



Balanced neural architecture and the idling brain

Brent Doiron^{1,2*} and Ashok Litwin-Kumar^{2,3}

¹ Department of Mathematics, University of Pittsburgh, Pittsburgh, PA, USA

² Center for the Neural Basis of Cognition, University of Pittsburgh and Carnegie Mellon University, Pittsburgh, PA, USA

³ Program for Neural Computation, University of Pittsburgh and Carnegie Mellon University, Pittsburgh, PA, USA

Edited by:

Arnulf Graf, California Institute of Technology, USA

Reviewed by:

Jaime de la Rocha, Institut D'Investigacions Biomèdiques August Pi i Sunyer, Spain
Jorge F. Mejias, New York University, USA

*Correspondence:

Brent Doiron, Department of Mathematics, University of Pittsburgh, Thackeray Hall, Pittsburgh, PA 15260, USA
e-mail: bdoiron@pitt.edu

A signature feature of cortical spike trains is their trial-to-trial variability. This variability is large in the spontaneous state and is reduced when cortex is driven by a stimulus or task. Models of recurrent cortical networks with unstructured, yet balanced, excitation and inhibition generate variability consistent with evoked conditions. However, these models produce spike trains which lack the long timescale fluctuations and large variability exhibited during spontaneous cortical dynamics. We propose that global network architectures which support a large number of stable states (attractor networks) allow balanced networks to capture key features of neural variability in both spontaneous and evoked conditions. We illustrate this using balanced spiking networks with clustered assembly, feedforward chain, and ring structures. By assuming that global network structure is related to stimulus preference, we show that signal correlations are related to the magnitude of correlations in the spontaneous state. Finally, we contrast the impact of stimulation on the trial-to-trial variability in attractor networks with that of strongly coupled spiking networks with chaotic firing rate instabilities, recently investigated by Ostojic (2014). We find that only attractor networks replicate an experimentally observed stimulus-induced quenching of trial-to-trial variability. In total, the comparison of the trial-variable dynamics of single neurons or neuron pairs during spontaneous and evoked activity can be a window into the global structure of balanced cortical networks.

Keywords: balanced cortical networks, neural variability, spontaneous cortical activity, cortical circuits, spiking models

1. INTRODUCTION

The rich structure of neural firing patterns provides ample challenges for the systems neuroscience community. The complexity of *in vivo* neural responses depends, in part, on detailed single neuron biophysics, synaptic dynamics, and network interactions. To make sense of such complexity, it is often necessary to treat brain activity as probabilistic, similar to how statistical physics treats large ensembles of particles. In this spirit, collecting the responses of a neuron (or a population of neurons) over many trials of an experiment permits a statistical characterization of neural activity. A widespread observation is that the trial averaged response of a neuron's spike output is stimulus or action tuned (Decharms and Zador, 2000), and this has grounded many theories of neural coding (Dayan and Abbott, 2001). However, only considering the trial averaged response of a neuron glosses over many complex response dynamics that may give further insight into neural function.

Large trial-to-trial variability of the neural response is a general characteristic of cortical dynamics (Britten et al., 1992; Shadlen and Newsome, 1994; Averbeck et al., 2006; Cohen and Kohn, 2011; Ponce-Alvarez et al., 2013). In stimulus or task evoked states, spiking responses are well modeled as an inhomogeneous Poisson process, whose firing rate dynamics can be obtained from trial averaged responses. In this case, the response variability is equal to the trial averaged evoked response, and little is learned about cortical processing by considering neural variability. Neural

activity during spontaneous dynamics (dynamics in the absence of a driving stimulus or action) is more complex (Ringach, 2009). In particular, the trial-to-trial spike variability during spontaneous states is larger than that predicted by a Poisson process (Churchland et al., 2006, 2010), and exhibits fluctuations over a range of timescales (Smith and Kohn, 2008). The state of cortex during spontaneous conditions is both an influence (Arieli et al., 1996; Fukushima et al., 2012) and a reflection of stimulus evoked activity (Tsodyks et al., 1999; Luczak et al., 2009; Luczak and MacLean, 2012). Thus, understanding the mechanics behind spontaneous activity may give new insight into the structure and function of cortical circuits.

Recent theoretical studies have shown that cortical models with balanced excitation and inhibition and clustered excitatory connectivity capture the trial-to-trial variability of both spontaneous and evoked conditions (Deco and Hugues, 2012; Litwin-Kumar and Doiron, 2012). In this article, we present a unified view of spontaneous dynamics in a variety of neural architectures. The core mechanic is that networks with a large number of stable states allow spontaneous fluctuations to stochastically "sample" the various states (Goldberg et al., 2004), leading to high spiking variability in spontaneous conditions. For stimuli to quench spiking variability, such models require that the recurrent architecture that supports network metastability be coherent with feedforward afferent projections. Under this assumption, we predict that pairwise spiking correlations as measured in the spontaneous state

will be related to signal correlations in evoked states. Our article extends the framework of balanced excitation and inhibition (van Vreeswijk and Sompolinsky, 1996, 1998; Renart et al., 2010) to more structured network architectures, leading to a theory of variability that is consistent with spontaneously active cortex and its relation to the evoked cortical response.

2. RESULTS

2.1. NEURAL VARIABILITY IN BALANCED EXCITATORY AND INHIBITORY NETWORKS

Cortical circuits have recurrent excitatory and inhibitory connections that are sparse and random (Figure 1A). As a first model, we consider no underlying global cortical architecture and rather assume wiring that obeys a pairwise independent model (meaning that a single connection between two neurons occurs with a probability that is independent of the rest of the network). Despite sparse wiring, the high number of cortical neurons ensures that a postsynaptic cell receives a large number ($\sim 10^3$) of synaptic inputs from other cortical neurons (Binzegger et al., 2004). Furthermore, the membrane potential deflection from an individual excitatory or inhibitory input can be on the order of 1 mV (Lefort et al., 2009). The convergence of a high number of large amplitude excitatory or inhibitory inputs could overwhelm a postsynaptic cell. However, in practice, the firing rates of cortical neurons are neither negligible nor driven to saturation (Hromádka et al., 2008). To explain this seeming contradiction between anatomy and physiology, Shadlen and Newsome (1994, 1998) proposed that large excitation and inhibition effectively balance one another (on average), so that the net mean input is small and output rates can be moderate (Figure 1B). This conjecture has experimental support from *in vivo* whole cell recordings from cortical neurons, in which the magnitudes of isolated excitatory and inhibitory inputs are large yet balanced when neurons are near their resting state (Haider et al., 2006).

A key feature of balanced cortical networks is that spiking activity is dominated by the *fluctuations* in synaptic input

(Shadlen and Newsome, 1994). These fluctuations can be sizable, since synaptic strengths are relatively large (only a small fraction of afferent inputs are necessary to drive the voltage to its threshold value). Synaptic input fluctuations provide a mechanism for the large dynamic and trial-to-trial variability of spiking activity that is characteristic of cortex (Figure 1B, top). Using techniques from statistical physics, van Vreeswijk and Sompolinsky (1996, 1998) provided a formal link between balance in recurrent networks and spiking variability. Specifically, they considered the statistics of the synaptic input into a representative cell in the network that received on average K inputs from recurrent excitatory and inhibitory wiring. With minimal assumptions, the mean M and variance V of the input obey:

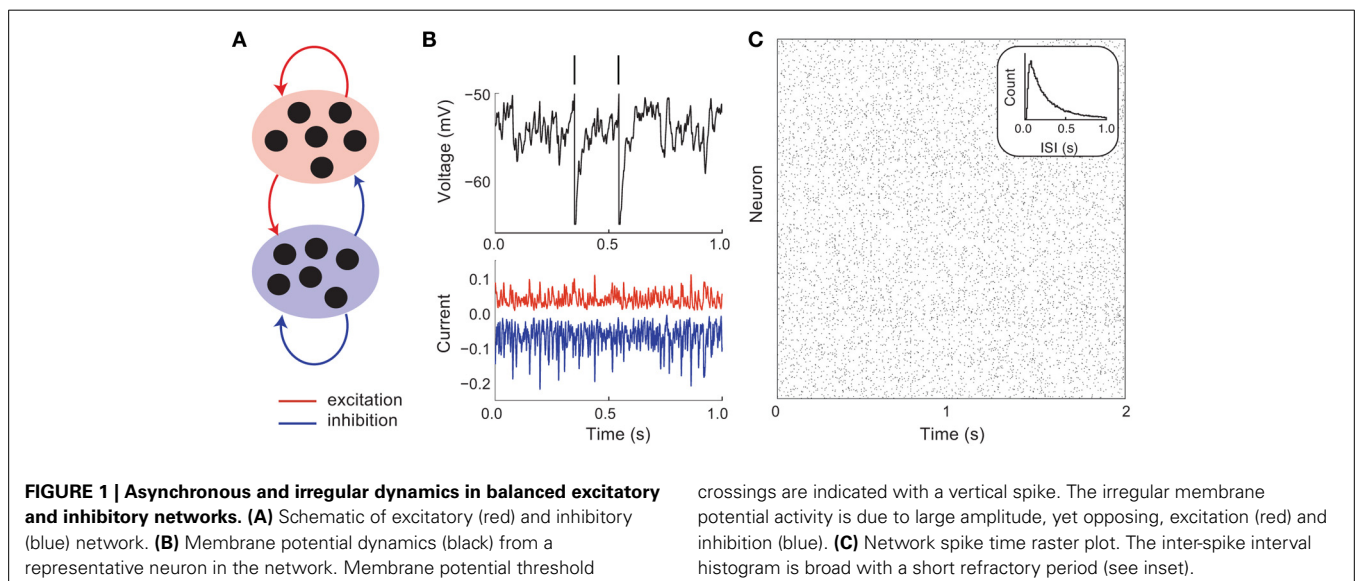
$$M = J_E r_E K - J_I r_I K + I_{\text{ext}},$$

$$V = (J_E)^2 r_E K + (J_I)^2 r_I K.$$

Here J_α and r_α are the synaptic strength and firing rates of the $\alpha \in \{E, I\}$ population and I_{ext} is an external drive to the network. For a reasonable theory we require that M and V to be independent of system size for large K . For the input variance V to neither vanish nor explode as $K \rightarrow \infty$, we must scale $J_\alpha = j_\alpha / \sqrt{K}$ where the parameter $j_\alpha \sim O(1)$. If we let the external input scale as $I_{\text{ext}} = i_{\text{ext}} \sqrt{K}$, then we have the following expression for M :

$$M = \sqrt{K} \underbrace{(j_E r_E - j_I r_I + i_{\text{ext}})}_{O(1/\sqrt{K})} \sim O(1).$$

For the mean input M to neither vanish nor explode as $K \rightarrow \infty$, the interior of the bracketed term must, as a whole, scale as $O(1/\sqrt{K})$. This only occurs when the recurrent and external excitation are balanced by recurrent inhibition, and with $i_{\text{ext}} > 0$ this can occur robustly over a wide range of network parameters.



Using this mean field approach van Vreeswijk and Sompolinsky showed that balanced networks were dynamically stable over a wide range of parameters, with asynchronous spiking dynamics (**Figure 1C**), and single neuron statistics consistent with Poisson-like variability (**Figure 1C**, inset). An important component of this variability is that it is self-generated through network interactions, making it akin to high dimensional chaos. Such chaotic network-based variability has *in vivo* experimental support (London et al., 2010), and been the focus of recent theoretical investigations (Banerjee et al., 2008; Monteforte and Wolf, 2010, 2012). Furthermore, the fact that the variability is self-generated shows that the high degree of variability of cortical dynamics may be an intrinsic property of a cortical circuit, rather than being inherited from external sources to the network.

In these analyses it is often assumed that connectivity is sparse. However, later work by Renart et al. (2010) relaxed this assumption and showed that asynchronous activity could be achieved in densely coupled balanced networks. More recent work has discussed the impact of balance in working memory (Boerlin and Denève, 2011; Lim and Goldman, 2013)¹, and the formation of orientation selectivity in visual cortices that lack functional topography (Hansel and van Vreeswijk, 2012). In total, balanced networks have been a successful framework to probe the mechanisms through which cortical dynamics are irregular, both in time and over repeated presentations of a fixed input.

2.2. SPONTANEOUS AND EVOKED DYNAMICS IN METASTABLE POPULATION DYNAMICS

The spike patterns from model neurons in balanced networks capture the irregular and asynchronous spike dynamics of evoked cortical response. However, balanced networks fail to capture the dynamics of cortical networks in spontaneous conditions (i.e., when the cortex is “idling”). In both sensory and motor cortices, the variability of spiking activity during spontaneous conditions is larger than that of a homogeneous Poisson process of the same firing rate (Churchland et al., 2006, 2010). A useful simplification is to treat a spike train in spontaneous conditions as a “doubly stochastic” process (Byron et al., 2009; Churchland and Abbott, 2012; Ponce-Alvarez et al., 2013), with one process modeling fast Poisson-like spike discharge and the other process capturing a slow fluctuation in the firing rate. Balanced networks of model neurons with independent wiring do not show the slow firing rate variability, and rather only capture the fast spiking variability (**Figure 1C**). In this section we show how higher order architectural structure in excitatory–excitatory connectivity allow balanced networks to replicate the variability reported in both spontaneous and evoked conditions.

2.2.1. Excitatory networks with clustered architectures

Over the past decade, several experimental groups have performed careful analysis of the microstructure in cortical networks. One clear finding is an architectural structure that is beyond that of a simple independent wiring model. In particular, small

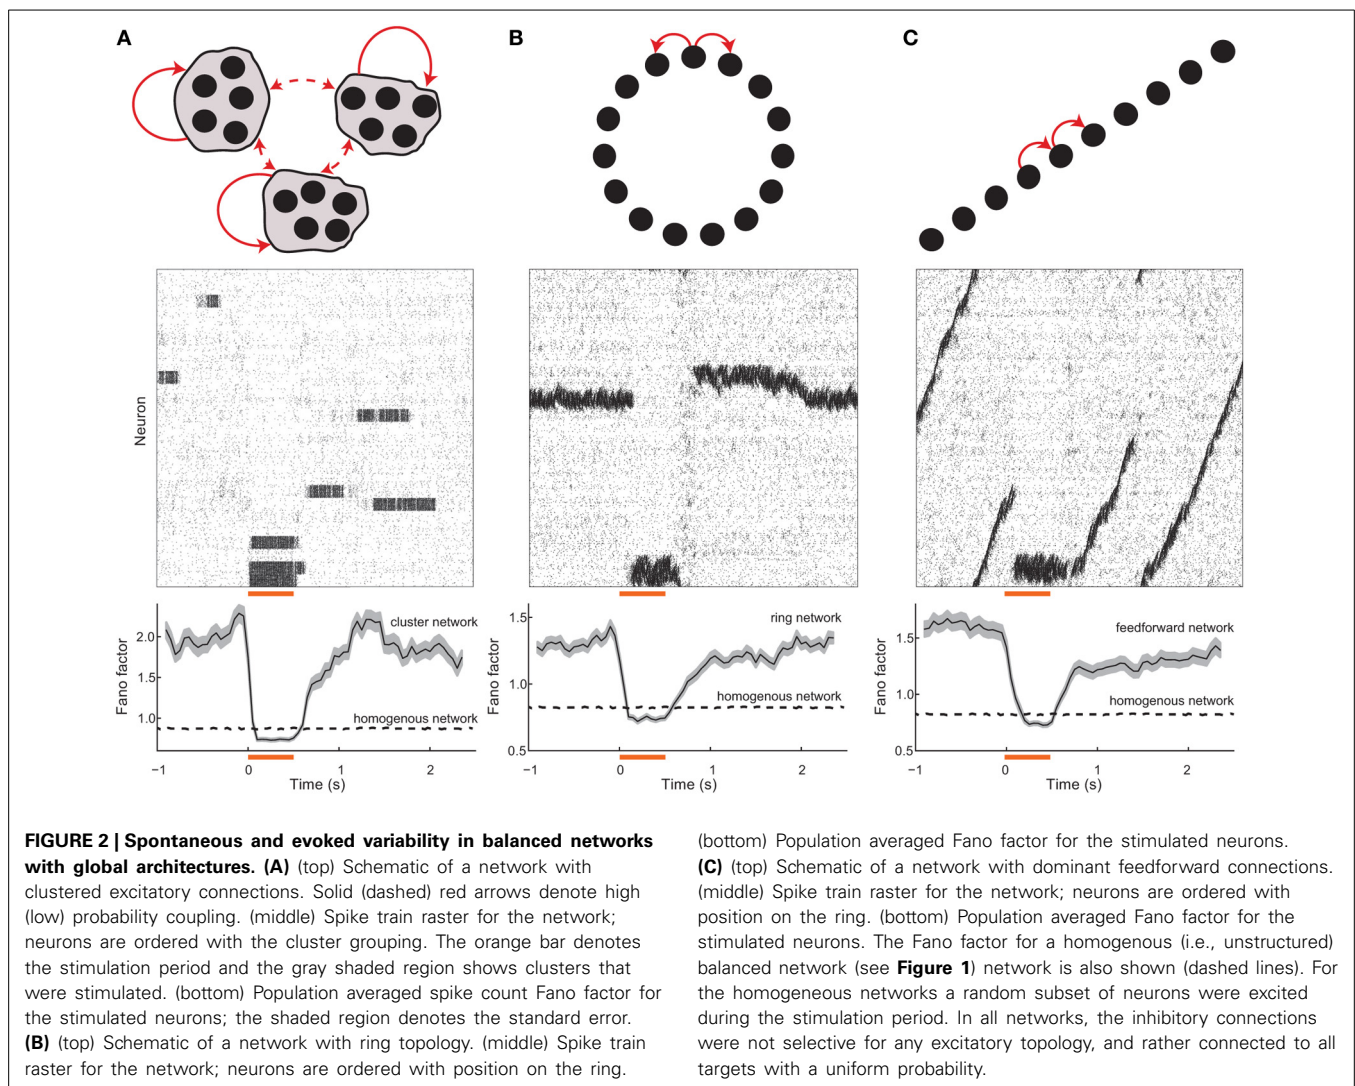
interconnected clusters of excitatory neurons are overrepresented in the cortex (Song et al., 2005; Perin et al., 2011). Cluster membership is often stimulus (Hofer et al., 2011; Ko et al., 2011), circuit input (Yoshimura and Callaway, 2005; Yoshimura et al., 2005), or activity (Yassin et al., 2010) dependent. Motivated by these anatomical findings, we subdivided excitatory cells into small clusters, and let the wiring probability be higher for two cells within the same cluster, compared to cells that are in distinct clusters (**Figure 2A** top, Materials and Methods). The chaotic dynamics of the network (**Figure 2A**, middle) is very distinct from that of an unclustered balanced network (**Figure 1C**). Specifically, cells within a cluster show periods of coordinated low and high firing activity, and the transitions between activity states occur over long timescales. These firing rate transitions are in addition to the spike time variability caused by the balanced state input fluctuations, so that spiking dynamics are “doubly stochastic.” Indeed, the spike count Fano factor over trials (normalized measure of variability, see Materials and Methods) is significantly above unity (**Figure 2A**, bottom), consistent with experimental measurements across cortex (Churchland et al., 2010).

When the neurons in a cluster are stimulated together they have a higher likelihood of being in a high firing rate state when compared to spontaneous conditions. Furthermore, recurrent inhibition permits only a fraction of the clusters to have high firing rates at any time, meaning that neurons in unstimulated clusters have a higher likelihood of being in a low firing rate state compared to spontaneous conditions. Thus, by stimulating a subset of the clusters, the network is biased to be in a specific firing rate configuration. This removes the trial-to-trial rate variability characteristic of spontaneous dynamics (**Figure 2A**, middle), leaving only the Poisson-like spike dynamics due to balanced excitation and inhibition. As a result the Fano factor is reduced to near one in evoked conditions (**Figure 2A**, bottom), again matching experiments (Churchland et al., 2010). In our analysis we have corrected for reductions in Fano factor that would be caused by an increase in firing rate. This was done through a mean-matching algorithm, originally presented in Churchland et al. (2010), that discards data in evoked and spontaneous conditions so that the firing rate distributions overlap (see Methods). Thus, the stimulus-induced reduction in Fano factor cannot be trivially due to an increase in the mean spike count that exceeds an increase in spike count variance, and is rather because of an active reduction in spike count variance.

The mean field analysis that was applied to unclustered balanced networks can be applied to clustered networks. Consider the mean synaptic input M given to a representative neuron. Split the recurrent excitation into the number of connections coming from within the neuron’s cluster K_{in} , and the remaining connections from outside the cluster $K_{out} = K - K_{in}$. This gives the mean input as:

$$M = \sqrt{K} \underbrace{(j_E^{\text{out}} r_E^{\text{out}} - j_I r_I + i_{\text{ext}})}_{O(1/\sqrt{K})} + \underbrace{(K^{\text{in}}/\sqrt{K})}_{O(1)} \underbrace{(j_E^{\text{in}} r_E^{\text{in}} - j_E^{\text{out}} r_E^{\text{out}})}_{O(1)} \sim O(1).$$

¹However, in these studies a fine balance between excitation and inhibition is required, necessitating some fine tuning of network parameters.



For excitatory cells that are within and outside the cluster we distinguish their connection strengths, J_E^{in} and J_E^{out} , and firing rates, r_E^{in} and r_E^{out} . The first term on the right hand side is $O(1)$ when the large recurrent excitation from outside the cluster and external drive are balanced by large recurrent inhibition. This is identical to the balance condition of the unclustered network. The second term is a perturbation from the balanced state that is induced by the clustered architecture. For this term to neither vanish nor overwhelm the network we require that the number of inputs that drive a cell from within its cluster K^{in} scale as \sqrt{K} (van Vreeswijk and Sompolinsky, 2005). For our network we set $K \approx 800$ so that $K^{\text{in}} \approx 30$, meaning that the excitation from within the cluster is a small fraction of the total recurrent excitation. Equivalently, for each neuron to be in a single cluster, we require the network to be divided into $Np_{EE}^{\text{in}}/K^{\text{in}} \approx 50$ clusters. Overall, this means that the network will have a large number of stable states that provide a rich repertoire of dynamics in spontaneous conditions. The interested reader can see a full exposition of the trial-to-trial variability of balanced networks with clustered connections in Litwin-Kumar and Doiron (2012).

The above cluster-based mechanism for a stimulus-induced reduction in spiking variability suggests a general framework. Highly variable spontaneous dynamics should occur in balanced networks with a large number of symmetric stable firing rate configurations. When population driven fluctuations cause stochastic transitions between distinct rate configurations, individual neurons will inherit these transitions as a new source of long timescale firing variability. Stimuli can break the network symmetry and bias specific firing rate patterns, reducing trial-to-trial firing rate variability. As a proof of principle, we show stimulus-induced reductions in variability in balanced networks with two global architectures that are distinct from clustered networks, yet satisfy the core requirements of our framework.

2.2.2. Excitatory networks with ring and feedforward architectures

Networks with local excitation and global (or lateral) inhibition arranged on a “ring” (Figure 2B, top) have been well studied as models of cortical orientation tuning (Ben-Yishai et al., 1995) as well as working memory (Compte et al., 2000). Ring networks support neutrally stable solutions where elevated rates

are localized in space on the ring, often labeled a “bump” of activity (Ermentrout, 1998). In the absence of a stimulus, the network has a continuum of stable bump solutions, and population fluctuations allow a slow timescale wandering of the bump (Figure 2B, middle). Over realizations of the network, the random initial position and stochastic wandering of the bump drive slow firing rate variability so that single neuron Fano factors are large (Figure 2B, bottom). When a specific region on the ring is stimulated it becomes globally stable and other regions lose their stability, so that a bump of activity appears in that region and is pinned in space. Consequently, the firing rate variability is removed and single unit Fano factors drop to levels below spontaneous conditions (Figure 2B, bottom).

Another well studied architecture in theoretical neuroscience is networks with dominant feedforward coupling (Figure 2C top)². Feedforward networks support the creation and propagation of synchronous waves of activity along the network (Kumar et al., 2010), as well as provide a mechanism for selective amplification of activity that produces long timescale dynamics (Goldman, 2009; Murphy and Miller, 2009). The symmetry in the network supports wave initiation at any position in the chain, so that population fluctuations spuriously cause and destroy wave dynamics over time (Figure 2C, middle). Over trials of the network, these wave events drive single unit firing rate variability so that Fano factors are elevated (Figure 2C, bottom). Stimuli at a fixed point along the chain break the translational symmetry along the network and effectively pin activity in that region since global network inhibition prevents wave propagation. Thus,

²In our networks we did not force a strict feedforward structure, instead we introduced an anisotropy along the chain so that backward connections were less likely than forward connections.

over repeated trials the wave induced dynamics are absent and the evoked Fano factor are below what is observed in spontaneous conditions (Figure 2C, bottom).

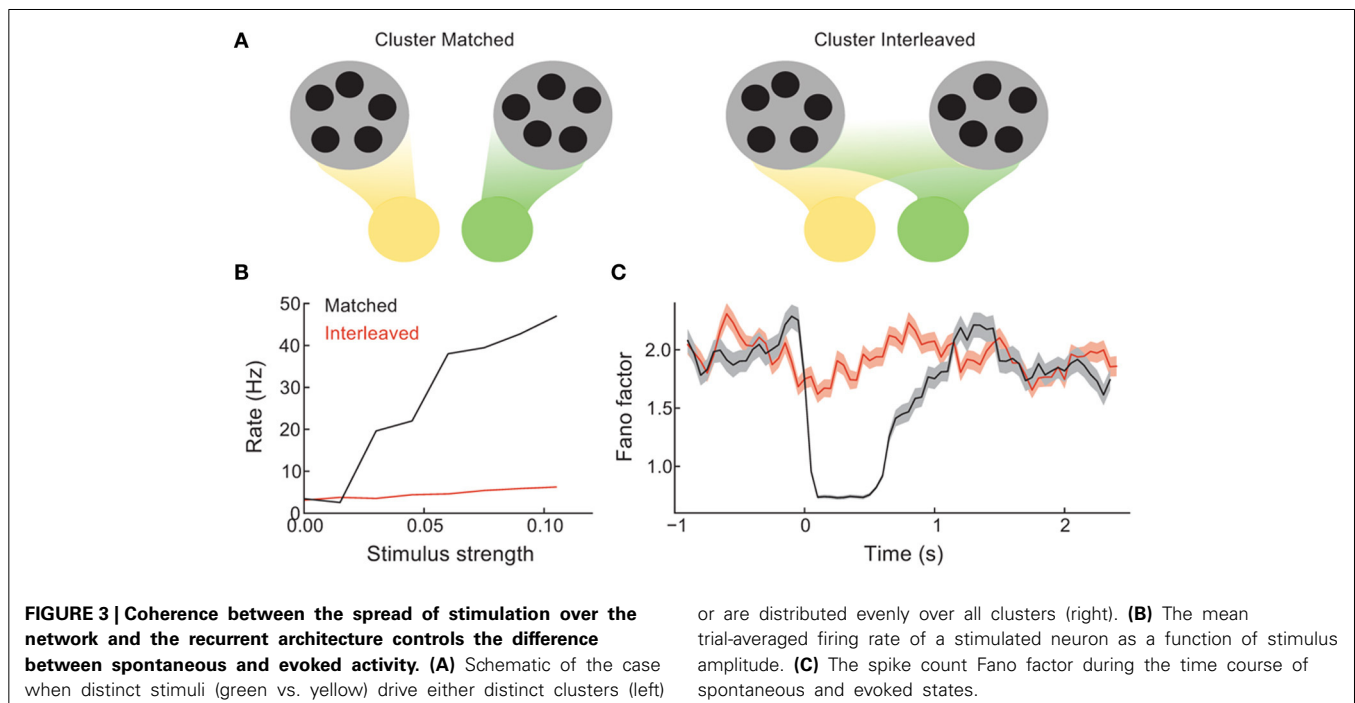
In total, networks with balanced excitation and inhibition with ring or feedforward architectures also capture high variability in spontaneous conditions and the reduction of variability in evoked states.

2.3. COHERENCE BETWEEN STIMULUS DRIVE AND RECURRENT ARCHITECTURES

In the balanced network with clustered excitation, we chose feedforward stimulus drive to be aligned with cluster membership (Figure 2). In other words, neuron pairs that belonged to the same cluster received coordinated stimulation, while neuron pairs in different clusters did not. In this section we compare the recruitment of clustered network activity when stimulation respects cluster membership (cluster matched, Figure 3A, left) to when a random assortment of neurons are stimulated, with no relation to the underlying cluster membership (cluster interleaved, Figure 3A, right).

The stimulus-response gain of the trial-averaged single neuron firing rate is higher for the cluster matched than for the cluster interleaved protocol (Figure 3B). Furthermore, the cluster interleaved stimulation did not cause a reduction in the spike count Fano factor compared to the spontaneous state, in contrast to the cluster matched case (Figure 3C). Thus, for the stimulus intensities studied, only cluster matched stimulation cause an evoked response that differs from the spiking statistics of the spontaneous state. This result is expected, as we describe below.

By design, the cluster of neurons driven by cluster matched inputs share more recurrence with one another than the neurons driven by cluster interleaved inputs. Because of this the



positive feedback recruited by the stimulus is larger for the cluster matched protocol than for the interleaved protocol. Positive feedback increases response gain in spiking networks (Sutherland et al., 2009), allowing the response to weak inputs to be amplified. Furthermore, when the stimulation is interleaved, each cluster receives inputs to only a fraction of its neurons, distributed roughly evenly over the clusters. Thus, stimulus interleaved stimulation does not bias any one firing rate configuration more than any other, and the mechanism for firing rate variability is left unaffected. Hence so that the Fano factor remains high throughout stimulation. In total, these results show that a coherence between how the stimulus is distributed over the network and the global recurrent architecture is needed for evoked spiking variability to differ from those in the spontaneous state.

2.4. RELATING SPIKING CORRELATIONS IN THE SPONTANEOUS AND EVOKED STATES

Balanced networks with unstructured connectivity have a stable asynchronous state (Renart et al., 2010). As such, in these networks the spike count correlation coefficient from pairs of neurons is on average close to zero, and for large networks ($N > 1000$) the likelihood of correlation coefficients above 0.2 is negligible. The clustered network is also balanced and, similar to the unclustered network, has a density of pairwise correlation that is approximately Gaussian with a small spread about a near zero mean (Figure 4A, black curve). While asynchronous dynamics are consistent with some cortical recordings (Ecker et al., 2010; Renart et al., 2010), there are many studies that show an average positive spike count correlation, with some pairs having significant (>0.2) correlations (Cohen and Kohn, 2011). However, in these experimental studies the correlation is measured from evoked, and not spontaneous, states.

The trial-to-trial co-variability of pairwise response in evoked states is termed “noise” correlations (Averbeck et al., 2006), to be distinguished from correlations in the trial-averaged activity due to the signal preferences of the neuron pair, i.e., “signal” correlations. In the clustered network we have assumed that stimulus preference and cluster membership are related (Figure 3). When pairwise correlation is conditioned on each neuron of the pair

being a member of the same cluster, the mean correlation is far from zero, and the density has a heavy tail at positive correlations (Figure 4A, green curve). This is due to neurons that are members of the same cluster being subject to the same firing rate fluctuations in spontaneous conditions, manifesting as a positive correlation coefficient. Thus, clustered architectures in balanced networks provide a mechanism for positive spike count correlations between select pairs. Further, because of the large number of clusters in the network, the number of pairs that are in the same cluster is far less than the number of total pairs in the network. This fact allows the clustered architecture to support a rough asynchronous state across the entire network, a requirement of balanced solutions (Renart et al., 2010).

The combination of these results suggests that a pair of neurons in the same cluster will have both a net positive correlation and an assumed common feedforward stimulus drive. This produces a clear and testable prediction: the signal correlation between pairs of neurons should be positively related to their correlation in the spontaneous state. This is certainly the case in our network simulations (Figure 4B); however, this has yet to be tested from *in vivo* data (to our knowledge).

2.5. SPONTANEOUS ACTIVITY “SAMPLES” THE RANGE OF EVOKED STATES

Our analysis of spontaneous and evoked network activity has focused on its reflection in single neuron (Figures 2, 3) or pairwise (Figure 4) spike count statistics. However, the high correlation for pairs of neurons within the same cluster suggests that an analysis of the population dynamics of pairs of clusters will provide a network perspective on the dynamics of spontaneous and evoked activity.

In the spontaneous state, cluster activity randomly transitions between low and high firing rate regimes. The recurrent inhibition regulates the overall network activity such that if neurons in one cluster are firing at high rates then neurons in another cluster are likely firing at low rates. This is seen in the joint activity of a pair of clusters in the network, since their joint firing rate trajectory shows an anti-correlation (Figure 5, gray curve). However, when a stimulus is given that drives either one of the

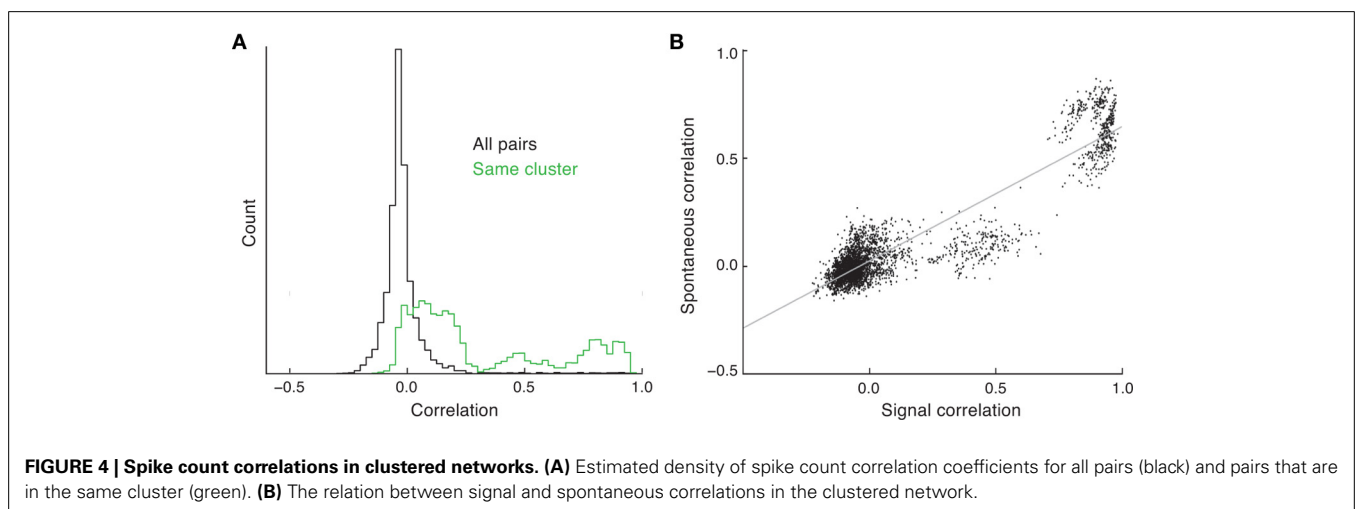
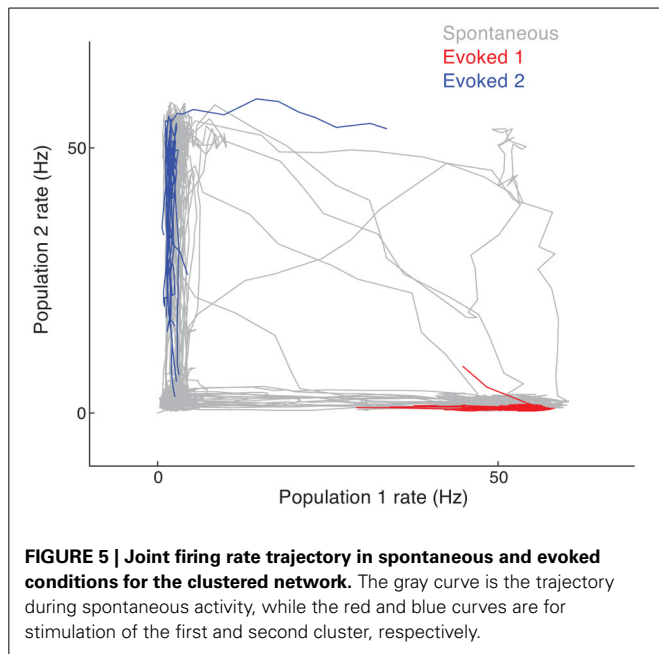


FIGURE 4 | Spike count correlations in clustered networks. (A) Estimated density of spike count correlation coefficients for all pairs (black) and pairs that are in the same cluster (green). **(B)** The relation between signal and spontaneous correlations in the clustered network.



clusters, the joint firing rate trajectory is effectively clamped so that the driven cluster's firing rate is high and the other cluster's is low. In the space of joint firing rates this collapse of network dynamics is such that the evoked state is a subset of the spontaneous state (Figure 5, red and blue curves compared to the gray curves).

2.6. STIMULUS QUENCHED VARIABILITY IN CHAOTIC FIRING RATE NETWORKS

We have presented a model for spontaneous cortical dynamics with high variability that is based on either a discrete (Figure 2A) or continuous (Figures 2B,C) attractor structure. Attractor dynamics require that recurrent cortical architecture provides a large set (or continuum) of stable network states. Such a constraint confers a fragility to the model dynamics, with unstructured heterogeneity in network architecture causing a collapse of network dynamics to only a few stable states. This criticism of attractor dynamics is often countered by assuming appropriate homeostatic mechanisms which provide a robustness to the attractor structure (Renart et al., 2003; Vogels et al., 2011). However, more research is needed to properly understand how attractor dynamics depends on the combination of homeostasis and rich spontaneous dynamics.

An alternative model of cortex is a recurrent excitatory–inhibitory network built from phenomenological firing rate models (Ermentrout and Terman, 2010). When recurrent coupling is strong, these models show chaotic solutions (Sompolinsky et al., 1988; Sussillo and Abbott, 2009; Rajan et al., 2010), where firing dynamics are trial-to-trial variable due to the extreme sensitivity of the state of the network at the time of stimulation. In contrast to attractor networks, these networks are not dependent on any fine structured architecture, so that the chaotic solutions are robust to parameter variations. However, by design these models consider only output firing rates, and we must assume

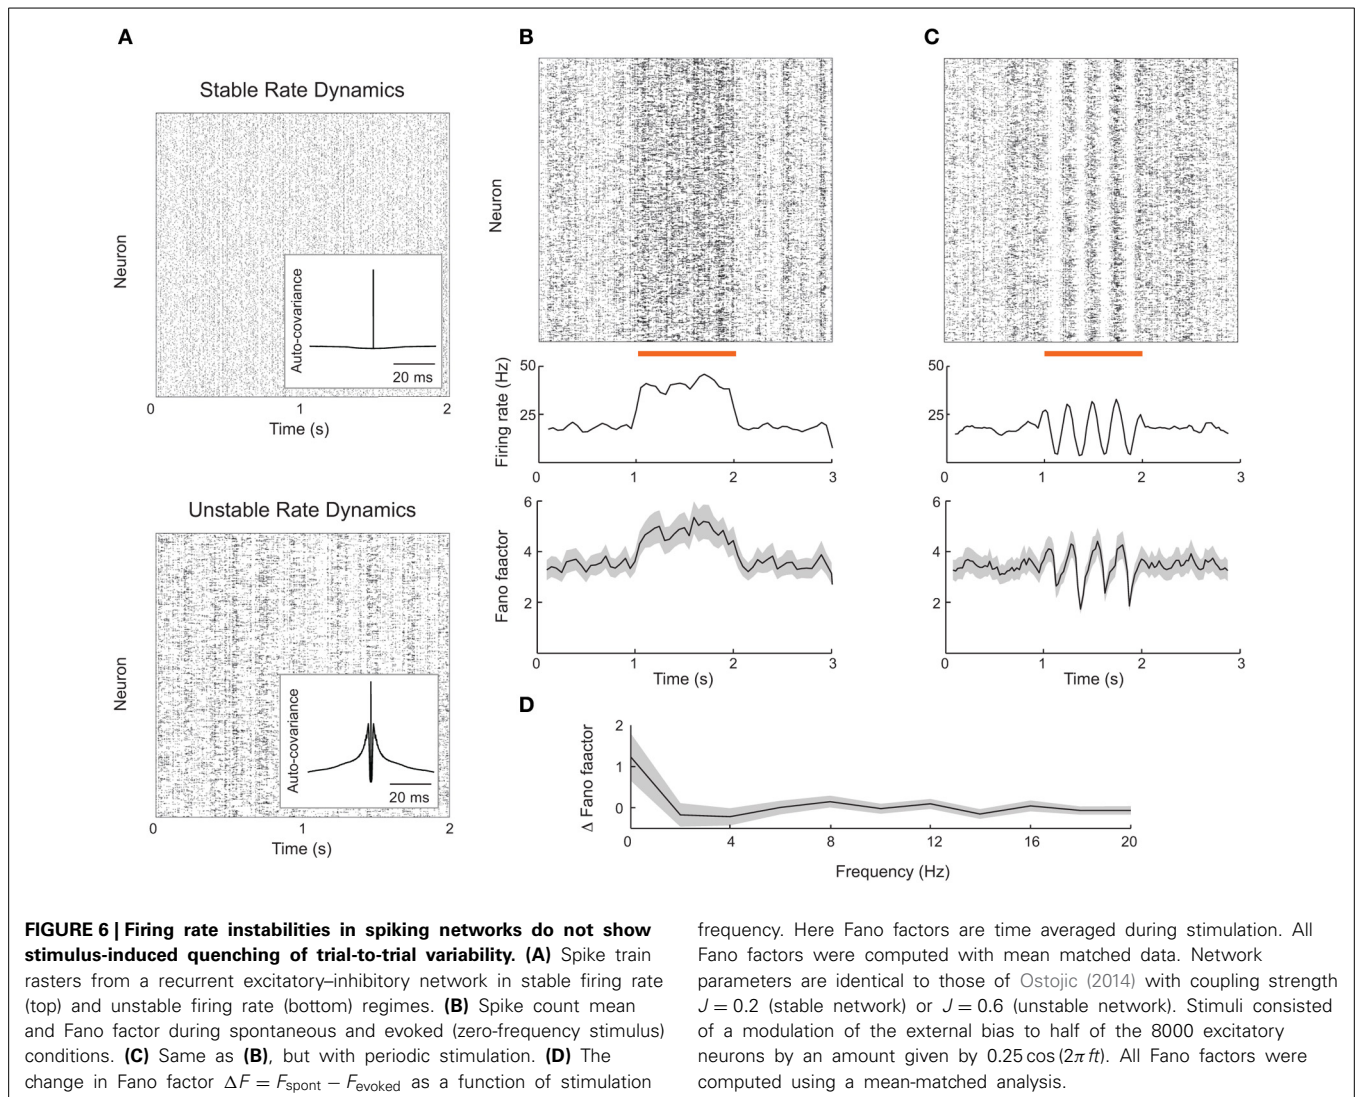
a suitable mechanism for how spike trains are generated from a rate dynamic. It is natural to consider the rate solutions as generating an inhomogeneous Poisson process. In this case, the trial variable chaotic dynamics of the rate model would produce spike trains that are “doubly stochastic” point processes, with an expected Fano factor larger than unity (Churchland and Abbott, 2012).

Rajan et al. (2010) observed that global, dynamic stimulation quenched internally generated chaotic dynamics in firing rate models. Thus, network firing rates in stimulated conditions lacked the dynamic variability characteristic of the spontaneous state. If a network produced spikes according to an inhomogeneous Poisson process with firing rates consistent with this model, then we expect a stimulus-induced reduction in the Fano factor, similar to that reported in our attractor model (Figure 2). However, it is still unclear if a network composed of actual spiking units, rather than interacting firing rates, is capable of reproducing the results of Rajan et al. (2010). This reproduction is not trivial because in most recurrent networks firing rates are stable (Figure 6A, top) (van Vreeswijk and Sompolinsky, 1996, 1998; Brunel, 2000; Renart et al., 2010), as opposed to chaotic.

Fortunately, Ostojic (2014) has recently shown how strong recurrent coupling in excitatory–inhibitory networks of spiking neuron models can destabilize firing rates, causing chaotic dynamics reminiscent of those predicted in firing rate networks (Figure 6A, bottom). This destabilization produces longer timescale spike train fluctuations reminiscent of a burst-like discharge that is absent in standard balanced networks (compare inset auto-covariance functions in Figure 6A). The network nevertheless shows roughly asynchronous spiking behavior and is a reasonable candidate model for rich spontaneous cortical dynamics. We implemented this spiking network in an effort to test the alternative theory for stimulus induced reduction in spike count variability proposed by Rajan et al. (2010).

The rate variability in the strongly coupled network produces high trial-to-trial variability in spiking dynamics in the spontaneous state, as characterized by a sizable Fano factor (~ 4 ; Figure 6B, bottom panel before stimulation). When a static depolarizing input stimulus was applied to a half of the excitatory neurons, the firing rates of the stimulated neurons naturally increased (Figure 6B, top and middle panels). However, the mean matched Fano factor also increased, because the elevated firing rates recruited more burst behavior (Figure 6B, bottom panels). This is in contrast to the drop in variability seen in the attractor networks presented earlier (Figure 2), as well recorded spike trains from sensory and motor cortex (Churchland et al., 2010).

We also tested if periodic stimulation could quench the spiking variability in strongly coupled networks, as suggested by chaotic rate networks (Rajan et al., 2010) as well as data from rodent visual cortex (White et al., 2012). Periodic input was reflected in a periodic modulation of both firing rate and spike count Fano factor (Figure 6C); however, the average Fano factor was not reduced by stimulation. This weak modulation of Fano factor, despite large changes in evoked rates, occurred over a range of stimulation frequencies (Figure 6D). Thus, further work is needed to determine whether certain stimulus or network structures can



frequency. Here Fano factors are time averaged during stimulation. All Fano factors were computed with mean matched data. Network parameters are identical to those of Ostojic (2014) with coupling strength $J = 0.2$ (stable network) or $J = 0.6$ (unstable network). Stimuli consisted of a modulation of the external bias to half of the 8000 excitatory neurons by an amount given by $0.25 \cos(2\pi ft)$. All Fano factors were computed using a mean-matched analysis.

reproduce the quenching of chaos seen in continuous firing rate networks.

3. DISCUSSION

The relation between spontaneous and evoked cortical dynamics is a popular topic of study (Ringach, 2009). Voltage sensitive dye imaging of visual cortex in cats has shown the spontaneous dynamics of large regions of cortex can mirror evoked dynamics (Arieli et al., 1996; Tsodyks et al., 1999; Luczak and MacLean, 2012). Population recordings in the auditory system of rodents (Luczak et al., 2009) and primates (Fukushima et al., 2012) show that spontaneous dynamics contain a large repertoire of activity, with an evoked state being merely a subset of that repertoire. Our study contributes to this literature by identifying cortical network architecture which supports rich spontaneous dynamics, that nevertheless has a clear relation to evoked dynamics. We expand on our previous study (Litwin-Kumar and Doiron, 2012) and argue that the excess variability in spontaneous conditions (Churchland et al., 2010) is a characteristic of balanced networks with a large number of stable firing rate states.

3.1. BALANCED ATTRACTOR NETWORKS AND SCALING

The networks discussed in this paper are a merging of ideas from the balanced network literature (van Vreeswijk and Sompolinsky, 1996, 1998; Renart et al., 2010), which proposes models of trial-to-trial variability, and the Hopfield network literature (Hopfield, 1982), which proposes models of neuronal assemblies. This connection has been made in several past studies (Amit and Brunel, 1997; van Vreeswijk and Sompolinsky, 2005; Renart et al., 2007; Roudi and Latham, 2007). Our contribution is specifically to study the implications of such attractor structures for trial-to-trial variability. Below, we discuss more thoroughly the plausibility of merging these ideas in the context of scaling arguments (see Renart et al., 2007 for an additional discussion of balance and scaling).

One requirement needed to merge these two lines of research was a specific scaling of the size of the neuronal assemblies. In balanced networks, if neurons receive K connections, then the connection strength J must be proportional to $1/\sqrt{K}$. We considered the case of *dense connectivity*, so that $K \propto N$. Further, the perturbation away from a homogeneous balanced network due

to the attractor structure was set to be $O(1)$, to prevent large attractors from dominating the network activity. Let the size of a neuronal assembly be K^{in} neurons and define the *coding fraction* as K^{in}/N . Then we must have:

$$K^{\text{in}}(J_E^{\text{in}}p^{\text{in}} - J_E^{\text{out}}p^{\text{out}}) \sim O(1),$$

where $J_E^{\text{in}} - J_E^{\text{out}}$ is how much stronger connections within the assembly are and $p^{\text{in}} - p^{\text{out}}$ is how much more probable they are. There are three scenarios:

1. $K^{\text{in}} \propto K$, that is, each assembly is composed of a macroscopic number of neurons and the coding fraction is constant. This is the case of a classic Hopfield network (Hopfield, 1982). However, to satisfy Equation (1), we must have $J_E^{\text{in}}p^{\text{in}} - J_E^{\text{out}}p^{\text{out}} \propto 1/K$. This presents a fine-tuning problem, as the J 's themselves are proportional to $1/\sqrt{K}$, much larger than the desired difference between J_E^{in} and J_E^{out} . Classic Hopfield networks avoided this problem as J was proportional to $1/K$.
2. $K^{\text{in}} \propto \sqrt{K}$, so the coding fraction is small. This was the case considered in our study and does not require fine-tuning as setting $J_E^{\text{in}} = aJ_E^{\text{out}}$ for some constant a ensures Equation (1) is satisfied. In such networks, the number of neurons in a local assembly is smaller by a factor of \sqrt{K} than the total number of inputs received. If stimuli are assumed to excite a few assemblies at once, this corresponds to a network exhibiting *sparse coding*, which is frequently encountered in cortex (Hromádka et al., 2008).
3. $K^{\text{in}} \propto 1$. This corresponds to a case of extremely sparse coding, as assembly size does not grow at all with K . Further, in this case we must have $J_E^{\text{in}} = O(1)$ while $J_E^{\text{out}} = O(1/\sqrt{K})$, corresponding to a few extremely strong local connections. For these reasons, we do not consider this case.

Typically, spiking implementations of attractor networks fall into the second category (Amit and Brunel, 1997; van Vreeswijk and Sompolinsky, 2005; Renart et al., 2007), as was true for our study. Roudi and Latham (2007) took an alternative approach and studied the first category, differentiating between “background weights” that were $O(\sqrt{K})$ and “foreground weights” (connections between neurons in the same assembly) that were $O(1/K)$. In order to obtain stable assembly activation, these foreground weights needed to be tuned to within 6% in a network of 10,000 neurons, and it is unclear that such fine-tuning is plausible.

It is therefore an open question how to robustly merge balanced and attractor networks with a non-vanishing coding fraction. Unlike primary sensory regions, higher-order associative regions such as prefrontal cortex have a large degree of “mixed selectivity,” with high coding fractions and assembly overlap (Rigotti et al., 2013). The fine-scale architecture of such regions has yet to be determined.

3.2. SYMMETRIC AND FINITE SIZE CLUSTERS

Despite their name, balanced networks are *robust* to changes in recurrent synaptic strengths J and admit a stable asynchronous solution over a wide range of network firing rates (van Vreeswijk and Sompolinsky, 1996, 1998; Renart et al., 2010). However, the

mechanism for spontaneous dynamics presented in this study requires a strong assumption of symmetry in the network architecture. More to the point, rich spontaneous dynamics requires that each firing rate configuration be roughly equally stable. If this is not the case then some stable configurations will be over-represented in the spontaneous state, and the population firing rates will tend to reside in those states. In the limit where one specific state is very stable, then the spiking variability in spontaneous and evoked dynamics would be similar. In our model, symmetric metastability is achieved by assuming that all clusters are of the same size, introducing a certain fragility into the framework. This is not an exact symmetry since quenched connection variability is present in any realization of the network. Nevertheless, our framework cannot support clusters of widely different sizes, as might be expected in real cortical networks. Future work should leverage powerful homeostatic mechanisms (Turrigiano and Nelson, 2000; Vogels et al., 2011) to prevent “winner-take-all” dynamics when the size and recurrence are heterogeneous across attractors.

The fluctuations in our model are an emergent feature of recurrent wiring, as opposed to being imposed by an external source. Balanced networks have a well defined solution as the number of inputs $K \rightarrow \infty$ (van Vreeswijk and Sompolinsky, 1996, 1998; Renart et al., 2010). However, in the clustered network the internal variability that promotes stochastic transitions between stable firing rate states requires $K < \infty$. Otherwise the time between firing rate transitions diverges, and the spontaneous state is no longer dynamically rich. Estimates for the number of neurons that are afferent to a cortical cell are imperfect, since anatomically and functionally defined architectures are often distinct, making a concrete value for K is difficult to obtain. Further, estimates for the size of a cluster are also difficult; however, (Perin et al., 2011) give an indirect number of ~ 100 neurons per cluster, consistent with our network model. For models of cortical networks to improve, techniques for functional microcircuit measurements of large numbers of cells must be developed.

3.3. CORTICAL ARCHITECTURE DETERMINES VARIABILITY IN SPONTANEOUS AND EVOKED STATES

In our model, feedforward afferent projections are assumed to be coherent with recurrent architecture (Figure 3A). Without this feature stimulus drive does not quench trial-to-trial variability (Figure 3C), and our model would then be at odds with experimental data across sensory and motor cortex (Churchland et al., 2010). Also, if stimulus input recruited distinct populations than those recruited during spontaneous activity, then the patterning of spontaneous activity would not resemble that in evoked states, again in disagreement with cortical data (Tsodyks et al., 1999; Luczak et al., 2009). This circuit assumption is well justified from recent circuit recordings in mouse visual cortex (Hofer et al., 2011; Ko et al., 2011), where the stimulus correlation and connection probability of pairs of neurons are positively correlated with one another. Thus, a specific stimulus should drive neurons that are members of a specific cluster to a larger degree than other neurons in the network. Finally, such a wiring rule is consistent with standard Hebbian plasticity rules, where repeated co-activation of neuron pairs

would engage plasticity rules that strengthen recurrent coupling between them.

Understanding the trial-to-trial variability in evoked conditions is important for a complete description of a population code (Dayan and Abbott, 2001; Averbeck et al., 2006; Josić et al., 2009; Ponce-Alvarez et al., 2013). For this reason, a large number of studies focus on the trial-to-trial “noise” correlations of neuron pairs in evoked states (Cohen and Kohn, 2011). In both visual (Smith and Kohn, 2008; Ponce-Alvarez et al., 2013) and auditory (Rothschild et al., 2010) cortex, signal and noise correlations are themselves correlated, so that neurons that have similar stimulus preference, i.e., positive signal correlation, also have positive noise correlation. This suggests that the underlying circuits that establish stimulus preference also contribute to common fluctuations. Our model extends this idea to the spontaneous state, and predicts a positive relation between signal correlations and the correlations measured during spontaneous conditions. However, balanced networks (including the ones used in this study) fail to capture the positive average pairwise correlation that is consistent with some (Cohen and Kohn, 2011), yet not all (Ecker et al., 2010), data sets (Figure 4A, black curve). A spiking network framework that internally generates non-zero mean correlation consistent with data is, to our knowledge, an open problem.

In sum, our model shows that a coherence between feedforward and recurrent architectures produces cortical dynamics that is consistent with a wide array of spontaneous and evoked data, and makes a clear and testable prediction for pairwise correlations in the spontaneous state.

4. MATERIALS AND METHODS

4.1. SPIKING NETWORK SIMULATIONS

Neurons were modeled as leaky integrate-and-fire units whose voltages obeyed:

$$\dot{V} = \frac{1}{\tau}(\mu - V) + I_{syn}(t)$$

When neurons reached a threshold $V_{th} = 1$, a spike was emitted and they were reset to $V_{re} = 0$ for an absolute refractory period of 5 ms. The membrane time constant τ was 15 and 10 ms for excitatory and inhibitory neurons, respectively. The bias μ was chosen according to a uniform random distribution between 1.1 and 1.2 for excitatory neurons and between 1 and 1.05 for inhibitory neurons. While these values are superthreshold, balanced dynamics ensured that the mean membrane potentials were subthreshold (van Vreeswijk and Sompolinsky, 1998). In Figure 1, the non-dimensionalized voltages were transformed so that $V_{th} = -50$ mV and $V_{re} = -65$ mV.

Synapses between neurons were modeled as differences of exponentials, and the total synaptic input to neuron i in population x was:

$$I_{i,syn}^x(t) = \sum_{jy} J_{ij}^{xy} F^y * s_j^y(t)$$

where $*$ denotes convolution. Here $x, y \in \{E, I\}$ denote excitatory or inhibitory populations of $N^E = 4000$ and $N^I = 1000$ neurons

each, J_{ij}^{xy} is the strength of synaptic connections from neuron j in population y to neuron i in population x , $F^y(t)$ is the synaptic filter for projections from neurons in population y , and $s_j^y(t)$ is the spike train of neuron j in population y , a series of delta-functions at the time points where the neuron emitted a spike. $F^y(t) = \frac{1}{\tau_2 - \tau_1} (e^{-t/\tau_1} - e^{-t/\tau_2})$, with $\tau_2 = 3$ ms for excitatory synapses and 2 ms for inhibitory synapses while $\tau_1 = 1$ ms.

Connection probabilities p^{xy} from neurons in population y to x were $p^{EI} = p^{IE} = p^{II} = 0.5$, while p^{EE} was on average 0.2. In clustered networks, excitatory neurons were partitioned into 50 assemblies of 80 neurons each and the connection probability was set to p_{in}^{EE} for neurons in the same assembly and p_{out}^{EE} for neurons in different assemblies. These were chosen so that p^{EE} remained on average 0.2, but $p_{in}^{EE}/p_{out}^{EE} = R^{EE} = 2.5$. For ring and feedforward networks, periodic boundary conditions were applied by identifying $i = 4000$ with $i = 0$. For ring networks, neurons i and j were said to belong to the same assembly if $|i - j| < 40$. For feedforward networks, connections from neuron i to neuron j were drawn with probability p_{in}^{EE} if $i - j \in [-35, 45]$ and p_{out}^{EE} otherwise, thus biasing connections in one direction along the ring. If a connection from neuron j in population y to neuron i in population x existed, $J_{ij}^{xy} = J^{xy}$ (unless $x, y = E$ and neurons were in the same assembly, then connection strength was multiplied by 1.9), otherwise $J_{ij}^{xy} = 0$. Synaptic strengths were $J^{EE} = 0.024$, $J^{EI} = -0.045$, $J^{IE} = 0.014$, and $J^{II} = -0.057$. These parameters, multiplied by 15 mV, would give the deflection of the membrane potential of the post-synaptic target, neglecting leak, in our dimensionalized units. When clusters of excitatory neurons were stimulated, stimulation was accomplished by increasing μ for neurons in those clusters by 0.07. Simulations were performed using Euler integration with a timestep of 0.1 ms.

4.2. SPIKE TRAIN STATISTICS

Spike train statistics were computed for excitatory neurons. We denote the spike times of neuron i as $\{t_{i1}, t_{i2}, t_{i3}, \dots\}$. We can then define neuron i 's spike train: $y_i(t) = \sum_k \delta(t - t_{ik})$. The number of spikes emitted by the neuron between times t and $t + \Delta t$ is

$$N_i(t, t + \Delta t) = \int_t^{t + \Delta t} y_i(t') dt'$$

The firing rate of a neuron over an interval $(t, t + \Delta t)$ was defined as

$$r_i(t, t + \Delta t) = \frac{1}{\Delta t} N_i(t, t + \Delta t)$$

For the networks studied, firing rates and other statistics for the spontaneous state were calculated with $t = 1.5$ s to prevent effects due to initial conditions and $\Delta t = 1.5$ s.

We also computed the Fano factor $F_i(t, t + \Delta t)$ for neuron i by evaluating

$$F_i(t, t + \Delta t) = \frac{\text{Var}(N_i(t, t + \Delta t))}{\langle N_i(t, t + \Delta t) \rangle}$$

where the expectations are over repeated trials of the same network with random initial conditions. When computing the Fano

factor as a function of time relative to stimulus onset, we computed the mean-matched Fano factor described in Churchland et al. (2010) to control for changes in firing rate. Fano factors were computed over 100 ms windows.

We computed correlation coefficients for the spike counts of neuron pairs. The correlation between neurons i and j was given by

$$\rho_{ij} = \frac{\text{Cov}(N_i(t, t + \Delta t), N_j(t, t + \Delta t))}{\sqrt{\text{Var}(N_i(t, t + \Delta t)) \text{Var}(N_j(t, t + \Delta t))}}$$

where the covariances and variances were computed over overlapping windows within each trial and then averaged across trials. For spontaneous activity trial averaging was replaced with temporal averaging.

FUNDING

NSF-DMS1313225 and NSF-DMS1121784.

ACKNOWLEDGMENT

The authors thank the members of the Doiron lab for useful conversations.

REFERENCES

- Amit, D., and Brunel, N. (1997). Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cereb. Cortex* 7, 237–252. doi: 10.1093/cercor/7.3.237
- Arieli, A., Sterkin, A., Grinvald, A., and Aertsen, A. (1996). Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science* 273, 1868–1871. doi: 10.1126/science.273.5283.1868
- Averbeck, B. B., Latham, P. E., and Pouget, A. (2006). Neural correlations, population coding and computation. *Nat. Rev. Neurosci.* 7, 358–366. doi: 10.1038/nrn1888
- Banerjee, A., Seriès, P., and Pouget, A. (2008). Dynamical constraints on using precise spike timing to compute in recurrent cortical networks. *Neural Comput.* 20, 974–993. doi: 10.1162/neco.2008.05-06-206
- Ben-Yishai, R., Bar-Or, R. L., and Sompolinsky, H. (1995). Theory of orientation tuning in visual cortex. *Proc. Natl. Acad. Sci. U.S.A.* 92, 3844–3848. doi: 10.1073/pnas.92.9.3844
- Binzegger, T., Douglas, R. J., and Martin, K. A. (2004). A quantitative map of the circuit of cat primary visual cortex. *J. Neurosci.* 24, 8441–8453. doi: 10.1523/JNEUROSCI.1400-04.2004
- Boerlin, M., and Denève, S. (2011). Spike-based population coding and working memory. *PLoS Comput. Biol.* 7:e1001080. doi: 10.1371/journal.pcbi.1001080
- Britten, K. H., Shadlen, M. N., Newsome, W. T., and Movshon, J. A. (1992). The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J. Neurosci.* 12, 4745–4765.
- Brunel, N. (2000). Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *J. Comput. Neurosci.* 8, 183–208. doi: 10.1023/A:1008925309027
- Byron, M. Y., Cunningham, J. P., Santhanam, G., Ryu, S. I., Shenoy, K. V., and Sahani, M. (2009). Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *J. Neurophysiol.* 102, 614–635. doi: 10.1152/jn.90941.2008
- Churchland, M., and Abbott, L. (2012). Two layers of neural variability. *Nat. Neurosci.* 15, 1472–1474. doi: 10.1038/nn.3247
- Churchland, M. M., Yu, B. M., Cunningham, J. P., Sugrue, L. P., Cohen, M. R., Corrado, G. S., et al. (2010). Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat. Neurosci.* 13, 369–378. doi: 10.1038/nn.2501
- Churchland, M. M., Yu, B. M., Ryu, S. I., Santhanam, G., and Shenoy, K. V. (2006). Neural variability in premotor cortex provides a signature of motor preparation. *J. Neurosci.* 26, 3697–3712. doi: 10.1523/JNEUROSCI.3762-05.2006
- Cohen, M. R., and Kohn, A. (2011). Measuring and interpreting neuronal correlations. *Nat. Neurosci.* 14, 811–819. doi: 10.1038/nn.2842
- Compte, A., Brunel, N., Goldman-Rakic, P. S., and Wang, X.-J. (2000). Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cereb. Cortex* 10, 910–923. doi: 10.1093/cercor/10.9.910
- Dayan, P., and Abbott, L. (2001). *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. Cambridge, MA: MIT Press. doi: 10.1016/S0306-4522(00)00552-2
- Decharms, R. C., and Zador, A. (2000). Neural representation and the cortical code. *Annu. Rev. Neurosci.* 23, 613–647. doi: 10.1146/annurev.neuro.23.1.613
- Deco, G., and Hugues, E. (2012). Neural network mechanisms underlying stimulus driven variability reduction. *PLoS Comput. Biol.* 8:e1002395. doi: 10.1371/journal.pcbi.1002395
- Ecker, A., Berens, P., Keliris, G., Bethge, M., Logothetis, N., and Tolias, A. (2010). Decorrelated neuronal firing in cortical microcircuits. *Science* 327, 584–587. doi: 10.1126/science.1179867
- Ermentrout, B. (1998). Neural networks as spatio-temporal pattern-forming systems. *Rep. Prog. Phys.* 61, 353. doi: 10.1088/0034-4885/61/4/002
- Ermentrout, G. B., and Terman, D. H. (2010). *Mathematical Foundations of Neuroscience*, Vol. 64. New York, NY: Springer. doi: 10.1007/978-0-387-87708-2
- Fukushima, M., Saunders, R. C., Leopold, D. A., Mishkin, M., and Averbeck, B. B. (2012). Spontaneous high-gamma band activity reflects functional organization of auditory cortex in the awake macaque. *Neuron* 74, 899–910. doi: 10.1016/j.neuron.2012.04.014
- Goldberg, J. A., Rokni, U., and Sompolinsky, H. (2004). Patterns of ongoing activity and the functional architecture of the primary visual cortex. *Neuron* 42, 489–500. doi: 10.1016/S0896-6273(04)00197-7
- Goldman, M. S. (2009). Memory without feedback in a neural network. *Neuron* 61, 621–634. doi: 10.1016/j.neuron.2008.12.012
- Haider, B., Duque, A., Hasenstaub, A. R., and McCormick, D. A. (2006). Neocortical network activity *in vivo* is generated through a dynamic balance of excitation and inhibition. *J. Neurosci.* 26, 4535–4545. doi: 10.1523/JNEUROSCI.5297-05.2006
- Hansel, D., and van Vreeswijk, C. (2012). The mechanism of orientation selectivity in primary visual cortex without a functional map. *J. Neurosci.* 32, 4049–4064. doi: 10.1523/JNEUROSCI.6284-11.2012
- Hofer, S. B., Ko, H., Pichler, B., Vogelstein, J., Ros, H., Zeng, H., et al. (2011). Differential connectivity and response dynamics of excitatory and inhibitory neurons in visual cortex. *Nat. Neurosci.* 14, 1045–1052. doi: 10.1038/nn.2876
- Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. U.S.A.* 79, 2554–2558. doi: 10.1073/pnas.79.8.2554
- Hromádka, T., DeWeese, M. R., and Zador, A. M. (2008). Sparse representation of sounds in the unanesthetized auditory cortex. *PLoS Biol.* 6:e16. doi: 10.1371/journal.pbio.0060016
- Josić, K., Shea-Brown, E., Doiron, B., and de La Rocha, J. (2009). Stimulus-dependent correlations and population codes. *Neural Comput.* 21, 2774–2804. doi: 10.1162/neco.2009.10-08-879
- Ko, H., Hofer, S. B., Pichler, B., Buchanan, K. A., Sjöström, P. J., and Mrsic-Flogel, T. D. (2011). Functional specificity of local synaptic connections in neocortical networks. *Nature* 473, 87–91. doi: 10.1038/nature09880
- Kumar, A., Rotter, S., and Aertsen, A. (2010). Spiking activity propagation in neuronal networks: reconciling different perspectives on neural coding. *Nat. Rev. Neurosci.* 11, 615–627. doi: 10.1038/nrn2886
- Lefort, S., Tomm, C., Floyd Sarria, J., and Petersen, C. C. (2009). The excitatory neuronal network of the c2 barrel column in mouse primary somatosensory cortex. *Neuron* 61, 301–310. doi: 10.1016/j.neuron.2008.12.020
- Lim, S., and Goldman, M. S. (2013). Balanced cortical microcircuitry for maintaining information in working memory. *Nat. Neurosci.* 16, 1306–1314. doi: 10.1038/nn.3492
- Litwin-Kumar, A., and Doiron, B. (2012). Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nat. Neurosci.* 15, 1498–1505. doi: 10.1038/nn.3220
- London, M., Roth, A., Beeren, L., Häusser, M., and Latham, P. E. (2010). Sensitivity to perturbations *in vivo* implies high noise and suggests rate coding in cortex. *Nature* 466, 123–127. doi: 10.1038/nature09086
- Luczak, A., Barthó, P., and Harris, K. D. (2009). Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron* 62, 413–425. doi: 10.1016/j.neuron.2009.03.014

- Luczak, A., and MacLean, J. N. (2012). Default activity patterns at the neocortical microcircuit level. *Front. Integr. Neurosci.* 6:30. doi: 10.3389/fnint.2012.00030
- Monteforte, M., and Wolf, F. (2010). Dynamical entropy production in spiking neuron networks in the balanced state. *Phys. Rev. Lett.* 105:268104. doi: 10.1103/PhysRevLett.105.268104
- Monteforte, M., and Wolf, F. (2012). Dynamic flux tubes form reservoirs of stability in neuronal circuits. *Phys. Rev. X* 2:041007. doi: 10.1103/PhysRevX.2.041007
- Murphy, B. K., and Miller, K. D. (2009). Balanced amplification: a new mechanism of selective amplification of neural activity patterns. *Neuron* 61, 635–648. doi: 10.1016/j.neuron.2009.02.005
- Ostojic, S. (2014). Two types of asynchronous activity in networks of excitatory and inhibitory spiking neurons. *Nat. Neurosci.* 17, 594–600. doi: 10.1038/nn.3658
- Perin, R., Berger, T. K., and Markram, H. (2011). A synaptic organizing principle for cortical neuronal groups. *Proc. Natl. Acad. Sci. U.S.A.* 108, 5419–5424. doi: 10.1073/pnas.1016051108
- Ponce-Alvarez, A., Thiele, A., Albright, T. D., Stoner, G. R., and Deco, G. (2013). Stimulus-dependent variability and noise correlations in cortical mt neurons. *Proc. Natl. Acad. Sci. U.S.A.* 110, 13162–13167. doi: 10.1073/pnas.1300098110
- Rajan, K., Abbott, L. F., and Sompolinsky, H. (2010). Stimulus-dependent suppression of chaos in recurrent neural networks. *Phys. Rev. E* 82, 011903. doi: 10.1103/PhysRevE.82.011903
- Renart, A., De la Rocha, J., Bartho, P., Hollender, L., Parga, N., Reyes, A., et al. (2010). The asynchronous state in cortical circuits. *Science* 327, 587. doi: 10.1126/science.1179850
- Renart, A., Moreno-Bote, R., Wang, X., and Parga, N. (2007). Mean-driven and fluctuation-driven persistent activity in recurrent networks. *Neural Comput.* 19, 1–46. doi: 10.1162/neco.2007.19.1.1
- Renart, A., Song, P., and Wang, X.-J. (2003). Robust spatial working memory through homeostatic synaptic scaling in heterogeneous cortical networks. *Neuron* 38, 473–485. doi: 10.1016/S0896-6273(03)00255-1
- Rigotti, M., Barak, O., Warden, M. R., Wang, X.-J., Daw, N. D., Miller, E. K., et al. (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature* 497, 585–590. doi: 10.1038/nature12160
- Ringach, D. L. (2009). Spontaneous and driven cortical activity: implications for computation. *Curr. Opin. Neurobiol.* 19, 439–444. doi: 10.1016/j.conb.2009.07.005
- Rothschild, G., Nelken, I., and Mizrahi, A. (2010). Functional organization and population dynamics in the mouse primary auditory cortex. *Nat. Neurosci.* 13, 353–360. doi: 10.1038/nn.2484
- Roudi, Y., and Latham, P. E. (2007). A balanced memory network. *PLoS Comput. Biol.* 3:e141. doi: 10.1371/journal.pcbi.0030141
- Shadlen, M. N., and Newsome, W. T. (1994). Noise, neural codes and cortical organization. *Curr. Opin. Neurobiol.* 4, 569–579. doi: 10.1016/0959-4388(94)90059-0
- Shadlen, M. N., and Newsome, W. T. (1998). The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J. Neurosci.* 18, 3870–3896.
- Smith, M., and Kohn, A. (2008). Spatial and temporal scales of neuronal correlation in primary visual cortex. *J. Neurosci.* 28, 12591–12603. doi: 10.1523/JNEUROSCI.2929-08.2008
- Sompolinsky, H., Crisanti, A., and Sommers, H. (1988). Chaos in random neural networks. *Phys. Rev. Lett.* 61, 259–262. doi: 10.1103/PhysRevLett.61.259
- Song, S., Sjöström, P., Reigl, M., Nelson, S., and Chklovskii, D. (2005). Highly non-random features of synaptic connectivity in local cortical circuits. *PLoS Biol.* 3:e68. doi: 10.1371/journal.pbio.0030068
- Sussillo, D., and Abbott, L. F. (2009). Generating coherent patterns of activity from chaotic neural networks. *Neuron* 63, 544–557. doi: 10.1016/j.neuron.2009.07.018
- Sutherland, C., Doiron, B., and Longtin, A. (2009). Feedback-induced gain control in stochastic spiking networks. *Biol. Cybern.* 100, 475–489. doi: 10.1007/s00422-009-0298-5
- Tsodyks, M., Kenet, T., Grinvald, A., and Arieli, A. (1999). Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science* 286, 1943–1946. doi: 10.1126/science.286.5446.1943
- Turrigiano, G. G., and Nelson, S. B. (2000). Hebb and homeostasis in neuronal plasticity. *Curr. Opin. Neurobiol.* 10, 358–364. doi: 10.1016/S0959-4388(00)00091-X
- van Vreeswijk, C., and Sompolinsky, H. (1996). Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* 274, 1724–1726. doi: 10.1126/science.274.5293.1724
- van Vreeswijk, C., and Sompolinsky, H. (1998). Chaotic balanced state in a model of cortical circuits. *Neural Comput.* 10, 1321–1371. doi: 10.1162/089976698300017214
- van Vreeswijk, C., and Sompolinsky, H. (2005). “Irregular activity in large networks of neurons,” in *Methods and Models in Neurophysics*, eds C. Chow, B. Gutkin, D. Hansel, and C. Meunier (Amsterdam: Elsevier).
- Vogels, T., Sprekeler, H., Zenke, F., Clopath, C., and Gerstner, W. (2011). Inhibitory plasticity balances excitation and inhibition in sensory pathways and memory networks. *Science* 334, 1569–1573. doi: 10.1126/science.1211095
- White, B., Abbott, L. F., and Fiser, J. (2012). Suppression of cortical neural variability is stimulus- and state-dependent. *J. Neurophysiol.* 108, 2383. doi: 10.1152/jn.00723.2011
- Yassin, L., Benedetti, B. L., Jouhanneau, J., Wen, J. A., Poulet, J. F., and Barth, A. L. (2010). An embedded subnetwork of highly active neurons in the neocortex. *Neuron* 68, 1043–1050. doi: 10.1016/j.neuron.2010.11.029
- Yoshimura, Y., and Callaway, E. M. (2005). Fine-scale specificity of cortical networks depends on inhibitory cell type and connectivity. *Nat. Neurosci.* 8, 1552–1559. doi: 10.1038/nn1565
- Yoshimura, Y., Dantzker, J. L. M., and Callaway, E. M. (2005). Excitatory cortical neurons form fine-scale functional networks. *Nature* 433, 868–873. doi: 10.1038/nature03252

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Received: 12 December 2013; accepted: 07 May 2014; published online: 27 May 2014.
Citation: Doiron B and Litwin-Kumar A (2014) Balanced neural architecture and the idling brain. *Front. Comput. Neurosci.* 8:56. doi: 10.3389/fncom.2014.00056
This article was submitted to the journal *Frontiers in Computational Neuroscience*. Copyright © 2014 Doiron and Litwin-Kumar. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.