synchronized Up-and-Down states[3,8,9]. The findings of Hahn *et al.*[4] therefore suggest that both neocortex and MEC can act as rhythm generator circuits for cortical slow oscillations. The mechanism underlying the synchronization of cortical slow oscillations has been an important question in the field of network physiology[3,10]. One idea is that transitions between Up and Down states propagate as traveling waves across partially coupled oscillators. However, the work by Hahn *et al.*[4] argues against this scenario, as it suggests that the MEC generator can gain independence from the neocortical generators by skipping Down states. It is also possible that other cortical generators may show similar decoupling, thereby making the temporal dynamics of slow oscillations more complex.

Whether the decoupling effect of the MEC on cortical slow oscillations is restricted to the hippocampus or occurs elsewhere remains unknown. Such MEC-related decoupling may also extend to other cortical structures, such as the postrhinal cortex. But what is the functional role of such a release? And how can a hippocampal dialog during neocortical Down states be beneficial? Many MEC neurons exhibit firing patterns that are spatially tuned to the location of the animal in its environment, whereas most neurons in the LEC display a weak spatial selectivity that indicates the possible influence of nonspatial sensory inputs[11–13]. Because of this, it has been suggested that MEC-postrhinal-retrosplenial cortices may represent a spatial information stream, whereas the LEC-perirhinal-anterior cingulate cortices may represent a nonspatial information pathway[14,15]. Thus, during active waking behavior, CA1 hippocampal neurons may encode a combination of spatial and nonspatial information through the integration of MEC and LEC inputs, respectively. However, during off-line periods, at times when memory traces are thought to be consolidated, the temporal dynamics between the LEC Up-Down states and the MEC persistent Up states might differentially channel the reactivation of spatial and nonspatial memory traces.

Hippocampal sharp wave ripple (150–250 Hz) events have been implicated in the reactivation of waking firing patterns and memory consolidation[5]. In light of the study from Hahn *et al.*[4], hippocampal sharp wave ripple events during neocortical Down states might favor the reactivation of spatial information, whereas those in the Up states might act to bind space and events. This would provide a mechanism that emphasizes spatial information at times when the hippocampus is decoupled from the neocortex and allows the selective communication of cortical areas involved in the spatial information stream. Such regions could include the postrhinal and the retrosplenial cortices, if these regions prove similarly capable of decoupling from neocortical Down states via MEC-moderated influence. In contrast, sharp wave ripple events that occur during neocortical Up states could enable the integrated binding of spatial and nonspatial information by simultaneously recruiting not only the spatial stream but also the nonspatial stream. This hypothesis is but one example of how the work of Hahn *et al.*[4] could stimulate pertinent follow-up work to test the influence of MEC Down-state skipping on other cortical areas and eventually reveal its broader function. Further work could also investigate whether reactivation patterns are different during neocortical Up and Down states to test whether the possible binding of spatial and nonspatial information occurs in the hippocampus or only in higher cortical areas.

COMPETING FINANCIAL INTERESTS
The authors declare no competing financial interests.

1. Buzsáki, G. *Cereb. Cortex* **6**, 81–92 (1996).
2. McClelland, J.L., McNaughton, B.L. & O'Reilly, R.C. *Psychol. Rev.* **102**, 419–457 (1995).
3. Steriade, M. & Timofeev, I. *Neuron* **37**, 563–576 (2003).
4. Hahn, T.T.G., McFarland, J.M., Berberich, S., Sakmann, B. & Mehta, M.R. *Nat. Neurosci.* **15**, 1531–1538 (2012).
5. O'Neill, J., Pleydell-Bouverie, B., Dupret, D. & Csicsvari, J. *Trends Neurosci.* **33**, 220–229 (2010).
6. Rasch, B. & Born, J. *Curr. Opin. Neurobiol.* **17**, 698–703 (2007).
7. Giocomo, L.M., Zilli, E.A., Fransen, E. & Hasselmo, M.E. *Science* **315**, 1719–1722 (2007).
8. Cunningham, M.O. *et al. Proc. Natl. Acad. Sci. USA* **103**, 5597–5601 (2006).
9. Mann, E.O., Kohl, M.M. & Paulsen, O. *J. Neurosci.* **29**, 7513–7518 (2009).
10. Crunelli, V. & Hughes, S.W. *Nat. Neurosci.* **13**, 9–17 (2010).
11. Fyhn, M., Molden, S., Witter, M.P., Moser, E.I. & Moser, M.B. *Science* **305**, 1258–1264 (2004).
12. Hafting, T., Fyhn, M., Molden, S., Moser, M.B. & Moser, E.I. *Nature* **436**, 801–806 (2005).
13. Hargreaves, E.L., Rao, G., Lee, I. & Knierim, J.J. *Science* **308**, 1792–1794 (2005).
14. Witter, M.P. & Moser, E.I. *Trends Neurosci.* **29**, 671–678 (2006).
15. van Strien, N.M., Cappaert, N.L. & Witter, M.P. *Nat. Rev. Neurosci.* **10**, 272–282 (2009).

Two layers of neural variability

Mark M Churchland & L F Abbott

Variability in neuronal firing rates and spike timing can be modeled as doubly stochastic. A study now suggests that these phenomena could arise from a network built of deterministic neurons with balanced excitation and inhibition.

The principal mode of neural communication is the action potential: a stereotyped spike in voltage across the membrane of a neuron. Neural circuits represent and process information through the temporal and spatial patterns of their spikes. However, these patterns of spikes are surprisingly variable, even across trials in which an animal is experiencing identical stimuli or performing a nominally identical action. A common practice is to treat such variability as arising from the stochastic (Poisson) generation of spikes on the basis of a firing rate. There is debate regarding the degree to which deviations from stochastic spiking (for example, synchrony and spike-timing effects) encode additional information beyond the rate. Putting that debate aside, it is widely agreed that the rate carries a good deal of information and that much of the recorded spiking variability is effectively noise. Many of our most successful models and hypotheses are couched entirely in terms of firing rates. Notably, such models often allow firing rates themselves to be variable across trials (for example, this might be necessary to explain behavioral variability). Thus, the observed data is often modeled as 'doubly stochastic'[1]: a variable firing rate gives rise to variable spiking[2] (**Fig. 1**). Although such models provide good descriptions of the data, they seem to be at variance with the known biophysics of neurons. Nowhere inside a neuron is there a biophysical quantity that corresponds to a

M.M. Churchland is in the Department of Neuroscience, Grossman Center for the Statistics of Mind, David Mahoney Center for Brain and Behavior Research, Kavli Institute for Brain Science, Columbia University Medical Center, New York, New York, USA. L.F. Abbott is in the Departments of Neuroscience and Physiology and Cellular Biophysics, Columbia University Medical Center, New York, New York, USA.
e-mail: markchurchland@columbia.edu

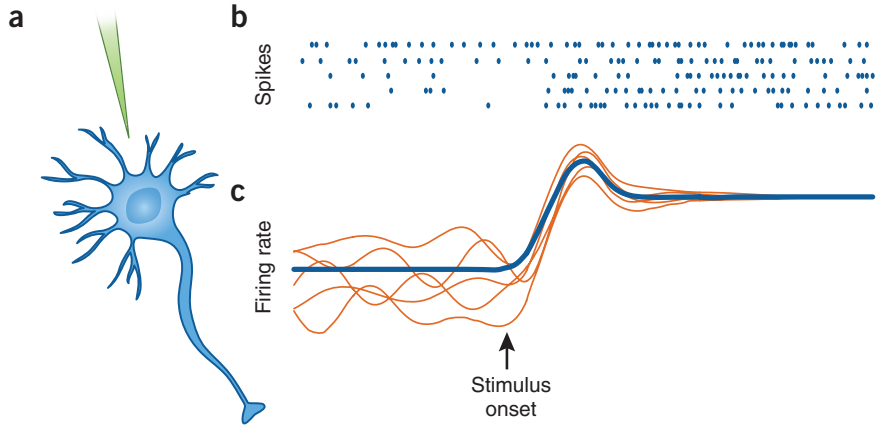

Figure 1 Schematic showing a typical presentation of extracellularly recorded data. (**a**) A recording electrode detects the occurrence of individual spikes from a neuron. (**b**) Spike times are plotted as a raster, with one tick per spike. The same neuron may be recorded for multiple repeats (trials) of the same stimulus, with each response plotted in one row. In this illustration, five trials are shown. Spiking appears to be stochastic, as each trial involves a different pattern of spikes. (**c**) Spikes are often said to occur with a frequency determined by an underlying firing rate (blue trace). If so, the firing rate can be estimated by filtering the data and averaging across trials. However, it is typically further assumed that the underlying firing rate may itself be variable across trials (orange traces). Thus, many models assume that the recorded data are doubly stochastic, with variable rates giving rise to variable spiking. A commonly observed effect is that the onset of a stimulus reduces firing-rate variability, but not spiking variability.

firing rate nor is there a random spike generator that might turn that rate into variable spikes. In this issue, Litwin-Kumar and Doiron[3] take a step toward resolving this apparent paradox by presenting a deterministic network model in which spiking activity appears as if it were doubly stochastic.

Litwin-Kumar and Doiron[3] show that a simulated network of spiking neurons can exhibit, as an emergent network property, both spontaneous firing rate fluctuations and spontaneous spiking variability. The simulated neurons are deterministic in that they do not employ a random spike generator nor are they directly endowed with an underlying firing rate. However, when placed in the proposed network architecture, these simple neurons act somewhat like real neurons and exhibit highly variable spiking at a rate that fluctuates over time. Consistent with experimental observations, spiking variability is 'private': one spike from one neuron implies little regarding whether another neuron will spike at the same instant[4]. In contrast, spontaneous firing-rate fluctuations are 'shared': groups of neurons collectively exhibit periods of high-frequency firing.

There is no true randomness in the network. The model is actually 'doubly chaotic' rather than doubly stochastic. But because chaotic dynamics are in practice unpredictable, they appear to be stochastic. The doubly chaotic character of the model is its novel feature. Prior network models constructed at the level of firing rates (without any spiking) have been shown to produce chaotic firing-rate fluctuations[5]. Similarly, simulated networks of spiking neurons have been shown to exhibit chaotic spiking dynamics[4,6–9]. The key advance in the network constructed by Litwin-Kumar and Doiron[3] is that it generates chaotic spiking along with irregular changes in firing rate. To do so, their network employs the principle of balanced excitation and inhibition.

Chaotic firing-rate and chaotic spiking network models typically rely on a balance of strong excitation and strong inhibition to generate highly variable activity. In a number of prior spiking networks[4,6–9], each neuron is bombarded by large amounts of excitation and inhibition that mostly cancel each other out, leaving behind fluctuations that occasionally and irregularly drive the neuron across the threshold for spiking[10]. Although spiking is deterministically chaotic, it appears to be roughly Poisson. A recent examination of synaptic strength argues that cortical networks are likely in a balanced regime[11]. The model of Litwin-Kumar and Doiron[3] starts with a balanced network of randomly connected spiking neurons and adds clustered connections to induce firing-rate fluctuations. Strong clustering is a brute-force way of causing a group of simulated spiking neurons to act like a firing rate unit (the average spiking of the whole cluster simply becomes the rate). The surprising feature of the Litwin-Kumar and Doiron[3] network is that firing-rate fluctuations emerge even with only modest clustering—when ~3% of the connections in an initially random network are rearranged. Although this is a small change from randomness at the level of the network, from the point of view of a given cluster, this modification is sufficient to produce a fairly large effect; the total strength of within-cluster excitation increases almost fivefold. This allows the neurons in each cluster to act synergistically, almost as if they formed one large 'rate unit'. Fluctuations in firing rate occur because the clusters are in competition: each cluster excites itself and inhibits the others. When a given cluster becomes active, it tends to stay active and to suppress the others. The active cluster therefore has an elevated overall firing rate. A little later, a different cluster may win the competition and become the most active. This competition produces extended fluctuations in firing rate. These rate fluctuations coexist with chaotic spiking that appears to be roughly Poisson. Indeed, variability in spiking drives the firing-rate transitions.

This network also reproduces a robust feature of the physiological data: the onset of a stimulus reduces firing-rate fluctuations while having little or no effect on spiking variability[2]. Previous simulations have addressed the possible sources of this effect[2,12,13], but Litwin-Kumar and Doiron[3] are the first to show that it exists in a deterministic network that internally generates both firing-rate and spiking variability. The reduction in firing-rate variability can be easily understood. If the stimulus drives clusters differentially, it resolves the competition among them, suppressing firing-rate variability and leaving only spiking variability.

The network developed by Litwin-Kumar and Doiron[3] elegantly illustrates how a single deterministic model can appear to be doubly stochastic, producing both firing-rate fluctuations and spiking variability. In doing so, it provides a solid conceptual framework that clarifies what we mean by terms such as firing-rate variability, spiking noise and underlying firing rate. That the network achieves its properties with only a modest deviation from random connectivity has important implications. For example, whether two neurons are connected says little about whether they share the same firing rate. Still, it should be stressed that the particular deviation from random connectivity (clustering) used in this study produces a very limited form of firing-rate dynamics: a competing set of attractor states that generate essentially two levels of firing, high and low. Those dynamics are not intended to compute anything or to perform any particular task; they are simply intended to emulate basic features of real data.

This class of model provides a potential explanation for features of neural data that we often take for granted and provides a theoretical justification for many standard analyses. For example, the single most

common analysis in systems neuroscience is to estimate the firing rate of a neuron by averaging the response across many stimulus repetitions. This analysis is appropriate if the underlying network supports both a time-evolving rate (which is effectively signal and is recovered by averaging) and spiking variability (which is effectively noise and is suppressed by averaging). Furthermore, experimentalists often attempt to segregate response variability into firing-rate variability and spiking noise[1,2]. It is generally assumed that the former can influence behavior[14] or perception[15], whereas the latter effectively acts as measurement noise. These are indeed fair assumptions for a network of the class simulated by Litwin-Kumar and Doiron[3]. Firing-rate fluctuations are shared among neurons (and can affect behavior), whereas spiking variability is local in time and mostly uncorrelated between neurons[4] (and would have difficulty affecting behavior).

A number of outstanding questions remain. Can the proposed class of spiking network instantiate the full range of dynamics that can be built from rate-based units, including dynamics that involve a continuous evolution of the rate (for example, integrators, limit cycles and chaotic attractors)? If so, should our models and hypotheses dispense with the nuisance of spiking variability and simulate dynamics directly at the level of rates? Or might structured details of spiking influence the evolution of the network at the level of rates, as has often been suggested? Such questions underscore that our field is not yet completely certain which aspects of a spike train are signal and which are noise. Still, we seem to have come to a fairly solid understanding that not every spike is sacred and that the dynamics of many networks are best described at the abstract level of rates.

COMPETING FINANCIAL INTERESTS
The authors declare no competing financial interests.

1. Churchland, A.K. *et al. Neuron* **69**, 818–831 (2011).
2. Churchland, M.M. *et al. Nat. Neurosci.* **13**, 369–378 (2010).
3. Litwin-Kumar, A. & Doiron, B. *Nat. Neurosci.* **15**, 1498–1505 (2012).
4. Renart, A. *et al. Science* **327**, 587–590 (2010).
5. Sompolinsky, H., Crisanti, A. & Sommers, H.J. *Phys. Rev. Lett.* **61**, 259–262 (1988).
6. Shadlen, M.N. & Newsome, W.T. *Curr. Opin. Neurobiol.* **4**, 569–579 (1994).
7. van Vreeswijk, C. & Sompolinsky, H. *Science* **274**, 1724–1726 (1996).
8. Vogels, T.P. & Abbott, L.F. *J. Neurosci.* **25**, 10786–10795 (2005).
9. Troyer, T.W. & Miller, K.D. *Neural Comput.* **9**, 971–983 (1997).
10. Carandini, M. *PLoS Biol.* **2**, e264 (2004).
11. London, M., Roth, A., Beeren, L., Hausser, M. & Latham, P.E. *Nature* **466**, 123–127 (2010).
12. Rajan, K., Abbott, L. & Sompolinsky, H. *Phys. Rev. E* **82**, (2010).
13. Deco, G. & Hugues, E. *PLOS Comput. Biol.* **8**, e1002395 (2012).
14. Churchland, M.M., Afshar, A. & Shenoy, K.V. *Neuron* **52**, 1085–1096 (2006).
15. Britten, K.H., Newsome, W.T., Shadlen, M.N., Celebrini, S. & Movshon, J.A. *Vis. Neurosci.* **13**, 87–100 (1996).

Cdk5 keeps memory on Trk

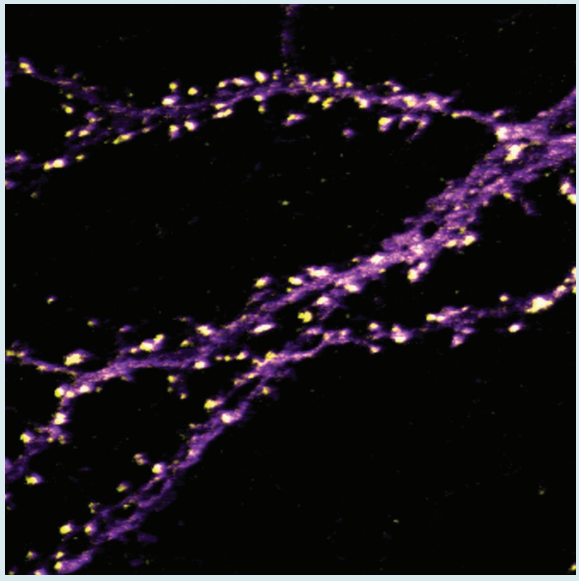

The structural and functional plasticity of synaptic connections is thought to contribute to the formation of new memories in the brain. There are a number of molecules that participate in the establishment of synaptic plasticity and in memory formation, including brain-derived neurotrophic factor (BDNF). Its receptor, TrkB, is a receptor tyrosine kinase that can, when activated, trigger autophosphorylation of its tyrosine residues. However, several other kinases can also phosphorylate this receptor on serine residues, and whether these forms of phosphorylation are actually critical for synaptic plasticity and memory remains an open question. On page 1506 of this issue, Lai and colleagues provide some important answers to this question by looking specifically at the role of the proline-directed serine and threonine kinase Cdk5 in the activation of TrkB and the subsequent effects on synaptic plasticity and memory.

The authors previously found that Cdk5 can phosphorylate TrkB on serine 478 to promote dendritic growth. To further determine the molecular mechanisms acting downstream of this signaling branch of the TrkB receptor and uncover the physiological relevance of S478 phosphorylation by Cdk5, Lai and colleagues generated a knock-in mouse in which TrkB was swapped with a phosphorylation-deficient S478A mutant TrkB receptor. They found that TrkB phosphorylated on its serine 478 colocalizes with PSD-95 in dendritic spines in hippocampal neurons in cultures (see image) and that this phosphorylation is crucial for BDNF-induced spine morphogenesis and glutamate-induced spine growth. Furthermore, S478 phosphorylation promoted the interaction of TrkB with the Rac-GEF TIAM1 and the activation of the Rac-PAK pathway, thereby linking phosphoS478-dependent signaling to the machinery that modulates actin dynamics and mediates spine remodeling. Crucially, S478 phosphorylation impaired CA3-CA1 long-term potentiation in hippocampal slices and S478A mutant mice exhibited severe deficits in spatial memory and novel object recognition tasks. However, their ability to learn contextual fear was intact.

The study by Lai and colleagues provides some new insights into to how Cdk5 and BDNF contribute to synaptic plasticity, learning and memory, and presents compelling evidence of the importance of serine phosphorylation of a receptor tyrosine kinase in this process.

Sebastien Thuault