

Mean-Driven and Fluctuation-Driven Persistent Activity in Recurrent Networks

Alfonso Renart*

arenart@andromeda.rutgers.edu

*Departamento de Física Teórica, Universidad Autónoma de Madrid,
Cantoblanco 28049, Madrid, Spain, and Volen Center for Complex Systems,
Brandeis University, Waltham, MA 02254, U.S.A.*

Rubén Moreno-Bote

rmoreno@cns.nyu.edu

*Department de Física Teórica, Universidad Autónoma de Madrid,
Cantoblanco 28049, Madrid, Spain*

Xiao-Jing Wang

xjwang@brandeis.edu

*Volen Center for Complex Systems, Brandeis University, Waltham, MA 02254,
U.S.A.*

Néstor Parga

nestor.parga@uam.es

*Departamento de Física Teórica, Universidad Autónoma de Madrid,
Cantoblanco 28049, Madrid, Spain*

Spike trains from cortical neurons show a high degree of irregularity, with coefficients of variation (CV) of their interspike interval (ISI) distribution close to or higher than one. It has been suggested that this irregularity might be a reflection of a particular dynamical state of the local cortical circuit in which excitation and inhibition balance each other. In this “balanced” state, the mean current to the neurons is below threshold, and firing is driven by current fluctuations, resulting in irregular Poisson-like spike trains. Recent data show that the degree of irregularity in neuronal spike trains recorded during the delay period of working memory experiments is the same for both low-activity states of a few Hz and for elevated, persistent activity states of a few tens of Hz. Since the difference between these persistent activity states cannot be due to external factors coming from sensory inputs, this suggests that the underlying

*Current address: Center for Molecular and Behavioral Neuroscience, Rutgers University, Newark, NJ 07102 USA.

network dynamics might support coexisting balanced states at different firing rates. We use mean field techniques to study the possible existence of multiple balanced steady states in recurrent networks of current-based leaky integrate-and-fire (LIF) neurons. To assess the degree of balance of a steady state, we extend existing mean-field theories so that not only the firing rate, but also the coefficient of variation of the interspike interval distribution of the neurons, are determined self-consistently. Depending on the connectivity parameters of the network, we find bistable solutions of different types. If the local recurrent connectivity is mainly excitatory, the two stable steady states differ mainly in the mean current to the neurons. In this case, the mean drive in the elevated persistent activity state is suprathreshold and typically characterized by low spiking irregularity. If the local recurrent excitatory and inhibitory drives are both large and nearly balanced, or even dominated by inhibition, two stable states coexist, both with subthreshold current drive. In this case, the spiking variability in both the resting state and the mnemonic persistent state is large, but the balance condition implies parameter fine-tuning. Since the degree of required fine-tuning increases with network size and, on the other hand, the size of the fluctuations in the afferent current to the cells increases for small networks, overall we find that fluctuation-driven persistent activity in the very simplified type of models we analyze is not a robust phenomenon. Possible implications of considering more realistic models are discussed.

1 Introduction

The spike trains of cortical neurons recorded *in vivo* are irregular and consistent, to a first approximation, to a Poisson process, possessing a roughly exponential interspike interval (ISI) distribution (except at very short intervals) and a coefficient variation (CV) of the ISI close to one (Softky & Koch, 1993). The possible implications of this fact on the basic principles of cortical organization have been the motivation for a large number of studies during the past 10 years (Softky & Koch, 1993; Shadlen & Newsome, 1994, 1998; Tsodyks & Sejnowski, 1995; van Vreeswijk & Sompolinsky, 1996, 1998; Zador & Stevens, 1998; Harsch & Robinson, 2000). An important idea that was analyzed by some of these studies was that a way out of the apparent inconsistency between the cortical neuron working as an integrator over the timescale of a relatively long time constant of the order of 10 to 20 ms of a very large number of inputs, and its irregular spiking, was to have similar amounts of excitatory and inhibitory drive. In this way, the mean drive to the cell was subthreshold, and spikes were the result of fluctuations, which occur irregularly, thus leading to a high CV (Gerstein & Mandelbrot, 1964). Although the implications of this result were first studied in a feedforward architecture (Shadlen & Newsome, 1994), it was soon discovered that a state

in which excitation and inhibition balance each other, resulting in irregular spiking, was a robust dynamical attractor in recurrent networks (Tsodyks & Sejnowski, 1995; van Vreeswijk & Sompolinsky, 1996, 1998); that is, under very general conditions, a recurrent network settles down into a state of this sort.

Although the original studies characterizing quantitatively the degree of spiking irregularity in the cortex were done using data from sensory cortices, it has since been shown that neurons in higher-order associative areas like the prefrontal cortex (PFC) also spike irregularly (Shinomoto, Sakai, & Funahashi, 1999; Compte et al., 2003) (see Figure 1). This is interesting because it is well known that cells in the PFC (Fuster & Alexander, 1971; Funahashi, Bruce, & Goldman-Rakic, 1989; Miller, Erickson, & Desimone, 1996; Romo, Brody, Hernández, & Lemus, 1999), as well as those in other associative cortices like the inferotemporal (Miyashita & Chang, 1988) or posterior parietal cortex (Gnadt & Andersen, 1988; Chafee & Goldman-Rakic, 1998), show activity patterns that are selective to stimuli no longer present to the animal and are thus being held in working memory. The activity of these neurons seems to be able to switch, on presentation of an appropriate brief, transient input, from a basal spontaneous activity level to a higher activity state. When the dimensionality of the stimulus to be remembered is low (e.g., the position of an LED on a computer screen or the frequency of a vibrotactile stimulus), the mnemonic activity during the delay period when the stimulus is absent seems to be graded (Funahashi et al., 1989; Romo et al., 1999), whereas when the dimensionality of the stimulus is high (e.g., a complex image), the single neurons seem to choose from a small number of discrete activity states (Miyashita & Chang, 1988; Miller et al., 1996). This last coding scheme is referred to as object working memory.

Since there is no explicit sensory input present during the delay period in a working memory task, the neuronal activity must be a result of the dynamics of the relevant neural circuit. There is a long tradition of modeling studies that have described delay-period activity as a reflection of dynamical attractors in multistable (usually bistable) networks presumed to represent the local cortical environment of the neurons recorded in the neurophysiological experiments (Hopfield, 1982; Amit & Tsodyks, 1991; Ben-Yishai, Lev Bar-Or, & Sompolinsky, 1995; Amit & Brunel, 1997b; Brunel, 2000a; Compte, Brunel, Goldman-Rakic, & Wang, 2000; Brunel & Wang, 2001; Hansel & Mato, 2001; Cai, Tao, Shelly, & McLaughlin, 2004). Originally inspired by models and techniques from the statistical mechanics of disordered systems, network models of persistent activity have progressively become more faithful to the biological circuits that they seek to describe. The landmark study (Amit & Brunel, 1997b) provided an extended meanfield description of the activity of a recurrent network of spiking current-based leaky integrate-and-fire neurons (LIF). One of its main achievements was to use the theory of diffusion processes to provide an intuitive, compact

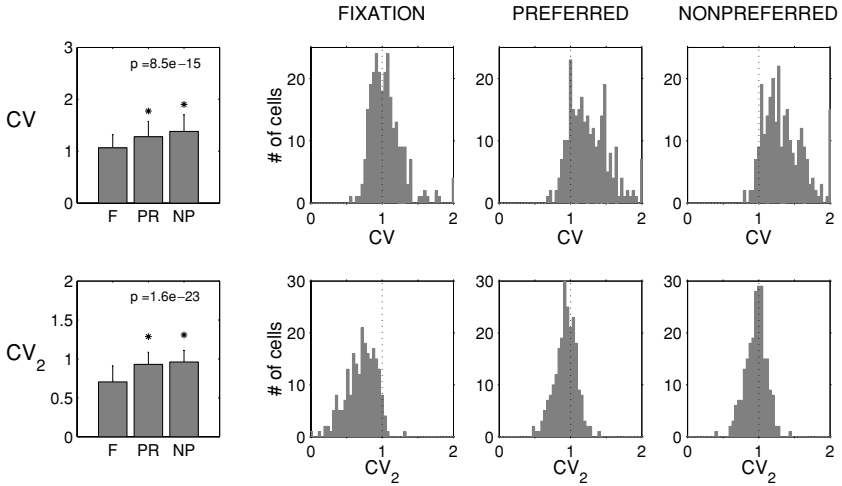


Figure 1: CV of the ISI of neurons in monkey prefrontal cortex during a spatial working memory task. The monkey made saccades to remembered locations on a computer screen after a delay period of a few seconds. On each trial, a dot of light (cue stimulus) was briefly shown in one of eight to-be-remembered locations, equidistant from the fixation point but at different angles. After the delay period, starting with the disappearance of the cue stimulus and terminating with the disappearance of the fixation point, the monkey made a saccade to the remembered location. Top and bottom rows correspond, respectively, to the CV and CV_2 (CV calculated using only consecutive ISIs to try to compensate from possible slow nonstationarities in the neurons instantaneous frequency) computed from spike trains of prefrontal cortical neurons recorded from monkeys performing an oculomotor spatial working memory task. Results shown correspond to analysis of the activity during the delay period of the task. The spike trains are irregular ($CV \sim 1$), and to a similar extent, both when the data correspond to trials in which the preferred (PR; middle column) positional cue for the cell was held in working memory (higher firing rate during the delay period) and when it corresponds to stimuli with the nonpreferred (NP; right column) positional cue for the particular neuron (lower firing rate during the delay period). See Compte et al. (2003) for details. Adapted with permission from Compte et al. (2003).

description of the spontaneous, low-rate, basal activity state of cortical cells in terms of self-consistency equations that included information about both the mean and the fluctuations of the afferent current to the cell. The theory proposed was both simple and accurate, and matched well the properties of simulated LIF networks.

The spontaneous activity state in Amit and Brunel (1997b) is effectively the balanced state described above, in which the recurrent connectivity is

dominated by inhibition and firing is due to the occurrence of positive fluctuations in the drive to the neurons. However, in Amit and Brunel (1997b), this same model was used to describe the coexistence of the spontaneous activity state with a persistent activity state with a physiologically plausible firing rate that would correspond to the spiking observed during the delay period in object working memory tasks, such as seen in, for example, Miyashita and Chang (1988).

Although the model, with its large number of subsequent improvements, has been successful in providing a fairly accurate description of simulated spiking networks, no effort has yet been made to study systematically the relationship between multistability and the irregularity of the spike trains, especially in the elevated activity state. As we will show below, the qualitative organization of the connectivity in the recurrent network not only determines the existence of a fluctuation-driven balanced spontaneous activity state in the network, but also the existence of bistability in the network, and whether the elevated activity states are fluctuation driven.

In order to perform a systematic analysis of the types of persistent activity that can be obtained in a network of current-based LIF neurons, two steps are important. First, we believe that the scaling of the synaptic connections with the number of afferent synapses per neuron should be made explicit. This approach was taken in the studies of the balanced state (Tsodyks & Sejnowski, 1995; van Vreeswijk & Sompolinsky, 1996), but is not present in the Amit and Brunel (1997b) framework. As we shall see, when the scaling is made explicit and the network is studied in the limit of a large number of connections per cell, the difference between the behavior of alternative circuit organizations (or architectures) becomes qualitative. Second, it would be desirable to be able to check for the spike train irregularity within the theory. In Amit and Brunel (1997b), spiking was assumed to be Poisson and, hence, to have a CV equal to 1. Poisson spike trains are completely characterized by a single number, the instantaneous firing probability, so there is nothing more to say about the spike train once its firing rate has been given. A general self-consistent description of the higher-order moments of spiking in a recurrent network of LIF neurons is extremely difficult, as the calculation of the moments of the ISI distribution becomes prohibitively complicated when the input current to a particular cell contains temporal correlations (although see Moreno-Bote & Parga, 2006). However, based on our study of the input-output properties of the LIF neuron under the influence of correlated inputs (Moreno, de la Rocha, Renart, & Parga, 2002), we have constructed a self-consistent description for the first two moments of the current to the neurons in the network, which relaxes the Poisson assumption and which we expect to be valid if the temporal correlations in the spike trains in the network are sufficiently short.

Some of the results presented here have already been published in abstract form.

2 Methods

We consider a network of current-based leaky integrate-and-fire neurons. The voltage difference across each neuron's membrane evolves in time according to the following equation,

$$\frac{dV(t)}{dt} = -\frac{V(t)}{\tau_m} + I(t),$$

with voltages being measured relative to the leak potential of the neuron. When the depolarization reaches a threshold voltage that we set at $V_{th} = 20$ mV, a spike is emitted, and the voltage is clamped at a reset potential $V_r = 10$ mV during a refractory period $\tau_{ref} = 2$ ms, after which the voltage continues to integrate the input current. The membrane time constant is $\tau_m = 10$ ms. When the neuron is inserted in a network, $I(t)$ represents the total synaptic current, which is assumed to be a linear sum of the contributions from each individual presynaptic cell.

We consider the simplest description of the synaptic interaction between the pre- and postsynaptic neurons, according to which each presynaptic action potential provokes an instantaneous "kick" in the depolarization of the postsynaptic cell. The network is composed of N_E excitatory and N_I inhibitory cells randomly connected so that each cell receives C_E excitatory and C_I inhibitory contacts, each with an efficacy ("kick" size) J_{E_j} and J_{I_k} , respectively ($j = 1, \dots, C_E$; $k = 1, \dots, C_I$). The total afferent current into the cell can be represented as

$$I(t) = \sum_{j=1}^{C_E} J_{E_j} s_j(t) - \sum_{k=1}^{C_I} J_{I_k} s_k(t),$$

where $s_{j(k)}(t)$ represents the spike train from the j th excitatory (k th inhibitory) neuron. Since according to this description, the effect of a presynaptic spike on the voltage of the postsynaptic neuron is instantaneous, $s(t)$ is a collection of Dirac delta functions, that is, $s(t) \equiv \sum_j \delta(t - t_j)$, where t_j are the spike times.

2.1 Mean-Field Description. Spike trains in the model are stochastic, with an instantaneous firing rate (i.e., a probability of measuring a spike in $(t, t + dt)$ per unit time) denoted by $\nu(t) = \nu$. The second-order statistics of the process is characterized by its connected two-point correlation function $C(t, t')$, giving the joint probability density (above chance) that two spikes happen at $(t, t + dt)$ and at $(t', t' + dt)$, that is, $C(t, t') \equiv \langle s(t)s(t') \rangle - \langle s(t) \rangle \langle s(t') \rangle$. Stochastic spiking in network models is usually assumed to follow Poisson statistics, which is both a fairly good approximation to what is commonly observed experimentally (see, e.g.,

Softky & Koch, 1993; Compte et al., 2003) and convenient technically since Poisson spike trains lack any temporal correlations. For Poisson spike trains, $C(t, t') = \nu\delta(t - t')$, where ν is the instantaneous firing probability.

We have previously analyzed the effect of temporal correlations in the afferents to a LIF neuron on its firing rate (Moreno et al., 2002). Temporal correlations measured *in vivo* are often well fitted by an exponential (Bair, Zohary, & Newsome, 2001). We considered exponential correlations of the form

$$C(t, t') = \nu \left[\delta(t' - t) + (F_\infty - 1) \frac{e^{-\frac{|t'-t|}{\tau_c}}}{2\tau_c} \right], \quad (2.1)$$

where F_∞ is the Fano factor of the spike train for infinitely long time windows. The Fano factor in a window of length T is defined as the ratio between the variance and the mean of the spike count on the window. It is illustrative to calculate it for our process,

$$F_T \equiv \frac{\sigma_{N(T)}^2}{\langle N(T) \rangle},$$

where $N(T)$ is the (stochastic) spike count in a window of length T ,

$$N(T) = \int_0^T dt s(t),$$

so that $\langle N(T) \rangle = \nu T$, and the spike count variance is given by

$$\sigma_{N(T)}^2 \equiv \int_0^T \int_0^T dt dt' C(t, t') = \nu T + \nu(F_\infty - 1)[T - \tau_c(1 - e^{-\frac{T}{\tau_c}})].$$

When the time window is long compared to the correlation time constant, that is, $T \gg \tau_c$, then $\sigma_{N(T)}^2 \propto F_\infty \nu T$; hence, our use of the factor F_∞ in the definition of the correlation function. An interesting point to note is that for time windows that are long compared to the correlation time constant, the variance of the spike count is linear in time, which is a signature of independence across time, that is, independent variances add up (for the Poisson process, $(\sigma_{N(T)}^2)_{\text{Poisson}} = \nu T$, so that $(F_T)_{\text{Poisson}} = 1$). If the characteristic time of the postsynaptic cell integrating this stochastic current (its membrane time constant) is very long compared with τ_c , we expect that the main effect of the deviation from Poisson of the input spike trains will be on the amplitude of the current variance, with the parameter τ_c playing only a marginal role, as it does not appear in $\sigma_{N(T)}^2$ when $T \gg \tau_c$. As we show

below, a rigorous analysis of the effect of correlations on the mean firing rate of a LIF neuron confirms this intuitive picture.

The postsynaptic cell receives many inputs. Recall that the total current is given (we consider for simplicity for this discussion that the cell receives C inputs from a single, large, homogeneous population composed of N neurons) by $I(t) = J \sum_j^C s_j(t)$. Thus, the mean and correlation function of the total afferent current to a given cell are

$$\begin{aligned} \langle I(t) \rangle &= C J v \\ C_I(t, t') &= \langle I(t')I(t) \rangle - \langle I(t) \rangle \langle I(t') \rangle \\ &= J^2 \sum_{ij} [\langle s_i(t) s_j(t') \rangle - \langle s_i(t) \rangle \langle s_j(t') \rangle] \\ &= C J^2 C(t, t') + C(C-1) J^2 C_{cc}(t, t'), \end{aligned}$$

where $C(t, t')$ is the (auto)correlation function in equation 2.1 and $C_{cc}(t, t')$ is the cross-correlogram between any two given cells of the pool of presynaptic inputs (which we have assumed to be the same for all pairs). We restrict our analysis to very sparse random networks—networks with $C \ll N$ —so that the fraction of synaptic inputs shared by any two given neurons can be assumed to be negligible. In this case, the cross-correlation between the spike trains of the two cells $C_{cc}(t, t')$ will be zero. This approximation simplifies the analysis of the network behavior significantly and allows for a self-consistent solution for the network's steady states. Thus, the temporal structure of the total current to the cell is described by

$$C_I(t, t') = \sigma_0^2 \left[\delta(t - t') + \frac{\alpha}{2\tau_c} e^{-\frac{|t-t'|}{\tau_c}} \right] \quad (2.2)$$

with

$$\sigma_0^2 = C J^2 v \quad \text{and} \quad \alpha = F_\infty - 1.$$

We have previously calculated the output firing rate of an LIF neuron subject to an exponentially correlated input (Moreno et al., 2002). The calculation is done using the diffusion approximation (Ricciardi, 1977) in which the discontinuous voltage trajectories are approximated by those obtained from an equivalent diffusion process. The diffusion approximation is expected to give accurate results when the overall rate of the input process is high, with the amplitude of each individual event being very small (Ricciardi, 1977).

For small but finite τ_c , the analytic calculation of the firing rate of the cell can be done only when the deviation of the input current from a white noise

process is small, that is, it has to be done assuming that $k \equiv \sqrt{\tau_c/\tau_m} \ll 1$ and that $\alpha \ll 1$. More specifically, we found that if $k = 0$, then the firing rate can be calculated for arbitrary values of α , but if k is small but finite, then an expression can be found for the case when both k and α are small (see Moreno et al., 2002, for details).

If $k = 0$, then the result one obtains is that the firing rate of the neuron is given by the same expression that one finds for the case of a white noise input, but with an effective variance that takes into account the change in amplitude of the fluctuations due to the non-Poisson nature of the inputs. The effective variance is equal to

$$\sigma_{\text{eff}}^2 = \sigma_0^2(1 + \alpha) = CJ^2vF_\infty,$$

which is exactly the slope of the linear increase with the size of the time window T of the variance in the spike count $N_I(T)$ of the total input current. This result can be understood in terms of the Fano factor calculation outlined above. Assuming $k = 0$ is equivalent to assuming an infinitely long time window for the calculation of the Fano factor, and in those conditions we also saw that the only effect of the temporal correlations is to renormalize the variance of the spike count with respect to the poisson case.

In order to set up a self-consistent scenario, we have to close the loop, by calculating a measure of the variability of the postsynaptic cell and relating it to the same property in the spike trains of its inputs. To do this, we note that if the spike trains in the model can be described as renewal processes, these processes have a property that relates their spike count variability and their ISI variability,

$$F_\infty = CV^2,$$

if a point process is renewal (Cox, 1962). Renewal point processes are characterized by having independent ISIs, which are not necessarily exponentially distributed. Since we are assuming that the temporal correlations in the spike trains are short anyway, and the firing rates of the cells in the persistent activity states that we are interested in are not very high, then we expect the renewal assumption to be appropriate. The final step is to make sure that the result for the firing rate (the inverse of the mean ISI) in terms of the effective variance also holds for higher moments of the postsynaptic ISI, not only for the first, and this is indeed the case (Renart, 2000); that is, the CV of the ISI when $k = 0$ is given by the same expression as when the input is a white noise process, but with a renormalized variance equal to σ_{eff}^2 .

Thus, under the assumptions described above, there is a way of computing the output rate and CV of an LIF neuron solely in terms of the rate and CV of its presynaptic inputs. In the steady states, both input and output

firing rate and CV will be the same, and this provides a couple of equations that determine these quantities self-consistently. In the remainder of the letter, we thus use the common expressions for the mean and CV of the first passage time of the Ornstein-Uhlenbeck (OU) process,

$$\nu^{-1} = \tau_{\text{ref}} + \tau_m \sqrt{\pi} \int_{\frac{V_{\text{res}} - \mu_V}{\sigma_V \sqrt{2}}}^{\frac{V_{\text{th}} - \mu_V}{\sigma_V \sqrt{2}}} dx e^{x^2} [1 + \text{erf}(x)] \quad (2.3)$$

$$\text{CV}^2 = 2\pi \nu^2 \int_{\frac{V_{\text{res}} - \mu_V}{\sigma_V \sqrt{2}}}^{\frac{V_{\text{th}} - \mu_V}{\sigma_V \sqrt{2}}} dx e^{x^2} \int_{-\infty}^x dy e^{y^2} [1 + \text{erf}(y)]^2, \quad (2.4)$$

where μ_V and σ_V^2 are the mean and variance of the depolarization of the postsynaptic cell (in the absence of threshold; Ricciardi, 1977). In a stationary situation, they are related to the mean μ and variance σ^2 of the afferent current to the cell by

$$\mu_V = \tau_m \mu; \quad \sigma_V^2 = \frac{1}{2} \tau_m \sigma^2.$$

Following the arguments above, the mean and (effective) variance of the current to the cells are given by

$$\begin{aligned} \mu &= C J \nu \\ \sigma^2 &\equiv \sigma_{\text{eff}}^2 = C J^2 \nu \text{CV}^2 \end{aligned}$$

for the mean and variance of the white noise input current. Finally, it is easy to show that if the presynaptic afferents to the cell come from a set of different statistically homogeneous subpopulations, the previous expressions generalize readily to

$$\mu_i = \sum_j C_{ij} J_{ij} \nu_j \quad (2.5)$$

$$\sigma_i^2 \equiv \sigma_{i\text{eff}}^2 = \sum_j C_{ij} J_{ij}^2 \nu_j \text{CV}_j^2 \quad (2.6)$$

as long as the timescales of the correlations in the spike trains of the neurons in the different subpopulations are all of the same order. Inhibitory subpopulations are characterized by negative connection strengths.

2.2 Dynamics. A detailed characterization of the dynamics of the activity of the network is beyond the scope of this work. Since our main interest is the steady states of the network, we use a simple, effective dynamical

scheme that is consistent with the self-consistent equations that determine the steady states. In particular, we use the subthreshold dynamics of the first two moments of the depolarization in terms of the first two moments of the current (Ricciardi, 1977; Gillespie, 1992):

$$\frac{d\mu_V}{dt} = -\frac{\mu_V}{\tau_m} + \mu; \quad \frac{d\sigma_V^2}{dt} = -\frac{\sigma_V^2}{\tau_m/2} + \sigma^2. \quad (2.7)$$

In using these equations, our assumption is that the activity of the population follows instantaneously the distribution of the depolarization. Thus, at every point in time, we use expressions 2.5 and 2.6 for μ and σ appearing in the right-hand side of equations 2.7, which depend on the rate $\nu(\mu_V, \sigma_V)$ and $\text{CV}(\mu_V, \sigma_V)$ as given in equations 2.3 and 2.4. The only dynamical variables are therefore μ_V and σ_V^2 (Amit & Brunel, 1997b; Mascaro & Amit, 1999).

2.3 Numerical Analysis of the Analytic Results. The phase plane analysis of the reduced network was done using both custom-made C++ code and the program XPPaut. The integrals in equations 2.3 and 2.4 were calculated analytically for very large and very small values of the limits of integration (using asymptotic expressions for the error function; Abramowitz & Stegun, 1970) and numerically for values of the integration limits of order one. The corresponding C++ code was incorporated into XPPaut through the use of dynamically linked libraries for phase plane analysis. Some of the cusp diagrams were calculated without the use of XPPaut by the direct approach of looking for values of the parameters at which the number of fixed points changed abruptly.

2.4 Numerical Simulations. We simulated an identical network to the one used in the mean-field description (see the captions of Figures 12 and 13 for parameters). In the simulation, on every time step $dt = 50 \mu\text{s}$, it is checked which neurons in the network receive any spikes. The membrane potential of cells that do not receive spikes is integrated analytically. The membrane potential of cells that receive spikes is integrated analytically within that dt , taking into account the synaptic postsynaptic potentials (PSPs) but assuming that there is no threshold. Only at the end of the time step is it checked whether the membrane potential is above threshold. If this is the case, the neuron is said to have produced a spike. This procedure effectively introduces a (very short but nonzero) synaptic integration time constant. Emitted spikes are fed back into the network using a system of queues to account for the synaptic delays (Mattia & Del Giudice, 2000).

3 Results

3.1 Network Architecture and Scaling Considerations. We first study the issue of how the architecture of the network determines the qualitative nature of the steady-state solutions for the firing rate of the cells in the network. In particular, we are interested in understanding under which conditions there are multiple steady-state solutions (bistability or, in general, multistability) in networks in which cells will fire with a high degree of irregularity. We consider a network for object working memory composed of a number of nonoverlapping subpopulations, or columns, defined by selectivity with respect to a given external stimulus. Each subpopulation contains both excitatory and inhibitory cells. The synaptic properties of neurons within the same column are assumed to be statistically identical. Thus, the column is, in our network, the minimal signaling unit at the average level, that is, all the neurons within a column behave identically on average. As will be shown below, the type of bistability realizable for the column depends critically on its size (more specifically, on the number of connections that a given cell receives from within its own column). Different architectures will thus be considered in which the total number of afferent connections per cell C is constant (and large) but the number of columns in the network varies, effectively varying the number of afferent connections from a given column to a cell. A multicolumnar architecture of this sort is inspired in the anatomical organization of the PFC, in which columnar organized putative excitatory cells and interneurons show similar response profiles during working memory tasks (Rao, Williams, & Goldman-Rakic, 1999).

As noted in section 1, many of the properties of the network can be inferred from a scaling analysis. In the limit in which the connectivity is very sparse, so that correlations between the spike trains of different cells can be neglected (see section 2), the relevant parameter is the number of afferent connections per neuron C . We will investigate the behavior of the network in the limit $C \rightarrow \infty$ (the “extensive” limit) since, in this case, the different scenarios become qualitatively different. Of course, physical neurons receive a finite number of connections, but the rationale is that the physical solution can be considered a small perturbation to the solution found in the $C = \infty$ case, which is much easier to characterize. One should keep in mind that even if C becomes very large, we still need to impose the sparseness condition for our analysis to be valid, which implies that it should always hold that $N \gg C$.

When considering current-based scenarios in the extensive limit, one is forced to normalize the connection strengths (the size of the unitary PSPs, which we denote generally by J) by (some power of) the number of connections per cell C , in order to keep the total afferent current within the (presynaptic) dynamic range of the neuron (whose order of magnitude is given by the distance between reset and threshold). As we show below,

different scaling schemes of J with C lead to different relative magnitudes of the mean and fluctuations of the afferent current into the cells in the extensive limit, and this in turn determines the type of steady-state solutions for the network. We thus proceed to analyze the expressions for the mean and variance of the afferent current (see equations 2.5 and 2.6) under different scaling assumptions.

We consider multicolumnar networks in which the C afferent connections to a given cell come from N_c different “columns” (each contributing C_c connections, so that $C = N_c C_c$). Each column is composed of an excitatory and an inhibitory subpopulation. The multicolumnar structure of the network is specified by the following scaling relationships,

$$\begin{aligned} N_c &\equiv n_c C^\alpha \\ C_c &\equiv n_c^{-1} C^{1-\alpha}, \end{aligned}$$

with $0 \leq \alpha \leq 1$ and n_c order one, that is, independent of C . The case $\alpha = 0$ corresponds to a finite number n_c of columns, each contributing a number of connections of order C . The case $\alpha = 1$ corresponds to an extensive number of columns, each contributing a number of connections of order one—that is, a fixed number as the total number of connections C grows.

Although connection strengths between the different subpopulations can all be different, we assume that they can be classified into two types according to their scaling with C : those between cells within the same column, of strength J_w , and those between cells belonging to different columns, of strength J_b (the scaling is assumed to be the same for excitatory and for inhibitory connections). We define

$$\begin{aligned} J_w &\equiv j_w C^{-\alpha_w} \\ J_b &\equiv j_b C^{-\alpha_b} \end{aligned}$$

where $\alpha_w, \alpha_b > 0$ and the j 's are all order one.

In these conditions, the afferent current to the excitatory or inhibitory cells (it does not matter which, for this analysis) from their own subpopulation is characterized by

$$\begin{aligned} \mu_{\text{in}} &= C_c [J_{E_w} \nu_{E_{\text{in}}} - J_{I_w} \nu_{I_{\text{in}}}] \\ &= C^{1-\alpha-\alpha_w} [j_{E_w} \nu_{E_{\text{in}}} - j_{I_w} \nu_{I_{\text{in}}}] n_c^{-1} \equiv C^{1-\alpha-\alpha_w} f_{\mu_{\text{in}}} \\ \sigma_{\text{in}}^2 &= C_c [J_{E_w}^2 \nu_{E_{\text{in}}} \text{CV}_{E_{\text{in}}}^2 + J_{I_w}^2 \nu_{I_{\text{in}}} \text{CV}_{I_{\text{in}}}^2] \\ &= C^{1-\alpha-2\alpha_w} [j_{E_w}^2 \nu_{E_{\text{in}}} \text{CV}_{E_{\text{in}}}^2 + j_{I_w}^2 \nu_{I_{\text{in}}} \text{CV}_{I_{\text{in}}}^2] n_c^{-1} \equiv C^{1-\alpha-2\alpha_w} f_{\sigma_{\text{in}}}, \end{aligned}$$

where the f 's are linear combinations of rates and CVs weighted by factors

of order one. We proceed by assuming that all other columns are in the same state $v_{E,I_{\text{out}}}$, $\text{CV}_{E,I_{\text{out}}}$, so that the current to the cells in the column under focus from the rest of the network is characterized by

$$\begin{aligned}\mu_{\text{out}} &= (N_c - 1)C_c [J_{E_b} v_{E_{\text{out}}} - J_{I_b} v_{I_{\text{out}}}] \\ &= C^{1-\alpha_b} (1 - n_c^{-1} C^{-\alpha}) [j_{E_b} v_{E_{\text{out}}} - j_{I_b} v_{I_{\text{out}}}] \equiv C^{1-\alpha_b} (1 - n_c^{-1} C^{-\alpha}) f_{\mu_{\text{out}}} \\ \sigma_{\text{out}}^2 &= (N_c - 1)C_c [J_{E_b}^2 v_{E_{\text{out}}} \text{CV}_{E_{\text{out}}}^2 + J_{I_b}^2 v_{I_{\text{out}}} \text{CV}_{I_{\text{out}}}^2] \\ &= C^{1-2\alpha_b} (1 - n_c^{-1} C^{-\alpha}) [j_{E_b}^2 v_{E_{\text{out}}} \text{CV}_{E_{\text{out}}}^2 + j_{I_b}^2 v_{I_{\text{out}}} \text{CV}_{I_{\text{out}}}^2] \\ &\equiv C^{1-2\alpha_b} (1 - n_c^{-1} C^{-\alpha}) f_{\sigma_{\text{out}}}.\end{aligned}$$

In the extensive limit, the terms $1 - n_c^{-1} C^{-\alpha}$ become equal to one if $\alpha > 0$ and are of order one if $\alpha = 0$, in which case they can be included as an extra multiplicative factor in the f terms. We thus omit them from now on.

In addition to their recurrent inputs, cells receive a similar number of external excitatory inputs as well, but since we are interested in the generation of irregularity by the network, we will assume this external drive to be deterministic, that is, characterized by

$$\mu^{\text{ext}} = C J_{\text{ext}} v_{\text{ext}} = C^{1-\alpha_{\text{ext}}} j_{\text{ext}} v_{\text{ext}} \equiv C^{1-\alpha_{\text{ext}}} f_{\text{ext}},$$

with $J_{\text{ext}} = j_{\text{ext}} C^{-\alpha_{\text{ext}}}$ and $\alpha_{\text{ext}} > 0$.

The scaling with C of the different components of the total afferent current is thus given by

$$\begin{aligned}\mu_{\text{in}} &= C^{1-\alpha-\alpha_w} f_{\mu_{\text{in}}} & \sigma_{\text{in}}^2 &= C^{1-\alpha-2\alpha_w} f_{\sigma_{\text{in}}} \\ \mu_{\text{out}} &= C^{1-\alpha_b} f_{\mu_{\text{out}}} & \sigma_{\text{out}}^2 &= C^{1-2\alpha_b} f_{\sigma_{\text{out}}} \\ \mu^{\text{ext}} &= C^{1-\alpha_{\text{ext}}} f_{\text{ext}}.\end{aligned}$$

If α , α_b , α_w are such that the variances vanish as $C \rightarrow \infty$, the corresponding networks will consist of regularly spiking neurons. Since we are interested in irregular spiking, we therefore look for solutions in which σ_{in}^2 , σ_{out}^2 or both remain order one in the $C \rightarrow \infty$ limit. There are several ways to achieve this.

3.1.1 Scenario 1: Homogeneous Balanced Network. This case is associated with the choice $\alpha = 0$. In this case, the size of the columns is of the same order as the size of the whole network (i.e., the number of columns, n_c , is order one), in which case the in and out quantities become equivalent. A finite variance is achieved by setting $\alpha_w = \alpha_b = 1/2$, that is, $J \propto 1/\sqrt{C}$. This scenario is equivalent to the network studied originally in Tsodyks and Sejnowski (1995) and van Vreeswijk and Sompolinsky (1996, 1998). In such

a network, the mean input from the recurrent network grows as the square root of the number of inputs,

$$\mu_{\text{in}} + \mu_{\text{out}} = \sqrt{C}(f_{\mu_{\text{in}}} + f_{\mu_{\text{out}}}).$$

This quantity can be positive or negative depending on the excitation-inhibition balance in the network. The overall mean input into the neurons is obtained by adding the external input: $\mu = \mu_{\text{in}} + \mu_{\text{out}} + \mu^{\text{ext}}$. In order not to saturate the dynamic range of the cell in the extensive limit, the overall mean current into the neurons should remain of order one as $C \rightarrow \infty$. Hence, it is needed that

$$\mu = \sqrt{C}[f_{\mu_{\text{in}}} + f_{\mu_{\text{out}}} + f_{\text{ext}}C^{\frac{1}{2}-\alpha_{\text{ext}}}] \sim O(1), \quad (3.1)$$

which is possible only if the term in square brackets vanishes as $1/\sqrt{C}$. If the synapses from the external inputs vanish like $1/C$ ($\alpha_{\text{ext}} = 1$), then the external input has a negligible contribution to the overall mean input to the cells. In this case, since both $f_{\mu_{\text{in}}}$ and $f_{\mu_{\text{out}}}$ are linear combinations of the firing rates of the neurons inside and outside the column under focus, equation 3.1 effectively becomes, in the extensive limit, a set of linear homogeneous equations for the activity of the different columns (note that although we have, for brevity, written only one, there are four equations like equation 3.1, for the excitatory and inhibitory subpopulations inside and outside the column under focus). Thus, unless the matrix of coefficients of the firing rates in $f_{\mu_{\text{in}}}$ and $f_{\mu_{\text{out}}}$ for the excitatory and inhibitory subpopulations is not full rank, the only solution of equation 3.1 in the extensive limit is given by a silent, zero rate, network (van Vreeswijk & Sompolinsky, 1998). On the other hand, if $\alpha_{\text{ext}} = 1/2$, the linear system defined by equations 3.1 in the extensive limit is not homogeneous anymore. Hence, in the general case (except for degeneracies), if $\alpha_{\text{ext}} = 1/2$, there is a single self-consistent solution for the firing rates in the network, in which the activity in each subpopulation is proportional to the external drive v_{ext} (van Vreeswijk & Sompolinsky, 1998). This highlights the importance of a powerful external excitatory drive. When $\alpha_{\text{ext}} = 1/2$, the external drive by itself would drive the cells to saturation if the recurrent connections were inactivated. In the presence of the recurrent connections, the activities in the excitatory and inhibitory subpopulations adjust themselves to compensate this massive external drive. The firing rates in the self-consistent solution correspond to the only way in which this compensation can occur for all the subpopulations at the same time. It follows that inasmuch as the different inputs to the neuron combine linearly, unless the connectivity matrix is degenerate, which requires some kind of fine-tuning mechanism, bistability in a large, homogeneous balanced network is not a robust property.

3.1.2 Scenario 2: Homogeneous Multicolumnar Balanced Network. Since the linearity condition that results in a unique solution follows from the mean current to the cells from within the column growing with C , we impose that $\mu_{\text{in}} \sim O(1)$, that is, $1 - \alpha - \alpha_w = 0$. If this is the case, the variance coming from within the column goes as $C^{\alpha-1}$. We consider first the case $\alpha < 1$. In these conditions, the variance coming from the activity inside the column vanishes for large C . To keep the total variance finite, we set $\alpha_b = 1/2$. If we also choose $\alpha = 1/2$, then $\alpha_w = \alpha_b = 1/2$, so the network is homogeneous in that all the connection strengths scale similarly with C regardless of whether they connect cells belonging to the same or different columns. Since $\alpha = 1/2$, there are many columns in the network, and the number of connections coming from inside a particular column is a very small fraction (which decreases like $1/\sqrt{C}$) of the total number of afferent connections to the cell.

The fact that $\alpha_b = 1/2$ implies that the mean current coming from outside the column will still grow like \sqrt{C} . Thus, in order for the cells not to saturate, the excitation and the inhibition from outside the column have to balance precisely; the steady state of the rest of the network becomes, again, a unique, linear function of the external input to the cells (where again we choose $\alpha_{\text{ext}} = 1/2$ to avoid the quiescent state). However, now the mean current coming from inside the column is independent of C , so the steady-state activity inside the column is not determined by a set of linear equations. Instead, it should be determined self-consistently using the nonlinear transfer function in equation 2.3, which, in principle, permits bistability. This scenario is, in fact, equivalent to the one studied in Brunel (2000b), where a systematic analysis of the conditions in which bistability in a network like this can exist has been performed. (Although no explicit scaling of the synaptic connection strengths with C was assumed in Brunel, 2000b, the essential fact that the total variance to the cells in the subpopulation that supports bistability is constant is considered in that article.) As will be shown in detail below, the fact that the potential multiple steady-state solutions in this scenario differ only in the mean current to the cells, not in their variance (which is fixed by the balance condition on the rates outside the column), leads necessarily to a lower (in general, significantly lower) CV in the activity in the cells in the elevated persistent activity state. Therefore, in a network with $J \propto 1/\sqrt{C}$ scaling, bistability is possible in small subsets of neurons comprising a fraction $\propto 1/\sqrt{C}$ of the total number of connections per cell, but the elevated persistent activity states are characterized by a change in the mean drive to the cells at constant variance, and, as we show below, this leads to a significant decrease in the spiking irregularity in the elevated persistent activity states.

3.1.3 Scenario 3: Heterogeneous Multicolumnar Network. In order for the CV in the elevated persistent activity state to remain close to one, the variance of the afferent current to the cells inside the column should depend

on their own activity. Thus, in addition to the condition $1 - \alpha - \alpha_w = 0$ necessary for bistability, we have to impose that $\sigma_{\text{in}}^2 \propto C^{\alpha-1}$ be independent of C , that is, $\alpha = 1$, which implies $\alpha_w = 0$. In these conditions, the extensive number of connections per cell come from an extensive number of columns, with the number of connections from each column remaining a finite number. The $\alpha_w = 0$ condition reflects the fact that since cells receive only a finite number of intracolumnar connections, the strength of these connections does not need to scale in any specific way with C . As for the activity outside the network, one could now, in principle, choose either $J^b \propto 1/C$ or $J^b \propto 1/\sqrt{C}$ (corresponding to $\alpha_b = 1, 1/2$, respectively), since there is already a finite amount of variance coming from within the column. In the first case, the rest of the network contributes only a noiseless deterministic current whose exact amount has to be determined self-consistently, and in the second it contributes both to the total mean and variance of the afferent current to the neurons in the column. In this last case, as in the previous two scenarios, the $J^b \propto 1/\sqrt{C}$ scaling results in the need for balance between the total excitation and inhibition outside the network, which (again, if $\alpha_{\text{ext}} = 1/2$) leads to a unique solution for the activity of the rest of the population linear in the external drive to these neurons.

In this scenario, the network is heterogeneous since the strength of the connections from neurons within the same column is larger than those from neurons in other columns. Since the rate and CV of the cells inside the column have to be determined self-consistently in this case, we proceed to do a systematic quantitative analysis of this scenario in the next section. From the scaling considerations described in this section, it is already clear, though, that a potential bistable solution with high CV is possible only in a small network.

3.2 Types of Bistable Solutions in a Reduced Network. In this section, we consider the network described in scenario 3 in the previous section, with the choice $\alpha_b = 1/2$, and analyze the types of steady-state solutions for the activity in a particular column of finite size. The rest of the network is in a balanced state, and its activity is completely decoupled from the activity of the column, which is too small in size to make a difference in the overall input to the rest of the cells. For our present purposes, all that matters about the afferents outside the column (from both the rest of the network and the external ones) is that they provide a finite net input to the cells in the column. We denote the mean and variance of that fixed external current by $\mu_{E,I}^{\text{ext}}/\tau_m$ and $2(\sigma_{E,I}^{\text{ext}})^2/\tau_m$, where the factors with the membrane time constant τ_m have been included so that $\mu_{E,I}^{\text{ext}}$ and $(\sigma_{E,I}^{\text{ext}})^2$ represent the contribution to the mean and variance of the postsynaptic depolarization (in the absence of threshold) in the steady states arising from outside the column.

The net input to the excitatory and inhibitory populations in the column under focus is thus characterized by

$$\begin{aligned}\mu_E &= c_{EE} j_{EE} v_E - c_{EI} j_{EI} v_I + \frac{\mu_E^{\text{ext}}}{\tau_m} \\ \mu_I &= c_{IE} j_{IE} v_E - c_{II} j_{II} v_I + \frac{\mu_I^{\text{ext}}}{\tau_m} \\ \sigma_E^2 &= c_{EE} j_{EE}^2 v_E^2 \text{CV}_E^2 + c_{EI} j_{EI}^2 v_I^2 \text{CV}_I^2 + \frac{(\sigma_E^{\text{ext}})^2}{\tau_m/2} \\ \sigma_I^2 &= c_{IE} j_{IE}^2 v_E^2 \text{CV}_E^2 + c_{II} j_{II}^2 v_I^2 \text{CV}_I^2 + \frac{(\sigma_I^{\text{ext}})^2}{\tau_m/2},\end{aligned}$$

where the number of connection and connection strength parameters c and j are all order one. We proceed by simplifying this scheme further in order to reduce the dimensionality of the system from four to two, which will allow a systematic exploration of the effect of all the parameters on the type of steady-state solutions of the network. In particular, we make the inputs to the excitatory and inhibitory populations identical,

$$\begin{aligned}c_{IE} = c_{EE} \equiv c_E \quad c_{II} = c_{EI} \equiv c_I \quad j_{IE} = j_{EE} \equiv j_E \quad j_{II} = j_{EI} \equiv j_I \\ \mu_E^{\text{ext}} = \mu_I^{\text{ext}} \equiv \mu^{\text{ext}} \quad (\sigma_E^{\text{ext}})^2 = (\sigma_I^{\text{ext}})^2 \equiv (\sigma^{\text{ext}})^2,\end{aligned}$$

so that the whole column becomes statistically identical: $v_E = v_I \equiv v$ and $\text{CV}_E = \text{CV}_I \equiv \text{CV}$. For simplicity, we also assume that the number of excitatory and inhibitory inputs is the same: $c_E = c_I = c$. Thus, we are left with a system with four parameters,

$$c_\mu = c(j_E - j_I); \quad c_\sigma = \sqrt{c(j_E^2 + j_I^2)}; \quad \mu^{\text{ext}}; \quad \sigma^{\text{ext}}, \quad (3.2)$$

all with units of mV, and two dynamical variables (from equation 2.7)

$$\frac{d\mu_V}{dt} = -\frac{\mu_V}{\tau_m} + \mu; \quad \frac{d\sigma_V^2}{dt} = -\frac{\sigma_V^2}{\tau_m/2} + \sigma^2, \quad (3.3)$$

where

$$\mu = c_\mu v + \frac{\mu^{\text{ext}}}{\tau_m}; \quad \sigma^2 = c_\sigma^2 v \text{CV}^2 + \frac{(\sigma^{\text{ext}})^2}{\tau_m/2}$$




Figure 2: Mean firing rate ν (left), CV (middle), and product νCV^2 (right) of the LIF neuron as a function of the mean and standard deviation of the depolarization. Parameters: $V_{th} = 20$ mV, $V_{res} = 10$ mV, $\tau_m = 10$ ms, and $\tau_{ref} = 2$ ms.

and

$$\nu(\mu_V, \sigma_V)^{-1} = \tau_{ref} + \tau_m \sqrt{\pi} \int_{\frac{V_{res} - \mu_V}{\sigma_V \sqrt{2}}}^{\frac{V_{th} - \mu_V}{\sigma_V \sqrt{2}}} dx e^{x^2} [1 + \operatorname{erf}(x)] \quad (3.4)$$

$$CV^2(\mu_V, \sigma_V) = 2\pi \nu^2 \int_{\frac{V_{res} - \mu_V}{\sigma_V \sqrt{2}}}^{\frac{V_{th} - \mu_V}{\sigma_V \sqrt{2}}} dx e^{x^2} \int_{-\infty}^x dy e^{y^2} [1 + \operatorname{erf}(y)]^2 \quad (3.5)$$

The parameters c_μ and c_σ^2 measure the strength of the feedback that the activity in the column produces on the mean and variance of the current to the cells. c_μ can be less than equal to, or greater than zero. A value larger (less) than zero implies that the activity in the column has a net excitatory (inhibitory) effect on the neurons. In general, we assume the positive parameter c_σ^2 to be independent of c_μ (implying the recurrent feedback on the mean and on the variance can be manipulated independently). Note, however, that since $j_I/j_E > 0$, c_σ^2 cannot be arbitrarily small, that is, $c_\sigma^2 > c_\mu^2/c$. Equations 3.4 and 3.5 are plotted as a function of μ_V and σ_V in Figure 2.

3.2.1 Mean- and Fluctuation-Driven Bistability in the Reduced System. The nullclines for the two equations 3.3 are given by

$$\nu(\mu_V, \sigma_V) = \frac{\mu_V - \mu^{ext}}{\tau_m c_\mu}; \quad \nu(\mu_V, \sigma_V) CV^2(\mu_V, \sigma_V) = \frac{\sigma_V^2 - (\sigma^{ext})^2}{(\tau_m/2)c_\sigma^2}.$$

The values of (μ_V, σ_V) that satisfy these equations are shown in Figure 3 for several values of the parameters c_μ, c_σ . The nullclines for the mean (see Figure 3, left) are the projection on the (μ_V, σ_V) plane of the intersection of the surface in Figure 2, left, with a plane parallel to the σ_V axes, shifted by μ^{ext} and tilted (i.e., with slope) at a rate $1/(\tau_m c_\mu)$. Since the mean firing

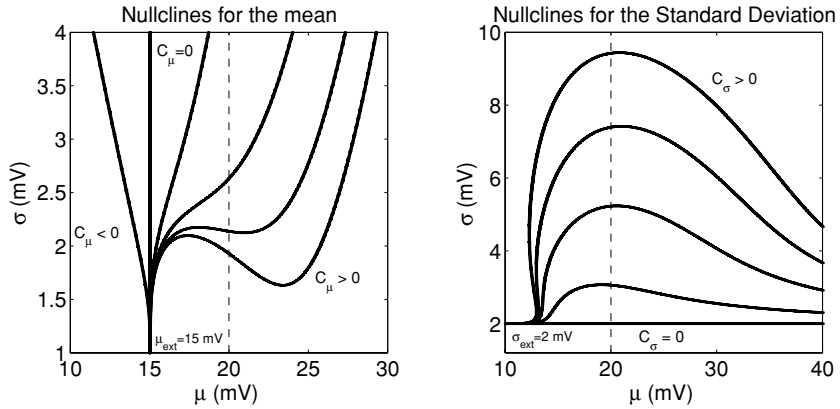


Figure 3: Nullclines of the equation for the mean μ_V (left) and standard deviation σ_V (right). The nullcline for the mean depends on only c_μ , and the one for the standard deviation depends on only c_σ .

rate as a function of the average depolarization changes curvature (it has an inflection point near threshold, where firing changes from being driven by fluctuations to being driven by the mean), the nullcline for the mean has a “knee” when the net feedback c_μ is large enough and excitatory.

Similarly, the nullclines for the standard deviation of the depolarization (see Figure 3, right) are the projection on the (μ_V, σ_V) plane of the intersection of the surface in Figure 2, right, with a parabolic surface parallel to the μ_V axes, shifted by $(\sigma^{\text{ext}})^2$ and with a curvature $2/\tau_m c_\sigma^2$. Again, this curve can display a “knee” for high enough values of the net strength of the feedback onto the variance c_σ^2 .

The fixed points of the system are given by the points of intersection of the two nullclines. We now show, through two particular examples, the main result of this letter: depending on the degree of balance between excitation and inhibition, two types of bistability can exist: mean driven and fluctuation driven.

Mean-Driven Bistability. Figure 4 shows a particular example of the type of bistability obtained for low-moderate values of c_σ and moderate-high values of c_μ . Figure 4a shows the time evolution of the firing rate (bottom) and CV (top) in the network when the external drive to the cells is transiently elevated. In response to the transient input, the network switches between a low-rate, high-CV basal state, into an elevated activity state. For this particular type of bistability, the CV in this state is low. The nullclines for this example are shown in Figure 4b. The σ_V nullcline is essentially parallel to the μ_V axis, and it intercepts the μ_V nullcline (which has a pronounced “knee”) at three points: one below (stable), one around (unstable), and one above




Figure 4: Example of mean-driven bistability. (a) CV (top) and firing rate (bottom) in the network as a function of time. Between $t = 0$ s and $t = 0.5$ s (dashed lines), the mean of the external drive to the neurons was elevated from 18 mV to 19 mV, causing the network to switch to its elevated activity fixed point. In this fixed point, the CV is low. (b) Nullclines for this example. The two stable fixed points differ primarily in the mean current that the cells are receiving, with an essentially constant variance. Hence, the CV in the elevated persistent-activity fixed point is low. Parameters: $\mu^{\text{ext}} = 18$ mV, $\sigma^{\text{ext}} = 0.65$ mV, $c_\mu = 7.2$ mV, $c_\sigma = 1$ mV. Dotted line: neuronal threshold.

(stable) threshold. The stable fixed point below threshold corresponds to the state of the system before the external input is transiently elevated in Figure 4a. It is typically characterized by a low rate and a relatively high CV, as subthreshold spiking is fluctuation driven and thus irregular. However, since the CV drops fast for $\mu_V > V_{\text{th}}$ at the values of σ_V at which the μ_V

nullcline bends (see Figure 2, middle), the CV in the suprathreshold fixed point is typically low (but see section 4). Fluctuations in the current to the cells play little role in determining the spike times of the cells in this elevated persistent activity state.

Qualitatively, this is the type of excitation-driven bistability that has been thoroughly analyzed over the past few years. It is expected to be present in networks in which small populations of excitatory cells can be bistable in the presence of global inhibition by virtue of selective synaptic potentiation. It is also expected to be present in larger populations if the fluctuations arising from the recurrent connections are weak compared to those coming from external afferents, for instance, due to synaptic filtering (Wang, 1999).

Fluctuation-Driven Bistability. If the connectivity is such that the mean drive to the neurons is only weakly dependent on the activity, that is, c_μ is small, but at the same time the activity has a strong effect on the variance, that is, c_σ is large, the system can also have two stable fixed points, as shown in the example in Figure 5 (same format as in the previous figure). In this situation, however, the two fixed points are subthreshold, and they differ primarily in the variance of the current to the cells. Hence, spiking in both fixed points is fluctuation driven, and the CV is high in both of them; in particular, it is slightly higher in the elevated activity fixed point (see Figure 5a).

This type of bistability can be realized only if there is a precise balance between the net excitatory and inhibitory drive to the cells. Since c_σ must be large in order for the σ_V nullcline to display a “knee,” both the net excitation and inhibition should be large, and in these conditions, a small c_μ can be achieved only if the balance between the two is precise. This suggests that this regime will be sensitive to changes in the parameters determining the connectivity; that is, it will require fine-tuning, a conclusion that is supported by the analysis below.

Mean- and fluctuation-driven bistability are not discrete phenomena. Depending on the values of the parameters, the elevated activity fixed point can rotate in the (μ_V, σ_V) plane, spanning intermediate values from those shown in the examples in Figures 4 and 5. We thus now proceed to a systematic characterization of all possible behaviors of the reduced system as a function of its four parameters.

3.2.2 Effect of the External Current. Since $c_\sigma^2 > 0$, the σ_V nullcline always bends upward (see Figure 3), that is, the values of σ_V in the nullcline are always larger than σ^{ext} . Assuming for simplicity that c_σ can be arbitrarily low, this implies that no bistability can exist unless the external variance is low enough. In particular, for every value of the external mean μ^{ext} , there is a critical value σ_{c1}^{ext} defined as the value of σ_V at which the first two derivatives of the μ_V nullcline vanish (see Figure 6, middle). For $\sigma^{\text{ext}} > \sigma_{c1}^{\text{ext}}$, the two nullclines can cross only once, and therefore no bistability of any

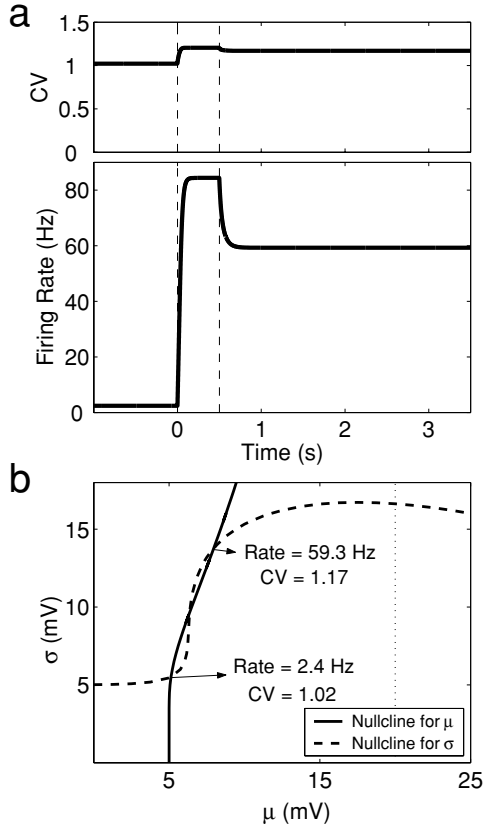


Figure 5: Example of fluctuation-driven bistability. (a) CV (top) and firing rate (bottom) in the network as a function of time. Between $t = 0$ s and $t = 0.5$ s (dashed lines), the standard deviation of the external drive to the neurons was elevated from 5 mV to 7 mV, causing the network to switch to its elevated activity fixed point. In this fixed point, the CV is slightly higher than in the basal state. (b) Nullclines for this example. The two stable fixed points differ primarily in the variance of the current that the cells are receiving, with little change in the mean. Hence, the CV in the elevated persistent-activity fixed point is slightly higher than in the low-activity fixed point; that is, the CV increases with the rate. Parameters: $\mu^{\text{ext}} = 5$ mV, $\sigma^{\text{ext}} = 5$ mV, $c_\mu = 5$ mV, $c_\sigma = 20.2$ mV. Dotted line: Neuronal threshold.

kind is possible in the reduced system (see Figure 6, left). For values of σ^{ext} only slightly lower than this critical value, the jump between the low- and the high-activity stable fixed points in the (μ_V, σ_V) plane is approximately horizontal, so the type of bistability obtained is mean driven. For lower




Figure 6: The external variance determines the existence and types of bistability possible in the network. For $\sigma^{\text{ext}} > \sigma_{c1}^{\text{ext}}$ (left), no bistability is possible. $\sigma^{\text{ext}} = \sigma_{c1}^{\text{ext}}$ marks the onset of bistability (middle). At $\sigma^{\text{ext}} = \sigma_{c2}^{\text{ext}}$ bistability becomes possible in a perfectly balanced network (a network with $c_\mu = 0$) (right).

values of the external variance, a point is eventually reached at which bistability becomes possible in a perfectly balanced network. Again, for each μ^{ext} , one can define a second critical value σ_{c2}^{ext} in the following way: σ_{c2}^{ext} is the value of σ^{ext} at which the point where the derivative at the inflection point of the σ_V nullcline is infinite occurs at a value of μ_V equal to μ^{ext} (see Figure 6, right). For values of $\sigma^{\text{ext}} < \sigma_{c2}^{\text{ext}}$, bistability is possible in networks in which the net recurrent feedback is inhibitory.

Since both critical values of σ^{ext} are functions of the external mean, they define curves in the $(\mu^{\text{ext}}, \sigma^{\text{ext}})$ plane. These curves are plotted in Figure 7. Both are decreasing functions of μ^{ext} and meet at threshold, implying that bistability in the reduced network is possible only for subthreshold mean external inputs (see section 4).

3.2.3 Phase Diagrams of the Reduced Network. For each point in the $(\mu^{\text{ext}}, \sigma^{\text{ext}})$ plane, the external current is completely characterized, and the only two parameters left to be specified are c_μ, c_σ . In particular, in the regions where bistability is possible, it will exist for only appropriate values of c_μ and c_σ . The two insets in Figure 7 show phase diagrams in the (c_μ, c_σ) plane showing the regions of bistability in two representative points: one in which $\sigma_{c2}^{\text{ext}} < \sigma^{\text{ext}} < \sigma_{c1}^{\text{ext}}$, in which bistability is possible only in excitation-dominated networks (top-right inset), and one in which $\sigma^{\text{ext}} < \sigma_{c2}^{\text{ext}}$, in which bistability is possible in both excitation- and inhibition-dominated networks (bottom-left inset). In this latter case, the region enclosed by the curve in which bistability can exist stretches to the left, including the region with $c_\mu \lesssim 0$.

We have further characterized the nature of the fixed-point solutions in these two cases by plotting the rate and CV on each point in the (c_μ, c_σ) plane on which bistability is possible, as well as the ratio between the rate and CV in the high- versus low-activity states. Instead of showing this in the

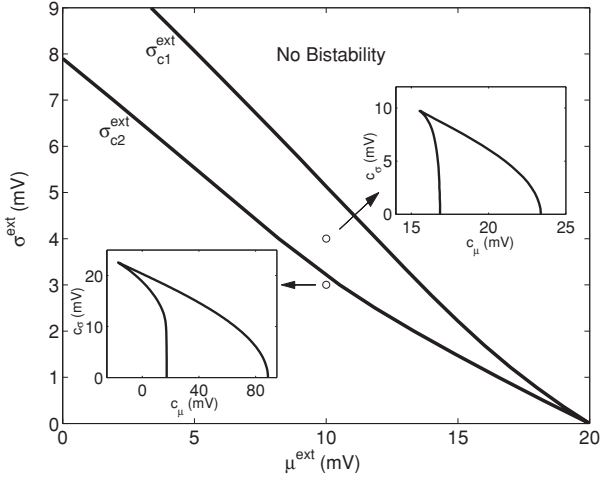


Figure 7: Phase diagram with the effect of the mean and variance of the external current on the existence and types of bistability in the network. The two insets represent the regions of bistability in the (c_μ, c_σ) plane at the corresponding points in the $(\mu^{\text{ext}}, \sigma^{\text{ext}})$ plane. Fluctuation-driven bistability is possible only near and below the lower critical line $\sigma^{\text{ext}} = \sigma_{c2}^{\text{ext}}(\mu^{\text{ext}})$. Top-right and bottom-left insets correspond to $\mu^{\text{ext}} = 10$ mV; $\sigma^{\text{ext}} = 4$ mV and $\mu^{\text{ext}} = 10$ mV; $\sigma^{\text{ext}} = 3$ mV, respectively.

(c_μ, c_σ) plane, we have inverted Equations 3.2 to show (assuming a constant $c = 100$) the results as a function of the unitary EPSP j_E and of the ratio of the unitary inhibitory to excitatory PSPs j_I/j_E , which measures the degree of balance in the network, that is, $j_I/j_E = 1$ implies a perfect balance:

$$j_E = \frac{c_\mu + \sqrt{2cc_\sigma^2 - c_\mu^2}}{2c}; \quad j_I/j_E = \frac{\sqrt{2cc_\sigma^2 - c_\mu^2} - c_\mu}{\sqrt{2cc_\sigma^2 - c_\mu^2} + c_\mu}. \quad (3.6)$$

In Figure 8 we show the results for the case where the external variance is higher than σ_{c2}^{ext} , so that bistability is possible only if the net recurrent connectivity is excitatory. Overall, the shape of the bistable region in this space is a diagonal to the right. This means that closer to the balanced region, the net excitatory drive (proportional to j_E) has to be higher in order for bistable solutions to exist. The low-activity fixed point (left column) is subthreshold, and thus spiking is fluctuation driven, characterized by a high CV. In this case, the high-activity fixed point is suprathreshold, so the CV in this fixed point is in general small (see the bottom right). Of course,




Figure 8: Bistability in an excitation-dominated network in the $(j_E, j_I/j_E)$ plane. (Top and bottom left) Firing rate and CV in the low-activity fixed point. (Top and bottom right) Ratio of the firing rate and CV between the high- and low-activity fixed points. (Inset) same as bottom-right panel in the (c_μ, c_σ) plane. Parameters: $c = 100$, $\mu^{\text{ext}} = 10$ mV, and $\sigma^{\text{ext}} = 4$ mV.

very close to the cusp, at the onset of bistability, the CV (and rate) in both fixed points is similar.

In Figure 9 the same information is shown for the case where the external variance is lower than $\sigma_{c_2}^{\text{ext}}$ so that bistability is possible when the recurrent connectivity is dominated by excitation or inhibition. In this case, the region where the CV in the high- and low-activity fixed points is similar is larger, corresponding to situations in which $j_I/j_E \sim 1$, that is, excitation and inhibition in the network are roughly balanced. Only in this region is the firing rate in the high-activity state not too high, $\lesssim 100$ Hz. In this regime, when excitation dominates, the rate in the high-activity state becomes very high. Note also that the transition between the relatively low-rate, high-CV regime and the very high-rate, low-CV regime at $j_I/j_E \sim 0.9$ is relatively abrupt.

Finally, we can use the relations 3.6 to study quantitatively the effect of the number of connections c on the regions of bistability, something we




Figure 9: Mean and fluctuation-driven bistability in the $(j_E, j_I/j_E)$ plane. Panels as in Figure 8. (Inset) Portion of the bistability region with fluctuation-driven fixed points in the (c_μ, c_σ) plane. When the network is approximately balanced, that is, $j_I/j_E \sim 1$, the CV in the high-activity state is high. Parameters: $c = 100$, $\mu^{\text{ext}} = 10$ mV, and $\sigma^{\text{ext}} = 3$ mV.

did in section 3.1 at a qualitative level based on scaling considerations. Figure 10 shows the effect of increasing the number of afferent connections per cell on the shape of the region of bistability for $\sigma^{\text{ext}} = 4$ mV (left) and for $\sigma^{\text{ext}} = 3$ mV (right). The results are clearer when shown in the plane $(j_{EC}, j_I/j_E)$, where j_{EC} represents the net mean excitatory drive to the neurons. The range of values of the net excitatory drive in which bistability is allowed in the excitation-dominated regime, where $j_I/j_E < 1$, does not depend very strongly on c . However, for both $\sigma^{\text{ext}} = 4$ mV and $\sigma^{\text{ext}} = 3$ mV, when inhibition and excitation become more balanced, a higher net excitatory drive is needed. In particular, when $\sigma^{\text{ext}} = 3$, the bistable region always includes the balanced network, $j_I/j_E = 1$, but the range of values of $j_I/j_E \sim 1$ where bistability is possible (being in this case fluctuation driven) considerably shrinks. Thus, as noted in section 3.1, bistability in a large, balanced network requires a precise balance of excitation and inhibition.




Figure 10: Effect of number of connections c on the phase diagram in the $(j_E c, j_I/j_E)$ plane for the case where $\mu^{\text{ext}} = 10$ mV and $\sigma^{\text{ext}} = 4$ mV (left) and $\sigma^{\text{ext}} = 3$ mV (right). Note that in the right panel, $j_I/j_E = 1$ is always in the region of bistability, but the range of values of $j_I/j_E \sim 1$ in the bistable region decreases significantly with c .

The precise balance between excitation and inhibition required to obtain solutions with fluctuation-driven bistability is also evident when one analyzes the effect of changing the net input to the excitatory and inhibitory subpopulations within a column. In this case, the excitatory and inhibitory mean firing rate and CV become different. We have chosen to study the effect of the different ratios of excitation and inhibition to the excitatory and inhibitory populations. In particular, defining

$$\gamma_E \equiv \frac{j_{EI}}{j_{EE}} \quad \text{and} \quad \gamma_I \equiv \frac{j_{II}}{j_{IE}},$$

we have considered the effect of having $\gamma_E \neq \gamma_I$ while still considering that the excitatory connections to the excitatory and inhibitory populations are equal, that is, $j_{EE} = j_{IE} \equiv j_E$. To proceed with the analysis, we started by specifying a point in the parameter space of the symmetric network in which excitation and inhibition were identical by choosing a value for $(\mu^{\text{ext}}, \sigma^{\text{ext}}, c_\mu, c_\sigma)$. Then, fixing $c = 200$, we used the relationships 3.6 to solve for j_E and γ and, defining $\gamma_E \equiv \gamma$, we found which values of γ_I resulted in bistable solutions. Correspondingly, when $\gamma_I = \gamma_E$, the two subpopulations within the column become identical again.

We performed this analysis for two initial sets of parameters of the symmetric network: one corresponding to mean-driven and the other to fluctuation-driven bi-stability. The results of this analysis are shown in Figure 11. The type of bistability does not change qualitatively depending on the value of γ_I/γ_E in the mean-driven case (left column). For the right

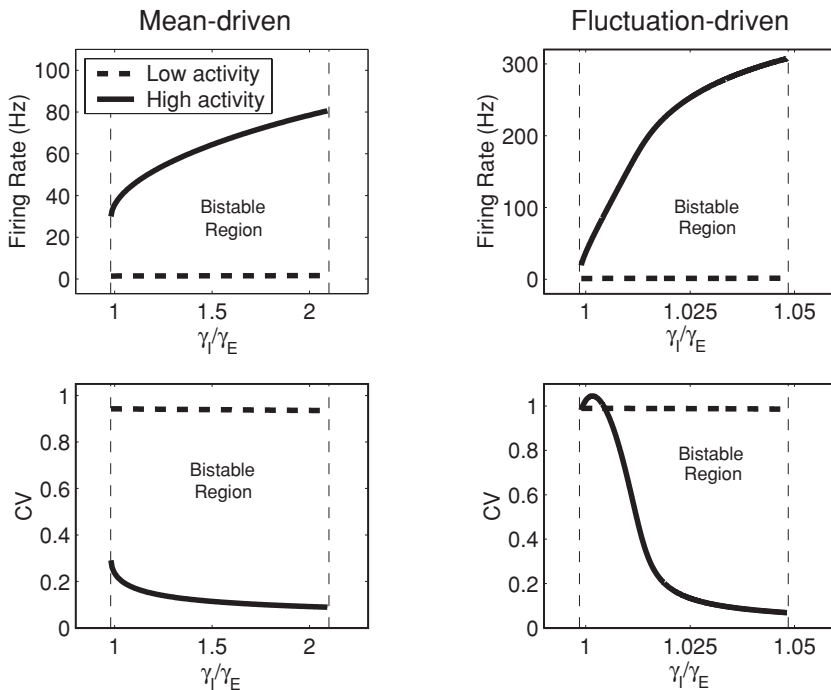


Figure 11: Effect of different levels of input to the excitatory and inhibitory subpopulations. The ratio between the inhibitory and excitatory connection strengths γ was allowed to be different for each subpopulation. Left and right columns correspond to mean- and fluctuation-driven bistability in the corresponding situation for the symmetric network. The network is bistable for values of γ_I/γ_E within the dashed lines. Parameters on the left column: $\mu^{\text{ext}} = 18$ mV, $\sigma^{\text{ext}} = 0.65$ mV, $c_\mu = 7.2$ mV, $c_\sigma = 1$ mV. Parameters on the right column: $\mu^{\text{ext}} = 10$ mV, $\sigma^{\text{ext}} = 3$ mV, $c_\mu = 0.5$ mV, $c_\sigma = 19$ mV.

column, however, the original fluctuation-driven regime is quickly abolished as γ_I/γ_E increases, leading to very high activity and low CV in the high-activity fixed point. Note that the size of the bistable region is also much smaller in this case.

3.3 Numerical Simulations. We conducted numerical simulations of our network to investigate whether the two types of bistable states that the mean-field analysis predicts, the mean-driven and the fluctuation-driven regimes, can be realized. In addition to the approximations that we are forced to make in order to be able to construct the mean-field theory itself, the more qualitative and robust result that fluctuation-driven bistable points require large fluctuations in relatively small networks with relatively large

synapses leads to the question of whether potential fixed points in this very noisy network will indeed be stable. We found that under certain conditions, both types of bistable points can be realized in numerical simulations.

In Figure 12 we show an example of a network supporting bistability on the mean-driven regime. In this network, the recurrent connections are dominated by excitation, and the mean of the external drive leaves the neurons very close to threshold, with the fluctuations of this external drive being small. As expected, the irregularity in the elevated-activity fixed point is low, with a CV ~ 0.2 . The mean μ_V in this fixed point is above threshold. The mean-field theory predicts for the same network a rate of 46.7 Hz and a CV of 0.21.

An example of another network showing bistability, this time in the fluctuation-driven regime, is shown in Figure 13. In this network, the recurrent connections are dominated by inhibition, and the external drive leaves the membrane potential relatively far from threshold on average but has large fluctuations. Taking into account only consecutive ISIs, the temporal irregularity is still large: $CV_2 \sim 0.8$. The spike trains in the elevated activity state are quite irregular, partly, but not only, due to large, temporal fluctuations in the instantaneous network activity. The mean-field prediction for these parameters gives a rate of 91.5 Hz and a CV of 1.6.

In order for the elevated activity states to be stable, in both the mean-driven and, especially, the fluctuation-driven regimes, we needed to use a wide distribution of synaptic delays in the recurrent connections: between 1 and 10 ms for Figure 12 and 1 and 50 ms for Figure 13. Narrow distributions of delays lead to oscillations, which destabilize the stationary activity states. The emergence and properties of these oscillations in a network similar to the one we study here have been described in Brunel (2000a).

Although such long synaptic delays are not expected to be found in connections between neighboring local cortical neurons, our network is extremely simple and lacks many elements of biological realism that would work in the same direction as the wide distributions of delays. Among these are slow and saturating synaptic interactions (NMDA-mediated excitation; (Wang, 1999) and heterogeneity in cellular and synaptic properties.

The large and slow temporal fluctuations in the instantaneous rate in Figure 13 are due to the large fluctuations in the nearly balanced external and recurrent drive to the cells and the wide distribution of synaptic delays. These fluctuations lead to high trial-to-trial variability in the activity of the network, as shown in Figure 14. In this figure, we show nine trials with identical parameters as in Figure 13, and only different seeds for the random number generator. On each panel, the mean instantaneous activity across all nine trials (the thick line) is shown along with the activity in the individual trial. Sometimes the large fluctuations lead to the activity returning to the basal spontaneous state. Other times they provoke long-lasting periods of elevated firing (above average). Nevertheless, on a large fraction of the trials, a memory of the stimulus persists for several seconds.

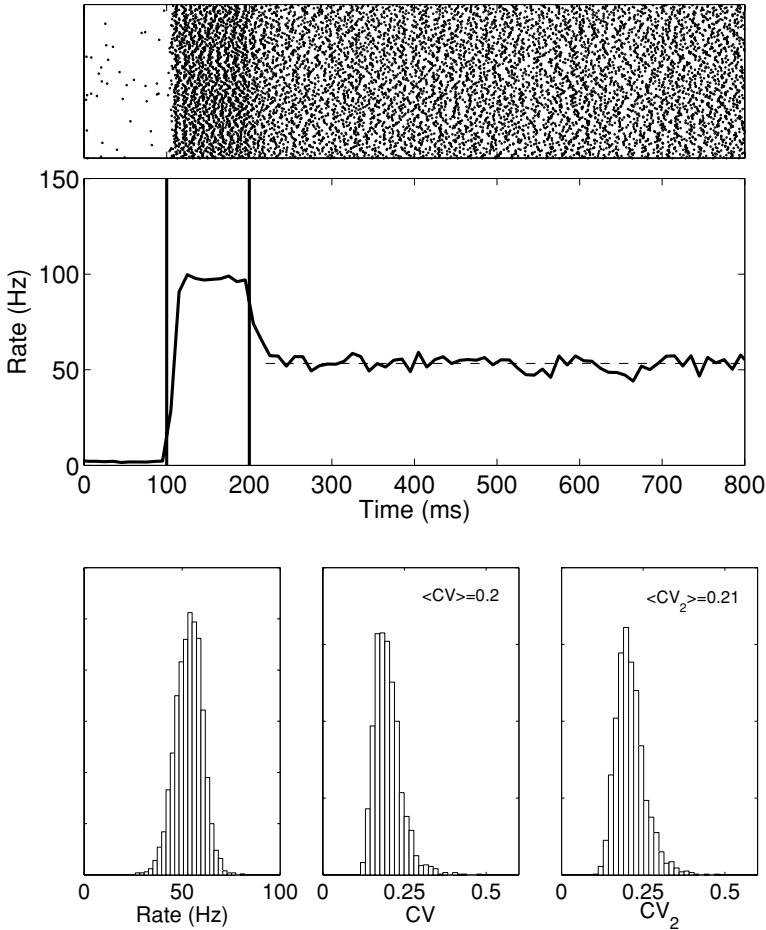


Figure 12: Numerical simulations of a bistable network in the mean-driven regime. The rate of the external afferents was raised between 200 and 300 ms (vertical bars). (Top) Raster display of the activity of 200 neurons in the network. (Middle) Instantaneous network activity (temporal bin of 10 ms). The dashed line represents the average network activity during the delay period, 53.4 Hz. (Lower panels) Distribution across cells of the rate (left), CV (middle), and CV_2 (right) during the delay period. The fact that the CV and CV_2 are very similar reflects the stationarity of the instantaneous activity. Single-cell parameters as in the caption to Figure 2. The network consists of two populations of excitatory and inhibitory cells (1000 neurons each) connected at random with 0.1 probability. Delays are uniformly distributed between 1 and 10 ms. External spikes are all excitatory, with PSP size 0.09 mV. The external rate is 19.25 KHz. This leads to $\mu^{\text{ext}} = 17.325$ mV and $\sigma^{\text{ext}} = 0.883$ mV. Recurrent EPSPs and IPSPs are 0.138 mV and -0.05 mV, respectively, leading to $c_\mu = 8.8$ mV and $c_\sigma = 1.468$ mV.

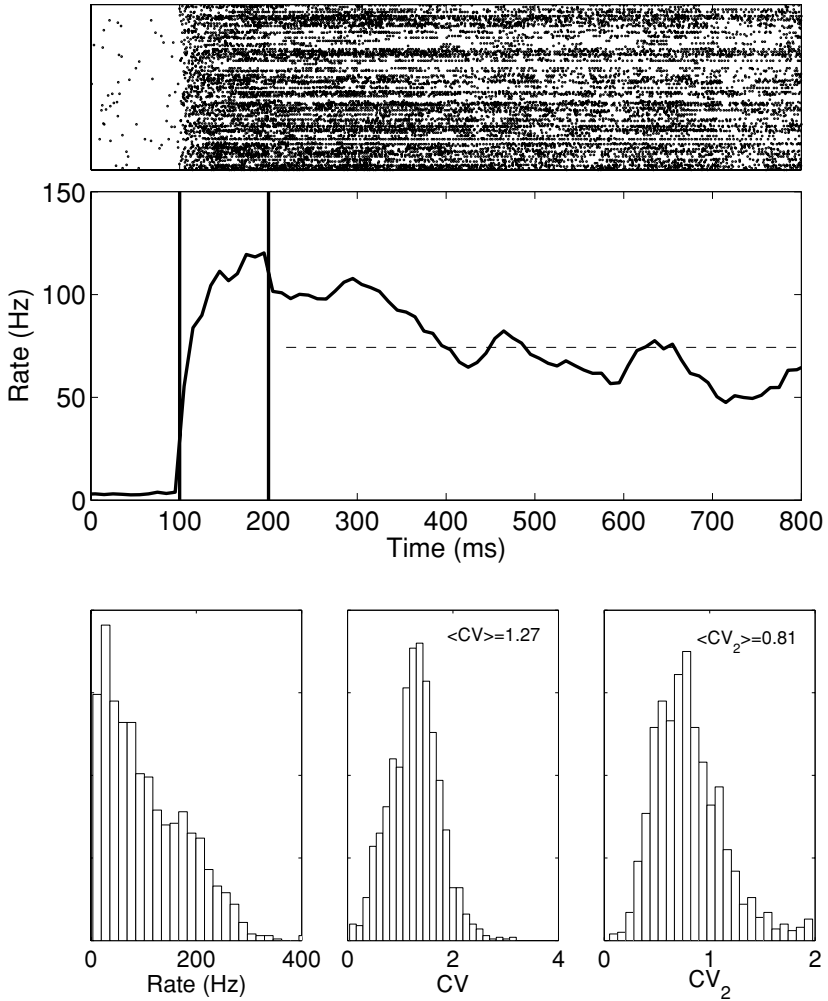


Figure 13: Numerical simulations of a bistable network in the fluctuation-driven regime. Panels as in Figure 12. Parameters as in Figure 12 except distribution of delays, which is uniform between 1 and 50 ms. External spikes are excitatory, with PSP size 1.85 mV and rate 0.78 KHz and inhibitory, with PSP size -1.85 mV and rate 0.5 KHz. This leads to $\mu^{\text{ext}} = 5.18$ mV and $\sigma^{\text{ext}} = 4.68$ mV. Recurrent EPSPs and IPSPs are 1.85 mV and -1.98 mV, respectively, leading to $c_\mu = -13$ mV and $c_\sigma = 27.1$ mV.




Figure 14: Trial-to-trial variability in the fluctuation-driven regime. Each panel is a different repetition of the same trial, in a network identical to the one described in Figure 13. The thick line represents the average across all nine trials, and the thin line is the instantaneous network activity in the given trial. Vertical bars mark the time during which the rate of the external inputs is elevated.

In the mean-driven regime, the trial-to-trial variability is very low (not shown).

We conclude that despite quantitative differences in the rate and CV between the mean-field theory and the simulations, it is possible, albeit difficult, to find both mean-driven and fluctuation-driven bistability in small networks of LIF neurons.

4 Discussion

In this letter, we have aimed at an understanding of the different ways in which a simple network of current-based LIF neurons can be organized in order to support bistability, the coexistence of two steady-state solutions for the activity of the network that can be selected by transient external stimulation. We have shown that in addition to the well-known case in which strong excitatory feedback can lead to bistability, bistability can also be obtained when the recurrent connectivity is nearly balanced, or even when its net effect is inhibitory, provided that an increase in the activity in

the network provides a large enough increase in the size of the fluctuations of the current afferent to the cells. When bistability is obtained in this fashion, the CV in both steady states is close to one, as found experimentally (see Figure 1; Compte et al., 2003). We have done a systematic analysis at the mean-field level (and a partial one through numerical simulations) of a reduced network where the activity in the excitatory and the inhibitory subpopulations was equal by construction (implying balance at the level of the output activity) and studied which types of bistable solutions are obtained depending on the level of balance in the currents (the parameter c_μ), that is, balance at the level of the inputs. This simple model allows for a complete understanding of the role of the different model parameters.

The first phenomenon, which we have termed *mean-driven bistability*, can essentially be traced back to the shape of the curve relating the mean firing rate of the cell to the average current it is receiving (at a constant noise level; Brunel, 2000b), that is, the $f - I$ curve. In order for bistability to exist, this curve should be a sigmoid, for which it is enough that the neurons possess a threshold and a refractory period. If, in addition, the low-activity fixed point is to have a nonzero activity (consistent with the fact that cells in the cortex fire spontaneously), then the neuron should display nonzero activity for subthreshold mean currents. This can be achieved if the current is noisy, where the noise is due to the spiking irregularity of the inputs to the cell. When this type of bistability is considered in a network of LIF neurons, the mean current to the cells in the high-activity fixed point is above threshold. Under general assumptions, this leads invariably to fairly regular spiking in this high-activity fixed point. Of course, tuning the parameters of the current in such a way that the mean current in the high-activity fixed point is only very marginally suprathreshold will result in only a small decrease of the CV with respect to the low-activity fixed point (e.g., Figure 2 in Brunel & Wang, 2001). On the other hand, in this scenario, it is relatively easy (it does not take much tuning) for the firing rate in the elevated persistent activity state not to be very much higher than that in the low-activity state, for example, below 100 Hz (see Figure 8).

When the recurrent connectivity is balanced, bistable solutions can exist in which both fixed points are subthreshold, so that spiking in both fixed points is fluctuation driven and thus fairly irregular. This can be the case if the fluctuations in the depolarization due to current from outside the column are low enough (see Figure 6). However, in order for these solutions to exist, first, the overall inputs to the excitatory and the inhibitory subpopulations should be close enough (ensuring balance at the level of the firing activity in the network); second, both of these inputs, the one to the excitatory and the one to the inhibitory subpopulation, should themselves also be balanced (be composed of similar amounts of excitation and inhibition); and third, both the net excitatory drive and the inhibitory drive to the cells should be large. This third condition, if the first two are satisfied, results in a high, effective fluctuation feedback gain: an increase in the activity of

the cells results in a large increase in the size of the fluctuations in the afferent current to the neurons (a large value of the parameter c_σ). However, it also implies that the excitation-inhibition balance condition will be quite stringent; it will require tuning, especially when the network is large. In addition, if this balance is slightly broken, since both excitation and inhibition are large, the corresponding firing rate in the elevated persistent activity state becomes very large, for example, significantly higher than 100 Hz (see Figure 9). In fact, based just on scaling considerations (see section 3.1), one can conclude that this type of bistability can be present only (unless one allows for perfect tuning) in small networks. If the network is large, the excitation-inhibition balance condition has, in general, a single (albeit very robust) solution (van Vreeswijk & Sompolinsky, 1996, 1998). It is intriguing, however, that several lines of evidence in fact suggest a fairly precise balance of local excitation and inhibition in cortical circuits, at both the output level (Rao et al., 1999) and the input level (Anderson, Carandini, & Ferster, 2000; Shu, Hasenstaub, & McCormick, 2003; Marino et al., 2005).

4.1 Limitations of the Present Approach. Most of the results we have presented are based on an exhaustive analysis of the stationary states of a mean-field description of a simple network of LIF neurons. Several limitations of our approach should be noted. First, in order to be able to go beyond the Poisson assumption, we have had to make a number of approximations (discussed in section 2) that are expected to be valid only on limited regions of the large parameter space. Second, we have focused only on the stationary fixed points of the system, neglecting an examination of any oscillatory solutions. Oscillations in networks of LIF neurons in the high-noise regime have been extensively studied by Brunel and collaborators (see, e.g., Brunel & Hakim, 1999; Brunel, 2000a; Brunel & Wang, 2003). Third, in order to be able to provide an analytical description, we have considered a very simplified network lacking many aspects of biological realism known to affect network dynamics, most important, a more realistic description of synaptic dynamics (Wang, 1999; Brunel & Wang, 2001). Finally, the use of a mean-field description based on the diffusion approximation to study small networks with big synaptic PSPs might lead to problems, since the diffusion approximation assumes these PSPs are (infinitely) small. Large PSPs might lead to fluctuations that are too strong, which would destabilize the analytically predicted fixed points.

In order to check that the main qualitative conclusion of this study was not an artifact due to the mean-field approach, we simulated the network of LIF neurons used in the mean-field description, adding synaptic delays. Provided the distribution of delays was wide, we observed both types of bistable solutions. However, as expected, the fluctuation-driven persistent activity states show large, temporal fluctuations that sometimes are enough to destabilize them.

The evidence we have provided is suggestive, but given the limitations listed above, it is not conclusive. Addressing the limitations of this work will involve using recently developed analytical techniques (Moreno-Bote & Parga, 2006), along with a systematic exploration of the behavior of more realistic recurrent networks through numerical simulations.

4.2 Self-Consistent Second-Order Statistics. The extended mean-field theory we have used builds on the landmark study by Amit and Brunel (1997b), which provided a general-purpose theory for the study of the different types of steady states in recurrent networks of spiking neurons in the presence of noise, while leaving room for different degrees of biophysical realism. Our contribution has been to try to go beyond the Poisson assumption in order to allow a self-consistent solution to the second-order statistics of the spike trains in the network. If the spike trains are assumed to be Poisson, there is only one parameter to be determined self-consistently: the firing rate. Under the approximations made in this letter, the statistics are characterized by two parameters, the firing rate and the CV, which provides information about the degree of irregularity of the spiking activity. In order to go beyond the Poisson assumption, we have assumed the spike trains in the network can be described as renewal processes with a very short correlation time. In these conditions, for time windows large compared to this correlation time, the Fano factor of the process is constant, but instead of being one, as for a Poisson process, it is equal to CV^2 . This motivates our strategy of neglecting the temporal aspect of the deviation from Poisson, which is extremely complicated to deal with analytically, and keep only its effect on the amplitude of the correlations. We have done this by using the expressions for the rate and CV of the first passage time of the OU process with a renormalized variance that takes into account the CV of the inputs. If the time constant of the autocorrelation of the process is exactly zero, this approximation becomes exact (Moreno et al., 2002), so we have assumed it will still be qualitatively valid if the correlation time constant is small. In this way, we have been able to solve for the CV of the neurons in the steady states self-consistently.

It has to be stressed that the fact that the individual inputs to a neuron are considered independent does not imply that the overall input process, made of the sum of each individual component, is Poisson. Informally, in order for the superposition to converge to a (homogeneous) Poisson process of rate λ , two conditions have to be met: given any set S on the time axis (say, any time interval), calling N_i^1 the probability of observing one spike in S from process i , and $N_i^{>2}$ the probability of observing two or more spikes in S from process i , then the superposition of the $i = 1, \dots, N$ processes will converge to a Poisson process if $\lim_{N \rightarrow \infty} \sum_i^N N_i^1 = \lambda S$ (with $\max\{N_i^1\} = 0$ as $N \rightarrow \infty$) and if $\lim_{N \rightarrow \infty} \sum_i^N N_i^{>2} = 0$ (see, e.g., Daley & Vere-Jones, 1988). The autocorrelated renewal processes that we consider in this letter do

not meet the second condition, which can also be seen in the fact that the superposition process has an autocorrelation given by equation 2.2, not by a Dirac delta function, as would be the case for a Poisson process. Despite this, it might be the case that if instead of any set S , one considers only a given time window T , both conditions could approximately be met in T , and we could say that the superposition process is locally Poisson in T . Whether this locally Poisson train will have the same effect on the postsynaptic cell as a real Poisson train of the same rate depends on a number of factors and has been studied in detail in Moreno et al. (2002) for the case of exponential autocorrelations. Other types of autocorrelation structures, for instance, regular spike trains, could lead to different results. This is an open problem.

4.3 Current-Based versus Conductance-Based Descriptions. We have analyzed a network of current-based LIF neurons. The motivations for this choice are that current-based LIF neurons are simpler to analyze, especially in the presence of noise, than conductance-based LIF neurons and also that there were a number of unresolved issues raised in the current-based framework that we have made an attempt to clarify. In particular, we were interested in understanding whether the framework of Amit and Brunel (1997b) could be used to produce bistable solutions in balanced networks like those studied in Tsodyks and Sejnowski (1995) and van Vreeswijk and Sompolinsky (1996, 1998) outside the context of multistability in recurrent networks. An important issue has been the relationship between different scaling relationships between the connection strengths and the number of afferents and the possible types of bistability attainable in large networks, when the number of afferents per cell tends to infinity. This analysis shows that large, homogeneous networks using the $J \sim 1/\sqrt{C}$ scaling needed to retain a significant amount of fluctuation at large C do not support bistability in a robust manner, a result already implicitly present in van Vreeswijk and Sompolinsky (1996, 1998). Reasonably robust bistability in homogeneous balanced networks requires that they are small.

Does one expect these conclusions to hold qualitatively if one considers the more realistic case of conductance-based synaptic inputs to the cells? The answer to this question is uncertain. In particular, scaling relationships between J and C , absolutely unavoidable in current-based scenarios to keep a finite input to the cells in the extensive C limit, are not necessary when synaptic inputs are assumed to induce a transient change in conductance. In the presence of conductances, the steady-state voltage is automatically independent of C for large C , regardless of the value of the unitary synaptic conductances. In fact, assuming a simple model for a cell having only leak, excitatory, and inhibitory synaptic conductances, the steady-state voltage in the absence of threshold is given by

$$V_{ss} = \frac{g_L}{g_{Tot}} V_L + \frac{g_E}{g_{Tot}} V_E + \frac{g_I}{g_{Tot}} V_I,$$

where V_L , V_E , V_I are the reversal potentials of the respective currents; g_L , g_E , g_I are the total leak, excitatory, and inhibitory conductance in the cell and

$$g_{\text{Tot}} = g_L + g_E + g_I.$$

(This expression ignores the effect of temporal fluctuations in the conductances, but it is a good approximation, since the mean conductance, being by definition positive, is expected to be much larger than its fluctuations.) The steady-state voltage is just the average of the different reversal potentials, each weighted by the relative amount of conductance that the respective current is carrying. Of course, each of the total synaptic conductances is proportional to the number of inputs, but since the steady-state voltage is just a weighted sum, it does not explode even if C tends to infinity.

It might seem that in fact, the infinite- C limit leads to an ill-defined model, as the membrane time constant vanishes in this case as $1/C$. If C_m is the membrane capacitance, then

$$\tau_m = \frac{C_m}{g_{\text{Tot}}} \sim 1/C,$$

assuming that $g_{E,I} \sim C$. We believe, however, that this is an artifact due to an incorrect way of defining the model in the large C limit. It implicitly assumes that the number of synaptic inputs grows at a constant membrane area, thus increasing indefinitely the local channel density. A more appropriate way of taking the large C limit is to fix the relative densities of the different channels per unit area and then assume the area becomes large. In this case, both C_m and the total leak conductance of the cell (proportional to the number of leak channels) will grow with the total area. This way of taking the limit respects the well-known decrease in membrane time constant as the synaptic input to the cell grows, but retaining a well-defined, nonzero membrane time constant in the extensive C limit (in this case, the range of values that τ_m can take is determined by the local differences in channel density, which is independent of the total channel number).

A crucial difference with the current-based cell is the behavior of the variance of the depolarization in the large C limit. A simple estimate of this variance can be obtained by ignoring threshold and considering only the low-pass filtering effect of the membrane (with time constant τ_m) on a gaussian noise current of variance σ_I^2 and time constant τ_s . It is straightforward to calculate the variance of the depolarization in these conditions, resulting in

$$\sigma_V^2 = \frac{\sigma_I^2}{g_{\text{Tot}}^2} \left(\frac{\tau_s}{\tau_s + \tau_m} \right).$$

If the inputs are independent, both the variance of the current and the total conductance of the cell are proportional to C , which implies that $\sigma_V^2 \sim 1/C$ for large C .

Therefore, the statistics of the depolarization in conductance-based and current-based neurons show a very different dependence with the number of inputs to the cell. In particular, it is unclear whether the main organizational principle behind the balanced state in the current-based framework, that is, the $J \sim 1/\sqrt{C}$ scaling that is needed to retain a finite variance in the $C \rightarrow \infty$ limit and that leads to the set of linear equations that specify the single solution for the activity in the balanced network, is relevant in a conductance-based framework. A rigorous study of this problem is beyond the scope of this work, but is one of the outstanding challenges for understanding the basic principles of cortical organization.

4.4 Correlations and Synaptic Time Constants. Our mean-field description assumes that the network is in the extreme sparse limit, $N \gg C$, in which the fraction of common input shared by two neurons is negligible, leading to vanishing correlations between the afferent current to different cells in the large C , large N limit. This is a crucial assumption, since it causes the variance of the depolarization in the network to be the sum of the variances of the individual spike trains, that is, proportional to C . If the correlation coefficient is finite as $C \rightarrow \infty$, the variance is proportional to C^2 (see, e.g., Moreno et al., 2002). In a current-based network, $J \sim 1/C$ scaling would lead to a nonvanishing variance in the large C limit without a stringent balance condition, and in a conductance-based network, it would lead to a C -independent variance for large C . This suggests that correlations between the cells in the recurrent network should have a large effect on both their input-output properties (Zohary, Shadlen, & Newsome, 1994; Salinas & Sejnowski, 2000; Moreno et al., 2002) and the network dynamics. The issue is, however, not straightforward, as simulations of irregular spiking networks with realistic connectivity parameters, which do show weak but significant cross-correlations between neurons (Amit & Brunel, 1997a), seem to be well described by the mean-field theory in which correlations are neglected (Amit & Brunel, 1997a; Brunel & Wang, 2001).

Noise correlations measured experimentally are small but significant, with normalized correlation coefficients on the range of a few percent to a few tenths for a review (see, e.g., Salinas & Sejnowski, 2001). It would thus be desirable to be able to extend the current mean-field theory to incorporate the effect of cross-correlations and to understand under which conditions their effect is important. The first steps in this direction have already been taken (Moreno et al., 2002; Moreno-Bote & Parga, 2006). The arguments of the previous section suggest that a correct treatment of correlations might be especially important in large networks of conductance-based neurons.

An issue of similar importance for the understanding of the high irregularity of cortical spike trains is the relationship between the time constant (or time constants) of the inputs to the neuron and its (effective) membrane time constant. Indeed, the need for a balance between excitation and inhibition in order to have a high-output spiking variability when receiving many irregular inputs exists only if the membrane time constant is relatively long—in particular, long enough that if the afferents are all excitatory, the input fluctuations are averaged out. If the membrane time constant is very short, large, short-lived fluctuations are needed to make the postsynaptic cell fire, and these occur only irregularly, even if all the afferents are excitatory (see, e.g., Figure 1 in Shadlen & Newsome, 1994). These considerations seem relevant since cortical cells receive large numbers of inputs that have spontaneous activity, thus putting the cell into a high-conductance state (see, e.g., Destexhe, Rudolph, & Pare, 2003) in which its effective membrane time constant is short—on the order of only a few milliseconds (Bernander, Douglas, Martin, & Koch, 1991; Softky, 1994). It has also been recognized that when the synaptic time constant is large compared to the membrane time constant, spiking in the subthreshold regime becomes very irregular, and in particular, the distribution of firing rates becomes bimodal. Qualitatively, in these conditions, the depolarization follows the current instead of integrating it. Relative to the timescale of the membrane, fluctuations are long-lived, and this separates two different timescales for spiking (which result in bimodality of the firing-rate distribution) depending on whether the size of a fluctuation is such that the total current is subthreshold (i.e., no spiking leading to a large peak of the firing rate histogram at zero) or suprathreshold (leading to a nonzero peak in the firing rate distribution) (Moreno-Bote & Parga, 2005). In these conditions, neurons seem “bursty,” and the CV of the ISI is high. Interestingly, recent evidence confirms this bimodality of the firing-rate distribution in spiking activity recorded *in vivo* in the visual cortex (Carandini, 2004).

Increases in the synaptic-to-membrane time constant ratio leading to more irregular spiking can be due to a number of factors: a very short membrane time constant if the neuron is a high-conductance state, relatively long excitatory synaptic drive if there is a substantial NMDA component in the excitatory EPSPs, or even long-lived dendrosomatic current sources, for instance, due to the existence of “calcium spikes” generated in the dendrites. There is evidence that irregular current applied to the dendrites of pyramidal cells results in higher CVs than the same current applied to the soma (Larkum, Senn, & Luscher, 2004).

4.5 Parameter Fine-Tuning. In order for both stable firing rate states of the networks we have studied to display significant spiking irregularity, the afferent current to the cells in both states needs to be subthreshold. We have shown that this requires a significant amount of parameter fine-tuning, especially when the number of connections per neuron is large. Parameter

fine-tuning is a problem, since biological networks are heterogeneous and cellular and synaptic properties change in time.

Regarding this issue, though, some considerations are in order. First, the model we have considered is extremely simple, especially at the single-neuron level. We have already pointed out possible consequences of considering properties such as longer synaptic time constants or some degree of correlations between the spiking activity of different neurons. Another biophysical property that we expect to have a large impact is short-term synaptic plasticity. In the presence of depressing synapses, the postsynaptic current is no longer linear in the presynaptic firing rate, thus acting as an activity-dependent gain control mechanism (Tsodyks & Markram, 1997; Abbott, Varela, Sen, & Nelson, 1997). It remains to be explored to what extent balanced bistability in networks of neurons exhibiting these properties becomes a more robust phenomenon.

Second, synaptic weights (as well as intrinsic properties; Desai, Rutherford, & Turrigiano, 1999) can adapt in an activity-dependent manner to keep the overall activity in a recurrent network within an appropriate operational range (Turrigiano, Leslie, Desai, Rutherford, & Nelson, 1998). Delicate computational tasks, which seem to require fine-tuning, can be rendered robust through the use of these types of activity-dependent homeostatic rules (Renart, Song, & Wang, 2003). It will be interesting to study whether homeostatic plasticity (Turrigiano, 1999) can be used to relax some of the fine-tuning constraints described in this letter.

4.6 Multicolumnar Networks and Hierarchical Organization. The fact that bistability is not a robust property of large, homogeneous balanced networks suggests that the functional units of working memory could correspond to small subpopulations (Rao et al., 1999). In addition, we have shown that bistability in a small, reduced network is possible only for subthreshold external inputs (see section 3.2.2). At the same time, it is known that a nonzero activity balanced state requires a very large (suprathreshold) excitatory drive (see section 3.1 and van Vreeswijk & Sompolinsky, 1998). This seems to point to a hierarchical organization: large networks receive massive excitation from long-distance projections, and this external excitation sets up a balanced state in the network. Globally, the activity in the large, balanced network follows the external input linearly. This large, balanced network then provides an already balanced (subthreshold) input to smaller subcomponents, which, in these conditions (in particular, if the variance of this subthreshold input is small enough; see figure 7), can display more complex nonlinear behavior such as bistability. From the point of view of the smaller subnetworks, the balanced subthreshold input can be considered external, since the size of this network is too small to make a difference in the global activity of the larger network (despite being recurrently connected, the activities in the large and small networks effectively

decouple). In the cortex, the larger balanced network could correspond to a whole cortical column. Indeed, long-range projections between columns are mostly excitatory (see, e.g., Douglas & Martin, 2004). Within a column, the smaller networks that interact through both excitation and inhibition could anatomically correspond to microcolumns (Rao et al., 1999) or, more generally, define functional assemblies (Hebb, 1949).

5 Summary

General principles of cortical organization (large numbers of active synaptic inputs per neuron) and function (irregular spiking statistics) put strong constraints on working memory models of spiking neurons. We have provided evidence that a network of current-based LIF neurons can exhibit bistability with the high persistent activity driven by either the mean or the fluctuations in the input to the cells. The fluctuation-driven bistability regime requires a strict excitation-inhibition balance that needs parameter tuning. It remains a challenge in future research to analyze systematically what the conditions are under which nonlinear phenomena such as bistability can exist robustly in large networks of more biophysically plausible conductance-based and correlated spiking neurons. It is also conceivable that additional biological mechanisms, such as homeostatic regulation, are important for solving the fine-tuning problem and ensuring a desired excitation-inhibition balance in cortical circuits. Progress in this direction will provide insight into the microcircuit mechanisms of working memory, such as found in the prefrontal cortex.

Acknowledgments

We are indebted to Jaime de la Rocha for providing the code for the numerical simulations and to Albert Compte for providing the data for Figure 1. A.R. thanks N. Brunel for pointing out previous related work, and A. Amarasingham for discussions on point processes. Support was provided by the National Institute of Mental Health (MH62349, DA016455), the A. P. Sloan Foundation and the Swartz Foundation, and the Spanish Ministry of Education and Science (BFM 2003-06242).

References

- Abbott, L. F., Varela, J. A., Sen, K., & Nelson, S. B. (1997). Synaptic depression and cortical gain control. *Science*, *275*, 220–224.
- Abramowitz, M., & Stegun, I. A. (1970). *Tables of mathematical functions*. New York: Dover.
- Amit, D. J., & Brunel, N. (1997a). Dynamics of a recurrent network of spiking neurons before and following learning. *Network*, *8*, 373–404.

- Amit, D. J., & Brunel, N. (1997b). Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cerebral Cortex*, *7*, 237–252.
- Amit, D. J., & Tsodyks, M. V. (1991). Quantitative study of attractor neural network retrieving at low spike rates II: Low-rate retrieval in symmetric networks. *Network*, *2*, 275–294.
- Anderson, J. S., Carandini, M., & Ferster, D. (2000). Orientation tuning of input conductance, excitation, and inhibition in cat primary visual cortex. *J. Neurophysiol.*, *84*, 909–926.
- Bair, W., Zohary, E., & Newsome, W. T. (2001). Correlated firing in macaque visual area MT: Time scales and relationship to behavior. *J. Neurosci.*, *21*, 1676–1697.
- Ben-Yishai, R., Lev Bar-Or, R., & Sompolinsky, H. (1995). Theory of orientation tuning in visual cortex. *Proc. Natl. Acad. Sci. USA*, *92*, 3844–3848.
- Bernander, O., Douglas, R. J., Martin, K. A., & Koch, C. (1991). Synaptic background activity influences spatiotemporal integration in single pyramidal cells. *Proc. Natl. Acad. Sci. USA*, *88*, 11569–11573.
- Brunel, N. (2000a). Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *J. Comput. Neurosci.*, *8*, 183–208.
- Brunel, N. (2000b). Persistent activity and the single cell f-I curve in a cortical network model. *Network*, *11*, 261–280.
- Brunel, N., & Hakim, V. (1999). Fast global oscillations in networks of integrate-and-fire neurons with low firing rates. *Neural Computation*, *11*, 1621–1671.
- Brunel, N., & Wang, X.-J. (2001). Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition. *J. Comput. Neurosci.*, *11*, 63–85.
- Brunel, N., & Wang, X.-J. (2003). What determines the frequency of fast network oscillations with irregular neural discharges? I. Synaptic dynamics and excitation-inhibition balance. *J. Neurophysiol.*, *90*, 415–430.
- Cai, D., Tao, L., Shelley, M., & McLaughlin, D. W. (2004). An effective kinetic representation of fluctuation-driven networks with application to simple and complex cells in visual cortex. *Proc. Natl. Acad. Sci.*, *101*, 7757–7762.
- Carandini, M. (2004). Amplification of trial-to-trial response variability by neurons in visual cortex. *PLoS Biol.*, *2*, E264.
- Chafee, M. V., & Goldman-Rakic, P. S. (1998). Matching patterns of activity in primate prefrontal area 8a and parietal area 7ip neurons during a spatial working memory task. *J. Neurophysiol.*, *79*, 2919–2940.
- Compte, A., Brunel, N., Goldman-Rakic, P. S., & Wang, X.-J. (2000). Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cerebral Cortex*, *10*, 910–923.
- Compte, A., Constantinidis, C., Tegnér, J., Raghavachari, S., Chafee, M., Goldman-Rakic, P. S., & Wang, X.-J. (2003). Temporally irregular mnemonic persistent activity in prefrontal neurons of monkeys during a delayed response task. *J. Neurophysiol.*, *90*, 3441–3454.
- Cox, D. R. (1962). *Renewal theory*. New York: Wiley.
- Daley, D. J., & Vere-Jones, D. (1988). *An introduction to the theory of point processes*. New York: Springer.

- Desai, N. S., Rutherford, L. C., & Turrigiano, G. G. (1999). Plasticity in the intrinsic excitability of cortical pyramidal neurons. *Nat. Neurosci.*, *2*, 515–520.
- Destexhe, A., Rudolph, M., & Pare, D. (2003). The high-conductance state of neocortical neurons in vivo. *Nat. Rev. Neurosci.*, *4*, 739–751.
- Douglas, R. J., & Martin, K. A. (2004). Neuronal circuits of the neocortex. *Ann. Rev. Neurosci.*, *27*, 419–451.
- Funahashi, S., Bruce, C. J., & Goldman-Rakic, P. S. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *J. Neurophysiol.*, *61*, 331–349.
- Fuster, J. M., & Alexander, G. (1971). Neuron activity related to short-term memory. *Science*, *173*, 652–654.
- Gerstein, G. L., & Mandelbrot, B. (1964). Random walk models for the spike activity of a single neuron. *Biophys. J.*, *4*, 41–68.
- Gillespie, D. T. (1992). *Markov processes: An introduction for physical scientists*. Orlando, FL: Academic Press.
- Gnadt, J. W., & Andersen, R. A. (1988). Memory related motor planning activity in posterior parietal cortex of macaque. *Exp. Brain Res.*, *70*, 216–220.
- Hansel, D., & Mato, G. (2001). Existence and stability of persistent states in large neuronal networks. *Phys. Rev. Lett.*, *86*, 4175–4178.
- Harsch, A., & Robinson, H. P. (2000). Postsynaptic variability of firing in rat cortical neurons: The roles of input synchronization and synaptic NMDA receptor conductance. *J. Neurosci.*, *20*, 6181–6192.
- Hebb, D. O. (1949). *Organization of behavior*. New York: Wiley.
- Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. USA*, *79*, 2554–2558.
- Larkum, M. E., Senn, W., & Luscher, H. M. (2004). Top-down dendritic input increases the gain of layer 5 pyramidal neurons. *Cereb. Cortex*, *14*, 1059–1070.
- Marino, J., Schummers, J., Lyon, D. C., Schwabe, L., Beck, O., Wiesing, P., & Obermayer, K. (2005). Invariant computations in local cortical networks with balanced excitation and inhibition. *Nat. Neurosci.*, *8*, 194–201.
- Mascaro, M., & Amit, D. J. (1999). Effective neural response function for collective population states. *Network*, *10*, 351–373.
- Mattia, M., & Del Giudice, P. (2000). Efficient event-driven simulation of large networks of spiking neurons and dynamical synapses. *Neural Computation*, *12*, 2305–2329.
- Miller, E. K., Erickson, C. A., & Desimone, R. (1996). Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *J. Neurosci.*, *16*, 5154–5167.
- Miyashita, Y., & Chang, H. S. (1988). Neuronal correlate of pictorial short-term memory in the primate temporal cortex. *Nature*, *331*, 68–70.
- Moreno, R., de la Rocha, J., Renart, A., & Parga, N. (2002). Response of spiking neurons to correlated inputs. *Phys. Rev. Lett.*, *89*, 288101.
- Moreno-Bote, R., & Parga, N. (2005). Membrane potential and response properties of populations of cortical neurons in the high conductance state. *Phys. Rev. Lett.*, *94*, 088103.
- Moreno-Bote, R., & Parga, N. (2006). Auto- and cross-correlograms for the spike response of lif neurons with slow synapses. *Phys. Rev. Lett.*, *96*, 028101.

- Rao, S. G., Williams, G. V., & Goldman-Rakic, P. S. (1999). Isodirectional tuning of adjacent interneurons and pyramidal cells during working memory: Evidence for microcolumnar organization in PFC. *J. Neurophysiol.*, *81*, 1903–1916.
- Renart, A. (2000). *Multi-modular memory systems*. Unpublished doctoral dissertation, Universidad Autónoma de Madrid.
- Renart, A., Song, P., & Wang, X.-J. (2003). Robust spatial working memory through homeostatic synaptic scaling in heterogeneous cortical networks. *Neuron*, *38*, 473–485.
- Ricciardi, L. M. (1977). *Diffusion processes and related topics on biology*. Berlin: Springer-Verlag.
- Romo, R., Brody, C. D., Hernández, A., & Lemus, L. (1999). Neuronal correlates of parametric working memory in the prefrontal cortex. *Nature*, *399*, 470–474.
- Salinas, E., & Sejnowski, T. J. (2000). Impact of correlated synaptic input on output firing rate and variability in simple neuronal models. *Journal of Neuroscience*, *20*, 6193–6209.
- Salinas, E., & Sejnowski, T. J. (2001). Correlated neuronal activity and the flow of neural information. *Nat. Rev. Neurosci.*, *2*, 539–550.
- Shadlen, M. N., & Newsome, W. T. (1994). Noise, neural codes and cortical organization. *Current Opinion in Neurobiol.*, *4*, 569–579.
- Shadlen, M. N., & Newsome, W. T. (1998). The variable discharge of cortical neurons: Implications for connectivity, computation, and information coding. *J. Neurosci.*, *18*, 3870–3896.
- Shinomoto, S., Sakai, Y., & Funahashi, S. (1999). The Ornstein-Uhlenbeck process does not reproduce spiking statistics of neurons in prefrontal cortex. *Neural Comput.*, *11*, 935–951.
- Shu, Y., Hasenstaub, A., & McCormick, D. A. (2003). Turning on and off recurrent balanced cortical activity. *Nature*, *432*, 288–293.
- Softky, W. R. (1994). Sub-millisecond coincidence detection in active dendritic trees. *Neuroscience*, *58*, 13–41.
- Softky, W. R., & Koch, C. (1993). The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J. Neurosci.*, *13*, 334–350.
- Tsodyks, M. V., & Markram, H. (1997). The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proc. Natl. Acad. Sci. USA*, *94*, 719–723.
- Tsodyks, M. V., & Sejnowski, T. (1995). Rapid state switching in balanced cortical network models. *Network*, *6*, 111–124.
- Turrigiano, G. G. (1999). Homeostatic plasticity in neuronal networks: The more things change, the more they stay the same. *Trends in Neurosci.*, *22*, 221–227.
- Turrigiano, G. G., Leslie, K. R., Desai, N. S., Rutherford, L. C., & Nelson, S. B. (1998). Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature*, *391*, 892–896.
- van Vreeswijk, C., & Sompolinsky, H. (1996). Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science*, *274*, 1724–1726.
- van Vreeswijk, C., & Sompolinsky, H. (1998). Chaotic balanced state in a model of cortical circuits. *Neural Computation*, *10*, 1321–1371.
- Wang, X.-J. (1999). Synaptic basis of cortical persistent activity: The importance of NMDA receptors to working memory. *J. Neurosci.*, *19*, 9587–9603.

Zador, A. M., & Stevens, C. F. (1998). Input synchrony and the irregular firing of cortical neurons. *Nat. Neurosci.*, *1*, 210–217.

Zohary, E., Shadlen, M. N., & Newsome, W. T. (1994). Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature*, *370*, 140–143.

Received August 22, 2005; accepted May 17, 2006.