

SUPPLEMENTARY INFORMATION

Contents	Page
1. Extracellular spike detection	1
2. Two experimental paradigms predict the same probability of an extra spike	3
3. Detecting up and down states	4
4. Biophysically realistic model	6
5. The effect of feedforward inhibition	10
6. The level of intrinsic noise at equilibrium	14
7. Network simulations	28
8. Error analysis	34
9. The log normal distribution	37
10. Difference of exponentials – normalization and integrals	38
11. List of parameters	39

1. Extracellular spike detection

This section describes pre-processing of the data recorded on the silicon probe (described in main text and Online Methods) prior to analysis. The extracellular signal was separated online into an LFP component (0-0.3 kHz) sampled at 3 kHz, and a high frequency component (0.3-5 kHz) sampled at 25 kHz. The high frequency component was used to detect spikes offline. Before looking for spikes on the extracellular electrodes, we removed stimulus-induced and spike-induced artifacts. For the former, we identified null-stimuli – stimuli that did not trigger intracellular spikes (see point 1 in Fig. S1a for an example and Fig. S1b for an average) – and used them to compute the null-stimulus triggered average on each channel. We then subtracted each channel's null-stimuli triggered average from the channel signal (aligned, of course, on each stimulus). To remove spike-induced artifacts, we followed essentially the same procedure, except that we computed an average voltage triggered on spontaneous intracellular spikes (which induced a very small capacitive signal; see Fig. S1a point 2 for an example and Fig. S1c for an average). We then subtracted each channel's spontaneous-spike triggered average from the channel signal, this time aligned on all intracellular spikes (spontaneous as well as triggered).

To improve the quality of the signal, we applied a low-pass filter based on the noise spectrum. The noise spectrum was estimated from the intracellular recordings during the

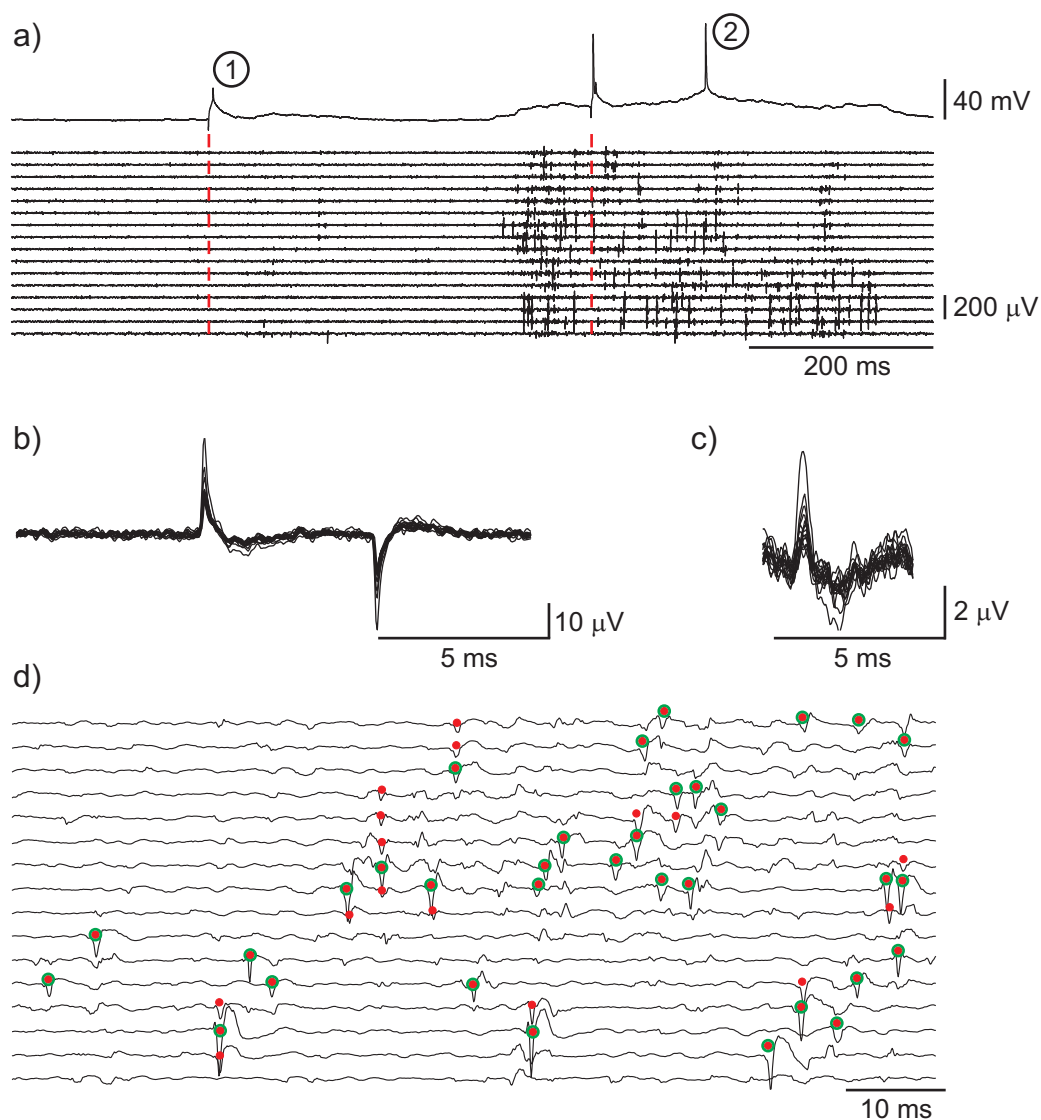


Figure S1: Simultaneous intracellular and extracellular recordings. **a.** A one second example of intracellular voltage (top) and simultaneous extracellular voltage on a 16 channel electrode (bottom, band-passed 0.3-5 kHz). Dashed red lines indicate stimulus onset. **b.** Average extracellular voltage triggered on stimuli that did not produce a spike (point 1 in panel a). This average is a capacitive artifact, and was subtracted from the voltage on each channel. **c.** Average extracellular voltage triggered on intracellular spikes (point 2 in panel a). This average is also a capacitive artifact (although a much smaller one), and was also subtracted from the voltage on each channel. **d.** After subtracting the artifacts, the signal on each channel was normalized to the noise (so the vertical scale bar has units of standard deviation) and smoothed using a Wiener filter (see text). Fast, large events were considered candidate spikes (red dots). These often occurred simultaneously on neighboring channels, in which case only the largest spike (over a range of 5 channels) was included in the analysis. Included spikes are marked with a green circle; events without a green circle were discarded.

down states (during which there were never any spikes on the extracellular electrodes). The spectrum was then used to construct a Wiener filter (using Matlab `wiener2` imaging toolbox function). Filtering the data resulted in a 15-30% improvement in signal to noise ratio for spike detection, with negligible effect on spike amplitude or shape.

Finally, to simplify the analysis, we normalized the voltage on each channel to the standard deviation of the noise on that channel. To estimate the standard deviation, we used $\text{median}(|V|)/0.6745$ where V is the voltage. This estimator avoids the upward bias associated with excursions in voltage caused by action potentials.¹ After normalization, candidate spikes were identified based on two criteria:

- There was a local minimum with amplitude greater than six standard deviations,
- There were zero-crossings less than 0.5 ms before the minimum and 1.0 ms after it.

Multiple spikes – spikes occurring within 0.5 ms of each other – were often detected on neighboring channels. When this happened, only one spike was selected (the one with the largest amplitude); the rest were discarded. This is illustrated in Fig S1d, where red dots are used to indicate all candidate spikes and red dots surrounded by green circles denote spikes used in the analysis.

2. Two experimental paradigms predict the same probability of an extra spike

Our results were primarily derived from two sets of experiments. The first measured the effect of current injections into single neurons *in vivo* (Fig. 3, main text); the second measured the effect of current injections on the surrounding networks (Fig. 4, main text). The question we address now is: do these two experiments give consistent results? To answer this, we compute the local increase in firing rate predicted by the first set of experiments and compare that to the increase measured in the second.

In our current injection experiments, we found that a single extra presynaptic spike increased the probability of a postsynaptic spike in a connected neuron by approximately 2% for 5 ms (Fig. 3, main text). This corresponds to 2 extra spikes per 100 trials per 5 ms, which in turn corresponds to an increase in firing rate of 4 Hz. To determine what this implies for the extracellular recordings, the first observation is that connectivity in barrel cortex is sparse, with a connection probability of about 4%.² Thus, if the extra spikes on the

extracellular electrodes all occurred within 5 ms, the expected increase in firing rate would be 0.16 Hz (0.04×4 Hz). It is unlikely that they do, however, because spikes were triggered using current pulses lasting 2.5-10 ms, causing a range of timings of the extra presynaptic spike. In addition, axonal delays, dendritic filtering, and latency to spike can further spread out the time of the postsynaptic spikes, adding another 1-5 ms. Thus, the extra postsynaptic spikes should occur primarily within 10-20 ms from the start of the stimulus, which implies an increase in firing rate of 0.04-0.08 Hz (i.e., a factor of 2-4 smaller than the 0.16 Hz increase expected if all the spikes were to occur within 5 ms). Assuming that the average firing rate in somatosensory cortex is 1 Hz,³ this is very close to the 0.03-0.065 Hz increase we see in the first 10-20 ms in the inset of Fig. 4c. Thus, not only does a single extra spike introduced into somatosensory cortex produce a measurable effect on the network – one that lasts for more than 50 ms – it produces an effect whose size is predicted by our single-neuron current injection experiments (Fig. 3, main text).

3. Detecting up and down states

Two key parameters in our analysis are the standard deviation of the membrane potential and the rate of rapidly depolarizing events. Both need to be calculated only during the up state, since it is the up state that resembles the *in vivo* membrane potential trajectory.^{4,5} Here we describe our procedure for separating the up states from the down states.

A typical voltage trace from an *in vivo*, whole cell recording is shown in Fig. S2a (see Sec. 1). Recordings were made in 11 second traces. To suppress spiking, we injected steady hyperpolarizing current between -250 and -450 pA; typically the amount of current required to suppress spiking was the amount required to hold the cell at -75 mV during voltage clamp. No current was injected during the first and last 250 ms. The first 500 ms of the hyperpolarized trace were discarded to avoid the sag associated with the I_h current, since the sag affects input resistance and complicates the detection of up and down states. We recorded until access resistance increased to 60 M Ω .

Up and down states are clearly visible, although they are somewhat irregular. The procedure for finding which portions of the voltage trace correspond to the up state is, briefly, as follows: find the voltage that corresponds to the down state, label as up all portions of the voltage trace that are more than 3 mV above the down state voltage, and fill

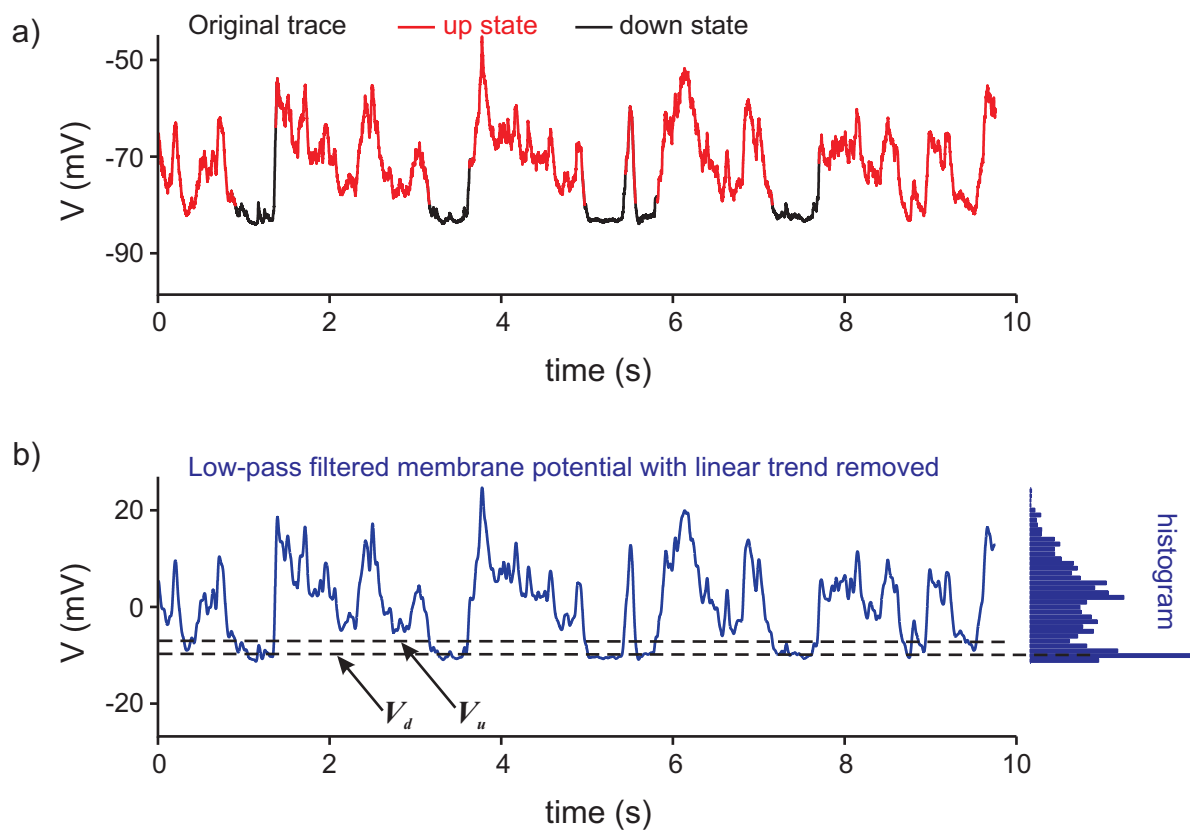


Figure S2: Detecting up and down states. **a.** A typical voltage trace from an *in vivo*, whole cell recording. Red labels up states; black labels down states. To suppress spiking, we hyperpolarized the cell by injecting negative current between -250 and -450 pA. **b.** Low pass filtered (25 Hz cutoff) membrane potential. The linear trend was subtracted to provide a flat baseline. The histogram of voltages is shown at the right. The highest peak in the histogram corresponds to the down state (V_d ; bottom dashed line). The threshold for up states, V_u , is 3 mV higher.

in any remaining short segments.

The hard part is finding the voltage that corresponds to the down state. This is done in several sub-steps. First, we low pass filter the voltage at 25 Hz to remove noise, and then subtract any linear trend in the data to give us a flat baseline. The resulting voltage trace is shown in Fig. S2b. Second, we construct a histogram of membrane potentials, binned at 1 mV (right side of Fig. S2b). Third, and last, we find the peak of the histogram and use that for the membrane potential of the down state, which we refer to as V_d (bottom dashed line in Fig. S2b). These sub-steps are carried out separately for each trace.

The next step is to set a threshold for up states, V_u , at $V_u = V_d + 3$ mV (top dashed line in Fig. S2b). All portions of the trace with voltages above V_u are candidates for the up

state. To avoid labeling brief hyperpolarizing transients as down states, any gaps between up states that are less than 150 ms are also candidates for the up state. Then, to eliminate rapid transients associated with down to up transitions, we eliminate the first 20 ms of all the up states.

The up and down states produced by this procedure are labeled red and black, respectively, in Fig. S2a. The up states were used to compute three quantities: the rate of rapidly depolarizing events (Fig. 5b, main text); the distribution of intrinsic fluctuations, $p(\sigma_V)$ (Eq. (S21) and Fig. S7); and the fraction of time the neurons spent in the down state (Fig S3e).

4. Biophysically realistic model

The experiments summarized in Fig. 3b of the main text tell us the probability of an extra spike as a function of the amplitude of a current pulse injected at the soma. What we want, however, is the average probability of a spike in response to realistic PSCs *in vivo*. Fortunately, it turns out (as we show below using a biophysically realistic neuron, and as has been shown theoretically in the limit of small PSC size^{6,7}) that the probability of an extra spike depends almost exclusively on the total charge injected into a neuron, and not the time course of that charge. Therefore, because probability is approximately linear in charge (see Fig. 3b, main text), the average probability of an extra spike *in vivo* is a function only of the average charge in a PSC *in vivo*.

This section is divided into two parts. In the first, we show that the probability of an extra spike depends almost exclusively on the total charge. In the second, we combine published data on PSP size with the electrophysiological properties of our cells to estimate the average charge in a PSC *in vivo*.

Probability depends only on injected charge

To investigate the relationship between PSCs and the probability of an extra spike, we constructed a biophysically realistic model of a neuron, including its active dendritic tree, with passive properties that matched those of the neurons we record from *in vivo*. It was based on a published model of a layer 5 pyramidal neuron,⁸ but with three changes: 1) the temperature of the model was increased from 23° C to 37° C, 2) as this speeds up action potential kinetics, all membrane conductances were increased by a factor of three, and 3) the passive reversal potential was set to -75 mV rather than -70 mV. We mimicked *in vivo*

conditions^{9,10} by bombarding the model neuron with synaptic input sufficiently large to generate voltage fluctuations on the order of 3 mV at the soma and a firing rate of 2.6 Hz.

To determine how realistic PSCs affected our simulated neuron, we computed numerically the probability of an extra spike in response to conductance changes at various points along the dendritic tree and current injection at the soma (see schematic in Fig. S3a). The probability was found by comparing two long trials (at least 10,000 s) in which the neurons received exactly the same synaptic input, but on the second trial received a conductance change every 100 ms. We then counted the number of spikes on the two trials in the 100 ms that followed the conductance change and took the difference; the probability of an extra spike was the average difference per conductance change. A typical 200 ms segment containing the two trials is shown in Fig. S3b.

Note that in our simulations we counted spikes for 100 ms after the conductance change, whereas in our experiments we counted spikes for only 5 ms (Figs. 2b and 3a, main text). That's because in our simulations the background synaptic input was identical on the two trials, so counting for long times didn't introduce noise. This can be seen in Fig. S3c, which shows the cumulative probability of an extra spike versus time since the conductance change: after a transient period lasting 30-40 ms, the cumulative probability is completely flat. This validates our procedure, and also shows that we really are seeing extra spikes, not simply shifted ones, since if the spikes were merely shifted the cumulative probability would return to baseline.

The conductance change on the dendrites was calibrated so that the total charge at the soma (computed from the time integral of the PSP; see Eq. (S1) below) took on discrete values, either ± 0.062 , ± 0.123 , or ± 0.247 pC, corresponding to current amplitudes of ± 25 , ± 50 or ± 100 pA. In Fig. S3d we plot the probability of an extra spike versus distance along the dendritic tree (with zero distance corresponding to somatic current injection) for three different somatic charges: 0.062 pC (blue dots), 0.123 pC (green dots) and 0.247 pC (red dots). In all cases the plots of probability versus distance are essentially flat – even though the shape of the somatic PSP has a strong dependence on distance (inset, to the right of the green dots).

These results suggest that what determines the probability of an extra spike is total charge at the soma. This conclusion could, of course, be specific to our particular simulated




Figure S3: The probability of an extra spike depends on total charge at the soma, and not the shape of the PSP. **a.** Schematic showing our model cell, which consists of a biophysically realistic compartmental model of a layer 5 pyramidal neuron. Charge was injected via conductance changes (except at the soma, where current was injected), and spikes were recorded at the soma. **b.** Typical traces showing two trials: one (black) on which the neuron did not receive a conductance change, another (green) on which it received a conductance change corresponding to a total charge of 0.123 pC. Inset shows a conductance change that did not cause an extra spike. **c.** Cumulative probability of an extra spike in response to a conductance change (which occurred at time $t = 0$). Blue, green and red traces correspond to charges of 0.062, 0.123 and 0.247 pC, respectively. **d.** Probability of an extra spike in response to a somatic current injection (distance = 0 μm) or conductance changes at various points along the dendritic tree. As in panel c, blue, green and red dots correspond to charges of 0.062, 0.123 and 0.247 pC, respectively. Inset: somatic PSP waveforms for conductance changes along the dendritic tree. **e.** Red and green: probability of an extra spike (for positive current changes) or a missed spike (for negative ones) versus charge, computed from our model neuron and scaled by a factor of 0.49 to take into account the fact that our model did not exhibit up and down states; see text. Red points correspond to current injected at the soma; green points to current injected at the distal dendrites, 403 μm from the soma. Black: probability measured experimentally; see Fig. 3b, main text.

neuron. However, that is unlikely, for two reasons. First, our model neuron was chosen more or less at random, from many possible models. Second, our results are consistent with theoretical analysis^{6,7} which showed rigorously that, for sufficiently small PSCs, it is total charge – and not anything else – that determines the probability of an extra spike.

Nevertheless, it is reasonable to ask that our model provides results that are close to those of our experimental observations (Fig. 3b, main text). In Fig. S3e we combine the data from our actual experiments (black points) with those from the simulations (red and green points, corresponding to injections at the soma and on the distal dendrites, 403 μm from the soma; error bars were dropped for clarity). For the latter, we corrected for up and down states by reducing the probability of an extra spike. This is because in our experiments the neurons spent 49% of their time in up states (computed from a random sample of 377 10-second traces from 5 cells), whereas in the simulations the neurons were always in the up state. Because extra spikes cannot occur in the down state, our simulations provided an overestimate of the probability of an extra spike. Thus, in Fig. S3e, for the model (red and green points) we multiply the probability of an extra spike by 0.49. When we do this, agreement between the simulations and experiments is very good.

Average charge in vivo

To determine the somatic charge that corresponds to the average PSC evoked by a single connection *in vivo*, we turn to published data. Because it is typically voltage, rather than charge, that is measured in experimental studies of synaptic connections, we need to translate between the two. The injected charge is related to the time integral of the PSP by the inverse of the membrane resistance, so we have

$$Q = \frac{1}{R} \int dt V_{\text{PSP}}(t) \quad (\text{S1})$$

where Q is the total charge at the soma, $V_{\text{PSP}}(t)$ is the time-dependent PSP at the soma, and R is the membrane resistance. To estimate Q , we need two quantities: the membrane resistance and the time integral of the PSPs. For the membrane resistance, we use the input resistance of our recorded cells, which was $42 \pm 4.3 \text{ M}\Omega$ ($n=28$), consistent with published values.^{9,11,12} For the time integral of the PSPs, we use published values of the connection strength between layer 5 pyramidal cells measured *in vitro*. These range from 0.85 to 1.3

mV (including failures).^{13–16} To estimate the PSP time course, we approximate it by a difference of exponentials. For the decay time, we use our measurements of the membrane time constant *in vivo* (8 ms). For the rise time, we use the fact that the membrane potential cannot rise faster than the synaptic conductance decays. Since the latter is 1.7 ms,¹⁷ we use that for the rise time. The integral in Eq. (S1) thus ranges from 10.4–15.8 mV-ms, or 13.1 ± 2.7 mV-ms (see Eq. (S38)). Dividing this by the effective input resistance, 42 ± 4.3 M Ω , and using log normal statistics, we find that the average charge per connection is 0.31 ± 0.07 pC. Note, therefore, that 0.31 pC corresponds to about 1 mV.

5. The effect of feedforward inhibition

In our analysis of the effects of missed and extra spikes, we made the simplifying assumption that time is discretized into finite steps, and there is no interaction within a time step (see Fig. 1a of the main text). Far more realistic, however, is asynchronous production of missed and extra spikes. This would allow, at least in principle, a scenario that is prohibited by the no-interaction assumption (see Fig. S4a): an extra spike occurs on an excitatory neuron, the first extra spike it produces is on an inhibitory neuron, that inhibitory neuron sends extra spikes to the postsynaptic targets of the original neuron, and the extra spikes that would have occurred are canceled (green arrows in Fig. S4a, and compare to Fig. 1a of the main text). If this, and this alone, were to happen, perturbations would die out very quickly.

In fact, this kind of feedforward inhibition¹⁸ is likely to reduce the number of extra spikes on excitatory neurons. However, that's not its only effect; it also leads to missed spikes, by canceling spikes that were not extra (i.e., by canceling spikes that would have existed whether or not there was an extra presynaptic spike; see black arrows in Fig. S4a). If every canceled spike is balanced by at least one missed spike, then perturbations grow just as rapidly as they would have in the absence of feedforward inhibition. That's because, by symmetry, a missed spike on one trial corresponds exactly to an extra spike on another.

To determine quantitatively the relationship between canceled and missed spikes, consider what happens in response to a single extra spike on an excitatory neuron. That extra spike delivers charge to each of its postsynaptic targets; say average charge q_i^E to excitatory neuron i . Charge is also delivered to inhibitory neurons, which causes them to produce extra spikes. Those extra spikes in turn deliver negative charge to their postsynaptic targets; say average




Figure S4: Feedforward inhibition. **a.** Same setup as in Fig. 1a of the main text, except that the first extra postsynaptic spike is on an inhibitory neuron (red arrow at the bottom of the panel). The extra inhibitory spike then cancels the two extra excitatory spikes that would have occurred without the inhibition (green arrows, and see Fig. 1a of the main text). The extra inhibitory spike also causes missed spikes (black arrows). **b.** Overlapping and non-overlapping neurons. Neurons that receive disynaptic input via an inhibitory neuron and a direct connection from the reference neuron are considered to be overlapping; those that receive disynaptic input via an inhibitory neuron and *no* direct connection from the reference neuron are considered to be non-overlapping. Note that whether a neuron is overlapping or non-overlapping depends on the reference neuron.

charge $-q_i^I$ to excitatory neuron i .

The negative charge has two effects. One is to reduce the probability of an extra spike on “overlapping” neurons – on neurons that receive both positive charge from a presynaptic neuron and, indirectly, negative charge from the same presynaptic neuron, via a disynaptic connection involving an inhibitory neuron (see Fig. S4b). The reduction is proportional to

the total negative charge from the inhibitory neuron times the slope of the positive portion of the regression line in Fig. 3b of the main text. Using S_+ to denote that slope, the number of canceled spikes, denoted $n_{canceled}$, is given by

$$n_{canceled} = S_+ \sum_i \Theta(q_i^E) \min(q_i^I, q_i^E),$$

where Θ is the Heaviside step function: $\Theta(q) = 1$ if $q > 0$ and 0 otherwise. The step function picks out the postsynaptic targets of the original extra spike, and the “min” operation (which returns the minimum of its two arguments) enforces the fact that negative charge larger than q_i^E does not contribute to canceling extra spikes.

The second effect of negative charge is to cause missed spikes, which can happen in two ways: via negative charge on non-overlapping neurons (neurons that receive a disynaptic connection involving an inhibitory neuron but no direct connection; see Fig. S4b), and *excess* negative charge on overlapping neurons. Using S_- to denote the absolute value of the slope of the negative portion of the regression line in Fig. 3b of the main text, and n_{missed} to denote the number of missed spikes, we have

$$n_{missed} = S_- \sum_i [q_i^I - q_i^E]^+,$$

where $[\cdot]^+$ is the threshold linear function: $[x]^+ = \max(x, 0)$. (Neurons for which $q_i^E = 0$ are non-overlapping; neurons for which $q_i^E > 0$ are overlapping.)

We can now compute our quantity of interest, the ratio of missed to extra spikes averaged over trials and pre and postsynaptic neurons. That ratio is given by

$$\frac{n_{missed}}{n_{canceled}} = \frac{S_-}{S_+} \frac{Q_{nonoverlap}^I}{Q_{overlap}^I},$$

where $Q_{nonoverlap}^I$ and $Q_{overlap}^I$ are the average charges delivered to non-overlapping and overlapping neurons, respectively; using angle brackets to indicate an average over which presynaptic excitatory neuron produces the original extra spike and over the amount of charge associated with each extra spike on both excitatory and inhibitory neurons, these quantities are given by

$$Q_{nonoverlap}^I = \left\langle \sum_i [q_i^I - q_i^E]^+ \right\rangle$$

$$Q_{overlap}^I = \left\langle \sum_i \Theta(q_i^E) \min(q_i^I, q_i^E) \right\rangle.$$

The slopes S_- and S_+ were determined experimentally to be 0.018 and 0.061, respectively (see Fig. 3b caption, main text), so their ratio is 1/3.4. Thus, to determine how many missed spikes there are for every canceled one, all we need to know is how much charge is delivered to non-overlapping neurons compared to overlapping neurons. If the ratio of these two charges is greater than 3.4, the number of missed spikes will exceed the number of canceled ones.

To determine the charge ratio, we assume that the charge delivered to a postsynaptic neuron does not depend on whether that neuron is overlapping or non-overlapping. This assumption sounds strong, but we should keep in mind that which neurons are overlapping and which are non-overlapping depends on which neurons produce extra spikes, and so changes from one trial to the next. Thus, while we don't necessarily expect independence to hold on individual trials, we expect it to hold on average.

With the independence assumption, determining the ratio of the charge delivered to non-overlapping versus overlapping neurons becomes simply a matter of counting the number of each, and taking their ratio. These two numbers depend primarily on the projection patterns of the excitatory neurons. For example, if connectivity is all to all, then *all* neurons are overlapping, and inhibitory neurons could indeed cancel all extra spikes. However, connectivity in the brain is not all to all, but sparse: locally, the connection probability of excitatory neurons is 4-10%.^{2,14,16} With this degree of sparseness, if inhibitory neurons made random connections, the ratio of non-overlapping to overlapping neurons would be at least 9:1 (because the probability that an inhibitory neuron would project to the postsynaptic target of a particular excitatory neuron would be 1/10). Inhibitory neurons are not, of course, likely to make random projections. However, as shown by Holmgren and colleagues,² the connectivity patterns of excitatory and inhibitory neurons in barrel cortex are very different, with inhibitory neurons making very local connectivity and excitatory neuron making (relatively) long range connectivity. This rules out the possibility that inhibitory neurons make most

of their connections to overlapping neurons, because that would require, at a minimum, each inhibitory neuron to have approximately the same projections pattern as at least one excitatory neuron (see Fig. S4b).

With a 9:1 ratio of non-overlapping to overlapping neurons, for every spike canceled by feedforward inhibition, there are approximately 2.6 ($=9/3.4$) missed spikes. Consequently, feedforward inhibition can only *increase* the rate at which perturbations grow. Because we do not know if feedforward inhibition really is fast enough to have much effect, we do not include it in our analysis, and instead use our naive (and more conservative) estimate of 28 extra spikes per spike.

6. The level of intrinsic noise at equilibrium

In the main text we showed that small perturbations to rat barrel cortex are amplified, and we estimated, using very qualitative analysis, the level at which they saturate. Here we first expand on that qualitative analysis, then we provide a much more rigorous one.

Our main tool, both here and in the analysis that follows, is self-consistency between the probability of a missed or extra spike arriving at a synapse (the presynaptic probability, denoted p_{pre}) and the probability of a missed or extra spike leaving a neuron (the postsynaptic probability, denoted p_{post}). As mentioned in the main text, if we discretize time, the presynaptic probability is easy to compute: if a neuron receives m missed and extra spikes in a time window that corresponds to a typical neuron's mean integration time, and it has K presynaptic connections, then $p_{pre} = m/K$. The postsynaptic probability is slightly harder, as two steps are required. First, if the m missed and extra spikes received by a neuron are reasonably uncorrelated, then they produce voltage fluctuations whose standard deviation, denoted σ_V , is proportional to $m^{1/2}\bar{V}_{PSP}$ where \bar{V}_{PSP} is the PSP size averaged over excitatory and inhibitory unitary postsynaptic potentials (not necessarily with equal weighting), and the factor of $m^{1/2}$ comes from central limit type arguments. Second, because of the approximate linearity between injected charge and probability (Fig. 3b, main text), the factor of $m^{1/2}$ in voltage fluctuations translates into a factor of $m^{1/2}$ in the probability of missed and extra spikes. Thus, $p_{post} \propto m^{1/2}\bar{p}_1$ where \bar{p}_1 is a weighted average of the probability of missed or extra postsynaptic spikes given a single missed or extra presynaptic spike. We expect \bar{p}_1 to be on the same order as \bar{p}_e (where, recall from the main text, \bar{p}_e is

the probability of an extra spike), so to first approximation we may write $p_{post} \propto m^{1/2} p_e$. Self-consistency then leads to Eq. (1) of the main text.

While this approach leads to qualitatively correct estimates for the level of intrinsic fluctuations, it suffers from two problems. The first is that we can show only that σ_V is proportional to $K\bar{p}_e\bar{V}_{PSP}$ (Eq. (2), main text), but we cannot determine the constant of proportionality. The second is that the analysis relies on a particular discretization time, which is essentially arbitrary. Thus, in this section we take a much more rigorous approach and compute quantitatively the level of membrane potential fluctuations in steady state.

The idea is as follows. Consider a recurrent network that receives identical input on multiple trials, but on each trial has slightly different initial conditions. The different initial conditions produce both different patterns of spiking activity and different time courses for the membrane potential at each neuron. Those differences are, of course, related, and the relation could be calculated two ways: given differences in spiking activity we could compute the differences in membrane potential (Fig. S5a \rightarrow S5b), and given differences in membrane potential we could compute differences in spiking activity (Fig. S5b \rightarrow S4c). Importantly, both computations are tractable (with a few simplifying assumptions). Our strategy, then, is to assume a particular difference in spiking activity (Fig. S5a), use that to compute differences in membrane potential (Fig. S5b), and use those differences to recompute differences in spike trains (Fig. S5c). At equilibrium, the two must be the same, a fact we use to compute the equilibrium rate of missed and extra spikes, and from that the equilibrium difference in membrane potential.

In principle, carrying out this procedure is straightforward. In practice, however, there are technical difficulties associated with computing differences in membrane potential and spiking activity across two trials, as “difference” is a multi-dimensional quantity. We thus make two approximations, both of which involve simplified measures of difference. For membrane potential, we use the variance across trials. For spiking activity, we use a variant of the Victor and Purpura spike train metric¹⁹ to both define a distance between spike trains and to rearrange the spike trains slightly. Loosely, if spikes on two trials are separated by a sufficiently short time, they are rearranged so that they occur at *exactly* the same time (spikes connected by gray lines in Figs. S5a and c); otherwise, they are considered to be missed or extra spikes (arrows in Figs. S5a and c). Because this procedure underestimates

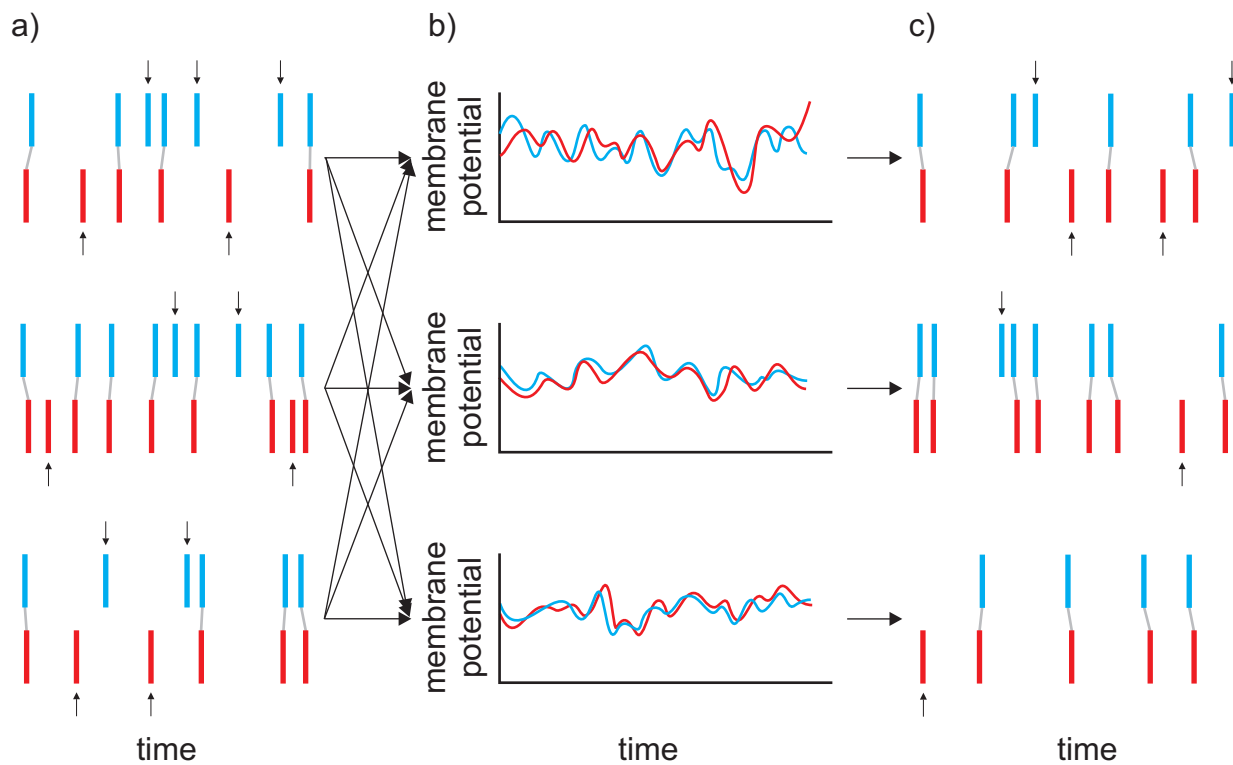


Figure S5: Computing membrane fluctuations self-consistently. Two trials of the same network, with the trials differing only by their initial conditions. Cyan refers to trial 1; red to trial 2. **a.** Activity patterns for three neurons. Vertical arrows indicated missed and extras spikes; gray lines link spikes that are considered to occur at identical times (see text). **b.** Synaptic drive for the same three neurons, with differences across trials computed from the spike trains in panel a. The arrows linking panels a and b indicate that each neuron receives input from many others. **c.** Activity patterns, with differences across trial coming from differences in membrane potential (indicated by the arrows linking panels b and c). Self-consistency demands that the degree of difference across the two trials is the same in panels a and c. Again, vertical arrows indicated missed and extra spikes and gray lines link spikes that are considered to occur at identical times.

the differences in the spiking activity (it explicitly makes spike trains more similar), it underestimates the trial-to-trial fluctuations. Consequently, we do not have to be especially careful in how we define “sufficiently short time”.

With these simplified distance measures, the self-consistency calculation becomes straightforward: First, we assume that missed and extra spikes occur at a particular rate, and compute the trial-to-trial variance in the membrane potential. Second, given that variance, we compute the rate at which missed and extra spikes occur. Self-consistency is enforced by equating the rates in the first and second steps.

Variability in membrane potential given missed and extra spikes

Consider two trials that differ only in their initial conditions. For simplicity, here we consider deterministic networks; below we generalize to networks that exhibit failures. Expressing the membrane potential on both trials as a sum of PSPs, and using a tilde to label the second trial, the difference in synaptic drive between the two trials, measured in voltage, is given by

$$\delta V(t) = \sum_{jk} v_{jk}(t - t_{jk}) - \sum_{jk} \tilde{v}_{jk}(t - \tilde{t}_{jk}). \quad (\text{S2})$$

Here j labels neuron and k labels spikes, so $v_{jk}(t - t_{jk})$ is the PSP produced by the k^{th} spike on neuron j , with t_{jk} denoting the time of the k^{th} spike. The subscript k on v_{jk} allows for trial-to-trial variability in PSP size, which arises both from variations in transmitter release and from the fact that a neuron's response to a presynaptic spike depends on its state. Note that there will, in general, be a different number of spikes on the second trial, so the first and second sums over k may have different numbers of elements. The sum over j runs over all neurons connected to the postsynaptic neuron of interest (which should also have a label, but we have dropped it for clarity). Finally, the v_{jk} can be either positive or negative, corresponding, respectively, to EPSPs and IPSPs.

The major approximation we make in writing Eq. (S2) is that the membrane potential adds linearly in response to successive spikes. This is undoubtedly violated at sufficiently high firing rate, but otherwise it is a reasonable approximation. Besides that, the model is very general; our formulation does not preclude v_{jk} that are sufficiently complicated to describe realistic neurons.

The variance of δV , averaged over a time T (which we ultimately take to infinity), is

$$\langle \delta V^2 \rangle = \frac{2}{T} \int dt \left[\sum_{jk, j'k'} v_{jk}(t - t_{jk}) v_{j'k'}(t - t_{j'k'}) - \sum_{jk, j'k'} v_{jk}(t - t_{jk}) \tilde{v}_{j'k'}(t - \tilde{t}_{j'k'}) \right]$$

where the integral over t runs over a range equal to T (e.g., from $-T/2$ to $T/2$). To derive this equation, we used the fact that the first and second trials have the same variance; that led to the factor of 2.

The first step is to separate the variability due to spike times from the variability in PSP size, which we do by introducing a second time-integral,

$$\langle \delta V^2 \rangle = \frac{2}{T} \int d\tau \left[\sum_{jk, j'k'} \int dt v_{jk}(t-\tau) v_{j'k'}(t) \delta(\tau - (t_{jk} - t_{j'k'})) \right. \\ \left. - \sum_{jk, j'k'} \int dt v_{jk}(t-\tau) \tilde{v}_{j'k'}(t) \delta(\tau - (t_{jk} - \tilde{t}_{j'k'})) \right]. \quad (\text{S3})$$

This expression is valid in the limit that $T \rightarrow \infty$, which is needed to shift the limits of integration on t .

The first sum in Eq. (S3) contains two kinds of spikes: those that occur on the same neuron and at the same time ($j = j'$ and $k = k'$), and all other spikes. Explicitly separating the two contributions to that term, we may write

$$\sum_{jk, j'k'} \int dt v_{jk}(t-\tau) v_{j'k'}(t) \delta(\tau - (t_{jk} - t_{j'k'})) = \quad (\text{S4}) \\ \sum_{jk} \int dt v_{jk}(t-\tau) v_{jk}(t) \delta(\tau) + \sum_{jk \neq j'k'} \int dt v_{jk}(t-\tau) v_{j'k'}(t) \delta(\tau - (t_{jk} - t_{j'k'}))$$

where the notation $jk \neq j'k'$ is shorthand for $j \neq j'$ or $k \neq k'$. The first sum on the right hand side reduces to the average sum of squares of PSPs,

$$\sum_{jk} \int dt v_{jk}(t-\tau) v_{jk}(t) \delta(\tau) = \left[K_E \nu_E T \int dt \overline{v_E^2}(t) + K_I \nu_I T \int dt \overline{v_I^2}(t) \right] \delta(\tau). \quad (\text{S5})$$

Here E and I refer to excitatory and inhibitory neurons, respectively, K is the average number of connections per neuron, ν is average firing rate, and the overline refers to averages over PSPs (i.e., an average over the index j in v_{jk}). The factors of K_E and K_I come from the sum on j ; the factors of $\nu_E T$ and $\nu_I T$ come from the sum on k .

For the second sum in Eq. (S4), we proceed in two steps. First we sum on k' . This converts the δ -function into a cross-correlogram. It is convenient to define cross-correlogram between neurons j and j' (or the auto-correlogram if $j = j'$), denoted $C_{jj'}(\tau)$, as

$$C_{jj'}(\tau) \equiv \nu_{j'}^{-1} \sum_{k'} \delta(\tau - (t_{jk} - t_{j'k'})),$$

where $\nu_{j'}$ is the firing rate of neuron j' . Note that we have chosen a slightly non-standard normalization in which $C_{jj'}(\tau)$ approaches 1 when $|\tau|$ is very large. Assuming that PSP size is independent of spike times, we have

$$\sum_{jk \neq j'k'} \int dt v_{jk}(t - \tau) v_{j'k'}(t) \delta(\tau - (t_{jk} - t_{j'k'})) = \sum_{j,j',k} \int dt v_{jk}(t - \tau) \bar{v}_{j'}(t) \nu_{j'} C_{jj'}(\tau).$$

The second step is to sum over k , which brings out a factor of $\nu_j T$, and we have

$$\sum_{jk \neq j'k'} \int dt v_{jk}(t - \tau) v_{j'k'}(t) \delta(\tau - (t_{jk} - t_{j'k'})) = \sum_{j,j'} \int dt \bar{v}_j(t - \tau) \bar{v}_{j'}(t) T \nu_j \nu_{j'} C_{jj'}(\tau). \quad (\text{S6})$$

Equations (S5) and (S6) take care of the first term in Eq. (S3). They can't, however, be directly used for the second, since for that term, in general there are no equal-time spikes (the PSPs occur on different trials). To get around this problem, we invoke the Victor and Purpura method,¹⁹ as discussed above, and take any spikes from the same neuron that differ by less than, say, 2.5 ms (half the integration time of extra spikes observed in our experiments; see Figs. 2b and 3a of the main text), and shift them so that they occur at equal times (Fig. S6a). Those spikes occur at a rate equal to the difference between the average firing rate and the rate of extra spikes, the latter denoted $\nu_{E,ex}$ for extra excitatory spikes and $\nu_{I,ex}$ for extra inhibitory ones. Noting also that the cross-correlograms for spikes on different trials are different from those for spikes on the same trial, we replace $C_{jj'}$ with $\tilde{C}_{jj'}$, and we have

$$\begin{aligned} \sum_{jk, j'k'} \int dt v_{jk}(t - \tau) \tilde{v}_{j'k'}(t) \delta(\tau - (t_{jk} - \tilde{t}_{j'k'})) = & \quad (\text{S7}) \\ & \left[K_E(\nu_E - \nu_{E,ex})T \int dt \bar{v}_E^2(t) + K_I(\nu_I - \nu_{I,ex})T \int dt \bar{v}_I^2(t) \right] \delta(\tau) \\ & + \sum_{j,j'} \int dt \bar{v}_j(t - \tau) \bar{\tilde{v}}_{j'}(t) T \nu_j \nu_{j'} \tilde{C}_{jj'}(\tau). \end{aligned}$$

Inserting Eqs. (S5-S7) into Eq. (S3), and using the fact that, on average, $\bar{\tilde{v}}_{j'}(t) = \bar{v}_j(t)$, we arrive at




Figure S6: Missed and extra spikes. **a.** Spikes come in two categories: those that are the same on two trials (or, more accurately, those that are sufficiently close on two trials; see Figs. S5a and c), and those that are extra. The former are black; the latter, which occur at rate $\nu_{E,ex,i}$ or $\nu_{I,ex,i}$ for neuron i (depending on whether the neuron is excitatory or inhibitory), are magenta. **b.** A single trial showing only the extra spikes. Neuron 1 is excitatory; neuron 2 is inhibitory. To the extent that the extra excitatory and inhibitory spikes line up perfectly (as shown for the three inhibitory spikes), excitation and inhibition can exactly cancel, something that reduces voltage fluctuations. However, because there are about four times fewer inhibitory neurons compared to excitatory ones, and thus about four times fewer spikes on inhibitory neurons, this cancellation can occur, at most, on only 1/4 of the excitatory spikes. The remaining 3/4 contribute to voltage fluctuations.

$$\begin{aligned} \langle \delta V^2 \rangle &= 2K_E \nu_{E,ex} \int dt \overline{v_E^2}(t) + 2K_I \nu_{I,ex} \int dt \overline{v_I^2}(t) \\ &+ 2 \sum_{j,j'} \int dt d\tau \nu_j \nu_{j'} \overline{v_j}(t - \tau) \overline{v_{j'}}(t) \left[C_{jj'}(\tau) - \tilde{C}_{jj'}(\tau) \right]. \end{aligned} \quad (\text{S8})$$

The last term on the right hand side of Eq. (S8) arises due to correlations among spikes, since if there were no correlations both $C_{jj'}$ and $\tilde{C}_{jj'}$ would be one. It is reasonable to assume that correlations are larger when considering spikes on the same trial versus spikes on different trials, which implies that $C_{jj'} > \tilde{C}_{jj'}$. This leaves us with a negative contribution from correlated excitatory and inhibitory spikes, since PSPs from neurons of different types have opposite signs ($v_j v_{j'} < 0$ if neuron j is excitatory and neuron j' is inhibitory or vice versa). This makes sense intuitively: if extra excitatory spikes tend to be accompanied by extra inhibitory spikes, that would reduce the voltage fluctuations.

Because of the negative contribution from correlations among excitatory and inhibitory

spikes, the last term in Eq. (S8) could, in principle, be sufficiently negative to cancel the first two. To what extent can this happen? To answer that, we return to Eq. (S2), and interpret it as a sum over the membrane potential associated with spikes that are sufficiently different on the two trials (see Fig. S5). We then consider the worst case scenario: the connectivity of inhibitory neurons perfectly matches that of excitatory neurons, and whenever there is an extra spike on an excitatory neuron there is an accompanying extra spike on an inhibitory neuron. Indeed, if there were as many inhibitory neurons as excitatory neurons, this scenario would eliminate the voltage fluctuations altogether. However, there are a factor of four more excitatory neurons than inhibitory ones, so such a scenario could never lead to complete cancellation. In fact, as can be seen in Fig. S6b, at worst correlations among excitatory and inhibitory spikes would eliminate the second term in Eq. (S8) (all inhibitory spikes canceled by excitatory ones) but it could reduce the first term by only 25% (25% of the excitatory spikes canceled by inhibitory ones). A conservative estimate of the voltage fluctuations would, therefore, consist of only the first term in Eq. (S8), reduced by an overall factor of 3/4. This, though, is extremely conservative, especially since the connectivity patterns of excitatory and inhibitory neurons in barrel cortex are very different.² We thus make a slightly less conservative assumption: inhibitory spikes cancel half the extra excitatory ones. With this assumption (which is still conservative) the first term in Eq. (S8) is reduced by 12.5% and the second cut in half, leaving the remaining half to restore the missing 12.5%. This effectively gives us the first term in Eq. (S8), leading to a very simple expression for the voltage fluctuation,

$$\langle \delta V^2 \rangle = 2K_E \nu_{E,ex} \int dt \overline{v_E^2}(t). \quad (\text{S9})$$

This expression represents a lower bound on the trial-trial variability in synaptic drive in a deterministic network receiving identical input on multiple trials, but differing in initial conditions.

The probability of missed and extra spikes given the variability in the synaptic drive

The next step is to compute the probability of a missed or extra spike. Here we use the results of our simulations, which show that the probability of missed or extra spikes is proportional to the total charge delivered to the soma (see Sec. 4 and Fig. 3c of the main text). Thus,

the mean firing rate of extra spikes at time t , computed only for excitatory neurons, may be written

$$\nu_{E,ex,raw}(t) = \frac{\eta}{R} [\delta V(t)]^+, \quad (\text{S10})$$

where η is the slope of the positive regression line in Fig. 3b of the main text, in units of probability/charge, and R is the average input resistance of the cells in our dataset. The threshold linear operator is needed because $\nu_{E,ex,raw}$ is the number of *extra* spikes, so we don't want to include missed ones. The extra subscript "*raw*" refers to the fact that Eq. (S10) gives us the raw rate of extra spikes; below we will see that it is reduced by an effective refractory period.

To find the mean value of the firing rate of extra spikes, we average Eq. (S10) over time, and then turn that into an average over the distribution of δV (for which we use angle brackets). This leads to

$$\nu_{E,ex,raw} = \frac{\eta}{R} \frac{1}{T} \int dt [\delta V]^+ = \frac{\eta}{R} \langle [\delta V]^+ \rangle .$$

As discussed above, we assume that δV is zero mean and Gaussian, for which the average is straightforward,

$$\nu_{E,ex,raw} = \frac{\eta}{R} \frac{\langle \delta V^2 \rangle^{1/2}}{(2\pi)^{1/2}} . \quad (\text{S11})$$

The rate of extra spikes in Eq. (S11) is the raw rate, and ignores the fact that after one extra spike the probability of a second one drops. This leads to an effective refractory period, and thus reduces the actual rate of extra spikes relative to $\nu_{E,ex,raw}$. We don't know the precise form of this refractory period, but, fortunately, our final result will be insensitive to the form we use. That's because the refractory period only becomes important in the very high noise level, where precise timing has already essentially been eliminated. We thus use an average refractory period of $\tau_r \equiv 1/\nu_E$, where, as above ν_E is the network averaged firing rate of excitatory neurons. Consequently, the average interval between extra spikes (which is the inverse of the actual rate of extra spikes, denoted $\nu_{E,ex}$) is increased by τ_r ; i.e.,

$1/\nu_{E,ex} = \tau_r + 1/\nu_{E,ex,raw}$. The actual rate of extra spikes is, then, given in terms of the raw rate as

$$\nu_{E,ex} = \frac{\nu_{E,ex,raw}}{1 + \tau_r \nu_{E,ex,raw}} = \frac{(\eta/R)\langle\delta V^2\rangle^{1/2}/(2\pi)^{1/2}}{1 + \tau_r(\eta/R)\langle\delta V^2\rangle^{1/2}/(2\pi)^{1/2}}. \quad (\text{S12})$$

Self-consistency

We now have two expressions for the variance of the membrane potential: Eqs. (S9) and (S12). These can be combined to eliminate $\nu_{E,ex}$, giving us a single equation for $\langle\delta V^2\rangle$. Note, however, that δV is the voltage *difference*; what we are interested in is the single trial variance, which is half of $\langle\delta V^2\rangle$. We thus define

$$\sigma_V^2 \equiv \frac{\langle\delta V^2\rangle}{2}. \quad (\text{S13})$$

Then, combining Eqs. (S9) and (S12), we have

$$\sigma_V^2 = \frac{\sigma_V \eta/R\pi^{1/2}}{1 + \tau_r \sigma_V \eta/R\pi^{1/2}} K_E \int dt \overline{v_E^2(t)}. \quad (\text{S14})$$

It is convenient to define

$$\Delta V_{\max}^2 \equiv \frac{K_E}{\tau_r} \int dt \overline{v_E^2(t)},$$

which is the maximum value of the right hand side of Eq. (S14) with respect to σ_V . Note that ΔV_{\max}^2 is the maximum variance in membrane potential, a quantity we can estimate from data (see Sec. 3, ‘‘Detecting up and down states’’, and Fig. S7). With this definition, the expression for σ_V^2 simplifies to

$$\sigma_V^2 = \Delta V_{\max}^2 \frac{\sigma_V \tau_r \eta/R\pi^{1/2}}{1 + \sigma_V \tau_r \eta/R\pi^{1/2}}. \quad (\text{S15})$$

Solving this for σ_V yields

$$\sigma_V = \Delta V_{\max} \frac{[1 + 4\xi^2]^{1/2} - 1}{2\xi} \quad (\text{S16})$$

where

$$\xi \equiv \frac{1}{\Delta V_{\max}} \frac{\eta}{R\pi^{1/2}} K_E \int dt \overline{v_E^2}(t). \quad (\text{S17})$$

This expression produces sensible results in the two natural limiting cases: When ξ is small, $\sigma_V \rightarrow \Delta V_{\max}\xi$, which is independent of ΔV_{\max} (see Eq. (S17)). This is as it should be: at small σ_V , the refractory period plays no role, and σ_V is given by its raw value. When ξ is large, on the other hand, the refractory period becomes important, and the voltage fluctuations saturate at ΔV_{\max} .

To compute ξ (which tells us, via Eq. (S16), σ_V), we need to determine the quantities on the right hand side of Eq. (S17). We do that by converting them into experimentally accessible quantities. We start with the last term, the time integral of the square of the membrane potential. For that we assume that the PSPs are described by a difference of exponentials with rise time τ_{rise} , decay time τ_{decay} , and amplitude V_{EPSP} (which is a random variable). The integral on the right hand side of Eq. (S17) is then straightforward (see Eq. (S39)), and we find that

$$\int dt \overline{v_E^2}(t) = \frac{\tau_{\text{decay}} \exp[2(\kappa \log \kappa)/(\kappa - 1)]}{2(1 + \kappa)} \overline{V_{EPSP}^2} \quad (\text{S18})$$

where $\kappa \equiv \tau_{\text{rise}}/\tau_{\text{decay}}$ is the ratio of the rise time to the decay time.

To compute $\overline{V_{EPSP}^2}$, we follow Song et al.,¹⁴ who found that EPSP amplitudes are well approximated by a log normal distribution with variance 0.94. For such a distribution, the ratio of the second moment to the square of the first obeys a simple relationship (see Eq. (S36b) and rearrange terms slightly),

$$\overline{V_{EPSP}^2} = e^{0.94} \overline{V_{EPSP}}^2. \quad (\text{S19})$$

Inserting Eqs. (S18) and (S19) into (S17), we arrive at

$$\xi = \frac{e^{0.94}}{\pi^{1/2}} \frac{\exp[2(\kappa \log \kappa)/(\kappa - 1)]}{2(1 + \kappa)} \frac{\tau_{\text{decay}}}{R\Delta V_{\text{max}}} \eta K_E \bar{V}_{EPSP}^2.$$

The variables η , K_E , and \bar{V}_{EPSP} all have uncertainty associated with them. Thus, ξ has both a mean and variance. These are computed in Sec. 8 (see Eq. (S34b)), and we find

$$\xi = \frac{21.42 \pm 12.95}{\Delta V_{\text{max}}}.$$

We are assuming that all variables follow a log normal distribution, so the relevant variables are the mean and variance of $\log \xi$; these are given by (see Eq. (S35))

$$\text{mean}[\log \xi] \equiv \mu_\xi = 3.04 - \log \Delta V_{\text{max}} \quad (\text{S20a})$$

$$\text{Var}[\log \xi] \equiv \sigma_\xi^2 = 0.31. \quad (\text{S20b})$$

Given the distribution of ξ , we can now compute the distribution of σ_V ; this is given by

$$p(\sigma_V) = \frac{1}{\Delta V_{\text{max}}} \frac{\xi(\sigma_V) \exp[-(\log[\xi(\sigma_V)] - \mu_\xi)^2/2\sigma_\xi^2]}{(2\pi\sigma_\xi^2)^{1/2}} \frac{\sigma_V^2 + \Delta V_{\text{max}}^2}{\sigma_V^2} \quad (\text{S21})$$

where $\xi(\sigma_V)$ is given in terms of σ_V , via Eq. (S16), as

$$\xi(\sigma_V) = \frac{\sigma_V/\Delta V_{\text{max}}}{1 - \sigma_V^2/\Delta V_{\text{max}}^2}.$$

The distribution in Eq. (S21) depends on ΔV_{max} , which varies from cell to cell. We computed ΔV_{max} for five cells, and $p(\sigma_V)$ is plotted for each of them in Fig. S5.

The distribution $p(\sigma_V)$ is important because it is used to compute the rate at which precisely timed events occur (Fig. 5b of the main text). This is done by measuring event rates at a particular σ_V and $\delta\tau$, as described in Fig. 5a of the main text, and then averaging that over $p(\sigma_V)$. Specifically, if we use $\nu_e(\sigma_V, \delta\tau)$ to denote the rate of event that rise by at least $2\sigma_V$ mV in $\delta\tau$ ms, then the average rate of events with precision $\delta\tau$ is $\int d\sigma_V p(\sigma_V) \nu_e(\sigma_V, \delta\tau)$. Note that both $\nu_e(\sigma_V, \delta\tau)$ and $p(\sigma_V)$ vary from cell to cell (the latter because ΔV_{max} varies from cell to cell). The error bars in Fig. 5b of the main text were computed assuming Poisson




Figure S7: Distribution of σ_V , computed from Eq. (S21), for nine cells. Dashed and solid lines are the distributions with and without whisker stimulation, respectively; each had a different value of ΔV_{\max} , as shown. For the black traces, whisker stimulation was not applied. For the green, red and cyan traces, curves that are the same color came from the same cell.

statistics, so that the variance in any measurement is $(\text{event count})/T^2$ where T is the total time used to count events.

Failures

In the above analysis, we assumed that all action potentials produce PSPs. This, however, is unrealistic: a propagating action potential may not arrive at a synapse, or, if it does, it may fail to cause release of neurotransmitter. In this section, we ask how failures affect the level of trial-to-trial voltage fluctuations. The answer, not surprisingly, is that failures increase them.

Failures clearly have no effect on the rate of extra spikes given voltage fluctuations; i.e., they do not alter Eq. (S12). They do, however, affect how voltage fluctuations depend on the rate of extra spikes. To see how, note that the factor $\nu_{E,ex}$ that appears in Eq. (S9) is the difference of two quantities: ν_E and $(\nu_E - \nu_{E,ex})$, as can be seen by examining Eqs. (S5) and (S7), respectively. The first quantity, ν_E , is the rate of spikes, as observed postsynaptically. Assuming, for simplicity, that failures occur with probability p_f , then this term is reduced by a factor of $1 - p_f$. The second quantity, $(\nu_E - \nu_{E,ex})$, is the rate of coincident postsynaptic spikes on two different trials. Since the trials are independent, at least with respect to

failures, this term is reduced by a factor of $(1 - p_f)^2$. Putting these together, we see that the factor of $\nu_{E,ex}$ that appears in Eq. (S9) is modified according to

$$\nu_{E,ex} \rightarrow (1 - p_f)\nu_E - (1 - p_f)^2(\nu_E - \nu_{E,ex}) = p_f(1 - p_f)\nu_E + (1 - p_f)^2\nu_{E,ex}.$$

Consequently, Eq. (S9) becomes

$$\sigma_V^2 = [p_f(1 - p_f)\nu_E + (1 - p_f)^2\nu_{E,ex}]K_E \int dt \overline{v_E^2}(t), \quad (\text{S22})$$

where we used Eq. (S13) to replace $\langle \delta V^2 \rangle / 2$ with σ_V^2 .

We now proceed in three steps: first, we use Eq. (S22) to express $\nu_{E,ex}$ in terms of σ_V^2 ; second, we insert that into Eq. (S12); and third, we solve the resulting equation for σ_V . The first step yields

$$\nu_{E,ex} = \frac{\sigma_V^2}{(1 - p_f)^2 K_E \int dt \overline{v_E^2}(t)} - \frac{p_f \nu_E}{1 - p_f}. \quad (\text{S23})$$

The next step is to insert Eq. (S23) into (S12), with $\langle \delta V^2 \rangle / 2$ replaced by σ_V^2 , which gives us

$$\sigma_V^2 = \left[(1 - p_f)^2 \frac{\sigma_V \eta / R \pi^{1/2}}{1 + \sigma_V \tau_r \eta / R \pi^{1/2}} + p_f(1 - p_f)\nu_E \right] K_E \int dt \overline{v_E^2}(t). \quad (\text{S24})$$

In the limit $p_f \rightarrow 0$, this expression reduces to Eq. (S14), as it should.

As in the previous analysis, we define ΔV_{\max}^2 to be the maximum value of the right hand side of Eq. (S24) with respect to σ_V ; this gives us

$$\Delta V_{\max}^2 = \left[\frac{(1 - p_f)^2}{\tau_r} + p_f(1 - p_f)\nu_E \right] K_E \int dt \overline{v_E^2}(t) = \frac{1 - p_f}{\tau_r} K_E \int dt \overline{v_E^2}(t) \quad (\text{S25})$$

where the second equality came from the fact that $\tau_r = 1/\nu_E$. The third step is to insert Eq. (S25) into (S24). Again using $\tau_r \nu_E = 1$, and performing a small amount of algebra, and we arrive at

$$\sigma_V^2 = \Delta V_{\max}^2 \frac{p_f + \sigma_V \tau_r \eta / R \pi^{1/2}}{1 + \sigma_V \tau_r \eta / R \pi^{1/2}}. \quad (\text{S26})$$

Comparing Eq. (S26) to (S15), we see that a nonzero failure rate increases the right hand side. As is easy to show, this increases σ_V . Thus, the effect of failures is to increase trial-to-trial variability.

7. Network simulations

Our analysis here and in the main text produced two observations. The first is that there is a relationship between extra spikes and chaotic dynamics: if one extra presynaptic spike causes more than one extra postsynaptic spike in a local network, then small perturbations grow, and network dynamics is chaotic. The second is that the number of extra spikes predicts, quantitatively, via Eq. (S16), a lower bound on the intrinsic noise at equilibrium, where intrinsic noise refers to the level of trial-to-trial fluctuations in the presence of identical input.

An important component of this work is that the above observations are general, and so do not depend on the details of the network (other than the existence of strong, local, recurrent connections). Nevertheless, the second observation in particular required a rather detailed calculation, so it makes sense to check it in at least one network. We therefore simulated a large, recurrent network of quadratic integrate and fire neurons, and tested two things. The first was whether more than one extra spike/spike predicted chaotic dynamics; the second was whether Eq. (S16) was satisfied. Note that numerous simulations have already shown that large networks of spiking neurons tend to be highly chaotic.^{20–24} What is new here is the test of the relationship between the number of extra spikes/spike and the trial-to-trial fluctuations.

The network we use has been describe previously²⁵ (although this version is slightly simpler in that we did not include spike-frequency adaptation). The parameters of the network are given in Table S1. The key quantities are the number of neurons (10,000), of which 80% are excitatory, and the connection probability between any two neurons (0.25), producing, on average, 2,500 connections per neuron. We simulated 8 networks that differed only in their PSP size, but were otherwise identical; these are labeled 1-8, and their PSP

Table S1. Network parameters. The neurons also received external Poisson input from excitatory and inhibitory neurons firing at 1,000 Hz and generating 2 mV PSPs. We performed one set of runs with a step size of 0.1 ms, but it had virtually no effect on our results. See Latham et al.²⁵ for a complete description of the network.

number of neurons	10,000
fraction of inhibitory neurons	0.2
connection probability	0.25
membrane time constant	4.0 ms
synaptic time constant	0.5 ms
resting membrane potential	-65 mV
threshold	-50 mV
excitatory reversal potential	0 mV
inhibitory reversal potential	-80 mV
step size	0.5 ms

sizes are given in Table S2. For the different networks the PSP sizes were chosen so that the firing rates for both excitatory and inhibitory neurons was about 5 Hz and there were no strong oscillations.

Spike rasters and a representative membrane potential for network 8 are shown in Fig. S8. These are typical; in fact, the plots for all the other networks look virtually identical. In these plots an extra spike was added at $t = 0$, but, because of the high connectivity, it had no visible effect on either the spike rasters or the membrane potential.

Our first step is to compute the probability of an extra postsynaptic spike given an extra presynaptic spike on a connected neuron, denoted p_1 , from network simulations. Here we

Table S2. The probability (p_1) and number (N_{extra}) of extra spikes versus PSP sizes for the eight networks. “E” and “I” refer to excitatory and inhibitory, respectively. PSP size is in mV. There was no spread in connection strength, so all PSPs were the same size. For all networks, the average firing rates of both excitatory and inhibitory neurons was approximately 5 Hz.

Network	E → E	E → I	I → E	I → I	p_1	N_{extra}
1	0.11	0.11	-0.17	-0.16	0.0010	2.5
2	0.15	0.16	-0.25	-0.23	0.0011	2.7
3	0.20	0.22	-0.35	-0.35	0.0015	3.6
4	0.35	0.39	-0.63	-0.65	0.0018	4.4
5	0.60	0.66	-1.05	-1.10	0.0021	5.2
6	0.75	0.80	-1.10	-1.10	0.0024	5.9
7	1.30	1.50	-1.70	-1.80	0.0023	5.7
8	1.70	1.60	-2.20	-2.00	0.0023	5.8




Figure S8: Simulations of network 8 from Table S2. **a.** Spike rasters for 100 randomly chosen neurons (out of 10,000). The bottom 20 are inhibitory; the top 80 are excitatory. An extra spike was added at $t = 0$ s, but its effects are not visible in this plot. **b.** Membrane potential of a randomly chosen excitatory neuron.

have a big advantage over the *in vivo* experiments: we can generate two runs with identical parameters and initial conditions, add a spike on the second one, and then count how many extra spikes were produced. For each network, we did this for 300 pairs of runs, each time with different initial conditions (the neurons had different membrane potentials at the start of the run), and with an extra spike on a different neuron. We counted spikes on all cells postsynaptic to the one that produced the extra spike, so, given that the average number of connections per neuron is 2,500, 300 pairs of runs gave us approximately 750,000 samples.

The results are shown in Fig. S9 for the eight networks listed in Table S2. Plotted is the cumulative number of extra postsynaptic spikes per neuron versus time relative to the presynaptic spike. Note that the rise time is relatively slow. This is because quadratic integrate and fire neurons (which we used in our simulations) are relatively slow to fire: during an action potential, it takes 3-5 ms for the membrane potential to go from rest to peak. This also explains the decrease in rise time with decreasing PSP size: spike generation in the quadratic integrate and fire neuron is much like a ball rolling off the top of a hill, and smaller PSPs correspond to smaller initial kicks to the ball.

The slow rise time made it difficult to identify exactly where the curves saturate. However, as we will see in Fig. S11 below, our results are insensitive to the precise value we choose.




Figure S9: Cumulative probability of an extra spike for the eight networks given in Table S2. The lowest curve corresponds to network 1, the next lowest to network 2, and so on. The intersections of the dashed line with the cumulative probability curves constitute our estimates for the probability of an extra spike.

Thus, we drew a line through the saturation point of the highest and lowest curves, as determined by eye (dashed line in Fig. S9), and used the intersection of that line with the cumulative probability curves as our estimate of the probability of an extra spike.

For all networks, the number of extra postsynaptic spikes per presynaptic spike is greater than 1 (last column in Table S2). Thus, they should all exhibit chaotic dynamics. To see if they do, in Fig. S10 we plot spike rasters immediately before and after a spike, and, below that, the standard deviation of the difference in membrane potentials between trials that did and did not have an extra spike at time $t = 0$. We show only networks 1 and 8, the ones with the smallest and largest number of extra spikes/spike, respectively.

There are two features to take away from Fig. S10. First, because each spike in network 1 (panels a and c) produced only about 2.5 extra spikes/spike (Table S2), the network is not very chaotic. Thus, after the extra spike at $t = 0$, the spike times on the two trials were closely related, if not identical (to within the step size) Nevertheless, close inspection indicates that there are some missed and extra spikes, and, as can be seen from panel c, perturbations grow rapidly: within about 100 ms the difference between the membrane potentials on the two trials has asymptoted to about 1 mV. Second, even though network 8 (panels b and d) produced only about a factor of two more extra spikes/spike than network 1 (5.8 versus 2.5; Table S2), shortly after the extra spikes, spike times on the two trials were completely unrelated. This is reflected in the difference in membrane potentials between the two trials, which asymptoted to a value close to the total average standard deviation.

So far, none of this is surprising, as numerous studies have already shown that spiking networks are chaotic.^{20–24} What our networks allow us to do that is new is to verify our

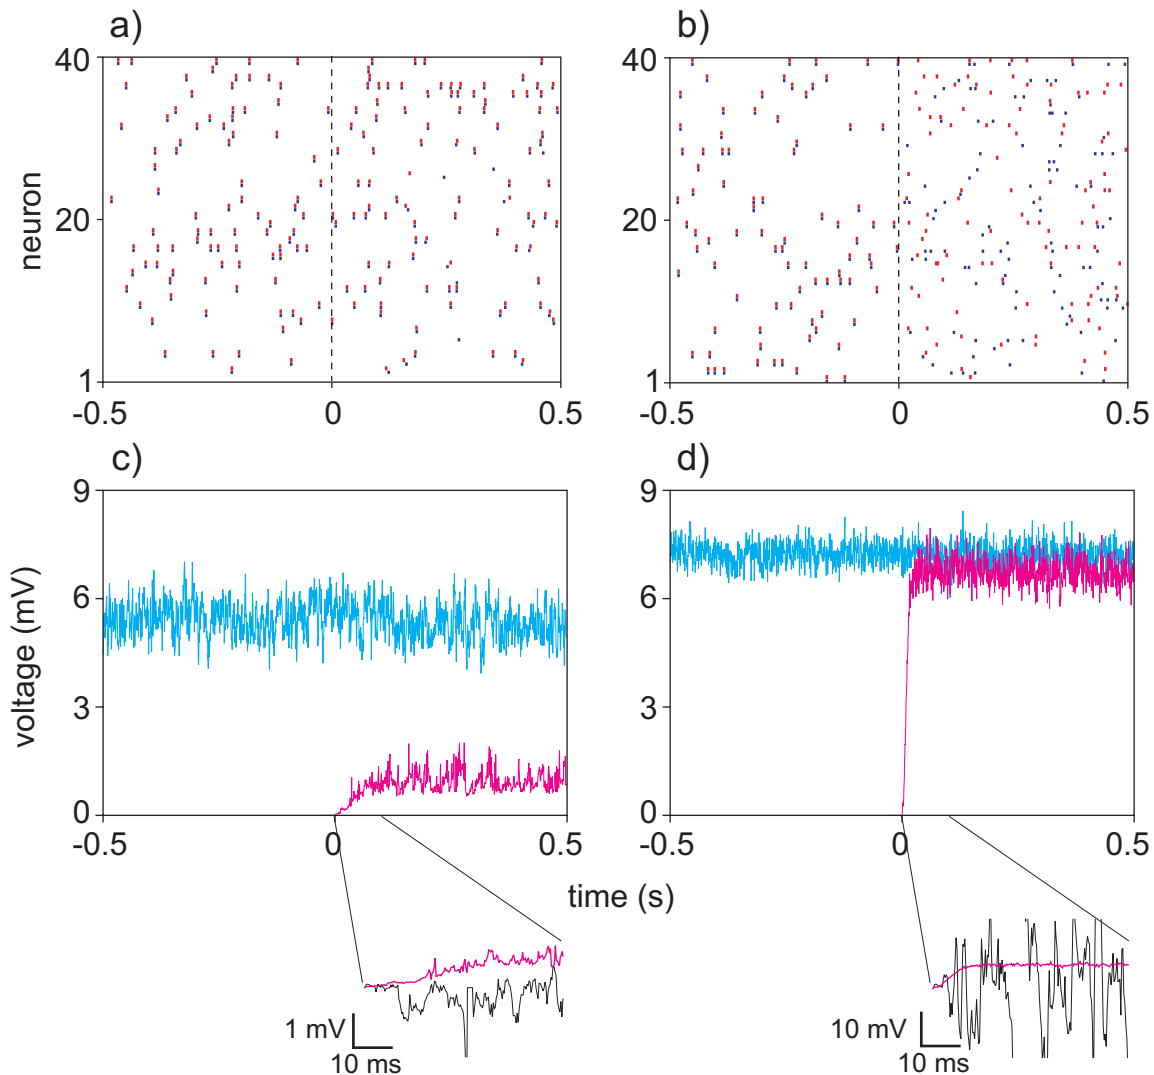


Figure S10: The effect of one extra spike. **a, b.** Spike rasters without (red) and with (blue) an extra spike for 40 randomly chosen neurons, the bottom 10 of which are inhibitory. Panels a and b correspond to networks 1 and 8 from Table S2, respectively. The extra spike was added at time $t = 0$. **c, d.** Standard deviation of the membrane potential (cyan) and the difference between the membrane potential with an extra spike and without one (magenta), averaged over the 300 pairs of runs, again for networks 1 and 8. Insets show the voltage difference on a single neuron (black traces), along with a blowup of the magenta trace (note the different voltage scales in panels c and d). To reduce the effect of spikes when computing standard deviation, voltage was truncated above -30 mV and below -60 mV. Averages were taken over 300 trials.

prediction of the lower bound on the trial-to-trial fluctuations at fixed input (the magenta traces in Figs. S10c and d). This lower bound, denoted (as above) σ_V , is given by Eq. (S16) in terms of two quantities: ΔV_{\max} and ξ . The first is just the trial averaged standard




Figure S11: Lower bound on the ratio $\sigma_V/\Delta V_{\max}$ (red line) along with the simulated values for 12 randomly chosen neurons in each of the 8 networks (black dots). Larger ξ corresponds to larger network index (e.g., network 1 corresponds to the smallest value of ξ , network 2 to the next smallest, and so on).

deviation of the membrane potential (cyan traces in Figs. S10c and d). The second is given in Eq. (S17). Because we used identical connection strengths for all neurons, we can replace the term $\overline{v_E^2}(t)$ that appears in Eq. (S17) with $v_E^2(t)$. Then, combining the relationships $\eta = p_1 \times (\text{charge per spike})$ and $(\text{charge per spike}) = R^{-1} \int dt v_E(t)$, Eq. (S17) becomes

$$\xi \equiv \frac{\exp[(\kappa \log \kappa)/(\kappa - 1)]}{2\pi^{1/2}(1 + \kappa)} \frac{p_1 K_E V_{EPSP}}{\Delta V_{\max}}$$

where V_{EPSP} is the EPSP size between excitatory neurons. To derive this expression, we used Eqs. (S38) and (S39) to convert the ratio $\int dt v_E^2(t)/\int dt v_E(t)$ into a function of κ . We take κ to be the ratio of the synaptic to the membrane time constant; from Table S1, we see that $\kappa = 0.125$.

In Fig. S11 we plot the ratio $\sigma_V/\Delta V_{\max}$ versus ξ for 12 randomly chosen neurons from each of our 8 networks (black dots). To compute $\sigma_V/\Delta V_{\max}$, we made plots like the one shown in Figs. S10c and d, computed the average heights of the magenta and cyan traces, and from that computed their ratio. We also plot, in red, the lower bound on $\sigma_V/\Delta V_{\max}$, which is given in Eq. (S16). In all cases, the network simulations produce ratios well above the predicted lower bound, thus verifying our calculations. The bound turns out to be not very tight, which is not surprising given the number of inequalities that went into its derivation. Note that we stopped the simulations at a relatively low value of ξ . We did this because even with $\xi < 0.5$, $\sigma_V/\Delta V_{\max}$ was close to its maximum value of 1. Increasing ξ would not, therefore, have provided much extra information.

Finally, in Fig. S12 we plot the network averaged probability of an extra spike (blue traces). These traces were computed by counting the total, network-wide, number of extra postsynaptic spikes produce by an extra presynaptic spike. The salient feature of these plots




Figure S12: Effect of extra spikes on the network. Blue: cumulative probability of an extra spike, averaged over all neurons in the network, for the 8 networks given in Table S2. Red: Cumulative probability of an extra spike; these curves are identical to the ones in Fig. S9.

is that, when extra spikes are averaged over the whole network, the effect is very small – the probability for network 8 is about 1 in 5,000. This number does not apply directly to our experiments, as the average probability of an extra spike was about 5 times higher in our experiments than it was for network 8 (28 versus 5.8). However, it indicates that detection of the increase in firing induced by a single spike is highly nontrivial.

8. Error analysis

Many of our quantities have large error bars, so we must consider full distributions rather than using a linear approximation; this makes error analysis nontrivial. In this section we provide details of exactly how errors are computed, and we derive confidence limits on the number of extra spikes and the voltage fluctuations.

The number of extra spikes, which we denote here N_{ex} , is given by

$$N_{ex} = \eta Q K \quad (\text{S27})$$

where η is slope of the regression line (Fig. 3b, main text) in units of probability/charge, Q is the average charge in an EPSP, and K is the average number of connections per neuron. In fact, Q is not a fundamental quantity; it depends on EPSP amplitude and membrane resistance, both of which carry uncertainty. We thus use Eqs. (S1) and (S38) to write

$$Q = \frac{1}{R} \bar{V}_{EPSP} \tau_{\text{decay}} \exp[(\kappa \log \kappa)/(\kappa - 1)]$$

where τ_{decay} is the decay time of EPSPs and κ is the ratio of rise to decay times of EPSPs (1.7 and 8 ms, respectively; see Methods). Inserting this expression into Eq. (S27) and using

$\tau_{\text{decay}} = 8$ ms and $\kappa = 1.7/8$, we have

$$N_{ex} = \tau_{\text{decay}} \exp[(\kappa \log \kappa)/(\kappa - 1)] \frac{\eta \bar{V}_{EPSP} K}{R} = 12.15 \frac{\eta \bar{V}_{EPSP} K}{R} \quad (\text{S28})$$

where the factor of 12.15 has units of ms.

The variance of the voltage fluctuations, σ_V^2 , depends on the parameter ξ (see Eq. (S16)). Combining Eqs. (S17-S19), ξ is written

$$\xi = \frac{0.8 e^{0.94}}{\pi^{1/2}} \frac{\exp[2(\kappa \log \kappa)/(\kappa - 1)]}{2(1 + \kappa)} \frac{\tau_{\text{decay}}}{\Delta V_{\text{max}}} \frac{\eta \bar{V}_{EPSP}^2 K}{R} = 8.31 \frac{\eta \bar{V}_{EPSP}^2 K}{R \Delta V_{\text{max}}}. \quad (\text{S29})$$

The factor of 0.8 translates from total number of connections, K , to the number of excitatory connection, K_E , and the factor of $e^{0.94}$ translates from \bar{V}_{EPSP}^2 to \bar{V}_{EPSP}^2 (see Eq. (S19)). The factor of 8.31 has units of ms.

Our goal now is to determine the distribution of N_{ex} and ξ given the distributions of the variables they depend on. Since all quantities are non-negative, we assume that they follow a log normal distribution. In the next section, we show that if a set of variables is related via

$$z \equiv \prod_i x_i^{n_i}, \quad (\text{S30}) \quad \{\text{zprod}\}$$

where the x_i are log normal and the n_i are integers, then z is also log normal. Furthermore, if $\log x_i$ has mean μ_i and variance σ_i^2 , then the mean and variance of $\log z$, denoted μ_z and σ_z^2 , are given by

$$\mu_z = \sum_i n_i \mu_i \quad (\text{S31a})$$

$$\sigma_z^2 = \sum_i n_i^2 \sigma_i^2. \quad (\text{S31b})$$

$\{\text{mu_sigma_z}\}$

We also show that the mean and variance of z , denoted $\langle z \rangle$ and $\langle \delta z^2 \rangle$, are related to the mean and variance of the x_i , denoted $\langle x_i \rangle$ and $\langle \delta x_i^2 \rangle$, via

Table S3. Statistics of η , \bar{V}_{EPSP} , K , and R . “mean” and “sd” are the mean and standard deviations of the parameters, δ_i^2 is the ratio of the variance to the square of the mean, as defined in Eq. (S33), and μ_i and σ_i are the mean and standard deviation of the logs of the parameters, computed from Eq. (S37). See Table S4 (Sec. 11) for additional details on these parameters.

variable	mean	sd	δ_i^2	μ_i	σ_i
η (prob/nC)	0.0608	0.0096	0.157	-2.81	0.156
\bar{V}_{EPSP} (mV)	1.075	0.225	0.209	0.0509	0.207
K	1500	500	0.333	7.26	0.325
R (M Ω)	42	4.3	0.102	3.73	0.102

$$\langle z \rangle = \prod_i \langle x_i \rangle^{n_i} (1 + \delta_i^2)^{(n_i^2 - n_i)/2} \quad (\text{S32a})$$

$$\langle \delta z^2 \rangle = \langle z \rangle^2 \left[\prod_i (1 + \delta_i^2)^{n_i^2} - 1 \right] \quad (\text{S32b})$$

{zbar}

where

$$\delta_i^2 \equiv \frac{\langle \delta x_i^2 \rangle}{\langle x_i \rangle^2}. \quad (\text{S33}) \quad \{\text{de12}\}$$

Let us apply these relationships to the expressions for N_{ex} and ξ , Eqs. (S28) and (S29). Both of these expressions depend on four variables: η , \bar{V}_{EPSP} , K , and R . Their relevant parameters – mean, variance, etc. – are listed in Table S3. Using this table and applying Eq. (S32), we see that

$$\begin{aligned} \text{mean}[N_{ex}] &= 12.15 \frac{0.0608 \times 1.075 \times 1500}{42} (1 + 0.102^2) \\ \text{Var}[N_{ex}] &= \langle N_{ex} \rangle^2 [(1 + 0.157^2)(1 + 0.209^2)(1 + 0.333^2)(1 + 0.102^2) - 1] \\ \text{mean}[\xi] &= 8.31 \frac{0.0608 \times 1.075^2 \times 1500}{42 \Delta V_{\max}} (1 + 0.209^2)^{3/2} (1 + 0.102^2) \\ \text{Var}[\xi] &= \text{mean}[\xi]^2 [(1 + 0.157^2)(1 + 0.209^2)^4 (1 + 0.333^2)(1 + 0.102^2) - 1]. \end{aligned}$$

Carrying out the arithmetic, we have

$$N_{ex} = 28.4 \pm 12.7 \quad (\text{S34a})$$

$$\xi = \frac{21.42 \pm 12.95}{\Delta V_{\max}}. \quad (\text{S34b})$$

{nex_xi}

Finally, we need the mean and variance of $\log \xi$. Using Eqs. (S29) and (S31) and Table S3, these are given by

$$\text{mean}[\log \xi] = \log 8.31 + \log 0.0608 + 2 \log 1.075 + \log 1500 - \log 42 - \log \Delta V_{\max} \quad (\text{S35a})$$

$$\text{Var}[\log \xi] = 0.156^2 + 4 \times 0.207^2 + 0.325^2 + 0.102^2. \quad (\text{S35b})$$

{logxi}

Performing the arithmetic gives us Eq. (S20).

9. The log normal distribution

In this section we derive Eqs. (S31) and (S32). We start by writing down the log normal distribution of the x_i . Using, as above, μ_i and σ_i^2 for the mean and variance of $\log x_i$, the distribution of $\log x_i$ is

$$p(\log x_i) = \frac{\exp[-(\log x_i - \mu_i)^2 / 2\sigma_i^2]}{(2\pi\sigma_i^2)^{1/2}}.$$

Taking logs of both sides of Eq. (S30), we have

$$\log z = \sum_i n_i \log x_i.$$

Since $\log z$ is linear in the $\log x_i$, z must also have a log normal distribution,

$$p(\log z) = \frac{\exp[-(\log z - \mu_z)^2 / 2\sigma_z^2]}{(2\pi\sigma_z^2)^{1/2}},$$

with μ_z and σ_z^2 given by Eq. (S31).

Equation (S31) tells us the relationship between the mean and variance of $\log z$ and the mean and variance of $\log x_i$. What we typically know, however, is the mean and variance of

x_i , and we want to use that to determine the mean and variance of z . To translate between the two, we use the fact that (as is easy to show)

$$\langle z \rangle = e^{\mu_z + \sigma_z^2/2} \quad (\text{S36a})$$

$$\langle \delta z^2 \rangle = \langle z \rangle^2 \left[e^{\sigma_z^2} - 1 \right] \quad (\text{S36b})$$

{zmoments}

where, as above, $\langle z \rangle$ and $\langle \delta z^2 \rangle$ are the mean and variance of z , respectively.

We can use Eq. (S31) to express μ_z and σ_z in terms of the μ_i and σ_i . It is convenient to take one further step and express the right hand side of Eq. (S36) in terms of the mean and variance of x_i . For this we invert Eq. (S36) (and replace the subscript z with i), leading to

$$e^{\mu_i} = \frac{\langle x_i \rangle}{(1 + \delta_i^2)^{1/2}} \quad (\text{S37a})$$

$$e^{\sigma_i^2} = 1 + \delta_i^2 \quad (\text{S37b})$$

{exps}

where δ_i^2 is given in Eq. (S33). Then, inserting Eq. (S31) into (S36), using (S37), and performing a small amount of algebra, we arrive at Eq. (S32).

10. Difference of exponentials – normalization and integrals

In our analysis we assume a difference of exponentials for both PSPs and current. This produces a relatively complicated relationship between the rise and decay times and the peak amplitude. Here we write down that relationship, and also compute the time integral of both the difference of exponential and its square.

Consider a function $f(t)$ that is a difference of exponentials with rise time τ_{rise} , decay time τ_{decay} , and amplitude f_0 . This function has the form

$$f(t) = f_0 \Theta(t) \frac{e^{-t/\tau_{\text{decay}}} - e^{-t/\tau_{\text{rise}}}}{(1 - \kappa) \exp[-(\kappa \log \kappa)/(1 - \kappa)]}$$

where

$$\kappa \equiv \frac{\tau_{\text{rise}}}{\tau_{\text{decay}}}.$$

With this normalization, it is straightforward to show that the maximum value of $f(t)$ is f_0 .

We need two quantities: the time integral of $f(t)$ and $f^2(t)$. These are given by, respectively,

$$\int dt f(t) = f_0 \tau_{\text{decay}} \exp[(\kappa \log \kappa)/(\kappa - 1)] \quad (\text{S38}) \quad \{\text{intf1}\}$$

and

$$\int dt f^2(t) = f_0^2 \tau_{\text{decay}} \frac{\exp[2(\kappa \log \kappa)/(\kappa - 1)]}{2(1 + \kappa)}. \quad (\text{S39}) \quad \{\text{intf2}\}$$

11. List of parameters

Table S4. Mean and standard deviation (sd) of parameters used to construct both the number of extra postsynaptic spikes per presynaptic spikes and the distribution of voltage fluctuations.

variable	mean	sd	source
number of connections, K	1500	500	literature ^{26,27}
EPSP amplitude (mV)	1.075	0.225	literature ¹³⁻¹⁶
EPSP rise time (ms)	1.7	–	experiment – Sec. 4
EPSP decay time (ms)	8.0	–	experiment – Sec. 4
average membrane resistance (M Ω)	42	4.3	experiment – Sec. 4
slope of regression line (prob/pC)	0.0608	0.0096	experiment – Fig. 3b, main text
ΔV_{max} (mV)	2.25-4.57	–	experiment – Fig. S7
current rise time (ms)	0.3	–	experiment – Sec. 4
current decay time (ms)	1.7	–	experiment – Sec. 4
charge/EPSP (nC)	0.31	0.061	derived – Eq. (S1)

References

- [1] D. Donoho and I.M. Johnstone. Ideal spatial adaptation by wavelet shrinkage. *Biometrika*, 81:425–455, 1994.
- [2] C Holmgren, T Harkany, B Svennenfors, and Y Zilberter. Pyramidal cell communication within local networks in layer 2/3 of rat neocortex. *J. Physiol.*, 551(Pt 1):139–153, 2003.
- [3] C P de Kock and B Sakmann. Spiking in primary somatosensory cortex during natural whisking in awake head-restrained rats is cell-type specific. *Proc. Natl. Acad. Sci. USA*, 106:16446–16450, 2009.
- [4] A.K. Lee, I.D. Manns, B. Sakmann, and M. Brecht. Whole-cell recordings in freely moving rats. *Neuron*, 51:399–407, 2006.
- [5] S. Crochet and C.C. Petersen. Correlating whisker behavior with membrane potential in barrel cortex of awake mice. *Nat. Neurosci.*, 9:608–610, 2006.
- [6] A. Herrmann and W. Gerstner. Noise and the psth response to current transients: I. general theory and application to the integrate-and-fire neuron. *J. Comput. Neurosci.*, 11:135–151, 2001.
- [7] M.J. Richardson. Firing-rate response of linear and nonlinear integrate-and-fire neurons to modulated current-based and conductance-based synaptic drive. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.*, 76:021919–021919, 2007.
- [8] Z.F. Mainen, J. Joerges, J.R. Huguenard, and T.J. Sejnowski. A model of spike initiation in neocortical pyramidal neurons. *Neuron*, 15:1427–1439, 1995.
- [9] A. Destexhe, M. Rudolph, J.M. Fellous, and T.J. Sejnowski. Fluctuating synaptic conductances recreate in vivo-like activity in neocortical neurons. *Neuroscience*, 107:13–24, 2001.
- [10] J. Waters and F. Helmchen. Background synaptic activity is sparse in neocortex. *J. Neurosci.*, 26:8267–8277, 2006.
- [11] M.E. Larkum and J.J. Zhu. Signaling of layer 1 and whisker-evoked ca^{2+} and na^{+} action potentials in distal and terminal dendrites of rat neocortical pyramidal neurons in

- [12] I.D. Manns, B. Sakmann, and M. Brecht. Sub- and suprathreshold receptive field properties of pyramidal neurones in layers 5a and 5b of rat somatosensory barrel cortex. *J. Physiol. (Lond.)*, 556 (Part 2):601–622, 2004.
- [13] H. Markram, J. Lübke, M. Frotscher, A. Roth, and B. Sakmann. Physiology and anatomy of synaptic connections between thick tufted pyramidal neurones in the developing rat neocortex. *J. Physiol.*, 500 (Part 2):409–440, 1997.
- [14] S. Song, P.J. Sjöström, M. Reigl, S. Nelson, and D. B. Chklovskii. Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biol.*, 3:e68, 2005.
- [15] B. Barbour, N. Brunel, V. Hakim, and J.P. Nadal. What can we learn from synaptic weight distributions? *Trends Neurosci.*, 30:622–629, 2007.
- [16] A. Thomson and C. Lamy. Functional maps of neocortical local circuitry. *Front. Neurosci.*, 1:19–42, 2007.
- [17] M Häusser and A. Roth. Estimating the time course of the excitatory synaptic conductance in neocortical pyramidal cells using a novel voltage jump method. *J. Neurosci.*, 17:7606–7625, 1997.
- [18] Q.Q. Sun, J.R. Huguenard, and D.A. Prince. Barrel cortex microcircuits: thalamocortical feedforward inhibition in spiny stellate cells is mediated by a small number of fast-spiking interneurons. *J. Neurosci.*, 26:1219–1230, 2006.
- [19] J.D. Victor and K.P. Purpura. Nature and precision of temporal coding in visual cortex: a metric-space analysis. *J. Neurophysiol.*, 76:1310–1326, 1996.
- [20] A. Banerjee. On the phase-space dynamics of systems of spiking neurons. I: model and experiments. *Neural Comput.*, 13:161–193, 2001.
- [21] A. Banerjee. On the phase-space dynamics of systems of spiking neurons. II: formal analysis. *Neural Comput.*, 13:195–225, 2001.
- [22] A. Banerjee. On the sensitive dependence on initial conditions of the dynamics of networks of spiking neurons. *J. Comput. Neurosci.*, 20:321–348, 2006.
- [23] A. Banerjee, P. Seriès, and A. Pouget. Dynamical constraints on using precise spike timing to compute in recurrent cortical networks. *Neural Comput.*, 20:974–993, 2008.

- [24] E.M. Izhikevich and G.M. Edelman. Large-scale model of mammalian thalamocortical systems. *Proc. Natl. Acad. Sci. USA*, 105:3593–3598, 2008.
- [25] P.E. Latham, B.J. Richmond, P.G. Nelson, and S.N. Nirenberg. Intrinsic dynamics in neuronal networks. I. theory. *J. Neurophysiol.*, 83:808–827, 2000.
- [26] V. Braitenberg and A. Schüz. *Anatomy of the Cortex*. Springer-Verlag, Berlin, 1991.
- [27] T. Binzegger, R.J. Douglas, and K.A. Martin. A quantitative map of the circuit of cat primary visual cortex. *J. Neurosci.*, 24:8441–8453, 2004.