

Clinical Applications of Stochastic Dynamic Models of the Brain, Part I: A Primer

James A Roberts^{1,2}, Karl J. Friston³, Michael Breakspear^{1,4}

¹ Systems Neuroscience Group, QIMR Berghofer Medical Research Institute, Herston, Brisbane, QLD 4006, Australia

² Centre for Integrative Brain Function, QIMR Berghofer Medical Research Institute, Herston, Brisbane, QLD 4006, Australia

³ The Wellcome Trust Centre for Neuroimaging, Institute of Neurology, University College London, London, UK

⁴ Metro North Mental Health Service, Royal Brisbane and Women's Hospital, Brisbane, QLD, Australia

* Corresponding author: Michael Breakspear, Systems Neuroscience Group, QIMR Berghofer Medical Research Institute, 300 Herston Rd, Herston, QLD 4006, Australia. Email: michael.breakspear@qimrberghofer.edu.au, phone: +61 7 3845 3692, fax: +61 7 3362 0111.

Keywords: stochastic, mathematical modeling, epilepsy, melancholia, computational psychiatry

Word counts:

Abstract: 194

Article body: 3982

Number of figures: 4

Number of tables: 0

Number of supplemental items: 1

Short title: Stochastic Dynamic Models

Abstract:

Biological phenomena arise through interactions between an organism's intrinsic dynamics and stochastic forces – random fluctuations due to external inputs, thermal energy, or other exogenous influences. Dynamic processes in the brain derive from neurophysiology and anatomical connectivity; stochastic effects arise through sensory fluctuations, brainstem discharges, and random microscopic states such as thermal noise. The dynamic evolution of systems composed of both dynamic and random effects can be studied with stochastic dynamic models (SDMs). This paper, Part I of a two-part series, offers a primer of SDMs and their application to large-scale neural systems in health and disease. The companion paper, Part II, reviews the application of SDMs to brain disorders. SDMs generate a distribution of dynamic states, which (we argue) represent ideal candidates for modeling how the brain represents states of the world. When augmented with variational methods for model inversion, SDMs represent a powerful means of inferring neuronal dynamics from functional neuroimaging data in health and disease. Together with deeper theoretical considerations, this work suggests that SDMs will play a unique and influential role in computational psychiatry, unifying empirical observations with models of perception and behavior.

1. Introduction

Biological organisms balance a tendency toward internal order and control against the need to pre-empt and adapt to a changing environment (1). Their functioning reflects a dynamic interplay of nonlinearity and feedback with stochastic fluctuations: This exchange of order and entropy yields complexity in its various guises (2). In the setting of the brain, innate dynamics arise from neurophysiological processes such as ion channels and dendritic filtering, while feedback occurs through short- and long-range axonal connections (3); stochastic inputs arise through sensory fluctuations, brainstem discharges, and thermal energy (random fluctuations at the microscopic level, such as the Brownian motion of ions). There is a well-established field devoted to modeling this interplay of dynamic processes and stochastic effects through a melding of calculus and statistical physics. The field is anchored by dynamic equations that govern the temporal (and possibly spatial) behavior of the system's state variables. These equations are derived from the biophysical properties of the system of interest and studied using analysis and simulation. Since the actual state variables (such as firing rates) cannot be directly observed in non-invasive human studies, measurement functions that map neuronal states onto observables (such as scalp electroencephalography (EEG) or the blood-oxygen-level dependent (BOLD) signal) are required to enable empirical predictions. Experimental data then allow models to be tested, compared, refined, or refuted.

In this paper, we introduce the equations that arise at the intersection of calculus and statistical physics, namely stochastic differential equations (SDEs). These lie at the heart of stochastic dynamic models of the brain, for which we offer micro- and mesoscopic examples. We also showcase the potential of stochastic differential models to unify observations of functional neuroimaging data with models of behavior. This forms the background for Part II (4), in which we review existing applications of SDEs to clinical disorders in neurology and psychiatry and consider future perspectives.

2. Stochastic dynamic models of the brain: A brief primer

2.1 Stochastic differential equations: Fundamentals

We first introduce the modeling of neural systems with SDEs. Key to this approach is the notion of a system’s *state* – the core dynamical variables that describe the system at any instant in time, such as firing rates, membrane potentials, channel conductances, etc. Models describe how states evolve in time – the dynamics. Given the current state and the dynamical rules of a particular model, it is possible to project (“integrate”) the state dynamics forward in time (i.e., solve the equations). Suppose we model a neuronal system with N state variables. We can represent these variables as a vector, $\mathbf{x} = [x_1, x_2, \dots, x_N]$ where, for example, x_1 is the cell membrane potential, x_2 is the firing rate, and x_3 is the conductance of a particular class of membrane channels, etc. Then, in the absence of any random fluctuations, the dynamic evolution of the state variables obeys a set of ordinary differential equations (ODEs),

$$\begin{aligned} \frac{dx_1}{dt} &= F_1(x_1(t), x_2(t), \dots, x_N(t)), \\ \frac{dx_2}{dt} &= F_2(x_1(t), x_2(t), \dots, x_N(t)), \\ &\vdots \\ \frac{dx_N}{dt} &= F_N(x_1(t), x_2(t), \dots, x_N(t)), \end{aligned} \tag{1}$$

where $\frac{d}{dt}$ is the usual differential operator (with respect to time) and the F_j are functions (possibly nonlinear) that embody the properties and interactions of the system. These equations can also be represented in a simpler vector form,

$$\frac{dx_i}{dt} = F_i(\mathbf{x}(t)). \tag{2}$$

An example of simple neural dynamics is given by the Morris-Lecar model (5), a two dimensional simplification of the Hodgkin-Huxley equations for the excitable membrane potential of a neuron. The membrane potential V is determined by the net current through all transmembrane ion channels (Na^+ , K^+ , Ca^{2+} , and leaky currents). The change in membrane potential at the cell soma is given by the sum of all ion channel currents plus any current I entering from the dendritic tree. In the full (four-dimensional) Hodgkin-Huxley model, the

voltages and temporal behavior of all the major ion channels are modeled explicitly. A reduction to the 2-D Morris-Lecar model is achieved by exploiting the fact that calcium and sodium channels respond more quickly to changes in membrane voltage than slower potassium channels. This means the calcium and sodium channels follow their voltage-dependent conductances instantaneously, while the potassium channels ‘relax’ to that value on a slower time scale – allowing us to focus on the slower dynamics, which ‘enslave’ faster dynamics. That is, because the fast variables reach equilibrium quickly after a perturbation, it is convenient to treat them as always being at equilibrium, such that the only remaining dynamics are in the slow variables – hence slow variables can be said to enslave fast ones (6). This is a common device in modeling dynamical systems known as an adiabatic approximation, which rests on the separation of time scales. Although widely used, it is of less value in systems where fast, microscopic fluctuations may drive slower, coarser subsystems (e.g. in turbulent fluid dynamics).

The dynamical states for the Morris-Lecar model are $\{V, n\}$, the membrane potential V and the proportion of open potassium ion channels n , which is proportional to the membrane conductance. The dynamic equations for the Morris-Lecar model are presented in Appendix A (Supplemental Material). An example time series is given in Figure 1A, showing the characteristic rapid spiking waveform of a supra-threshold cell. Figure 1B shows the corresponding limit cycle *attractor* in the *phase space* spanned by the dynamical variables $\{V, n\}$. The attractor is color coded to show the direction of the (clockwise) *flow* around the smooth limit cycle attractor.

The equations thus far capture the essence of a neuron that is kept spiking by a constant dendritic current. However, as discussed above, neural dynamics inevitably occur in the presence of noisy fluctuations. In the Morris-Lecar model, such fluctuations reflect stochastic effects at ion channels, thermal energy, the uneven distribution of channels in the membrane, and irregularities in dendritic currents – sources of so-called neural noise (7). Stochastic effects can be introduced by the addition of a random term to the state equation,

$$\frac{dx_i}{dt} = F_i(\mathbf{x}) + \mu_i \eta_i, \quad (3)$$

where η_i represents independent zero-mean, unit-variance Gaussian noise added incrementally at each time point (also called a *Wiener process*), and μ_i is a coefficient that scales the noise appropriately to each of the variables. Note that while it is apparently simple to add noise in this way (yielding the *Langevin equation*), neither classic calculus nor standard numerical integration schemes deal with the incorporation of incremental “rough” discontinuities in this way. An alternative way of writing Equation (3) is through the use of stochastic differential equations (SDEs). This formulation makes it clear that changes in the states x_i are governed by deterministic contributions to the temporal dynamics F of the system and the random fluctuations μ_i . For the interested reader, we provide the corresponding SDEs in Appendix B (Supplemental Material), together with a more formal treatment of stochastic fluctuations¹.

Figures 1C-D show example dynamics from the Morris-Lecar system with an additive stochastic term. Both the time series (Fig. 1C) and the attractor (Fig. 1D) show the impact of the irregular roughness of the additive noise term. The roughness is more apparent during the refractory phase of the firing cycle, but only because the underlying flow is slower during these periods, allowing more time for the noise to accrue. The impact of the noise is a relatively modest degradation in the regularity of the periodicity and amplitude of the spikes.

This is a simple example of a stochastic dynamic model with independent dynamic and noise terms. The noise acts to perturb the system as it traverses the limit cycle attractor. However, in many complex systems, noise does not enter as a simple state-independent (additive) term. An example is the well-known coupling between trade volume and volatility in financial markets: as the amount of trade increases, so do fluctuations in market value (9). Therefore, in many applications, the noise mixes with the states as it enters the system, yielding

$$\frac{dx_i}{dt} = F_i(\mathbf{x}) + G_i(\mathbf{x})\eta_i, \quad (4)$$

¹ All of the equations in this paper, together with an integration scheme for SDEs (8), can be downloaded in MATLAB form from sng.org.au/Downloads. Python code for integrating SDEs is also available.

where G is a function that captures the state-dependence of the stochastic influence. Equation (4) is a “generalized Langevin equation”. If G is linear in \mathbf{x} and each state mixes with only its own noise term, we have

$$\frac{dx_i}{dt} = F_i(\mathbf{x}) + \mu_i x_i \eta_i. \quad (5)$$

That is, the influence of the noise term is not constant, but scales in proportion to the states \mathbf{x} . Fluctuations at voltage-dependent ion channels are by definition state-dependent. An example of Morris-Lecar dynamics with state dependent noise is shown in Figures 1E and F. The stand-out feature of these panels is the increase in fluctuations in the sub-threshold regime (<-50 mV) compared to the preceding panels, and the contrasting smoother flow in the fast, suprathreshold phase (>-20 mV). This noisier subthreshold phase substantially increases the irregularity of interspike intervals.

Fluctuations induced by additive noise [Eq. (3)] accord to normal distributions, whereas those arising from state-dependent terms [Eq. (4) and (5)] often display a broad range of non-Gaussian, *heavy-tailed* forms (10). Thin tails are also possible, such as when noise decreases with the state values, though this can also counterintuitively yield heavy-tailed noise in other cases (10). It has previously been argued that this state-dependent noise equation is a more appropriate framework for modeling the brain, capturing the heavy-tailed, bimodal fluctuations in the alpha rhythm (11, 12). More generally, state-dependent fluctuations may be fundamental to human cognition, arising in perception (the “Weber-Fechner law”; (13)), movement (“Fitt’s law”; (14)) and even computation (“Hick’s law”; (15)).

Stochastic fluctuations are relevant to both large-scale neural models and microscopic spiking neural models, although the sources of fluctuations are different. Whereas noise at the microscopic level reflects stochastic effects of ion channels and heat, these may not be relevant at coarser scales. Unknown inputs from sensory systems, erratic discharges from brain stem nuclei, and inputs from other large-scale neural systems that have not been explicitly modeled can all be accommodated as random fluctuations in models of mesoscopic neural dynamics. However, systems with heavy-tailed statistics do not accord to the classic central limit theorem:

That is, small-scale fluctuations do not decay with increasing system size but persist as high-amplitude (anomalous) fluctuations (16). Hence state-dependent stochastic systems of the form of Eq. (5) can “transport” fluctuations between scales – from microscopic to mesoscopic.

Mesoscopic neural models – so-called *neural mass models* – represent local average neural states (such as the average firing rates or membrane potentials) of pyramidal and inhibitory neurons within a single cortical column (17). The process of averaging over a large number of neurons to reduce dimensionality is a *mean field approximation*. In addition to making the analysis more tractable, the mean field reduction allows one to study neuronal dynamics at scales accessible by non-invasive means and, arguably, at scales of most relevance to cognition and behavior. Figure 2 shows an example of a conductance-based neural mass model (the equations are in Appendix A, Supplemental Material). This model is basically an extension of the Morris-Lecar model with a simple mean field approximation (mean membrane potentials are converted to average firing rates through a sigmoid activation function) and incorporation of a simplified inhibitory population (which responds to pyramidal cell firing with a slow feedback inhibitory current) (18). In this example, the purely deterministic dynamics are smooth, aperiodic chaos (Figure 2A and B). The addition of a constant amplitude noise term “roughens” the flow, similar to the microscopic system, increasing the temporal and amplitude irregularity of the corresponding time series (Figure 2C and D), particularly at times when the underlying dynamics are slow. Incorporating a state-dependent noise term (Figure 2E and F) further changes the irregularity of the attractor’s slow phase: Although the variance of the fluctuations during the slow phase is approximately the same for additive and multiplicative noise (compare panels C and E), the dynamics are substantially more irregular in the case of state-dependent noise.

2.2 *From stochastic differential equations to stochastic differential models*

This concludes our survey of the fundamentals of stochastic neural equations. However, a “model” is arguably more than an equation for the hidden states, but should also include a framework for explaining empirical observations.

2.2.1. *Importance of a measurement function.* The underlying states of neuronal systems cannot usually be directly measured, particularly in non-invasive imaging modalities such as EEG and fMRI. To model empirical data, it is thus imperative to introduce a measurement function M that maps the states \mathbf{x} into a multivariate, observable data set \mathbf{y} composed of measurable quantities y_j ,

$$\mathbf{y} = M(\mathbf{x}) + \boldsymbol{\epsilon}, \quad (6)$$

where $\boldsymbol{\epsilon}$ is a random variable [a vector the same length as \mathbf{y}] representing observation noise in each channel/voxel. While a static measurement function M might be valid for EEG, the observation equation for the BOLD signal is in reality also a dynamical system (19, 20), coupled passively to the neural state,

$$\frac{d\mathbf{y}}{dt} = M(\mathbf{x}) + \boldsymbol{\epsilon}. \quad (7)$$

The absence of a $\dot{\mathbf{y}}$ term on the RHS implies that the observables passively reflect the neural state, with filtering and superposition plus noise.

This mapping from the unknown (or hidden) states to the noisy observables is crucial for inferring the parameters of SDEs from empirical data (21). If Eq. (7) is linear, its solution can be expressed in terms of its impulse response function (e.g., the hemodynamic response function, HRF), enabling efficient computation of the BOLD in terms of a convolution of the neural signal with the HRF. This is a commonly-used method (22), though solving the full nonlinear hemodynamic equations allows richer dynamics in general (19, 23). More comprehensive models of the fMRI BOLD signal also include spatial terms enabling more realistic treatments of spatiotemporal cerebral hemodynamics (20, 24-26).

While in the “classic sense” the measurement function maps neuronal states into functional imaging data, it can also be conceptualized as mapping neuronal states into behaviors, such as the mapping from neocortical activity to changes in limb (27) or eye (28) movement through dendritic filters and musculoskeletal forward models.

2.2.2. *Attractors, bifurcations and multistability.* Figures 1B and 2B show the *attractors* of the spiking neuronal and neural mass model – the long-term solutions, following initial transients. As evident in the other panels, attractors continue to play a key organizing role of a system's solutions in the presence of stochastic inputs. An attractor is said to be structurally stable when its basic morphology is robust to changes in the underlying parameters (such as increasing or decreasing the excitatory self-feedback). If a small change suddenly changes the attractor – for example, from a fixed point to a limit cycle – then the attractor is structurally unstable: The sudden change in the attractor is called a *bifurcation* (Fig. 4A,B). While stochastic inputs are enslaved to the attractor when it is structurally stable, they grow in variance and correlation length close to the point of a bifurcation yielding slow, high amplitude fluctuations known as *critical slowing* (29). These fluctuations occur because the noise is able to drive larger and longer excursions from the attractor as it becomes progressively less stable near the bifurcation. These large fluctuations lead to long temporal autocorrelations (the reduced damping endows the dynamics with a longer "memory").

In many complex dynamical systems, the co-existence of two or more attractors is possible. When such *multistability* exists, noise of an appropriate amplitude is able to knock the dynamics from one attractor to the other and back, hence yielding itinerant, erratic time series that jump between the corresponding (periodic/chaotic) waveforms (Fig. 4C-F). If the noise is multiplicative, then the system typically becomes trapped for longer than expected on each of the attractors, giving heavy-tailed *dwelling times*: After transition to a low-amplitude attractor, the multiplicative noise term shrinks, decreasing the probability of an excursion sufficient to leave that attractor. Conversely, after a transition toward a high-amplitude attractor, the multiplicative noise becomes large, increasing the likelihood of a large excursion away from the origin. There is strong evidence that the dominant rhythm of the human cortex – the alpha rhythm – shows heavy-tailed bistability (11, 12).

2.2.3. *Probability densities and the Fokker-Planck equation*

The SDEs describing the temporal behavior of neuronal states yield noisy time series and phase space orbits. Visual inspection of lengthy solutions, however, shows that the solutions converge toward an underlying distribution (Figure 3A). That is, while the individual solutions are stochastic, a large ensemble of such solutions conforms to a fixed, deterministic probability density function (PDF) (Figure 3B). The SDE can be recast as a deterministic equation, known *the Fokker-Planck equation* (FPE), whose solution is the PDF of the solutions. Just as an SDE consists of deterministic and stochastic terms, the FPE broadly consists of drift (the deterministic flow) and diffusion (the dispersion of states) terms.

In principle, FPEs can be derived from SDEs using appropriate mathematical techniques (30). In practice, this is a non-trivial exercise unless the system has a slowly moving attractor and simple, additive fluctuations (31). Nonetheless, theoretical consideration of a FPE serves two useful functions: First, it allows us to distinguish stochastic model inversion schemes that deal directly with the noisy time series, such as dynamic expectation maximization (32) and Kalman filters, from those that maximize model likelihood using (the sufficient statistics of) stationary distributions of the data, like the spectral (33-36), or higher-order data features such as the parameters of heavy-tailed PDFs (37-39). Formal model inversion of SDEs is an important but complex problem that has been considered in detail elsewhere (21, 40).

Second, knowledge of the FPE representation is crucial for an important insight into stochastic neural models: relating ensemble dynamics to neural “coding”. Put simply, the proposal is to equate the ensemble mean with the likely value of some perceptual attribute being encoded by a population of neurons, while the variance captures the uncertainty of that representation. While there are alternative ways of describing how the brain encodes expected states of the world – and uncertainty about those states – equating of the distribution of states with the probabilistic representation of “belief” allows a bridge between stochastic models of neural dynamics and probabilistic neuronal codes in perception and cognition (41, 42). Cognitive processes thus map onto stochastic neuronal dynamics: As we review in Part II (4), it is reasonable to conjecture a mapping from unstable and noisy stochastic neuronal activity to cognitive disorganization (such as in schizophrenia), or conversely between excessively stable and recursive dynamics to obsessions and recurrent ruminations. Likewise, the precision of perceptual inference maps onto

the inverse variance of population dynamics: Imprecise beliefs reflect highly-variable states in the corresponding (coding) population, and sub-optimal belief updating corresponds to non-adaptive changes in the trial-to-trial expression of neuronal state variability.

Accordingly, SDMs become more than just a mechanistic model of observed data and behavior, but may ultimately provide a biophysical grounding for computational accounts of cognitive (dys)-function. What scale of description – microscopic spiking or mesoscopic mean field – will be the most suitable for this enterprise? Learning paradigms, such as reinforcement learning, typically emphasize the importance of spike timing and suggest a privileged role of stochastic models of spiking neurons. Probabilistic decision making schemes call upon the mutual shaping of ensemble probability distributions in cortical hierarchies and speak to mesoscopic population density approaches. We return to these fundamental issues in the Discussion and suggest a formal reconciliation.

3. Discussion

In sum, the fusion of dynamical systems theory with statistical physics provides a powerful and unifying framework for modeling the dynamics of neuronal systems across spatial and temporal scales. Noise interacts with the underlying dynamic flow to yield a variety of neural behaviors, which include slow and fast fluctuations, multistable switching, and simple attracting dynamics. Subtle disturbances in the balance of noise and order yield maladaptive dynamics, such as trapping, unstable fluctuations, and the erratic appearance of abnormal waveforms. While stochastic dynamical systems are clearly apt for describing neuronal processes, they may also have something profound to say about computational architectures and the neuronal code:

3.1. Integrating stochastic models of neuronal activity and cognition – the neural code

The nature of the neuronal code remains a fundamental issue in computational neuroscience. If the brain furnishes representations, expectations, predictions and beliefs, it has to perform some form of Bayesian inference (41), or more generally probabilistic encoding (42, 43). In other

words, neuronal activity must encode probability distributions (44, 45); though see (46). The two schools of thought on probabilistic coding can be broadly divided into *sampling schemes* and *sufficient statistics*. In sampling schemes, a probability distribution is encoded directly by the sample distribution over ensembles of neuronal activities, such that the most likely value of a belief is encoded in the ensemble mean and the precision or confidence associated with that belief is encoded by the dispersion of population activity. In machine learning, these approaches include importance sampling and particle filtering (47). Alternatively, encoding may be achieved with the moments of a probability distribution; for example, the mean and precision. Generally, this form of encoding considers a small number of sufficient statistics and an assumed form for the probability distribution (e.g., Gaussian) (48). In machine learning, this leads to approximate Bayesian inference and variational schemes such as Kalman filtering, predictive coding, and free energy minimization (49).

At first glance, these two encoding strategies appear fundamentally different: probabilistic sampling codes require large numbers of neurons but can encode probability distributions with arbitrary functional forms. In contrast, variational schemes call upon considerably fewer neurons (or populations) to approximate probabilistic beliefs but require a fixed form. But are they really so different? It can be shown that the variational schemes (and associated neuronal codes) are just a mean field description of underlying population dynamics (50). In other words, one can treat encoding in terms of sufficient statistics as a mean field description of a sampling scheme – such that both essentially describe the same thing. For example, the mean of a large neuronal population can be taken to encode the expectation of a probabilistic belief.

However, things get more interesting when we consider higher order moments: Are the second order moments of beliefs (i.e., inverse variance or *precision*) simply encoded by the corresponding dispersion of neuronal population activity, or by some other biophysical quantity that is updated in an approximately Bayes-optimal fashion? At first glance, one might argue that there is a fundamental difference between encoding uncertainty in terms of the variance of stochastic fluctuations in neuronal activity – as opposed to another neuronal state; e.g., activity in the classical neuromodulator systems (dopamine, acetylcholine, and serotonin) – or via synaptic gain.

In variational schemes such as predictive coding, precision is normally associated with the gain or postsynaptic excitability of units encoding prediction errors. This allows precise predictions to have more influence on other neuronal populations encoding expectations (that generate predictions). Precision is then tuned according to the influence of neuromodulatory inputs on local gain control (51). It might therefore appear that there is an irreconcilable difference between representing confidence in terms of synaptic gain, as opposed to the sample precision of a neuronal ensemble. However, a key insight here is that in the mean field reduction underlying neural mass models, the dispersion of the neuronal ensemble determines the gain of sigmoid activation functions relating neuronal firing to depolarization (52). In other words, if we consider the population density dynamics (i.e., the Fokker-Planck formulation) of neuronal message passing, where individual neurons have threshold on-off activation functions, then the population's activation function is the cumulative sample distribution of the population. This means that the dispersion of neuronal activity corresponds to the inverse postsynaptic gain at the population level.

This insight essentially dissolves the distinction between codes based on sampling schemes and sufficient statistics (at the level of density dynamics). This enables one to either study stochastic activity in neuronal populations, or describe their density dynamics deterministically, internally consistent with variational approaches (i.e., neuronal codes based upon sufficient statistics like predictive coding). Such considerations form a crucial bridge between stochastic models of neuronal dynamics and phenomenological models of uncertainty in cognition and behavior. They may also help reconcile computational accounts of how neuromodulatory systems adapt decision making according to perceived uncertainty with biophysical models of population activity that incorporate – for example – dopaminergic neurons (53) or neuromodulatory gain control (54). As a corollary, aberrant encoding of precision – as in schizophrenia – could be accounted for by SDMs (55, 56). More generally, if SDM provides a principled formalism for describing the neuronal encoding of beliefs, it may also provide a formal description of abnormal beliefs (i.e., false inference) in psychiatry. This concludes our primer on the theory of SDMs. In Part II (4), we review applications of SDMs to large-scale neural systems in health and disease.

Acknowledgements

This work was supported by the National Health and Medical Research Council (Program Grant 1037196), the Australian Research Council (Centre of Excellence for Integrative Brain Function CE140100007), and the Wellcome Trust (Wellcome Trust Principal Research Fellowship: 088130/Z/09/Z).

Financial disclosures

The authors declare no conflicts of interest and have no financial disclosures.

References:

1. Friston K (2013): Life as we know it. *Journal of the Royal Society Interface*. 10:20130475.
2. Tononi G, Sporns O, Edelman GM (1994): A measure for brain complexity: relating functional segregation and integration in the nervous system. *Proc Natl Acad Sci U S A*. 91:5033-5037.
3. Honey CJ, Kötter R, Breakspear M, Sporns O (2007): Network structure of cerebral cortex shapes functional connectivity on multiple time scales. *Proc Natl Acad Sci U S A*. 104:10240-10245.
4. Roberts JA, Friston KJ, Breakspear M (2016): Clinical applications of stochastic dynamic models of the brain, Part II: A review. (*submitted*).
5. Morris C, Lecar H (1981): Voltage oscillations in the barnacle giant muscle fiber. *Biophys J*. 35:193.
6. Haken H (1977): *Synergetics: an introduction*. Berlin: Springer.
7. Faisal AA, Selen LPJ, Wolpert DM (2008): Noise in the nervous system. *Nature Reviews Neuroscience*. 9:292-303.
8. Aburn MJ, Holmes CA, Roberts JA, Boonstra TW, Breakspear M (2012): Critical fluctuations in cortical models near instability. *Front Physiol*. 3:331.
9. Bekaert G, Wu G (2000): Asymmetric volatility and risk in equity markets. *Review of Financial Studies*. 13:1-42.
10. Anteneodo C, Riera R (2005): Additive-multiplicative stochastic models of financial mean-reverting processes. *Physical Review E*. 72:026106.
11. Freyer F, Aquino K, Robinson PA, Ritter P, Breakspear M (2009): Bistability and non-Gaussian fluctuations in spontaneous cortical activity. *J Neurosci*. 29:8512-8524.
12. Freyer F, Roberts JA, Becker R, Robinson PA, Ritter P, Breakspear M (2011): Biophysical mechanisms of multistability in resting-state cortical rhythms. *J Neurosci*. 31:6353-6361.
13. Weber EH (1834): *De Pulsu, resorptione, auditu et tactu: Annotationes anatomicae et physiologicae*. Leipzig: Koehler.
14. Harris CM, Wolpert DM (1998): Signal-dependent noise determines motor planning. *Nature*. 394:780-784.
15. Cordes S, Gelman R, Gallistel CR, Whalen J (2001): Variability signatures distinguish verbal from nonverbal counting for both large and small numbers. *Psychonomic Bulletin & Review*. 8:698-707.
16. Roberts JA, Boonstra TW, Breakspear M (2015): The heavy tail of the human brain. *Curr Opin Neurobiol*. 31:164-172.
17. Jansen BH, Rit VG (1995): Electroencephalogram and visual evoked potential generation in a mathematical model of coupled cortical columns. *Biol Cybern*. 73:357-366.
18. Breakspear M, Terry JR, Friston KJ (2003): Modulation of excitatory synaptic coupling facilitates synchronization and complex dynamics in a biophysical model of neuronal dynamics. *Network*. 14:703-732.
19. Friston KJ, Mechelli A, Turner R, Price CJ (2000): Nonlinear responses in fMRI: the Balloon model, Volterra kernels, and other hemodynamics. *Neuroimage*. 12:466-477.

20. Aquino KM, Robinson PA, Drysdale PM (2014): Spatiotemporal hemodynamic response functions derived from physiology. *J Theor Biol.* 347:118-136.
21. Li B, Daunizeau J, Stephan KE, Penny W, Hu D, Friston K (2011): Generalised filtering and stochastic DCM for fMRI. *Neuroimage.* 58:442-457.
22. Gitelman DR, Penny WD, Ashburner J, Friston KJ (2003): Modeling regional and psychophysiological interactions in fMRI: the importance of hemodynamic deconvolution. *Neuroimage.* 19:200-207.
23. Drysdale PM, Huber JP, Robinson PA, Aquino KM (2010): Spatiotemporal BOLD dynamics from a poroelastic hemodynamic model. *J Theor Biol.* 265:524-534.
24. Aquino KM, Schira MM, Robinson PA, Drysdale PM, Breakspear M (2012): Hemodynamic Traveling Waves in Human Visual Cortex. *PLoS Comput Biol.* 8:e1002435.
25. Aquino KM, Robinson PA, Schira MM, Breakspear M (2014): Deconvolution of neural dynamics from fMRI data using a spatiotemporal hemodynamic response function. *Neuroimage.* 94:203-215.
26. Kriegeskorte N, Cusack R, Bandettini P (2010): How does an fMRI voxel sample the neuronal activity pattern: compact-kernel or complex spatiotemporal filter? *Neuroimage.* 49:1965-1976.
27. Heitmann S, Boonstra T, Gong P, Breakspear M, Ermentrout B (2015): The rhythms of steady posture: Motor commands as spatially organized oscillation patterns. *Neurocomputing.* 170:3-14.
28. Friston K, Breakspear M, Deco G (2012): Perception and self-organized instability. *Front Comput Neurosci.* 6:44.
29. Kelso JAS, Scholz JP, Schöner G (1986): Nonequilibrium phase transitions in coordinated biological motion: Critical fluctuations. *Physics Letters A.* 118:279-284.
30. Risken H (1989): *The Fokker-Planck Equation.* 2nd ed. Berlin: Springer.
31. Deco G, Jirsa VK, Robinson PA, Breakspear M, Friston K (2008): The dynamic brain: from spiking neurons to neural masses and cortical fields. *PLoS Comput Biol.* 4:e1000092.
32. Friston KJ, Trujillo-Barreto N, Daunizeau J (2008): DEM: a variational treatment of dynamic systems. *Neuroimage.* 41:849-885.
33. Robinson PA, Rennie CJ, Rowe DL, O'Connor SC (2004): Estimation of multiscale neurophysiologic parameters by electroencephalographic means. *Hum Brain Mapp.* 23:53-72.
34. Robinson PA, Rennie CJ, Wright JJ, Bahramali H, Gordon E, Rowe DL (2001): Prediction of electroencephalographic spectra from neurophysiology. *Physical Review E.* 63:021903.
35. Van Albada SJ, Kerr CC, Chiang AKI, Rennie CJ, Robinson PA (2010): Neurophysiological changes with age probed by inverse modeling of EEG spectra. *Clin Neurophysiol.* 121:21-38.
36. Moran R, Kiebel S, Stephan K, Reilly R, Daunizeau J, Friston K (2007): A neural mass model of spectral responses in electrophysiology. *Neuroimage.* 37:706-720.
37. Chapman SC, Rowlands G, Watkins NW (2002): Extremum statistics: a framework for data analysis. *Nonlinear Processes in Geophysics.* 9:409-418.
38. Freyer F, Roberts JA, Ritter P, Breakspear M (2012): A canonical model of multistability and scale-invariance in biological systems. *PLoS Comput Biol.* 8:e1002634.
39. Roberts JA, Iyer KK, Finnigan S, Vanhatalo S, Breakspear M (2014): Scale-Free Bursting in Human Cortex following Hypoxia at Birth. *J Neurosci.* 34:6557-6572.

40. Friston K, Mattout J, Trujillo-Barreto N, Ashburner J, Penny W (2007): Variational free energy and the Laplace approximation. *Neuroimage*. 34:220-234.
41. Knill DC, Pouget A (2004): The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends Neurosci*. 27:712-719.
42. Ma WJ, Beck JM, Latham PE, Pouget A (2006): Bayesian inference with probabilistic population codes. *Nat Neurosci*. 9:1432-1438.
43. Dayan P, Hinton GE, Neal RM, Zemel RS (1995): The Helmholtz machine. *Neural Comput*. 7:889-904.
44. Sanger TD (1996): Probability density estimation for the interpretation of neural population codes. *J Neurophysiol*. 76:2790-2793.
45. Zemel RS, Dayan P, Pouget A (1998): Probabilistic interpretation of population codes. *Neural Comput*. 10:403-430.
46. Sanborn AN, Chater N (2016): Bayesian brains without probabilities. *Trends in Cognitive Sciences*.
47. Shi L, Griffiths TL (2009): Neural implementation of hierarchical Bayesian inference by importance sampling. *Adv Neural Inf Process Syst*, pp 1669-1677.
48. Sahani M, Dayan P (2003): Doubly distributional population codes: simultaneous representation of uncertainty and multiplicity. *Neural Comput*. 15:2255-2279.
49. Friston KJ, Daunizeau J, Kilner J, Kiebel SJ (2010): Action and behavior: a free-energy formulation. *Biol Cybern*. 102:227-260.
50. Friston K (2009): The free-energy principle: a rough guide to the brain? *Trends in Cognitive Sciences*. 13:293-301.
51. Yu AJ, Dayan P (2005): Uncertainty, neuromodulation, and attention. *Neuron*. 46:681-692.
52. Marreiros AC, Daunizeau J, Kiebel SJ, Friston KJ (2008): Population dynamics: variance and the sigmoid activation function. *Neuroimage*. 42:147-157.
53. Schwartenbeck P, FitzGerald TH, Mathys C, Dolan R, Friston K (2014): The dopaminergic midbrain encodes the expected certainty about desired outcomes. *Cereb Cortex*.bhu159.
54. Li S-C, von Oertzen T, Lindenberger U (2006): A neurocomputational model of stochastic resonance and aging. *Neurocomputing*. 69:1553-1560.
55. Adams RA, Stephan KE, Brown HR, Frith CD, Friston KJ (2013): The Computational Anatomy of Psychosis. *Frontiers in Psychiatry*. 4:47.
56. Stephan KE, Binder EB, Breakspear M, Dayan P, Johnstone EC, Meyer-Lindenberg A, et al. (2016): Charting the landscape of priority problems in psychiatry, part 2: pathogenesis and aetiology. *The Lancet Psychiatry*. 3:84-90.
57. Lopes da Silva F, Blanes W, Kalitzin SN, Parra J, Suffczynski P, Velis DN (2003): Epilepsies as dynamical diseases of brain systems: basic models of the transition between normal and epileptic activity. *Epilepsia*. 44:72-83.
58. Benjamin O, Fitzgerald THB, Ashwin P, Tsaneva-Atanasova K, Chowdhury F, Richardson MP, et al. (2012): A phenomenological model of seizure initiation suggests network structure may explain seizure frequency in idiopathic generalised epilepsy. *The Journal of Mathematical Neuroscience*. 2:1-30.
59. Suffczynski P, Kalitzin S, Lopes Da Silva FH (2004): Dynamics of non-convulsive epileptic phenomena modeled by a bistable neuronal network. *Neuroscience*. 126:467-484.

Figure captions

Figure 1: Example dynamics from the Morris-Lecar model. Left column: Time series for the membrane potential V . Right column: phase portraits in the (n, V) plane. **A,B:** Deterministic case (no noise). **C,D:** Additive noise. **E,F:** Multiplicative noise.

Figure 2: Example dynamics from a neural mass model. Left column: Time series for the membrane potential v . Right column: phase portraits in the (V, W, Z) state space (see Appendix A, Supplemental Material). **A,B:** Deterministic case (no noise). **C,D:** Additive noise. **E,F:** Multiplicative noise.

Figure 3: Density of trajectories from the Morris-Lecar model with multiplicative noise. **A:** Phase portrait; same as Fig. 1F but for a long (7000 ms) run. **B:** Probability density function for the trajectory in panel A, warm colors denote higher densities. Note the relatively smooth nature of the resulting stationary distribution, despite the randomness of the trajectories.

Figure 4: Noise-induced transitions in a simple bistable model (the Hopf-bifurcation model used in Ref. (38); similar ideas have been explored in biophysical seizure models (57-59)). See Appendix A for further details of the model and definitions of technical terms. **A:** Bifurcation diagram for a supercritical Hopf bifurcation. For $\beta < 0$, the only attractor is a stable fixed point ($r = 0$, black solid line). For $\beta > 0$, the fixed point is unstable (black dotted line), and a stable limit cycle (green) emerges with radius $r > 0$. The green line depicts the radius of the limit cycle for the corresponding value of β . **B:** Bifurcation diagram for a subcritical Hopf bifurcation and bistability. The $r = 0$ fixed point behaves the same as in panel A. The difference here is that the limit cycle that emerges at $\beta = 0$ exists for $\beta < 0$ and is unstable (red). This unstable limit cycle then meets a stable limit cycle (green) at a “saddle-node” bifurcation of periodic orbits. There is thus a region of bistability where the fixed point (black) and the upper limit cycle (green) are

both stable, separated by an unstable limit cycle (red). **C-F:** Left column: time series for one of the two Cartesian coordinates in the phase space. Right column: trajectories in phase space (blue) with stable attractor (green) and unstable "separatrix" (red) that demarcates the boundary across which a sudden transition (e.g., seizure onset) occurs. The polar radial coordinate r in panels A and B is given by $r^2 = x^2 + y^2$. **C,D:** The low-amplitude case: the trajectory remains well within the transition boundary, exhibiting only low-amplitude fluctuations (e.g., healthy non-seizure dynamics). **E,F:** The high-amplitude case: random fluctuations cross the boundary triggering high-amplitude oscillations (e.g., a seizure).