Biol Philos (2017) 32:1205–1223
https://doi.org/10.1007/s10539-017-9601-3

Indeterminism in the brain

Bryce Gessell[1]

Received: 24 January 2017 / Accepted: 17 October 2017 / Published online: 27 October 2017

Abstract Does the brain behave indeterministically? I argue that accounting for ion channels, key functional units in the brain, requires indeterministic models. These models are probabilistic, so the brain does behave indeterministically in a weak sense. I explore the implications of this point for a stronger sense of indeterminism. Ultimately I argue that it is not possible, either empirically or through philosophical argument, to show that the brain is indeterministic in that stronger sense.

Keywords Indeterminism · Neuroscience · Neurobiology · Philosophy of neuroscience · Ion channels

Introduction

In this paper I argue that the brain is indeterministic, although only in a weak sense, and that nothing follows about more robust indeterminism.

What does it mean to say the brain is indeterministic? The question is surprisingly difficult to answer. Some theories link determinism to causation or predictability (van Strien 2014), but recent work avoids that connection. Philosophers usually study determinism in physics, and physics may not contain a notion like "cause" (Butterfield 2005; Hoefer 2016). A definition that doesn't commit us to causes begins with models: a model is deterministic if its state at one time fixes its

I wish to thank Robert Brandon, Felipe De Brigard, Paul Henne, Rafael Ventura, and Jing Hu for their comments on earlier drafts, as well as two anonymous reviewers, whose criticisms have greatly improved the article.

✉ Bryce Gessell
bgesslo7277@gmail.com

[1] Duke University, Durham, NC, USA

states at all times (Werndl 2009). To ask whether the brain is indeterministic, then, is first of all to ask whether some model of its behavior is indeterministic. Depending on the answer, we may then consider whether the brain itself, or one of its subsystems, is indeterministic. We therefore have two questions: one about the deterministic character of some model, and another about the real system the model represents.

In this paper I take up both questions. I start with a general discussion on determinism. I distinguish two senses in which a system may be indeterministic, corresponding to the two questions above. On one sense, the best models representing a system are indeterministic. On the other sense, the system itself is indeterministic. I explore the relationship between these two senses, and discuss how experimental and philosophical evidence might support them.

Next I turn to the brain. First I argue, contrary to Weber (2005), that the brain is indeterministic in the weak sense. That is, the best models of ion channel gating—a critical class of neural events—are indeterministic. I then evaluate three philosophical proposals for inferring the strong sense from the weak, and I conclude that none will work. Neither empirical arguments nor philosophical ones are capable of showing that the brain is strongly indeterministic. I close by considering some implications of this result for thinking about the brain. The main upshot is an end to worries about whether the brain is deterministic or not, since current methods are incapable of answering the question either way.

Determinism and two senses of indeterminism

In this section I discuss determinism and distinguish two senses of indeterminism for a given system. I examine the relationship between these senses and discuss how empirical evidence supports them. I imagine a hypothetical experiment which attempts to justify a move from weak to strong indeterminism. No matter the set-up, however, such an experiment cannot support strong indeterminism. The conclusion of this section is that we can only infer strong indeterminism with *philosophical* arguments, if we can infer it at all.

First we must understand the difference between the determinism of a model and that of an actual system. This difference parallels the weak and strong senses of indeterminism. A system is weakly indeterministic if the best models representing that system are indeterministic. A strongly indeterministic system is one that is itself indeterministic—its current states and properties, independent of our models, don't fix its past or future states.[1] Now a model of a system can be either deterministic or indeterministic; the system itself may be either deterministic or indeterministic as well. These properties are doubly dissociable. A model's being deterministic need not entail anything about the system it represents, nor vice-versa, and the same is true for indeterminism.

Defining determinism in terms of models and fixed states is a relatively recent approach. Departing from considerations of sufficient reasons, previous thinking

[1] Some authors, such as Butterfield (2005), speak of worlds rather than systems.

about determinism connected it to causation (van Strien 2014). For a phenomenon to be fully determined meant that it had reasons or causes sufficient to explain its existence. From this tradition we also get Laplace's demon (Laplace 1814), which joined determinism and predictability.

Others philosophers have discussed the reasons for abandoning these notions of determinism (Earman 1986), and I will not review them here. Suffice it to say that contemporary philosophers of science usually study determinism in terms of models (Werndl 2016) or theories (Butterfield 2005). Werndl gives an informal definition of determinism for a model: a model is deterministic if "the state of a model at one time fixes the state of the model at all times" (2016, 1). Similarly, Butterfield gives a formal definition for theories: "a theory is deterministic if, and only if: for any two of its models, if they have instantaneous slices that are isomorphic, then the corresponding final segments are also isomorphic" (2005).[2] The core insight is that a single state fixes all future and past states. These definitions do not require causation or epistemic properties like predictability.[3]

A model is not the same as the system it represents, and a model's deterministic properties are dissociable from those of its system. For an actual system to be indeterministic—that is, for that system to be strongly indeterministic—more is required than that the system's best model be indeterministic. We need an additional argument which goes beyond the model to show that the modeled system itself has indeterministic properties. In principle that argument could be empirical or philosophical, but regardless of its features, the inference to strong indeterminism requires it.

The distinction between senses of indeterminism respects the need for extra reasons which would bridge the gap between weak and strong. This gap between the indeterminism of a model and that of its system is an important consequence of Werndl's recent work (Werndl 2009, 2011, 2012, 2013). Building on results from Ornstein and Weiss (1991), Werndl showed that certain deterministic and indeterministic descriptions are observationally equivalent. In other words, in some cases, a deterministic model can replace an indeterministic one, and vice-versa, while still making the same predictions (Werndl 2011).

Suppose, for example, that one wonders whether some system with an indeterministic description is itself indeterministic. We might give the following argument: the indeterministic model represents reality; therefore, the represented system ought to have certain properties of the model. One property the model and the system share is indeterminism, and therefore the system itself is indeterministic as well. This is an argument for strong indeterminism on the basis of scientific

[2] Similar to Butterfield, Van Inwagen (2015) gives a definition in terms of alternative futures. Nothing depends here on whether determinism includes only future states or past states as well. Likewise, though Werndl (2013) restricts herself to talk of models, any difference between the determinism of theories and models is irrelevant in this paper.

[3] My discussion is also orthogonal to concerns about objective chances (Lewis 1980, 1994), or what Werndl (2016) calls "ontic probabilities." These are "probabilities that are real features of the world" (Werndl 2016, 14). Objective chances worried Lewis because, if they exist in a deterministic world, they seem to present a problem for the Principal Principle. That principle says that an agent should set their credence to the actual chance of an event, conditional on that agent's admissible information. See Frigg and Hoefer (2015).

realism.[4] The force behind it is a firm realist commitment, which underwrites the sharing of properties between the model and its system.

Werndl's results block this realist argument in some cases where a system has two observationally-equivalent descriptions. If a system has two such descriptions, one deterministic and the other indeterministic, the inference from weak to strong indeterminism is unwarranted. For we could just as well use the alternative description, along with the same commitment to realism, to infer that the system is strongly deterministic.

Not all systems are subject to Werndl's results, since some descriptions won't be observationally equivalent. In modeling empirical systems we will sometimes have reasons for preferring an indeterministic description over a deterministic one. When we do, we say that the system is weakly indeterministic. The reasons for our preference could be statistical superiority or fit with other evidence. But even in cases where we have justification for an indeterministic model, the gap between weak and strong indeterminism remains.

We could try to bridge the gap with an empirical argument, perhaps by doing experiments to address the question of indeterminism. What is interesting about such experiments is how clearly they reveal our intuitions about determinism. These intuitions are very powerful. To see them in action, consider the following example.

Suppose we have two models for some physical system, one deterministic and the other indeterministic. Our testing shows that the indeterministic model is more accurate, since it reproduces observational measurements more exactly. On that basis, we infer that our target system is weakly indeterministic. Now suppose we'd like to know more than that—suppose we want to know not just whether the best model is indeterministic, but whether the system itself is. In other words, we want to know whether the system is strongly indeterministic. How could we approach this question empirically?

We decide to return to our previous experiment. The first test of that experiment showed that the indeterministic model outperformed the deterministic one. But for this second test, we will do the experiment *twice*. We will make the greatest effort possible for everything to be the same between the two runs: the same set-up, external environment, conditions of measurement, and so on. Our plan is to run the experiment once and then run it again to see how the results compare. Note here that making conditions as identical as possible is all we can hope for, not just with an actual experiment but even with a hypothetical experiment. A hypothetical experiment in which we imagine the second set-up to be particle-for-particle identical to the first would be illegitimate, because the "outcome," or intuitive verdict, would beg the question for or against determinism. That is, our intuitive response to a case of particle-for-particle identical conditions exposes our a priori commitment to determinism or indeterminism. Rather than establishing that position, the intuition reveals it.

So after fixing the set-ups we run our two experiments. There are two possible outcomes: either the results between the two experiments differ, or they're the same.

[4] Brandon and Carson (1996) make this argument about evolutionary theory, and I review their approach below.

I will take these possibilities one at a time, focusing on how they relate to moving from weak to strong indeterminism for some target system.

The first outcome could be that the experimental results differ. We did the same experiment twice, and our measurements of the key variables are different between runs, despite our best efforts to make the set-ups the same. In this case, which of the following two conclusions is more plausible?

(A1) The target system is strongly indeterministic, since the results differed between runs.
(A2) Something about the experimental set-ups was different, which caused the results to differ.

Given the practical limitations on experimental conditions, (A2) is the more plausible conclusion. The number of possible influences on a given study is enormous, and we have no hope of controlling them all. At least some of the things we can't control will influence the system and our measurements. We could try a third run and be even more careful, but we would still face a choice between (A1) and (A2), and it's hard to see what could convince us that (A1) is a better conclusion.

Notice, however, that this way of thinking about (A1) and (A2) betrays a deep intuitive commitment to determinism. Our measurements differed (we think) *because* the set-ups differed; had they been *exactly* the same, we may feel, then the results would have been the same too. This intuitive commitment, born of experience with natural systems, is nevertheless difficult to budge through empirical evidence.

On the other hand, the outcome of our two experiments could be that the measurements are identical. At least four possible conclusions may interest us now:

(B1) Determinism is true, since we got the same results.
(B2) Indeterminism is true, and we just *happened* to get the same results.
(B3) If we measured more finely, we'd see that the results actually differ, and that therefore determinism is true.
(B4) If we measured more finely, we'd see that the results actually differ, and that therefore *indeterminism* is true, since the experimental set-ups were identical.

As before, each of (B1)–(B4) betrays a commitment regarding the very question at issue. All four conclusions are compatible with the evidence, so only one's metaphysical assumptions could make a difference between them. Deciding between (B1) and (B2) relies on one's a priori view of whether determinism is true, since both its truth and falsity could lead to the same results. On the other hand, the counterfactuals in (B3) and (B4) rely on that a priori view as well. Both possibilities could produce the observed results. The point, of course, is not about limits on what our intuitions can tell us. The point is about what it is possible to show about determinism through experiment. Even if we performed these experiments in real life, the justification for any conclusion seems to presuppose

that same conclusion. Experimentation simply does not provide the right kind of evidence.[5]

An experimental effort to argue for strong indeterminism is inadequate, because the clearest possible results support mutually incompatible conclusions. Thus the issue of determinism or indeterminism for an actual system does not seem to be decidable on the basis of experiment.[6] What we lack to cross the gap between weak and strong indeterminism is not more data but philosophical arguments.

In this section I have noted the differences between determinism and indeterminism, on the one hand, and determinism for models and determinism for real systems on the other. I distinguished between weak and strong senses of indeterminism. A system is weakly indeterministic if the best models of its behavior are indeterminsitic; strong indeterminism applies to systems which are indeterministic themselves. Since there is a gap between the weak and strong senses, we need additional arguments to infer the latter from the former. These arguments cannot essentially rely on empirical data. What we are looking for is a philosophical justification for strong indeterminism.

My remarks in this section concern physical systems generally, while the next two sections discuss the brain in particular. In "Ion channel gating is weakly indeterministic" section, I argue that the brain is indeterministic in the weak sense. Empirical evidence will be relevant to that argument. Then in "Philosophical arguments for strong indeterminism" section I return to the problem of the gap, taking up three philosophical approaches for going from weak to strong indeterminism.

Ion channel gating is weakly indeterministic

Brain indeterminism in the weak sense says that the best models for some brain process are indeterministic. I will argue that the brain is in fact weakly indeterministic. This is not a trivial point, as we shall see; a great deal of evidence bears on it. I begin by giving some background to ion channels. Then I argue that indeterministic models are the best ones we have for ion channel gating, a key class of neural events.

[5] Here we see the appeal of predictability and causation in explicating determinism. Physics may have expunged causes, but they permeate our own experience of the world; causes and prediction are the marks by which we would judge an indeterministic system. An indeterministic system could be one whose behavior is not predictable even in principle, and perhaps is it unpredictable because not all of the system's events are fully caused. These questions arise naturally when investigating deterministic properties. They are not likely to help us think about indeterminism, however, at least for the brain.

[6] I exclude problems of quantum measurement, which may be different; see Butterfield (2005). The Everett interpretation of quantum mechanics takes Schrödinger's equation to be deterministic, and the de Broglie–Bohm theory is explicitly deterministic. It is a separate question whether one of these interpretations may be preferable on the basis of experiment. In a sense the idea is already familiar in thinking about incompatible interpretations of quantum mechanics. What matters here is that we are in much the same position for other physical processes, at least when it comes to indeterminism.

Ion channel structure and function

Ion channels are proteins straddling cell membranes. A pore spans the middle of the channels, running from the extracellular matrix to the cell's interior. The pore allows charged particles to enter and exit, depending on factors like ambient voltage and concentration gradients. Voltage-gated sodium channels, for example, open in response to depolarizing currents which begin near the cell body. They allow sodium to cross the membrane, which further depolarizes the cell and creates more depolarizing currents. The result is an action potential, perhaps the most important means of neural information sharing. Other ion channels respond to other stimuli, like mechanical pressure or ligand binding.

The opening and closing of a channel is called "gating," and records of gating events are "dwell-time distributions." The goal of an ion channel model is to recreate these distributions. The best models are those which do so most accurately (Colquhoun and Hawkes 1995) and which seem most likely to correspond to actual molecular mechanisms (Craver 2014).

The classic approach to ion channels is the Markov model (Colquhoun and Hawkes 1981, 1982; Hille 1992; Colquhoun and Hawkes 1995). Markov models are paradigm cases of indeterministic, or stochastic, models (Werndl 2016). In a Markov model, channels move from open to closed states with a certain probability. Kinetic diagrams represent these transitions, like the following (over)simplified schema with one open and one closed state:

$$\text{Closed} \overset{\alpha}{\underset{\beta}{\rightleftharpoons}} \text{Open}$$

In this diagram α and β are "rate constants," or probabilities for state transitions. If α were .7, for example, a closed channel would have a .7 chance of opening before the next time step.

The critical feature of these models is the Markov property: the probability of changing states is independent of a system's history. Suppose a channel in the closed state has a .7 chance of opening in the next 1000 ms. If it hasn't opened 900 ms later, the probability of opening in the last 100 is still .7. In other words, the probability of a state change conditional on a system's past equals the unconditional probability of the same change. Markov systems are "memoryless" because the past doesn't affect the future. And since α and β are probabilities, Markov models are also indeterministic—the instantiation of one state fixes neither past nor future states.

Channel researchers didn't adopt Markov models arbitrarily. In fact, Hodgkin and Huxley's (HH) model of membrane conductance assumed Markovian or random channel gating (French and Horn 1983). When researchers succeeded in measuring individual ion channel currents for the first time (Neher and Sakmann 1976), they found that Markov models could accurately reproduce dwell-time distributions (Colquhoun and Hawkes 1981). The use of Markov models for channel behavior is fortuitous for claims of weak indeterminism about the brain. These models are indeterministic, and so if they are the best channel models, we could conclude that the brain is weakly indeterministic.

The story is not quite so simple, though, since there are also non-stochastic (deterministic) models of channel behavior (Weber 2005). Markov channel modeling arose in the 1980's, but during that time some proposed a new approach based on fractal mathematics (Liebovitch et al. 1987). Other non-Markov approaches followed (Millhauser et al. 1988). The first fractal models were indeterministic (Liebovitch et al. 1987), but the new approach culminated in a fully deterministic version (Liebovitch and Tóth 1990) and quickly spawned many others (Liebovitch and Czegledy 1991; Liebovitch and Toth 1991; Liebovitch 1993).[7] These fully-determined models mimic the behavior of random systems. They are chaotic, as their evolution over time is exquisitely sensitive to initial conditions. Since we only know those conditions with finite accuracy, we can't predict the future behavior of chaotic systems, even though each step in their evolution is fully determined by the one before it.[8]

In sum, we have two broad classes of ion channel models. The first are indeterministic Markov models, while the second are deterministic and based on fractal mathematics and chaos. Researchers have used both to explain channel behavior.

An argument for weak indeterminism for ion channels

Most models of channel behavior are indeterministic, but Weber (2005) says that deterministic models are just as good. In a deterministic model, fixing one state fixes all the states. Each step between states is fully determined by the one before it. Weber argues that, since both model types succeed in describing channel behavior, the choice between model types is underdetermined. In other words, we can't use them to support weak indeterminism for the brain.

Weber's claim of underdetermination is the same concern we saw earlier about observational equivalence. There are two ways to solve a problem of observational equivalence. The first is to argue directly that the underdetermination claim is false, by showing that one model type reproduces data distributions more accurately. The second way, discussed at length by Werndl (2012, 2013), is by indirect evidence. Indirect evidence refers to a model's coherence with other accepted facts or theories (Laudan and Leplin 1991; Moretti 2007). If two models are observationally equivalent, but one has more indirect support than the other, we have grounds for preferring the former.

Accordingly, my argument in the remainder of this section has two parts, following two different lines of evidence. Both strongly suggest that Markov models better capture gating behavior. The first part deals with distributions and answers the claim of underdetermination. The second shows that we have substantial indirect support for indeterministic models as well.

[7] See also Liebovitch and Czegledy (1992), Cavalcanti and Fontanazzi (1999) and Liebovitch and Krekora (2002).

[8] I have simplified things slightly—deterministic models are based on deterministic chaos, but they are not necessarily fractal. Though researchers often lump chaos and fractal models together (Liebovitch 1996, 170), "chaos does not imply fractal nor does fractal imply chaos" (Lowen and Teich 2005, 31).

We will begin with distributions. One direct comparison tested how well different model types could reproduce recordings from individual channels (McManus et al. 1988, 1989; Korn and Horn 1988; Millhauser et al. 1988; Sansom et al. 1989). This is the critical test for underdetermination, since it decides whether two models really are observationally equivalent. Using various criteria, researchers found that Markov models outperformed fractal models at describing many sequences of dwell-time distributions. For example, one test showed that Markov models gave accurate descriptions of a dataset for 72 of 72 distributions, while fractal models gave poor descriptions for 54 such distributions (McManus et al. 1988). They also found that Markov models were superior according to the likelihood ratio test and the asymptotic information criterion (Korn and Horn 1988; Akaike 1974). From the Markov model's superiority on these tests they concluded that there was no underdetermination among model types. Since then other investigators have taken their results as conclusive (Geng and Magleby 2015).

Hence it is not just that there are, as Werndl (2011) put it, "general differences between probability distributions of deterministic and stochastic descriptions." The deterministic models in these ion channel comparisons could not generate the correct distributions of data. These findings are the best evidence we can provide in favor of indeterministic models; for if two models don't produce equivalent distributions, they won't be observationally equivalent. Deterministic models may be able to mimic *some* properties of indeterministic models *some* of the time, but on balance the evidence shows that the indeterministic models are better at describing the data.

These results support the first part of my argument for weak indeterminism in the brain: indeterministic models recreate dwell-time distributions better than deterministic models.

The second part concerns indirect evidence. In reviewing this evidence for stochastic ion channel models, we look for coherence between the models and other aspects of related theories (Butterfield 2005; Werndl 2013).[9]. One way to put this is that the model has a "natural" interpretation: it involves "a realistic way of observing the system" (Werndl 2009). Many connections between a model and other points of theory could show the proper coherence and naturalness. I will review a few of them here.

We've seen that the Hodgkin–Huxley (HH) equations for membrance conductance assume random channel gating. Markov models are a natural way to describe the kinetic diagrams resulting from these equations (Lipscombe and Toro 2014). The HH view also assumes that a system has a finite number of discrete states, another feature of the Markov approach (Varanda et al. 2000; Austin 2008). After HH, researchers began to model channel conductance on classical chemical kinetics (Fitzhugh 1965; Qin 2014). The law of mass action, which is that the rate constant of a chemical reaction is proportional to the product of the chemical concentrations in the reaction (Colquhoun and Hawkes 1995), applies to such kinetics. This law

[9] Werndl's criterion is that *there is indirect evidence when the models (predictions all confirmed) are unified by a well confirmed theory and by confirmed similar additional assumptions about physical systems* (Werndl 2013, 2258, emphasis in original).

implies that a system is memoryless, which is just the Markov property itself (Colquhoun and Hawkes 1977). There are other connections as well between random gating, the HH model, and other aspects of chemical kinetics (Eyring 1935; Jones 2006; Austin 2008). These connections—to the HH model of membrane conductance and the foundations of chemical kinetics—are strong pieces of indirect evidence for indeterministic channel models. Deterministic gating approaches do not have such connections, making their relationship to other aspects of molecular and chemical theory unclear.

A second set of findings concerns the models' physical interpretations. Memoryless systems like Markov models produce distributions following a negative exponential function (Horn and Vandenberg 1984; Qin and Li 2004). This function is "more intimately associated with *physical realities*" than other possible model choices (Horrigan 2015, 86, emphasis in original). The reason is protein structure. Ion channels are proteins, and proteins are assemblies of amino acids which change their shape and orientation in order to participate in different reactions. Different shape and orientation combinations are called *conformational* states. Changes between conformational states occur in response to external stimuli, like variations in surrounding energy patterns.

Gating is a conformational change of the protein, since its parts act to allow or block ions through the pore. The Markov model assumes that each model state corresponds to a different conformational state, and that transitions between model states represent possible conformation transitions across local energy barriers.[10] A Markov model with five states, for example, attributes five conformational states to the channel protein, with the rate constants representing chances of crossing energy barriers. These states then connect to the parameters of a negative exponential function. Thus indeterministic channel models have a natural physical intepretation, with a simple correspondence between model states, functions, and their physical counterparts. All these factors make it more likely that the physical interpretation of a Markov model is both useful and accurate (Colquhoun and Hawkes 1995).

In contrast, it is difficult to find such theory-to-world fit for fractal and deterministic models. These models usually represent dwell times with a two-parameter power-law function (Lowen et al. 1999). Such equations and parameters have no ready physical interpretation or mechanism (Horrigan 2015). These approaches are therefore of limited scope, and no currently accepted gating model relies on them (ibid.). Furthermore, claiming deterministic chaos for a system in a noisy environment like a cell membrane carries a theoretical risk. A system's sensitivity to initial conditions amplifies noise as well as signal, making the concept of chaos less useful (Lowen and Teich 2005), independent of the problem of empirical interpretation.

These results—coherence with other aspects of neurobiology and chemistry, and ease of physical interpretation—are strong indirect evidence for indeterministic channel models. These models thus enjoy advantages in both their predicted

[10] Since shifts in amino acid structure happen on a micro-scale, a channel protein in fact has an enormous number of states. But there exist a smaller number of metastable states which comprise many substates, and these metastable states are those represented by channel models. See Colquhoun and Hawkes (1995, 400–401).

distributions and their connections to other theories. The first advantage blocks underdetermination; the second situates models in neighboring scientific fields.

The evidence overwhelmingly points to ion channels as weakly-indeterministic systems. From the ion channel data we can therefore conclude that the brain is indeterministic in the weak sense: the best models for one of its essential subprocesses are indeterministic. In the next section, I discuss what this might mean for strong indeterminism and the brain.

Philosophical arguments for strong indeterminism

Empirical evidence suggests that ion channels are weakly indeterministic. Important as this result is, however, nothing about strong indeterminism follows from it directly. Neither the evidence itself nor any devisable experiment helps us infer that some weakly-indeterministic system is also strongly indeterministic.

In this section, I evaluate three additional *philosophical* arguments which may help us make the leap from weak to strong indeterminism. Their goal is to move from properties of a model to conclude something about the world, though they take that step in different ways. The first is an argument from scientific realism, mentioned briefly above. The second is based on the best-systems account of laws, and the third on level-relative chances.

Two preliminary points will help clear the ground for my discussion. The first is that we cannot use quantum mechanics to make an argument about ion channels. Glymour (2001), for example, suggests that random-looking gating could be due to quantum effects "percolating up" to higher levels to cause openings and closings. Gating is a result of interactions in a channel's heat bath, however, and has nothing to do with lower-level particle physics. A related argument, that the thermodynamic forces behind channel gating are themselves strongly indeterministic, fails as well. This argument just passes the buck to thermodynamics and statistical mechanics, forcing us to justify strong indeterminism somewhere else.

The second point is a possible worry about weak indeterminism. Indeterministic channel models may be the best we have *now*, but what about the future? Perhaps a superior deterministic model will come along later, without needing fractals or deterministic chaos. Maybe it will even have a plausible physical interpretation.

This possibility is unlikely, given what we know about proteins, but we may still worry about a future deterministic model superseding current ones. Weak indeterminism is fairly non-committal, however. It says something about the statistical distributions a certain system creates, but says much less about the underlying nature of that system. We do not lose much *philosophically* if we find out later that some deterministic model is superior. Note that this does not mean weak indeterminism is unimportant. It's a useful property to know about for explanations, futher computational modeling, and predictions from interventions. And while it is not a trivial matter to show that a system is weakly indeterministic, the property's metaphysical commitments are light—which explains why we need extra arguments to get to strong indeterminism.

I will now evaluate the first of these arguments, each of which tries to bridge the gap between weak and strong indeterminism. The first we saw above. Brandon and Carson (1996)'s argument for strong indeterminism infers it directly from the weak sense.[11] Their concern is evolutionary theory. They reason that genetic drift, or gene frequency change due to sampling errors, is both irreducibly probabilistic and unavoidable in some evolutionary processes. That is, while drift in small populations is likely, in some circumstances a gene's frequency cannot remain the same from one generation to the next (Brandon and Carson 1996, 322). Thus the weak sense seems true for evolutionary theory—only a probabilistic or indeterministic model describes the system.

These authors then infer straightaway that evolutionary processes are strongly indeterministic as well, by means of a deep commitment to scientific realism. They claim that "if one is a realist in one's attitude towards science...then one should conclude that E[volutionary] T[heory] is fundamentally indeterminstic" [i.e., that the strong sense is true for evolutionary theory] (336). Since probabilites are unavoidable in our best theory, and our best theories truly describe reality, we should hold further that those probabilities are fundamental to evolutionary processes. The real systems involving those processes would therefore be indeterministic.

Another way of putting their point is that the state of an evolutionary system at one time does not fix the future states. Note here that we are not just talking about *models* of evolution; we are talking about the evolutionary processes themselves. On Brandon and Carson's view, the development of an evolutionary system through time could take multiple possible paths—not because we don't know enough to predict the actual path, but because the system itself is fundamentally indeterministic.

Turning to the brain, we could make the same argument about ion channels. Being scientific realists, we could infer that the rate constants describing channel behavior are fundamental features of the world. On this view, ion channels just are indeterministic, and their state at one time does not fix future or past states.

At first glance it may be tempting to make the following response. One could say that the argument from realism leads to impossible consequences, since it implies that under the same initial conditions with the same laws, we would observe different outcomes. One may think that this implication must be wrong, since the outcomes would clearly have to be the same. Therefore, this response says, Brandon and Carson's argument cannot be right. The problem with this response is that it begs the question, just as we saw in our earlier hypothetical experiments. Our intuitive reaction that the outcomes would be identical is not evidence for the determinist position; it is merely a statement of one's commitment to it. Whether we would get different outcomes is precisely the point at issue.

A better response to the realist argument would note, along with Graves et al. (1999), that no mechanism appears to underlie the indeterministic character of the evolutionary processes. Brandon and Carson (1996) cite quantum percolation as a

[11] Section 2 of Butterfield (2005) mentions a version of this argument as well but does so in the context of a complete theory for a possible world.

possibility (Sober 1984), but there is no demonstrable evidence that this actually occurs.[12] The same is true for ion channels. If we as scientific realists inferred strong indeterminism for ion channels directly, we would have to ask what gives rise to that behavior. And despite knowing a great deal about channel composition and function, we would have no mechanism to appeal to. Given the problems it would create in other parts of neurobiological theory, at best we should be agnostic about the argument from scientific realism. It cannot be the way to go from weak to strong indeterminism.

The remaining two arguments have a more a priori flavor, dealing less with empirical facts than with our framework for understanding them. The first is a strategy based on David Lewis's best-systems account of laws, and the second depends on level-relative chance.

I begin with Lewis. On his view empirical regularities are laws when they belong to the "best system" (Lewis 1973; Ramsey 2009). The best system is the set of generalizations striking the best balance among theoretical virtues like strength and simplicity. When we decide which generalizations should be our laws, these virtues compete against each other. Adding more regularities might offer more information while decreasing simplicity, or vice versa. Assembling the system requires finding the right balance between the virtues.

We could go from the weak sense to strong indeterminism if the best system itself included probabilities (Loewer 2001). The universe would evolve according to the best system's laws, and since some laws would be probabilistic, there would be multiple possible futures for some initial state. An indeterministic model might therefore indicate an indeterministic part of the real (best) system.

But why include probabilities in the best system? Chances bring vast gains in simplicity: a single chance rule can describe many disparate events. Such rules may increase strength as well, and so they might enter the best system if the corresponding gains were significant enough.

There are various ways to implement this strategy (Loewer 2001; Hoefer 2007; Frigg and Hoefer 2010, 2015). The specifics don't matter much here; what does is the connection to neuroscience. Suppose we were deciding whether to add probability rules for ion channels to our best system. What reasons could we give for including them? Such reasons would be difficult to come by. After all, why should we expect to find channel gating probabilities in the best system of laws for the whole universe? Perhaps including these probabilities would enhance the system's simplicity. In comparison with the rest of the universe, however, the probabilities would cover so few events that no gains in information or predictive power could make up for them. The best-systems account may work for strong indeterminism more generally, or on a larger scale, but it is not likely to help for the brain.

The third and last strategy depends on level-relative chances, as described in Glynn (2010).[13] Schaffer (2007) defines chance as a function $Ch_w(p)$ taking three

[12] But see Stamos (2001). Millstein (2000) notes that the percolation argument may be too vague to settle.

[13] Werndl (2012) mentions a similar view.

arguments: a time, a world, and a proposition describing an event. Glynn conceptualizes chance as the four-place function $Ch_{twl}(p)$. The additional argument l is the level of description for the event in p. With a fourth argument, it's possible for two chance functions describing one event to vary only along the l-dimension. If two such functions returned different values for the chance of a single event, then that event would have two distinct chances in the same world at the same time. One function might return a value of 1, for example, while a second (with different l) might return .43. Chances between 0 and 1, described at a higher level, would be compatible with deterministic chances at a lower level.

List and Pivato (2015) give much the same argument as Glynn. The ploy behind level-relativity is to isolate higher-level chances from intereference by lower-level facts. This is normally done by restricting higher-level chance functions such that they conditionalize only on higher-level information. For example, say we had a four-place chance function describing an ion channel's gating events. As a higher-level function it conditionalizes on information about neurobiological kinds: molecules and proteins, inorganic compounds, cell organelles, and so on. This function would return a value between 0 and 1, rather like a rate constant.

Now normally we would be able to conditionalize on lower-level information as well. The chance function might then return a value of either 0 or 1, in accordance with the laws governing the lower level. Glynn allows this low-level conditionalization but maintains that the two distinct chance values are both legitimate; ListPivato2015 disallow it, calling such a conditionalization a "category mistake" because it "mix[es] two different levels of description" (135). Either way the result is the same: higher-level chances end up compatible with lower-level determinism.

Level-relative chances then help us go from weak indeterminism to strong, since the strong sense follows trivially from the weak. This is because the level at which ion channels exist is far above any level for which we have deterministic laws. Given the laws or principled assumptions available at the neurobiological level, ion channels have multiple possible futures because those laws don't fix the future states—they evolve according to the level-relative chances in our models.

Glynn and List and Pivato give the most promising approach for getting strong indeterminism from weak. As long as we have weak indeterminism at one level, we appear to be justified in inferring strong indeterminism for that level, since we have a framework for making it compatible with determinism.[14]

Level-relative chances do not give us strong indeterminism for free; we cannot have it both ways. An ion channel described at one level may be indeterministic, as in a Markov model. Described at another it may be deterministic. Yet both cannot truly represent the real system, for that system either is or is not indeterministic. Its current state either does or does not fix all other states. In other words, arguing from level-relative chance actually allows *two* inferences: one at level n from an indeterministic model to an indeterministic system, and another at level $n \pm 1$ from a deterministic model to a deterministic system. But these inferences are not both

[14] List and Pivato (2015) target views like Schaffer (2007)'s, which seem to say that if we have indeterminism at one level, there must be indeterminism at every level below that; the indeterminism must go "all the way down."

compatible with the ground truth of a system. It cannot be that both are true representations, and indeed level-relative chances were not intended for this purpose, as they were meant to make objective chances compatible with epistemic probabilities (List and Pivato 2015). When used for the purpose of strong indeterminism, level-relative chances assume we know more than we actually do. They assume that we have enough information to decide which of the two contradictory inferences is better: the inference to a deterministic system, or the inference to an indeterministic one. Yet if we knew enough to make that decision, we wouldn't need level-relative chances to get strong indeterminism in the first place. We would just infer it in cases where we knew it already existed.

This approach tries to accommodate multiple incompatible descriptions for a single system. Perhaps we can describe an ion channel in several ways, but the system itself is only one way. The system's actual properties do not depend on how we describe them; either its current state fixes its past and future or it doesn't. Level-relative chances cannot tell us which is true. Like the previous two arguments, this view fails to push us from indeterminism for a model to indeterminism for the world.[15]

In this section I have discussed three approaches to infer the strong sense of indeterminism from the weak. All three fail to get us from a weakly indeterministic system to a strong one. We have already seen that empirical demonstrations are insufficient to show strong indeterminism; it appears that philosophical arguments are inadequate as well, at least for the brain. Even though the best models of ion channel function are indeterministic, we have no good reason to think that the channels themselves are. I discuss this result in the next section.

The brain and strong indeterminism

Philosophical discussions of indeterminism tend to focus on physics and low-level particle systems (Loewer 2001; Ismael 2009; Werndl 2009; Frigg and Hoefer 2015). Biology, save for a handful of interventions (Brandon and Carson 1996; Millstein 2003; Sansom 2003; Werndl 2012b), has proceeded in relative isolation from these debates. In this paper I have taken more general views on determinism and applied them to the brain. Ion channels make an excellent test case for these arguments, since they are essential for cognition and we already know a lot about their structure and behavior.

Applying general arguments to the case of determinism in biology is not always straightforward. Biology, and neurobiology in particular, counts on mechanisms. Central to the notion of a mechanism are causation and causal relationships (Machamer et al. 2000). Accordingly, our naïve view of determinism carries

[15] Level-relative chance views would have other unsavory implications if they could get us to strong indeterminism. For example, since the strong sense effectively collapses into the weak sense on level-relative chance, then real indeterministic systems would be everywhere—for any stochastic model, we could infer that the represented system itself is indeterministic. Level-relative chance would also blur the line between cases like ion channels and those like quantum mechanics, where (on some interpretations) the latter probabilities are more fundamental.

connections to causes and predictions which more sophisticated accounts eschew. When we bring these cause-less accounts into the context of the brain, we end up focusing on models and their properties. This focus introduces a gap between the model and the system it represents. That a system is weakly indeterministic—that our best model of it is stochastic—says little about the system itself.

Whether a system is weakly indeterministic is not a trivial matter, however. A great deal of evidence bears on the question for any given system. Previous results on observational equivalence underscore the importance of indirect evidence and broad coherence in theory-building and model choice. Such evidence and coherence mean that weak indeterminism can play an important role in explaining channel behavior, independent of whether we can reach strong indeterminism through some other argument.

We have already seen this role in physically interpreting the Markov model. The states of the model correspond to protein conformational states, which help explain why channels gate in certain ways. We could also ask further why an individual channel's dwell-time distribution doesn't seem to contain any particular patterns. A narrow answer to this question would cite the details of a particular exponential function and its parameters, but a more general answer says that the channel gates randomly. This more general explanation unifies the behavior of other ion channel types and helps us understand different phases of the action potential. From the perspective of neurobiology, an appeal to random behavior makes for a perfectly legitimate explanation. In these cases, the best explanations of certain neural phenomena essentially involve indeterminism. We cannot escape chance when explaining the brain.

On the other hand, because of its role in action, the brain attracts interest as a system whose behavior may not be fully determined. The unfortunate truth is that its complexity makes the brain a bad place to look for instances of strong indeterminism. So many factors influence action that we cannot be sure we are accounting for them all.

My argument assuages these worries, however, for it shows that neither empirical arguments nor philosophical ones are capable of showing that the brain is strongly indeterministic. Current scientific and analytical methods cannot answer the question either way. Even with ion channels, where we have an enormous amount of evidence in favor of indeterminism, we still cannot infer the strong sense from the weak. Since we cannot answer the question, there is no sense in worrying about it. The weak sense of indeterminism is all we have to work with, and the only sense that figures in our explanations and predictions of neural phenomena.

Neuroscience requires probabilities. Neuroscientific explanations depend on them, and neurobiological explanations of channel gating are incomplete without them. Stochasticity is an explanatory necessity for neuroscience. The empirical evidence tells us that, but the evidence stops there. Neither it nor philosophical arguments are capable of taking us farther.

To conclude, I have argued that the brain is indeterministic in the weak sense, which means that our best models of at least one key neural process are indeterministic. But neither the evidence itself nor philosophical argumentation can connect this weak sense to a strong sense of indeterminism. Yet despite the gap

between these two senses, indeterministic models are still essential for neurobiology.

References

Akaike H (1974) A new look at the statistical model identification. IEEE Trans Autom Control 19(6):716–723

Austin TD (2008) The emergence of the deterministic hodgkinhuxley equations as a limit from the underlying stochastic ion-channel mechanism. Ann Appl Probab 18(4):1279–1325

Brandon RN, Carson S (1996) The indeterministic character of evolutionary theory: no no hidden variables proof but no room for determinism either. Philos Sci 63:315–337

Butterfield J (2005) Determinism and indeterminism. Routledge encyclopedia of philosophy. Taylor and Francis

Cavalcanti S, Fontanazzi F (1999) Deterministic model of ion channel flipping with fractal scaling of kinetic rates. Ann Biomed Eng 27(5):682–695

Colquhoun D, Hawkes AG (1977) Relaxation and fluctuation of membrane currents that flow through drug-operated channels. Proc R Soc B Biol Sci 199:231–262

Colquhoun D, Hawkes AG (1981) On the stochastic properties of single ion channels. Proc R Soc B Biol Sci 1(211):205–235

Colquhoun D, Hawkes AG (1982) On the stochastic properties of bursts of single ion channel openings and of clusters of bursts. Philos Trans R Soc Lond Ser B Biol Sci 300:1–59

Colquhoun D, Hawkes AG (1995) The principles of the stochastic interpretation of ion-channel mechanisms. In: Neher E, Sakmann B (eds) Single-channel recording. Plenum Press, New York, pp 397–482

Craver CF (2014) The ontic account of scientific explanation. In: Kaiser MI, Scholz OR, Plenge D, Hüttemann A (eds) Explanation in the special sciences: the case of biology and history. Springer, New York, pp 27–52

Earman J (1986) A primer on determinism. University of Western Ontario series in philosophy of science, v. 32. Reidel

Eyring H (1935) The activated complex in chemical reactions. J Chem Phys 3(2):107–115

Fitzhugh R (1965) A kinetic model of the conductance changes in nerve membrane. J Cell Comp Physiol 66(S2):111–117

French RJ, Horn R (1983) Sodium channel gating: models, mimics, and modifiers. Annu Rev Biophys Bioeng 12:319–356

Frigg R, Hoefer C (2010) The present situation in the philosophy of science. Chapter determinism and chance from a humean perspective. Springer, Berlin, pp 351–371

Frigg R, Hoefer C (2015) The best humean system for statistical mechanics. Erkenntnis 80:551–574

Geng Y, Magleby KL (2015) Single-channel kinetics of BK (Slo1) channels. Front Physiol 5(January):1–24

Glymour B (2001) Selection, indeterminism, and evolutionary theory. Philos Sci 68(4):518–535

Glynn L (2010) Deterministic chance. Br J Philos Sci 61(1):51–80

Graves L, Horan BL, Rosenberg A (1999) Is indeterminism the source of the statistical character of evolutionary theory? Philos Sci 66(1):140–157

Hille B (1992) Ionic channels of excitable membranes. Sinauer Associates Inc, Sunderland, MA

Hoefer C (2007) The third way on objective probability: a sceptic's guide to objective chance. Mind 116(463):549–596

Hoefer C (2016) Causal determinism. In: Zalta EN (ed) The Stanford encyclopedia of philosophy, Spring 2016 ed. Metaphysics Research Lab, Stanford University

Horn R, Vandenberg CA (1984) Statistical properties of single sodium channels. J Gen Physiol 84(4):505–534

Horrigan FT, Hoshi T (2015) Models of ion channel gating, Chapter 7. CRC Press, Boca Raton, pp 83–102

Ismael J (2009) Probability in deterministic physics. J Philos 106(2):89–108

Jones SW (2006) Are rate constants constant? J Philos 571(Pt 3):502

Korn SJ, Horn R (1988) Statistical discrimination of fractal and markov models of single-channel gating. Biophys J 54(5):871–877
Laplace P-S (1814) Essai philosophique sur les probabilits. Courcier, Paris
Laudan L, Leplin J (1991) Empirical equivalence and underdetermination. J Philos 88(9):449–472
Lewis DK (1973) Counterfactuals. Blackwell, Oxford
Lewis D (1980) A subjectivist's guide to objective chance. In: Jeffrey RC (ed) Studies in inductive logic and probability. University of California Press, Berkeley, pp 83–132
Lewis D (1994) Humean supervenience debugged. Mind 103(412):473–490
Liebovitch LS (1993) Interpretation of protein structure and dynamics from the statistics of the open and closed times measured in a single ion-channel protein. J Stat Phys 70(1/2):329–337
Liebovitch LS (1996) Ion channel kinetics. In: Iannaccone PM, Khokha M (eds) Fractal geometry in biological systems: an analytical approach. CRC Press, Boca Raton, pp 31–56
Liebovitch LS, Czegledy FP (1991) Fractal, chaotic, and self-organizing critical system: descriptions of the kinetics of cell membrane ion channels. In: Mosekilde E, Mosekilde L (eds) Complexity, chaos, and biological evolution. Plenum Press, New York, pp 145–153
Liebovitch LS, Czegledy FP (1992) A model of ion channel kinetics based on deterministic, chaotic motion in a potential with two local minima. Ann Biomed Eng 20(5):517–531
Liebovitch LS, Krekora P (2002) The physical basis of ion channel kinetics: the importance of dynamics. In: Layton HE, Weinstein AM (eds) Institute for mathematics and its applications volumes in mathematics and its applications, membrane transport and renal physiology, vol 129. Springer, Berlin, pp 27–52
Liebovitch LS, Tőth TI (1990) Fractal activity in cell membrane ion channels. Ann N Y Acad Sci 591(1):375–391
Liebovitch LS, Toth TI (1991) A model of ion channel kinetics using deterministic chaotic rather than stochastic processes. J Theor Biol 148:243–267
Liebovitch LS, Fischbarg J, Koniarek JP, Todorova I, Wang M (1987) Fractal model of ion-channel kinetics. Biochim Biophys Acta 896:173–180
Lipscombe D, Toro CP (2014) Biophysics of voltage-gated ion channels, Chapter 13. Elsevier, Amsterdam, pp 377–407
List C, Pivato M (2015) Emergent chance. Philos Rev 124(1):119–152
Loewer B (2001) Determinism and chance. Stud Hist Philos Mod Phys 32(4):609–620
Lowen SB, Teich MC (2005) Fractal-based point processes. Wiley
Lowen SB, Liebovitch LS, White JA (1999) Fractal ion-channel behavior generates fractal firing patterns in neuronal models. Phys Rev E Stat Phys Plasmas Fluids Relat Interdiscip Top 59(5 Pt B):5970–5980
Machamer PK, Darden L, Craver CF (2000) Thinking about mechanisms. Philos Sci 67(1):1–25
McManus OB, Weiss DS, Spivak CE, Blatz AL, Magleby KL (1988) Fractal models are inadequate for the kinetics of four different ion channels. Biophys J 54:859–870
McManus OB, Spivak CE, Blatz AL, Weiss DS, Magleby KL (1989) Fractal models, Markov models, and channel kinetics. Biophys J 55(2):383–385
Millhauser GL, Salpeter EE, Oswald RE (1988) Rate-amplitude correlation from single-channel records. a hidden structure in ion channel gating kinetics? Biophys J 54(6):1165–1168
Millstein R (2000) Is the evolutionary process deterministic or indeterministic? An argument for agnosticism. Presented at the Biennial Meeting of the Philosophy of Science Association
Millstein R (2003) Determinism in physics and biology, chapter how not to argue for the indeterminism of evolution: a look at two recent attempts to settle the issue, Paderborn, pp 91–107
Moretti L (2007) Ways in which coherence is confirmation conducive. Synthese 157(3):309–319
Neher E, Sakmann B (1976) Single-channel currents recorded from membrane of denervated frog muscle fibres. Nature 260:799–802
Ornstein DS, Weiss B (1991) Statistical properties of chaotic systems. Bull Am Math Soc (NS) 24(1):11–116
Qin F (2014) Principles of single-channel kinetic analysis. Patch Clamp Methods Protoc 1183:371–399
Qin F, Li L (2004) Model-based fitting of single-channel dwell-time distributions. Biophys J 87(3):1657–1671
Ramsey FP (2009) Foundations of mathematics and other logical essays. Routledge, London
Sansom R (2003) Why evolution is really indeterministic. Synthese 136:263–280

Sansom MS, Ball FG, Kerry CJ, McGee R, Ramsey RL, Usherwood PNR (1989) Markov, fractal, diffusion, and related models of ion channel gating: a comparison with experimental data from two ion channels. Biophys J 56(6):1229–1243

Schaffer J (2007) Determinstic chance? Br J Philos Sci 58:113–140

Sober E (1984) The nature of selection. MIT Press, Cambridge

Stamos DH (2001) Quantum indeterminism and evolutionary biology. Philos Sci 68(2):164–184

Van Inwagen P (2015) Metaphysics. Westview Press, Boulder

van Strien M (2014) The norton dome and the nineteenth century foundations of determinism. J Gen Philos Sci/Z Allg Wiss 45(1):167–185

Varanda WA, Liebovitch LS, Figueiroa JN, Nogueira RA (2000) Hurst analysis applied to the study of single calcium-activated potassium channel kinetics. J Theor Biol 206(3):343–353

Weber M (2005) Indeterminism in neurobiology. Philos Sci 72(December):663–674

Werndl C (2009) Are deterministic descriptions and indeterministic descriptions observationally equivalent? Stud Hist Philos Sci Part B Stud Hist Philos Mod Phys 40(3):232–242

Werndl C (2011) On the observational equivalence of continuous-time deterministic and indeterministic descriptions. Eur J Philos Sci 1(2):193–225

Werndl C (2012a) Evidence for the deterministic or the indeterministic description? A critique of the literature about classical dynamical systems. J Gen Philos Sci/Z Allg Wiss 43(2):295–312

Werndl C (2012b) Probabilities, laws and structures, chapter probability, indeterminism and biological processes. Springer, Berlin, pp 263–277

Werndl C (2013) On choosing between deterministic and indeterministic models: Underdetermination and indirect evidence. Synthese 190(12):2243–2265

Werndl C (2016) Determinism and indeterminism. In: The oxford handbook of philosophy of science. Oxford University Press, Oxford