon, R. N. (1990). *Adaptation and Environment*, Princeton, N.J.: Princeton University Press.

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.

Cartwright, N. (1983). *How the Laws of Physics Lie*, Oxford: Oxford University Press.

Craver, C. F. (2007). *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*,

New York: Oxford University Press.

Cohen, J. & Callender, C. (2009). A better best system account of lawhood. *Philosophical Studies*. 145(1), 1-34.

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.

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.

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.

Glennan, S. (1992). *Mechanisms, Models, and Causation*. Ph.D. Dissertation. Chicago, IL: University of Chicago.

Glennan, S. (1996). Mechanisms and the nature of causation. *Erkenntnis*, 44(1), 49-71.

Glennan, S. (1997). Probable causes and the distinction between subjective and objective chance. *Noûs*, 31(4), 496–519.

Hajek, A. (1997). "Mises redux"—redux: fifteen arguments against finite frequentism. *Erkenntnis*,.45(2/3), 209-227.

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.

Hoefer, C. (1997). "On Lewis's objective chance: 'humean supervenience debugged'. *Mind*, 106(422), 322 – 334.

Jackson, F., & Pettit, P. (1990). Program explanation: a general perspective. *Analysis*, 50(2), 107-117.

Jacob, F. & Monod, J. (1961). Genetic regulatory mechanisms in the synthesis of proteins. *Journal of Molecular Biology* 3, 318-356.

Jeffery, R. C. (1969). Statistical explanation vs. statistical inference. In N. Rescher (Ed.), *Essays in Honor of Carl G. Hempel*. (pp. 104-113). Dordrecht: D. Reidel Publishing Co.

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.

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.

Lyon, A. (2011). Determinisitc probability: neither chance nor credence. *Synthese*, 182(3), 413-432.

Machamer, P., Darden, L., & Craver, C. (2000). Thinking about mechanisms. *Philosophy of Science*, 67(1), 1-25.

Matthen, M., & Ariew, A. (2002). Two ways of thinking about fitness and natural selection. *The Journal of Philosophy*, 99(2), 55–83.

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., & Beatty, J. H. (1979). The propensity interpretation of fitness. *Philosophy of Science*. 46(2), 263–286.

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. (2003b), Interpretations of probability in evolutionary theory. *Philosophy of Science,* 70(5), 1317-1328.

Pence, C., & 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.

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.

Rosenberg, A., & 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, 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.

Schaffer, J. (2007). Deterministic chance? *British Journal for the Philosophy of Science*, 58(2), 113-140.

Skipper, R. A., Jr. &. Millstein, R. L. (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(2), 327-347.

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.

von Mises R. (1957). *Probability, Statistics and Truth*, revised English ed., New York: Macmillan.

Venn, J., (1876). *The Logic of Chance*, (2nd ed.), London: Macmillan; reprinted, New York: Chelsea Publishing Co., 1962.

Walsh, D., Lewens, T., & Ariew, A. (2002). The trials of natural selection and drift. *Philosophy of Science*, 69(3), 452-473.

Woodward, J. (2003). *Making Things Happen: a Theory of Causal Explanation*. New York: Oxford University Press.