*LETTER* **Communicated by Paul Bressloff**

# Response Variability in Balanced Cortical Networks

**Alexander Lerchner***
*LerchnerA@mail.nih.gov*
*Technical University of Denmark, 2800 Lyngby, Denmark*

**Cristina Ursta**
*cristina@nordita.dk*
*Niels Bohr Institut, 2100 Copenhagen Ø, Denmark*

**John Hertz**
*hertz@nordita.dk*
*Nordita, 2100 Copenhagen Ø, Denmark*

**Mandana Ahmadi**
*ahmadi@nordita.dk*
*Nordita, 2100 Copenhagen Ø, Denmark*

**Pauline Ruffiot**
*pruffiot@ujf-grenoble.fr*
*Université Joseph Fourier, Grenoble, France*

**Søren Enemark**
*enemark@nbi.dk*
*Niels Bohr Institut, 2100 Copenhagen Ø, Denmark*

**We study the spike statistics of neurons in a network with dynamically balanced excitation and inhibition. Our model, intended to represent a generic cortical column, comprises randomly connected excitatory and inhibitory leaky integrate-and-fire neurons, driven by excitatory input from an external population. The high connectivity permits a mean field description in which synaptic currents can be treated as gaussian noise, the mean and autocorrelation function of which are calculated self-consistently from the firing statistics of single model neurons. Within this description, a wide range of Fano factors is possible. We find that the irregularity of spike trains is controlled mainly by the strength of the synapses relative to the difference between the firing threshold and the postfiring reset level of the membrane potential. For moderately strong**

* Current address: Laboratory of Neuropsychology, NIMH, NIH, Bethesda, MD 20893, USA

 

**synapses, we find spike statistics very similar to those observed in primary visual cortex.**

## 1 Introduction

The observed irregularity and relatively low rates of the firing of neocortical neurons suggest strongly that excitatory and inhibitory input are nearly balanced. Such a balance, in turn, finds an attractive explanation in the approximate, heuristic mean field description of Amit and Brunel (1997a, 1997b) and Brunel (2000). In this treatment, the balance does not have to be put in "by hand"; rather, it emerges self-consistently from the network dynamics. This success encourages us to study firing correlations and irregularity in models like theirs in greater detail. In particular, we would like to quantify the irregularity and identify the parameters of the network that control it. This is important because one cannot extract the signal in neuronal spike trains correctly without a good characterization of the noise. Indeed, an incorrect noise model can lead to spurious conclusions about the nature of the signal, as demonstrated by Oram, Wiener, Lestienne, and Richmond (1999).

Response variability has been studied for a long time in primary visual cortex (Heggelund & Albus, 1978; Dean, 1981; Tolhurst, Movshon, & Thompson, 1981; Tolhurst, Movshon, & Dean, 1983; Vogels, Spileers, & Orban, 1989; Snowden, Treue, & Andersen, 1992; Gur, Beylin, & Snodderly, 1997; Shadlen & Newsome, 1998; Gershon, Wiener, Latham, & Richmond, 1998; Kara, Reinagel, & Reid, 2000; Buracas, Zador, DeWeese, & Albright, 1998) and elsewhere (Lee, Port, Kruse, & Georgopoulos, 1998; Gershon et al., 1998; Kara et al., 2000; DeWeese, Wehr, & Zador, 2003). Most, though not all, of these studies found rather strong irregularity. As an example, we consider the findings of Gershon et al. (1998). In their experiments, monkeys were presented with flashed, stationary visual patterns for several hundred ms. Repeated presentations of a given stimulus evoked varying numbers of spikes in different trials, though the mean number (as well as the peristimulus time histogram) varied systematically from stimulus to stimulus. The statistical objects of interest to us here are the distributions of single-trial spike counts for given fixed stimuli. Often one compares the data with a Poisson model of the spike trains, for which the count distribution $P(n) = m^n \mathrm{e}^{-m}/n!$. This distribution has the property that its mean $\langle n \rangle = m$ is equal to its variance $\langle \delta n^2 \rangle = \langle (n - \langle n \rangle)^2 \rangle$. However, the experimental finding was that the measured distributions were quite generally wider than this: $\langle \delta n^2 \rangle > m$. Furthermore, when data were collected for many stimuli, the variance of the spike count was fit well by a power law function of the mean count: $\langle \delta n^2 \rangle \propto m^y$, with $y$ typically in the range 1.2 to 1.4, broadly consistent with the results of many of the other studies cited above.

Some of this observed variance could have a simple explanation: the condition of the animal might have changed between trials, so the intrinsic rate at which the neuron fires might differ from trial to trial, as suggested by Tolhurst et al. (1981). But it is far from clear whether all the variance can be accounted for in this way. Moreover, there is no special reason to take a Poisson process as the null hypothesis, so we do not even really know how much variance we are trying to explain.

In this article, we try to address the question of how much variability, or more generally, what firing correlations can be expected as a consequence of the intrinsic dynamics of cortical neuronal networks. The theories of Amit and Brunel (1997a, 1997b) and of van Vreeswijk and Sompolinsky (1996, 1998) do not permit a consistent study of firing correlations. The Amit-Brunel equations ignore firing correlations and variations in firing rate within neuronal populations; thus, they do not constitute a complete mean field theory. Although one can calculate the variability of the firing (Brunel, 2000), the calculation is not self-consistent. Van Vreeswijk and Sompolinsky use a binary neuron model with stochastic dynamics, which makes it difficult, if not impossible, to study temporal correlations that might occur in networks of spiking neurons. Therefore, in this article, we do a complete mean field theory for a network of leaky integrate-and-fire neurons, including, as self-consistently-determined order parameters, both firing rates and autocorrelation functions. This kind of theory is needed whenever the connections in the network are random. A general formalism for doing this was introduced by Fulvi Mari (2000) and used for an all-excitatory network; here we employ it for a network with both excitatory and inhibitory neurons. A preliminary study of this approach for an all-inhibitory network was presented previously (Hertz, Richmond, & Nilsen, 2003).

## 2 Model and Methods

The model network, indicated schematically in Figure 1, consists of $N_1$ excitatory neurons and $N_2$ inhibitory ones. In this work, we use leaky integrate-and-fire neurons, though the methods could be carried over directly to networks of other kinds of model neurons, such as conductance-based ones. They are randomly interconnected by synapses, both within and between populations, with the mean number of connections from population $b$ to population $a$ equal to $K_b$, independent of $a$. In specific calculations, we have used $K_1$ from 400 to 6400, and we take $K_2 = K_1/4$.

We scale the synaptic strengths in the way van Vreeswijk and Sompolinsky (1996, 1998) did, with each nonzero synapse from population $b$ to population $a$ having the value $J_{ab}/\sqrt{K_b}$. Thus, the mean value of a synapse $J_{ij}^{ab}$ is

$$\overline{J_{ij}^{ab}} = \frac{\sqrt{K_b} J_{ab}}{N_b}, \tag{2.1}$$

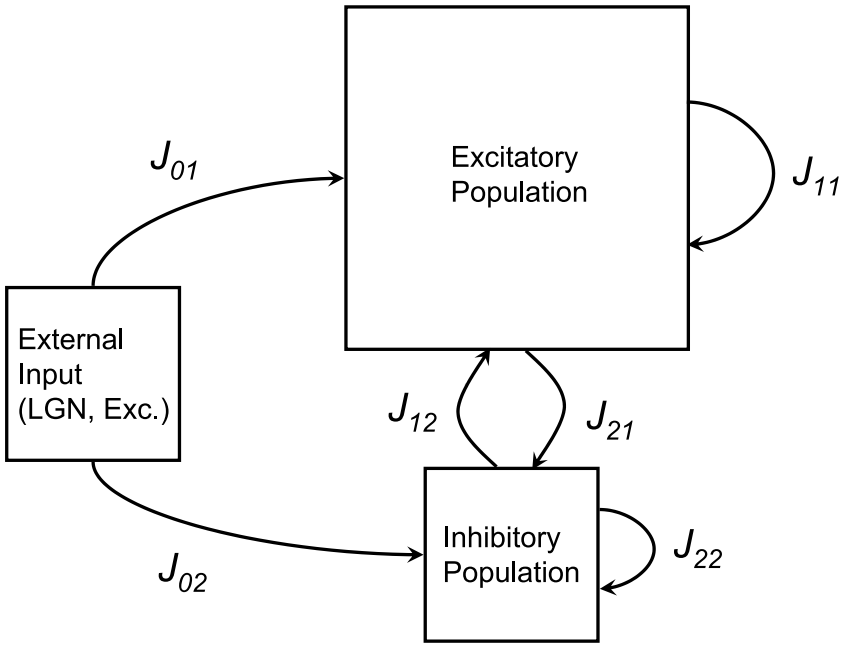


Figure 1: Structure of the Model Network.

and its variance is

$$\overline{\left(\delta J_{ij}^{ab}\right)^2} = \left(1 - \frac{K_b}{N_b}\right) \frac{J_{ab}^2}{N_b}. \tag{2.2}$$

The parameters $J_{ab}$ are taken to be of order 1, so the net input current to a neuron from the $K_b$ neurons in population $b$ connected to it is of order $\sqrt{K_b}$. With this scaling, the fluctuations in this current are of order 1.

Similarly, we assume that the external input to any neuron is the sum of $K_0 \gg 1$ contributions from individual neurons (in the lateral geniculate nucleus, if we are thinking about modeling V1), each of order $1/\sqrt{K_0}$, so the net input is of order $\sqrt{K_0}$. In our calculations, we have used $K_0 = K_1$.

We point out that this scaling is just for convenience in thinking about the problem. In the balanced asynchronous firing state, the large excitatory and inhibitory input currents nearly cancel, leaving a net input current of order 1. Thus, for this choice, both the net mean current and its typical fluctuations are of order 1, which is convenient for analysis. The physiologically relevant assumptions are only that excitatory and inhibitory inputs are separately much larger than their sum and that the latter is of the same order as its fluctuations.

Our synapses are not modeled as conductances. Our synaptic strength simply defines the amplitude of the postsynaptic current pulse produced by a single presynaptic spike.

The model is formally specified by the subthreshold equations of motion for the membrane potentials $u_i^a$ ($a = 1, 2$, $i = 1, \ldots, N_a$),

$$\frac{\mathrm{d}u_i^a}{\mathrm{d}t} = -\frac{u_i^a}{\tau} + \sum_{b=0}^{2} \sum_{j=1}^{N_b} J_{ij}^{ab} S_j^b(t), \tag{2.3}$$

together with the condition that when $u_i^a$ reaches the threshold $\theta_a$, the neuron spikes and the membrane potential is reset to a value $u_r^a$. The indices $a$ or $b = 0$, 1, or 2 label populations: $b = 0$ refers to the (excitatory) population providing the external input, $b = 1$ refers to the excitatory population, and $b = 2$ to the inhibitory population. In equation 2.3, $\tau$ is the membrane time constant (taken the same for all neurons, for convenience), and $S_j^b(t) = \sum_s \delta(t - t_{jb}^s)$ is the spike train of neuron $j$ in population $b$. We have ignored transmission delays, and we take the reset levels $u_r^a$ equal to the rest value of the membrane potential, 0. In our calculations, the thresholds are given a gaussian distribution with a standard deviation equal to 10% of the mean. We fix the mean threshold $\theta_a = 1$. Analogous variability in other single-cell parameters (such as membrane time constants) could also be included in the model, but for simplicity, we do not do so here.

We assume that the neurons in the external input population ($b = 0$) fire as independent Poisson processes. However, the neurons in the network ($b = 1, 2$) are not in general Poissonian; it is their correlations that we want to find in this investigation.

**2.1 Mean Field Theory: Stationary States.** We describe the mean field theory and its computational implementation first for the case of stationary rates. We will assume $N_a \gg K_a \gg 1$ (large but extremely dilute connectivity). Any mean field theory has to start with an ansatz for the structure of its order parameters. In words, our ansatz is that neurons fire noisily: thus, they are characterized by their rates (which can vary across a neuronal population) and autocorrelation functions. We assume the latter to contain a delta function spike at equal times of strength equal to the rate (because they are spiking neurons) plus a continuous part at unequal times. In our dilute limit, the theory is simplified by the fact that there are no cross-correlations between neurons. (Therefore, we generally drop the "auto" from "autocorrelation.")

Under these assumptions, each of the three terms in the sum on $b$ on the right-hand side of equation 2.3 can be treated as a gaussian random function with time-independent mean. (This can be proved formally using the generating-functional formalism of Fulvi Mari, 2000; it is a consequence of the fact that $K_b \gg 1$ and the independence of the $J_{ij}^{ab}$. Furthermore,

experiments (Destexhe, Rudolph, & Paré, 2003) show that a gaussian approximation is very good for real synaptic noise.) We write the contribution from population $b$ as

$$I_i^{ab}(t) = \sum_{j=1}^{N_b} \left(\overline{J_{ij}^{ab}} + \delta J_{ij}^{ab}\right)\left[r_b + \delta r_j^b + \delta S_j^b(t)\right], \tag{2.4}$$

where $r_b$ is the mean rate in population $b$, $\delta r_j^b$ is the deviation of the rate of neurons $j$ from its population mean, and $\delta S_j^b(t) = S_j^b(t) - r_j^b$ describes the fluctuations of the activity of neuron $j$ from its temporal mean $r_j^b$.

The mean (over both time and neurons in the receiving population $a$) comes from the product of the first terms:

$$\overline{\langle I_i^{ab}(t)\rangle} = \sqrt{K_b} J_{ab} r_b. \tag{2.5}$$

By $\langle \cdots \rangle$ we mean a time average or, equivalently, an average over "trials" (independent repetitions of the Poisson processes defining the input population neurons). We will generally use a bar over a quantity to indicate an average over the neuronal population or over the distribution of the $J_{ij}^{ab}$. (Note that these two kinds of averages are very different things.)

The fluctuations around this mean are of two kinds. One is the neuron-to-neuron rate variations in population $a$, obtained from the time-independent terms in equation 2.4:

$$\delta\langle I_i^{ab}\rangle = \sum_j \overline{J_{ij}^{ab}} \delta r_j^b + \sum_j \delta J_{ij}^{ab} r_j^b. \tag{2.6}$$

Using equations 2.1 and 2.2, their variance reduces, for $K_b \ll N_b$, to

$$\overline{\left(\delta\langle I_i^{ab}\rangle\right)^2} = \frac{J_{ab}^2}{N_b} \sum_j \left(r_j^b\right)^2 = J_{ab}^2 \overline{\left(r_j^b\right)^2}. \tag{2.7}$$

The second kind is the temporal fluctuations for single neurons, obtained from the terms in equation 2.4 involving $\delta S_j^b(t)$. Their population-averaged correlation function is proportional to the average correlation function in population $b$:

$$\overline{\langle \delta I_i^{ab}(t) \delta I_i^{ab}(t')\rangle} = \frac{J_{ab}^2}{N_b} \sum_j \langle \delta S_j^b(t) \delta S_j^b(t')\rangle \equiv J_{ab}^2 C_b(t - t'). \tag{2.8}$$

Thus, we can write this contribution to the input current for a single neuron as

$$I_i^{ab}(t) = J_{ab}\left[\sqrt{K_b}r_b + \sqrt{\overline{\left(r_j^b\right)^2}}x_i^{ab} + \xi_i^{ab}(t)\right], \tag{2.9}$$

where $x_i^{ab}$ is a unit-variance gaussian random number and

$$\left\langle \xi_i^{ab}(t)\xi_i^{ab}(t')\right\rangle = C_b(t-t'). \tag{2.10}$$

The $x_i^{ab}$ are time and trial independent, while the noise $\xi_i^{ab}(t)$ varies both in time within a trial and randomly from trial to trial. Note that for this model, a correct and complete mean field theory has to include rate variations, through $\overline{(r_j^b)^2}$, and the temporal firing correlations, given by $C_b(t-t')$, as well as the mean rates. In our treatment here, we will assume that the neurons in the external input population fire like Poisson processes, so $I_i^{a0}(t)$ is white noise. However, the neurons providing the source of the recurrent currents are not generally Poissonian, so their correlations appear in the statistics of the noise term.

The self-consistency equations of mean field theory are simply the conditions that the average output statistics of the neurons, $r_a$, $\overline{(r_j^a)^2}$ and $C_a(t-t')$, are the same as those used to generate the inputs for single neurons using integrate-and-fire neurons with synaptic input currents given by equation 2.9.

In an equivalent formulation, the second term in equation 2.9 can be omitted if the noise terms $\xi_i^{ab}(t)$ have correlations equal to the unsubtracted correlation function,

$$C_b^{\mathrm{tot}}(t-t') = \frac{1}{N_b}\sum_j \left\langle S_j^b(t)S_j^b(t')\right\rangle. \tag{2.11}$$

For $|t-t'|\to\infty$, $C_b^{\mathrm{tot}}(t-t') \to \overline{(r_j^b)^2}$, so $\xi_i^{ab}(t)$ acquires a random static component of mean square value $\overline{(r_j^b)^2}$.

In still another way to do it, one can use the average rate $r_b$ in place of its root mean square value in the second term on the right-hand side of equation 2.9 and employ noise with a correlation function:

$$\tilde{C}_b(t-t') = \frac{1}{N_b}\sum_j \left\langle \left(S_j^b(t)-r_b\right)\left(S_j^b(t')-r_b\right)\right\rangle. \tag{2.12}$$

For $|t - t'| \to \infty$,

$$\tilde{C}_b(t - t') \to \overline{\left(r_j^b - r_b\right)^2} \equiv \overline{\left(\delta r_j^b\right)^2}. \tag{2.13}$$

There are now two static random parts of $I_i^{ab}(t)$ in equation 2.9: one from the second term and one from the static component of the noise. Their sum is a gaussian random number with variance equal to $\overline{(r_j^b)^2}$. Thus, these three ways of generating the input currents are equivalent.

*2.1.1 The Balance Condition.* In a stationary, low-rate state, the mean membrane potential described by equation 2.3 has to be approximately stationary. If excitation dominates, we have $\mathrm{d}u_i^a/\mathrm{d}t \propto \sqrt{K_0}$, implying a firing rate of order $\sqrt{K_0}$ (or one limited only by the refractory period of the neuron). If inhibition dominates, the neuron will never fire. The only way to have a stationary state at a low rate (less than one spike per membrane time constant) is to have the excitation and inhibition nearly cancel. Then the mean membrane potential can lie a little below threshold, and the neuron can fire occasionally due to the input current fluctuations. This suggests the following heuristic theory, based on this approximate balance.

Using equations 2.3 and 2.5, we have

$$\sum_{b=0}^{2} J_{ab}\sqrt{K_b} r_b = \mathcal{O}(1), \tag{2.14}$$

or, up to corrections of $\mathrm{O}(1/\sqrt{K_0})$,

$$\sum_{b=0}^{2} \hat{J}_{ab} r_b = 0, \tag{2.15}$$

with $\hat{J}_{ab} = J_{ab}\sqrt{K_b/K_0}$. These are two linear equations in the two unknowns $r_a, a = 1, 2$, with the solution

$$r_a = \sum_{b=1}^{2} [\hat{\mathrm{J}}^{-1}]_{ab} J_{b0} r_0, \tag{2.16}$$

where $\hat{\mathrm{J}}^{-1}$ is the inverse of the $2 \times 2$ matrix with elements $\hat{J}_{ab}$, $a, b = 1, 2$. If there is a stationary balanced state, the average rates of the excitatory and inhibitory populations are given by the solutions of equation 2.16. This argument, given by Amit and Brunel and by Sompolinsky and van Vreeswijk, depends only on the rates, not on the correlations.

*2.1.2 Numerical Procedure.* For integrate-and-fire neurons in a stationary state, the mean field theory can be reduced to a set of analytic equations if neuron-to-neuron rate variations are neglected and a white-noise (Poisson firing) approximation is made (Amit & Brunel, 1997a, 1997b; Brunel, 2000). However, in a complete mean field theory for our model, the randomness in the connectivity forces these features to be taken into account, and it is necessary to resort to numerical methods. Thus, we simulate single neurons driven by gaussian synaptic currents; collect their firing statistics to compute the rates $r_a$, rate fluctuations $\overline{(\delta r_j^a)^2}$, and correlations $C_a(t-t')$; and then use these to generate improved input current statistics. The cycle is repeated until the input and output statistics are consistent. This algorithm was first used by Eisfeller and Opper (1992) to calculate the remanent magnetization of a mean field model for spin glasses.

Explicitly, we proceed as follows. We simulate single excitatory and inhibitory neurons over "trials" 100 integration time steps long. (We will call each time step a "millisecond." We have explored using smaller time steps and verified that there are no qualitative changes in the results.) We start from estimates of the rates given by the balance condition, which makes the net mean input current vanish. Then the sum over presynaptic populations of the $\mathcal{O}(\sqrt{K_b})$ terms in equation 2.9 vanishes, leaving only the rate variation and noise terms. We then run 10,000 trials of single excitatory and inhibitory neurons, selecting on each trial random values of $x_i^{ab}$ and $\xi_i^{ab}(t)$. Since at this point we do not have any estimates of either the rate fluctuations $\overline{(\delta r_j^b)^2}$ or the correlations $C_b(t-t')$, we use $r_b^2$ in place of $\overline{(r_j^b)^2}$ in equation 2.9 and use white noise for $\xi_i^{ab}(t)$: $C_b(t-t') \to r_b\delta(t-t')$.

The random choice of $x_i$ from trial to trial effectively samples across the neuronal populations, so we can then collect the statistics $r_a$, $\overline{(r_j^a)^2}$ (or, equivalently, $\overline{(\delta r_j^a)^2}$), and $C_a(t-t')$ from these trials. These can be used to generate an improved estimate of the input noise statistics to be used in equation 2.9 in a second set of trials, which yields new spike statistics again. This procedure is iterated until the input and output statistics agree. This may take up to several hundred iterations, depending on network parameters and how the computation is organized.

If one tries this procedure in its naive form, that is, using the output statistics directly to generate the input noise at the next step, it will lead to big oscillations and not converge. It is necessary to make small corrections (of relative order $1/\sqrt{K_0}$) to the previous input noise statistics to guarantee convergence.

When one computes statistics from the trials in any iteration, the simplest procedure involves calculating not the average correlation function $C_b(t-t')$ defined in equation 2.8 but, rather, $\tilde{C}_b(t-t')$ (see equation 2.12). From it, we can proceed in two ways. In one (the first of the three schemes described above for organizing the noise), from its $|t-t'| \to \infty$ limit we can obtain $\overline{(\delta r_j^b)^2}$, and thereby $\overline{(r_j^b)^2} = r_b^2 + \overline{(\delta r_j^b)^2}$ for use in equation 2.9. Subtracting this limiting value from $\tilde{C}_b(t-t')$ gives us $C_b(t-t')$ (which vanishes for

large $|t - t'|$) for use in generating the noise $\xi_i^{ab}(t)$. We will call this the *subtracted correlation method*.

Alternatively, as in the third of the schemes above, we can, at each step of our iterative procedure, generate noise directly with the correlations $\tilde{C}_b(t - t')$ (which have a long-range time dependence) and use $r_b^2$ in place of $(r_j^b)^2$ in equation 2.9. We call this the *unsubtracted correlation method*. We have verified that the two methods give the same results when carried out numerically, though the second one converges more slowly.

While the true rates in the stationary case are time independent and $C_a(t, t')$ is a function only of $t - t'$, the statistics collected over a finite set of noise-driven trials will not exactly have these stationarity properties. Therefore, we improve the statistics and impose time-translational invariance by averaging the measured $r_a(t)$ and $\overline{(\delta r_j^a(t))^2}$ over $t$ and averaging over the measured values $C_a(t, t')$ with a fixed $t - t'$.

After the iterative procedure converges, so that we have a good estimate of the statistics of the input, we want to run many trials on a single neuron and compute its firing statistics. This means that the numbers $x_i^{ab}$ ($b = 0, 1, 2$) should be held constant over these trials. In this case, it is necessary to subtract out the large $t - t'$ limit of $\tilde{C}_a(t - t')$ and use fixed $x_i^{ab}$ (constant in time and across trials) to generate the input noise. (If we did it the other way, without the subtraction, we would effectively be assuming that $x_i^{ab}$ changed randomly from trial to trial, which is not correct.)

In our calculations we have used 10,000 trials to calculate these single-neuron firing statistics. We perform the subtraction of the long-time limit of $\tilde{C}_a(t - t')$ at $|t - t'| = 50$, and we have checked that equation 2.12 is flat beyond this point in all the cases we have done.

If we perform this kind of measurement separately for many values of the $x_i^{ab}$, we will be able to see how the firing statistics vary across the population. Here, however, we will confine most of our attention to what we call the "average neuron": the one with the average value (0) of all three $x_i^{ab}$.

In particular, we calculate the mean spike count in the 100 ms trials and its variance across trials. From this we can get the Fano factor $F$ (the variance-to-mean ratio). We also compute the autocorrelation function, which offers a consistency check, since the Fano factor can also be obtained from

$$F = \frac{1}{r} \int_{-\infty}^{\infty} C(\tau) \mathrm{d}\tau. \tag{2.17}$$

(This formula is valid when the measurement period is much larger than the time over which $C(\tau)$ falls to zero.)

We will study how these firing statistics vary as we change various parameters of the model: the input rates $r_0$, parameters that control the balance of excitation and inhibition and the overall strength of the synapses.

This will give us some generic understanding of what controls the degree of irregularity of the neuronal firing.

**2.2 Nonstationary Case.** When the input population is not firing at a constant rate, almost the same calculational procedure can be followed, except that one does not average measured rates, their fluctuations, or correlation function over time. To start, we get initial instantaneous rate estimates from the balance condition, assuming that the time-dependent average input currents do not vary too quickly. (This condition is not very stringent; van Vreeswijk and Sompolinsky showed that the stability eigenvalues are proportional to $\sqrt{K_0}$, so if they have the right sign, the convergence to the balanced state is very rapid.)

To do the iterative procedure to satisfy the self-consistency conditions of the theory, it is simplest to use the second of the two ways described above (the unsubtracted correlation method). In this case, we get equations for the noise input currents just like equation 2.9, except that the second term is omitted and the $r_b$ are $t$-dependent and the correlation functions $C_b^{\mathrm{tot}}$ depend on both $t$ and $t'$, not just their difference. The only tricky part is the subtraction of the long-time limit of the correlation function, which is not simply defined.

We treat this problem in the following way. We examine the rate-normalized quantity,

$$\hat{C}_a(t,t') = \frac{C_a^{\mathrm{tot}}(t,t')}{r_a(t)r_a(t')}. \tag{2.18}$$

We find that this quantity is time-translation invariant (i.e., a function only of $t - t'$) to a very good approximation, so we perform the subtraction of the long-time limit on it. Then multiplying the subtracted $\hat{C}$ by $r_a(t)r_a(t')$ gives a good approximation to the true (subtracted) correlation function $C_a(t,t')$. The meaning of this finding is, loosely speaking, that when the rates vary (slowly enough) in time, the correlation functions just inherit these rates as overall factors without changing anything else about the problem.

We will use the this time-dependent formulation below to simulate experiments like those of Gershon et al. (1998), where the LGN input $r_0(t)$ to visual cortical cells is time dependent because of the flashing on and off of the stimulus.

## 3 Results

The results presented in this article were obtained from simulations with parameters $K_1 = 4444$ excitatory inputs and $K_2 = 1111$ inhibitory inputs to each neuron. The average number of external (excitatory) inputs $K_0$

was chosen to be equal to $K_2$. All neurons have the same membrane time constant $\tau$ of 10 ms.

To study the effect of various combinations in synaptic strength, we use the following generic form to define the intracortical weights $J_{ab}$:

$$\begin{pmatrix} J_{11} & J_{12} \\ J_{21} & J_{22} \end{pmatrix} = \begin{pmatrix} \epsilon & -2g \\ 1 & -2g \end{pmatrix}. \tag{3.1}$$

For the synaptic strengths from the external population, we use $J_{10} = 1$ and $J_{20} = \epsilon$. With this notation, $g$ determines the strength of inhibition relative to excitation within the network and $\epsilon$ the strength of intracortical excitation. Additionally, we scale the overall strength of the synapses with a multiplicative scaling factor denoted $J_s$ so that each synapse has an actual weight of $J_s \cdot J_{ab}$, regardless of $a$ and $b$.

Figure 2 summarizes how the firing statistics depend on all of the parameters $g$, $\epsilon$, and $J_s$. The irregularity of spiking, as measured by the Fano factor, depends most sensitively on the overall scaling of the synaptic strength, $J_s$. The Fano factor increases systematically as $J_s$ increases, and higher values of intracortical excitation $\epsilon$ also result in higher values of $F$. The same pattern holds for stronger intracortical inhibition, parameterized by $g$. For all of these cases, the mean firing rate remains virtually unchanged due to the dynamic balance of excitation and inhibition in the network, whereas the fluctuations increase with the increase of any of the synaptic weights.

Interspike interval (ISI) distributions are shown in Figure 3 for three different values of $J_s$, keeping $\epsilon$ and $g$ fixed at 0.5 and 1, respectively. For a Poisson spike train, the Fano factor $F = 1$, while $F > 1$ (which we term *super-Poissonian*) indicates a tendency of spikes occurring in clusters separated by accordingly longer empty intervals, and $F < 1$ (*sub-Poissonian*) indicates more regularity, reflected by a narrower distribution. We have adjusted the input rate $r_0$ so that the output rate is the same in all three cases.

The top panel of Figure 3 shows the ISI distribution of a super-Poissonian spike train, obtained for $J_s = 1.42$. Overlaied on the histogram of ISI counts is an exponential curve indicating a Poisson distribution with the same mean ISI length. Compared with the Poisson distribution, the super-Poissonian spike train contains more short intervals, as seen by the peak at short lengths, and also more long intervals, causing a long tail. Necessarily, the interval count around the average ISI length is lower than that for a Poisson spike train.

The ISI distribution in the middle panel of Figure 3 belongs to a spike train with a Fano factor close to one, obtained for $J_s = 0.714$. The overlaid exponential reveals a deviation from the ISI count: while intervals of diminishing length are the most likely ones for a real Poisson process, our neuronal spike trains always show some refractoriness reflected by a dip at

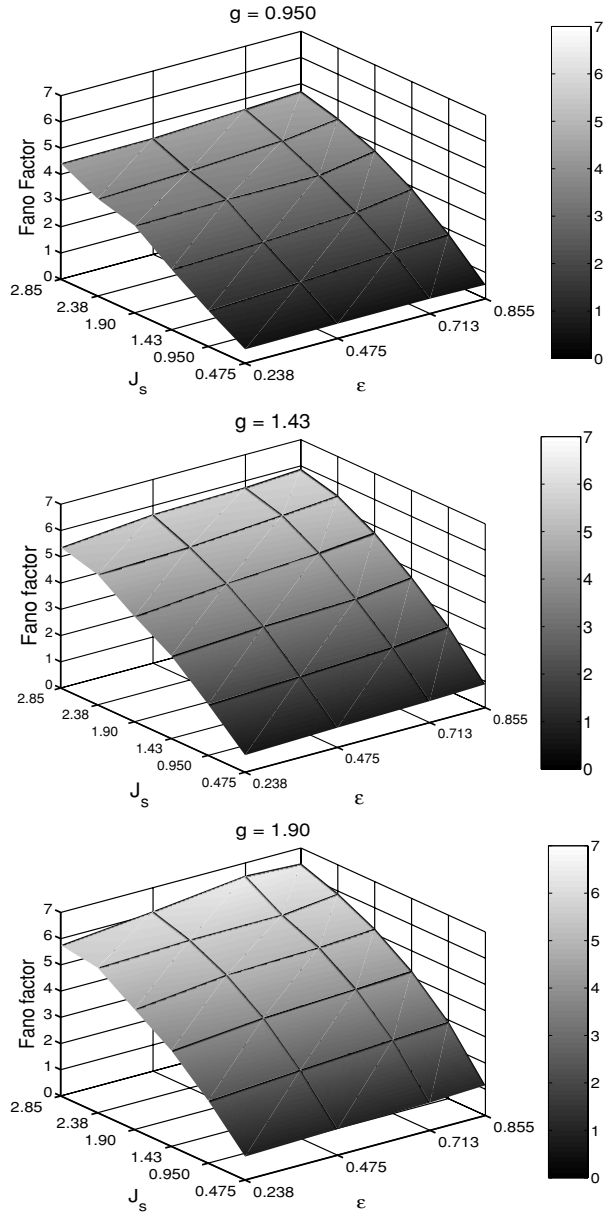


Figure 2: Fano factors as a function of overall synaptic strength $J_s$ and intracortical excitation strength $\epsilon$ for three different inhibition factors: $g = 1$, 1.5, and 2, respectively. The increase of any of these parameters results in more irregular firing statistics as measured by the Fano factor.

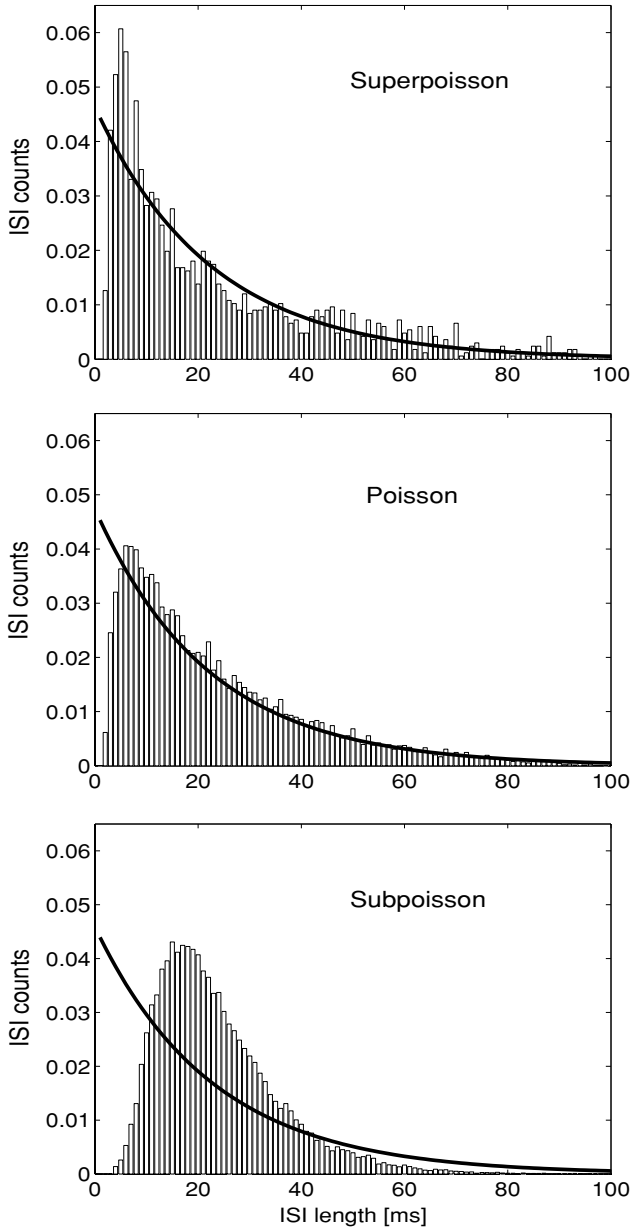


Figure 3: Interspike interval distributions for fixed $\epsilon = 0.5$ and $g = 1$, and three different values of overall synaptic strength $J_s$: 1.42 (super-Poissonian), 0.714 (Poissonian), and 0.357 (sub-Poissonian). Overlaid on each figure is the exponential fall-off of a true Poisson distribution with the same average rate as in all of the three cases.

the shortest intervals. (We have not used an explicit refractory period in our model. The dip seen here simply reflects the fact that it takes a little time for the membrane potential distribution to return to its steady-state form after reset.) Apart from this deviation, however, there is a close resemblance between the observed distribution and the "predicted" one.

Finally, the lower panel of Figure 3 depicts a case with $F < 1$, with weaker synapses, leading to a stronger refractory effect and (since the rate is fixed) an accordingly narrower distribution around the average ISI length, as compared to the overlaied Poisson distribution. This distribution was obtained with weak synapses produced by a small scaling factor of $J_s = 0.357$.

As mentioned in the previous section, the Fano factor can also be obtained by integrating over the spike train autocorrelation divided by the spike rate, equation 2.17. For a Poisson process, the autocorrelation vanishes for all lags different from zero. In contrast, $F > 1$ (super-Poissonian case) implies a positive integral over nonzero lags, whereas in the sub-Poissonian case, there must be a negative area under the curve. Figure 4 shows examples of autocorrelations for all of the three cases. For the super-Poissonian case (dashed line), there is a "hill" of positive correlations for short intervals,

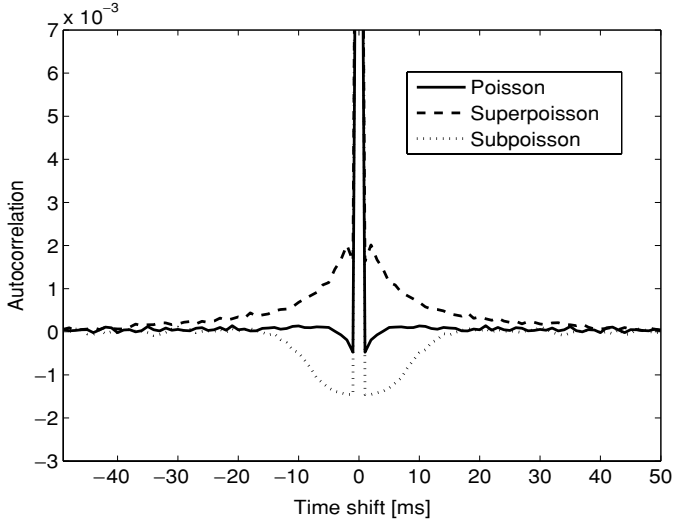


Figure 4: Three different spike train autocorrelations illustrating the relationship between the Fano factor $F$ and the area under the curve. For $F = 1$ (Poissonian, solid line), the autocorrelation is an almost perfect delta function. $F > 1$ (super-Poissonian, dashed line) is reflected by a hill generating a positive area, and $F < 1$ (sub-Poissonian, dotted line) is accompanied by a valley of negative correlations. (See the text for more details.)

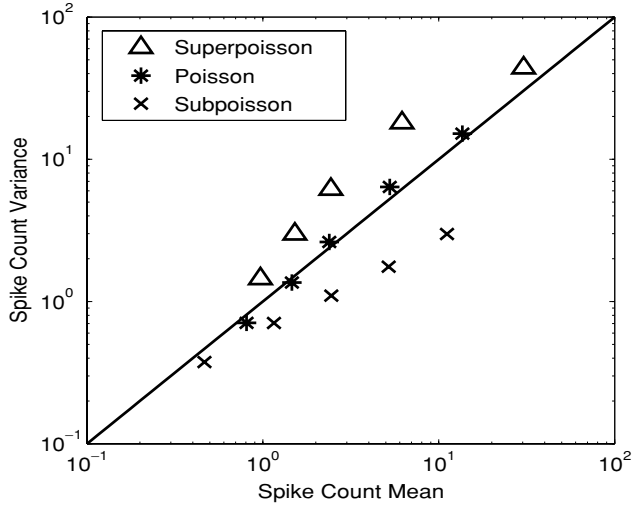


Figure 5: Spike count log(variance) versus log(mean) for three different values of overall synaptic strength $J_s$, varying the external input rate $r_0$. For $J_s = 1.19$ (super-Poissonian, triangles), the data look qualitatively like those from experiments. The other values for $J_s$ are 0.714 (Poisson, stars) and 0.357 (sub-Poissonian, crosses).

reflecting the tendency toward spike clustering. The sub-Poissonian autocorrelation (dotted line) shows a valley of negative correlations for short intervals, indicating well-separated spikes in a more regular spike train. The curve labeled as Poisson (solid line) does have a small valley around zero lag, which reflects once more the refractoriness of neurons to fire at extremely short intervals, unlike a completely random Poisson process. (Actually, the measured $F$ in this case is slightly greater than 1, implying that in this case, the integral of the very small positive tail for $t > 2$ ms is slightly larger than that of the (more obvious) negative short-time dip.)

Measurements on V1 neurons in awake monkeys (see, e.g., Gershon et al., 1998) suggest a linear relationship between the log variance and the log mean of stimulus-elicited spike counts. We find a similar dependence for neurons within our model network. Figure 5 shows results for three different values of $J_s$. In each case, five different values of the external input rate $r_0$ were tried, causing various mean spike counts and variances. The logarithm of the spike count variance is plotted as a function of the logarithm of the spike count mean, and a solid diagonal line indicates the identity, that is, a Fano factor of exactly 1. We see that for the largest value of $J_s$ used here, the data look qualitatively like those from experiments, with Fano factors in the range around 1.5 to 2.

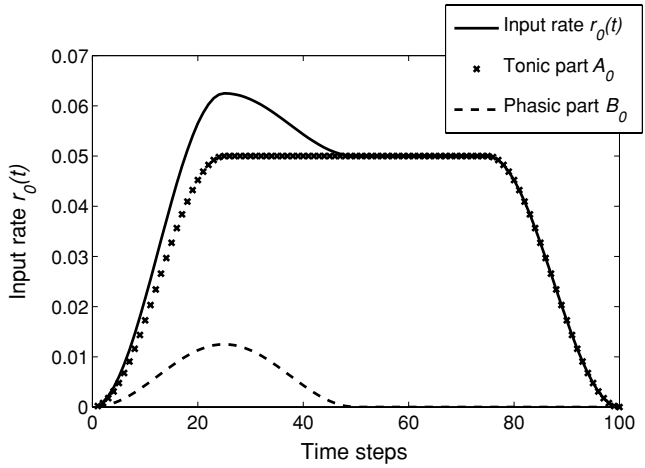


Figure 6: Parameterization of the time-dependent input rate $r_0(t)$. The input is modeled as the sum of three functions: (1) a stationary background rate (which is zero in this case); (2) a tonic part, which rises within the first 20 ms to a constant level of $A_0$, where it stays for 60 ms, falling back to zero within the last 20 ms; and (3) an initial phasic part, which is nonzero only in the first 50 ms, rising to a maximum value of $B_0$.

**3.1 Nonstationary Case.** The results presented in the previous section were obtained with stationary inputs, while experimental data like those from Gershon et al. (1998) were collected from visual neurons subject to time-dependent inputs. Therefore, we performed calculations of the spike statistics in which the input population rate $r_0$ was time dependent. The modeled temporal shape of $r_0(t)$ is depicted in Figure 6. It is the sum of three terms:

$$r_0(t) = R_0 + A(t) + B(t). \tag{3.2}$$

The first, $R_0$, is a constant, as in the preceding section. The second term, $A(t)$, rises to a maximum over a 25 ms interval, remains constant for 50 ms, and then falls off to zero over the final 25 ms,

$$A(t) = \begin{cases} 0.5A_0(1 - \cos(4t\pi/T)) & \text{for } 0 < t \le T/4 \\ A_0 & \text{for } T/4 < t \le 3T/4 \\ 0.5A_0(1 - \cos(4(T-t)\pi/T)) & \text{for } 3T/4 < t \le T, \end{cases} \tag{3.3}$$

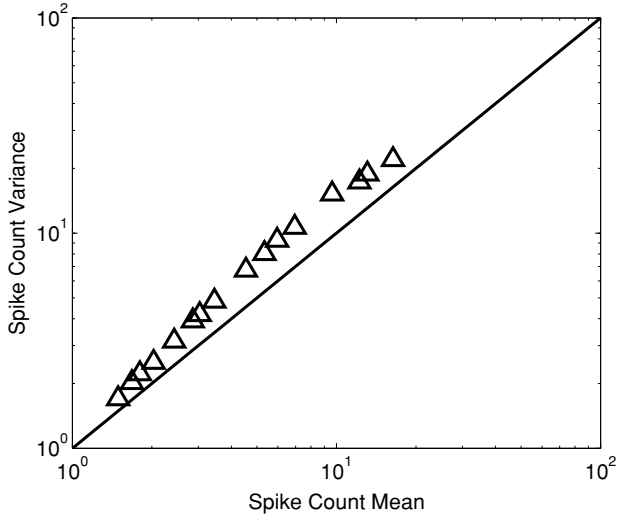


Figure 7: Spike count log(variance) versus log(mean) for time-varying external inputs with varying overall strength. The neuron in the simulated network (triangles) fires in a super-Poissonian regime, with an almost linear relationship for low spike rates between the log variance and the log mean, resembling closely data obtained from in vivo experiments. The diagonal solid line indicates the identity of variance and mean (Fano factor $F = 1$).

where $T$ is the total simulation interval of 100 ms. The third term, $B_0$, rises to a maximum in the first 25 ms and then falls back to zero in the next 25 ms, remaining zero thereafter:

$$B(t) = \begin{cases} 0.5B_0(1 - \cos(4t\pi/T)) & \text{for } 0 < t \le T/4 \\ 0.5B_0(1 - \cos(4(T/2 - t)\pi/T)) & \text{for } T/4 < t \le T/2 \\ 0 & \text{for } T/2 < t \le T. \end{cases} \tag{3.4}$$

Figure 7 shows the logarithm of the spike count variance plotted against the logarithm of the spike count mean for various nonstationary inputs characterized by different values of $A_0$ and $B_0$. The graph shows results for $J_s = 0.95$, $\epsilon = 0.5$, $g = 1$, and a background rate of $R_0 = 0.1$. Table 1 shows the choice of the 16 combinations of the stimulus parameters $A_0$ and $B_0$, together with the resulting Fano factors $F$ for the simulated neuron.

The data look qualitatively like those obtained from in vivo experiments by Gershon et al. (1998) and are similar to the super-Poissonian case in Figure 5. The neuron fires consistently in a super-Poissonian regime with Fano factors slightly higher than 1 and an almost linear relationship between the

Table 1: Stimulus Parameters $A_0$ and $B_0$ for the Results Depicted in Figure 7 and the Resulting Fano Factors $F$.

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| $A_0$ | 0.375 | 0.375 | 0.500 | 0.500 | 0.750 | 0.750 | 1.000 | 1.000 |
| $B_0$ | 0.125 | 0.375 | 0.125 | 0.375 | 0.250 | 0.750 | 0.250 | 0.750 |
| $F$ | 1.14 | 1.2 | 1.22 | 1.23 | 1.29 | 1.36 | 1.37 | 1.4 |
| $A_0$ | 1.500 | 1.500 | 2.000 | 2.000 | 3.000 | 3.000 | 4.000 | 4.000 |
| $B_0$ | 0.500 | 1.500 | 0.500 | 1.500 | 1.000 | 3.000 | 1.000 | 3.000 |
| $F$ | 1.48 | 1.5 | 1.55 | 1.53 | 1.57 | 1.41 | 1.43 | 1.34 |

log variance and the log mean for low spike counts. For higher spike counts, the curve bends toward values of lower Fano factors, just as for stationary inputs (see Figure 5). In both cases, this bend reflects the decrease in irregularity of firing caused by an increasingly prominent role of refractoriness for shorter interspike intervals.

**3.2 Comparison with Network Simulations.** An extensive exploration of the validity of mean field theory is beyond the scope of this letter. However, we have performed simulations of networks constructed according to the model of section 2 and compared their firing irregularity with that obtained in mean field theory. Specifically, we tested two main results from our mean field analysis, stating that Fano factors increase systematically with synaptic strength and that there is an approximate power law between the mean spike count and the spike count variance, similar to experimental findings (see Figures 2 and 5, respectively).

Figure 8 shows measured Fano factors for a typical neuron in a network with $K_0 = K_1 = 400$, $K_2 = 100$, and $N = 10000$, where we varied $J_s$ in the same range as in Figure 2 (other parameters were $g = 1$ and $\epsilon = 0.5$). The Fano factor increases systematically as $J_s$ increases, lying in the quantitative range predicted by mean field theory.

In addition (results not shown), we explored the lower and upper limits of Fano factors in our network. For $J_s = 0.1$, the average Fano factor of all neurons in the network was 0.034. Notwithstanding the very regular firing of all neurons in this network with very weak synapses, the overall activity remained asynchronous, as required in our mean field analysis. At the other extreme, for very strong synapses with $J_s = 32$, we observed an average Fano factor of 16.05, and individual neurons exhibited Fano factors of up to 30 and higher. These results show that networks of integrate-and-fire neurons exhibit a wide range of Fano factors in their balanced state, depending on synaptic strength.

In Figure 9, we show plots of the logarithm of the spike count variance as a function of the logarithm of the mean spike count for six individual

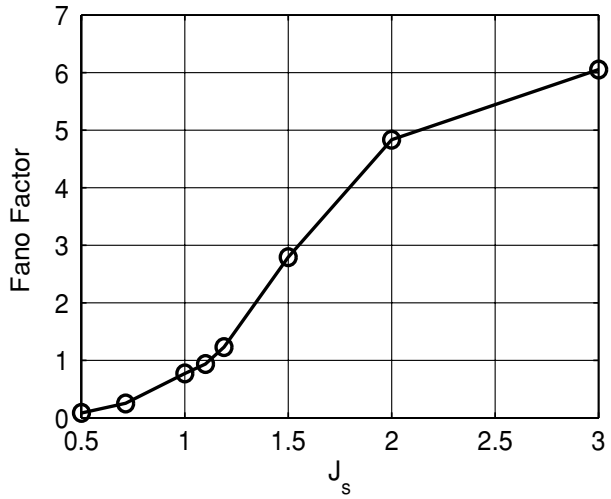


Figure 8: Fano factors as a function of overall synaptic strength $J_s$ obtained from network simulations for a randomly chosen neuron ($g = 1$, $\epsilon = 0.5$). A comparison with Figure 2 reveals that the mean field calculations correctly predict both the qualitative relationship between Fano factors and synaptic strength and the quantitative range of Fano factors for this range of $J_s$ values.

neurons in the network. Analogous to Figure 5, results for three different values of $J_s$ are shown (1.5, 1.1, and 0.714, indicated by triangles, stars, and crosses, respectively), each probed with five different strengths of external inputs. The neurons were chosen randomly from all 8000 excitatory neurons with nonzero firing rates. With the exception of neuron 4638 (in the lower middle panel), $J_s = 1.5$ resulted in super-Poissionian firing statistics, $J_s = 0.714$ in sub-Poissionian firing, and $J_s = 1.1$ in approximately Poisson statistics. There is a strong qualitative resemblance between the network simulation results in Figure 9 and the mean field results in Figure 5, with the latter showing spike count statistics of the hypothetical "average neuron" defined above.

Taken together, these results suggest that mean field theory provides a reliable way to estimate firing variability in balanced networks.

## 4 Discussion

Cortical neurons receive thousands of excitatory and inhibitory inputs, and despite the high number of inputs from nearby neurons with similar firing statistics and similar connectivity, their observed firing is very irregular (Heggelund & Albus, 1978; Dean, 1981; Tolhurst et al., 1981, 1983; Vogels, Spileers, & Orban, 1989; Snowden et al., 1992; Gur et al., 1997; Shadlen &

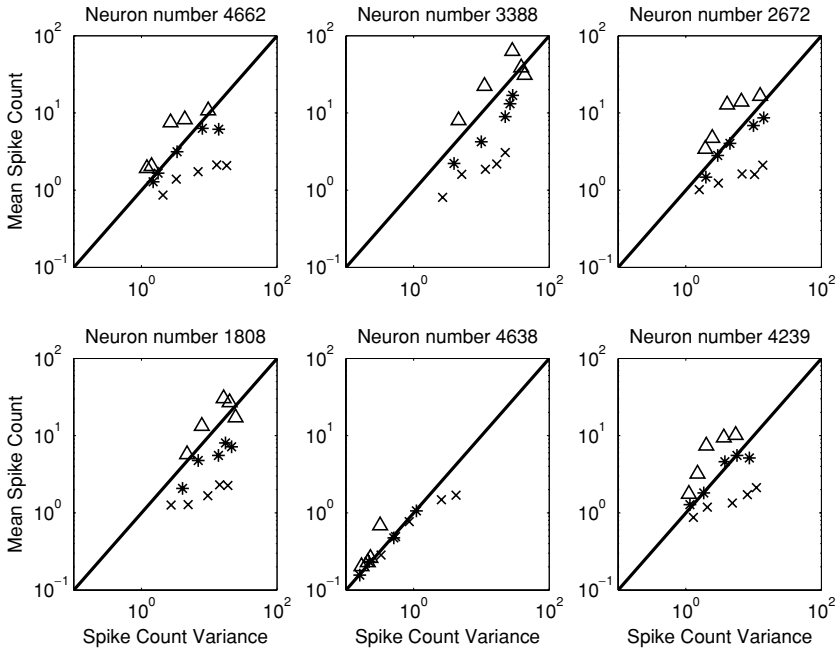


Figure 9: Spike count statistics obtained from network simulations for six randomly chosen neurons. Spike count log(variance) versus log(mean) are plotted as in Figure 5, for various external input rates at three different values of synaptic strength ($J_s = 1.5$, 1.1, and 0.714, represented by triangles, stars, and crosses, respectively). There is a close qualitative resemblance to the results from mean field calculations shown in Figure 5, where spike statistics of a hypothetical neuron with average overall input are shown.

Newsome, 1998; Gershon et al., 1998; Kara et al., 2000; Buracas et al., 1998; Lee et al., 1998; DeWeese et al., 2003). Dynamically balanced excitation and inhibition through a simple feedback mechanism provide an explanation that naturally accounts for this phenomenon without requiring fine-tuning of the parameters (Amit & Brunel, 1997a, 1997b; Brunel, 2000; van Vreeswijk and Sompolinsky, 1996, 1998). Moreover, neurons in such model networks show an almost linear input-output relationship (input current versus firing frequency), as do neurons in the neocortex.

Whenever one wants to understand a complex dynamical system, one asks whether there is an approximate theory, possibly exact in some interesting limit, that captures and affords insight into the observed properties. Here, the high connectivity of the cortical networks of interest suggests trying to obtain this insight from mean field theory, which becomes exact for an infinite, extensively connected system. In this article, we have formulated a

complete mean field description of the dynamically balanced asynchronous firing state in the dilute, high-connectivity limit $N_a \gg K_a \gg 1$. Because of the assumed random connection structure in the network, the mean field theory has to include autocorrelation functions and rate variations as well as population-mean rates as order parameters, as in spin glasses (Sompolinsky & Zippelius, 1982).

We used this mean field theory to analyze firing correlations. We found that the relationship between the observed irregularity of firing (spike count variance) and the firing rate (spike count mean) of the neurons resembles closely data collected from in vivo experiments (see Figures 5 and 7). To do this, we developed a complete mean field theory for a network of leaky integrate-and-fire neurons, in which both firing rates and correlation functions are determined self-consistently. Using an algorithm that allows us to find the solutions to the mean field equations numerically, we could elucidate how the strength of synapses within the network influences the expected firing statistics of cortical neurons in a systematic manner (see Figure 2).

We have shown that the irregularity of firing, as measured by the Fano factor, increases with increasing synaptic strengths (see Figure 2). Nearly Poisson statistics (with $F \approx 1$) are observed for moderately strong synapses, but the transition from sub-Poissonian to super-Poissonian statistics is smooth, without a special role for $F = 1$.

The higher irregularity in the spike counts is always accompanied by a tendency toward more "bursty" firing. (These bursts are a network effect because the model contains only leaky integrate-and-fire neurons, which do not burst on their own.) This burstiness can best be seen in the spike train autocorrelation function (see Figure 4), which acquires a hill of growing size and width around zero lag for increasing Fano factors. The interdependence between firing irregularity and bursting can be understood with the help of the ISI distributions depicted in Figure 3: when the rate, and thus the average ISI, is kept constant, then any higher count for shorter-than-average ISIs must be accompanied by an accordingly higher count for longer ISIs (indicating bursts), and vice versa. Thus, higher irregularity always goes hand in hand with a higher tendency toward temporal clustering of spikes.

Why do stronger synapses lead to higher irregularity in firing? The size of the input current fluctuations in equation 2.9 is controlled by the $J_{ab}$, and so, therefore, are the corresponding membrane potential fluctuations. Thus, for example, the width of the steady-state membrane potential distribution is proportional to $J_s$. We next have to consider where this distribution is centered. Remembering that, according to the balance condition, the firing rate is independent of $J_s$, and the center of the distribution has to move farther away from threshold as $J_s$ is increased in order to keep the rate fixed. Therefore, for very small $J_s$ almost the entire equilibrium membrane potential distribution will lie well above the postspike reset value, while for large $J_s$, it will be mostly below reset.

Immediately after a spike, the membrane potential distribution is a delta function centered at the reset (here 0). It then spreads, and its mean moves up or down toward its equilibrium value. This equilibration will take about a membrane time constant. If the equilibrium value is well above zero (the small-$J_s$ case), the probability of reaching threshold will be suppressed during this time, implying a refractory dip in the ISI distribution and the correlation function and a tendency toward a Fano factor less than 1.

In the large-$J_s$ case, where the membrane potential is reset much closer to the threshold than to its eventual equilibrium value, the initial rapid spread (with the width growing proportional to $J_s\sqrt{t}$) leads to an enhanced probability of early spikes. At short times, this diffusive spread dominates the downward drift of the mean (which is only linear in $t$). Thus, there is extra weight in the ISI distribution and a positive correlation function at these short times, leading to a Fano factor greater than 1.

Empirically, an approximate power-law relationship between the mean and variance of the spike count has frequently been observed for cortical neurons (see, e.g., Tolhurst et al., 1981; Vogels, Spileers, & Orban, 1989; Gershon et al., 1998; Lee et al., 1998). Our model shows the same qualitative feature (see Figures 5 and 6), though we have no argument that the relation should be an exact power law. However, this agreement suggests that the model captures at least part of physics underlying the firing statistics.

As already observed, not all of the variability in measured neuron responses has to be explained in the manner outlined above. Changing conditions during the run of a single experiment may introduce extra irregularity, caused by collecting statistics over trials with different mean firing rates. Our analysis shows why—and how much—irregularity can be expected due to intrinsic cortical dynamics.

Other authors have also studied firing irregularity, in phase-oscillator models (Bressloff & Coombes, 2000; Bressloff, Bressloff, & Cowan, 2000) and in a ring model with inhomogeneous excitation and inhibition (Lin, Pawelzik, Ernst, & Sejnowski, 1998). Both groups found that their models could produce highly irregular firing. In our work, we have tried to make a systematic study of how the irregularity (quantified by the Fano factor) depends on system parameters for a fairly simple model appropriate for describing local (intracolumn) neocortical networks.

We have used instantaneous synapses; that is, we have not included synaptic filtering of input spike trains in the calculations we have reported here. However, we have incorporated such filtering, with a simple exponential kernel, into our code and explored the effects of a nonzero synaptic current decay time $\tau_{\mathrm{syn}}$. We find that for small $\tau_{\mathrm{syn}}/\tau$, Fano factors grow proportional to this ratio,

$$F(\tau_{\mathrm{syn}}) = F(0)\left(1 + a\,\frac{\tau_{\mathrm{syn}}}{\tau}\right), \tag{4.1}$$

with $a = \mathcal{O}(1) > 0$. Since we are most interested in the limit $\tau_{\rm syn} \ll \tau$, we have not studied these corrections in detail. However, it should be noted that in a model where the synapses are modeled by conductances instead of current pulses (Lerchner, Ahmadi, & Hertz, 2004), the effective membrane time constant can become very small, so $\tau_{\rm syn}$ can be considerably larger than it (Destexhe et al., 2003). In this case, the dynamics become rather different.

Our formulation of the mean field theory is general enough to allow other straightforward extensions toward greater biological realism and more complicated network architectures. We have extended the model to include systematic structure in the connections, modeling an orientation hypercolumn in the primary visual cortex (Hertz & Sterner, 2003). Moreover, our algorithm for finding the mean field solutions is not restricted to networks of integrate-and-fire neurons. It can be applied to any kind of neuronal model. Furthermore, synaptic depression and facilitation can be incorporated by using synaptically filtered spike trains to compute the self-consistent solutions.

As we remarked earlier, if one ignores correlations in the synaptic input and neuron-to-neuron rate variations (Amit & Brunel, 1997a, 1997b), analytic self-consistent equations for the population rates can be derived. From these, one can calculate the steady-state Fano factor analytically in closed form (Brunel, 2000). Such a calculation is obviously not self-consistent, although it can give qualitative information about firing irregularity. We have done some calculations, using our single-neuron simulation methods, but in which we impose the Amit-Brunel approximations by hand when generating the input noise. As one could anticipate, we find that this procedure systematically underestimates Fano factors in the super-Poissonian regime, by factors of up to 2 or so at the largest values of $J_s$ studied here.

Since it is necessary to solve the full mean field theory numerically, one might ask: If it is necessary to resort to numerical solution anyway, why not just simulate the network directly? Our answer is that beyond the advantage of having to simulate only one neuron at a time, it is interesting to know what the predictions of mean field theory are. To the extent that they agree with network simulations, we can understand our findings in terms of single-neuron properties (albeit with self-consistent synaptic current statistics). Discrepancies would point to either finite-size or finite-concentration effects or more subtle correlation effects not included in the mean field ansatz. Identifying such effects, if they exist, would point the way toward future theoretical investigations, which could shed potentially useful light on the dynamics of these networks.

## References

Amit, D., & Brunel, N. (1997a). Dynamics of a recurrent network of spiking neurons before and following learning. *Network*, *8*, 373–404.

Amit, D., & Brunel, N. (1997b). Model of spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cerebral Cortex*, *7*, 237–252.

Bressloff, P. C., Bressloff, N. W., & Cowan, J. D. (2000). Dynamical mechanism for sharp orientation tuning in an integrate-and-fire model of a cortical hypercolumn. *Neural Comp.*, *12*, 2473–2511.

Bressloff, P. C., & Coombes, S. (2000). Dynamics of strongly coupled neurons. *Neural Comp.*, *12*, 91–129.

Brunel, N. (2000). Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *J. Comput. Neurosci.*, *8*, 183–208.

Buracas, G. T., Zador, A. M., DeWeese, M. R., & Albright, T. D. (1998). Efficient discrimination of temporal patterns by motion-sensitive neurons in primate visual cortex. *Neuron*, *20*, 959–969.

Dean, A. F. (1981). The variability of discharge of simple cells in the cat striate cortex. *Exp. Brain Res.*, *44*, 437–440.

Destexhe, A., Rudolph, M., & Paré, D. (2003). The high-conductance state of neocortical neurons in vivo. *Nature Rev. Neurosci.*, *4*, 739–761.

DeWeese, M. R., Wehr, M., & Zador, A. M. (2003). Binary spiking in auditory cortex. *J. Neurosci.*, *23*, 7940–7949.

Eisfeller, H., & Opper, M. (1992). New method for studying the dynamics of disordered spin systems without finite-size effects. *Phys. Rev. Lett.*, *68*, 2094–2097.

Fulvi Mari, C. (2000). Random networks of spiking neurons: Instability in the *Xenopus* tadpole moto-neuron pattern. *Phys. Rev. Lett.*, *85*, 210–213.

Gershon, E., Wiener, M. C., Latham, P. E., & Richmond, B. J. (1998). Coding strategies in monkey V1 and inferior temporal cortex. *J. Neurophysiol.*, *79*, 1135–1144.

Gur, M., Beylin, A., & Snodderly, D. M. (1997). Response variability of neurons in primary visual cortex (V1) of alert monkeys. *J. Neurosci.*, *17*, 2914–2920.

Heggelund, P., & Albus, K. (1978). Response variability and orientation discrimination of single cells in striate cortex of cat. *Exp. Brain Res.*, *32*, 197–211.

Hertz, J., Richmond, B., & Nilsen, K. (2003). Anomalous response variability in a balanced cortical network model. *Neurocomputing*, *52–54*, 787–792.

Hertz, J., & Sterner, G. (2003). Mean field model of an orientation hypercolumn. *Soc. for Neurosci. Abstract*, no. 911.19.

Kara, P., Reinagel, P., & Reid, R. C. (2000). Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons. *Neuron*, *27*, 635–646.

Lee, D., Port, N. L., Kruse, W., & Georgopoulos, A. P. (1998). Variability and correlated noise in the discharge of neurons in motor and parietal areas of primate cortex. *J. Neurosci.*, *18*, 1161–1170.

Lerchner, A., Ahmadi, M., & Hertz, J. (2004). High conductance states in a mean field cortical network model. *Neurocomputing*, *58–60*, 935–940.

Lin, J. K., Pawelzik, K., Ernst, U., & Sejnowski, T. J. (1998). Irregular synchronous activity in stochastically-coupled networks of integrate-and-fire neurons. *Network*, *9*, 333–344.

Oram, M. W., Wiener, M. C., Lestienne, R., & Richmond, B. J. (1999). Stochastic nature of precisely-timed spike patterns in visual system neural responses. *J. Neurophysiol.*, *81*, 3021–3033.

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.

Snowden, R. J., Treue, S., & Andersen, R. A. (1992). The response of neurons in areas V1 and MT of the alert rhesus monkey to moving random dot patterns. *Exp. Brain Res.*, *88*, 389–400.

Sompolinsky, H., & Zippelius, A. (1982). Relaxational dynamics of the Edwards-Anderson model and the mean-field theory of spin glasses. *Phys. Rev. B.*, *25*, 6860–6875.

Tolhurst, D. J., Movshon, J. A., & Dean, A. F. (1983). The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res.*, *23*, 775–785.

Tolhurst, D. J., Movshon, J. A., & Thompson, I. D. (1981). The dependence of response amplitude and variance of cat visual cortical neurones on stimulus contrast. *Exp. Brain Res.*, *41*, 414–419.

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 Comp.*, *10*, 1321–1371.

Vogels, R., Spileers, W., & Orban, G. A. (1989). The response variability of striate cortical neurons in the behaving monkey. *Exp. Brain Res.*, *77*, 432–436.

---

Received February 18, 2004; accepted August 23, 2005.