Cortical Reliability Amid Noise and Chaos

Max Nolte[1,*], Michael W. Reimann[1], James G. King[1], Henry Markram[1,2], Eilif B. Muller[1]

[1]Blue Brain Project, École Polytechnique Fédérale de Lausanne, 1202 Geneva, Switzerland
[2]Laboratory of Neural Microcircuitry, Brain Mind Institute, École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland

*Correspondence: *max.nolte@epfl.ch*

Summary

Typical responses of cortical neurons to identical sensory stimuli are highly variable. It has thus been proposed that the cortex primarily uses a rate code. However, other reports show spike-time coding under certain conditions. The potential role of spike-time coding is constrained by the variability arising directly from noise sources within local cortical circuits. Here, we quantified this internally generated variability using a detailed model of rat neocortical microcircuitry with biologically realistic noise sources. We found stochastic neurotransmitter release to be a critical component of this variability, which, amplified by recurrent connectivity, causes rapid chaotic divergence with a time constant on the order of 10-20 milliseconds. Surprisingly, however, relatively weak thalamocortical stimuli can transiently overcome the chaos, and induce reliable spike times with millisecond precision. We show that this effect relies on recurrent cortical connectivity and is not a simple result of feed-forward thalamocortical input. We conclude that recurrent cortical architecture simultaneously supports both chaotic network dynamics and millisecond spike-time reliability.

Keywords: *Stochastic synaptic transmission; neocortex; excitation-inhibition balance; biophysical model; variability; reliability; coding; information; noise; chaos*

Introduction

The typical electrical activity of cortical neurons is highly variable, in the sense that membrane potentials, spike times and interspike intervals vary during spontaneous activity as well as across trials with identical sensory stimuli (Mohajerani et al., 2013; Shadlen and Newsome, 1998; Stern et al., 1997; Tolhurst et al., 1983). A substantial part of this cortical variability is probably due to hidden variables such as unobserved stimuli, environmental parameters, or brain state (Pachitariu et al., 2015; Renart and Machens, 2014). For instance, it has been shown that, in the visual cortex, the act of running modulates responses of neurons to identical stimuli (Niell and Stryker, 2010). Furthermore, it has also been demonstrated that neocortical neurons respond reliably to somatic current injections in vitro (Mainen and Sejnowski, 1995). Moreover, while most cortical neurons display variable activity, some neurons in sensory cortices can encode sensory input with high spike-time precision (Hires et al., 2015; Kayser et al., 2010; Petersen et al., 2001). Taken together, these findings suggest that the variability simply encodes hidden or unobserved variables, and that cortical activity is essentially deterministic (Masquelier, 2013). However, there are two important reasons to believe that a large part of cortical variability is due to internally generated noise that carries no signal.

First, all cortical neurons are affected by well-established cellular noise sources, such as stochastic synaptic transmission and ion-channel noise (Faisal et al., 2008). These cellular noise sources ultimately originate from proteins susceptible to thermodynamic fluctuations, and are therefore intrinsic sources of noise (Faisal et al., 2008; Renart and Machens, 2014). In particular, synaptic transmission is based on a sequence of stochastic molecular events, where the low numbers of molecules involved do not allow stochastic properties to average out (Ribrault et al., 2011). In fact, in tightly controlled slice environments, the probability of vesicle release upon action potential arrival at a single cortical synapse is low (~50% between thick tufted layer 5 pyramidal neurons, see Markram et al., 1997), and estimated to be substantially lower in vivo (Borst, 2010) (~10% between same neurons, see Markram et al., 2015). The universal presence of synaptic noise suggests that cortical neurons respond far less reliably to presynaptic inputs than to current injections. It has been shown, moreover, that a simplified cortical network model with stochastic synapses can provide a sufficient explanation for variable spiking (Moreno-Bote, 2014). Furthermore, in vitro, some types of inhibitory neurons respond irregularly even to constant somatic current injections (Petilla Interneuron Nomenclature Group et al., 2008), unlike excitatory neurons which respond reliably (Mainen and Sejnowski, 1995). This is due to ion-channel noise that is amplified during action potential initiation (Mendonça et al., 2016). Moreover, even activity in regular firing excitatory neurons could be subject to ion-channel noise, for example during action potential propagation in thin axons (Faisal and Laughlin, 2007).

Second, models suggest (van Vreeswijk and Sompolinsky, 1996, 1998) and experiments show (London et al., 2010) that cortical networks have chaotic dynamics. This implies, by definition, that small perturbations, such as those due to intrinsic cellular noise, are amplified. Thus, extra or missing spikes in the network, for example, due to failed synaptic transmission, could change the

trajectories of electrical activity for all recurrently connected neurons in the network, and even the global firing rate. This could lead, in turn, to large steady-state fluctuations.

In spite of their potential importance, the separate and combined impacts of network dynamics and cellular noise sources on cortical neuronal variability remain largely unexplored. Nonetheless, there are several reasons why understanding what proportion of cortical neuronal variability is generated internally and how this variability arises is crucial for understanding the neural code.

First, strong internally generated variability due to chaotic network dynamics could prevent coding based on spike timing past the sensory periphery, and favor theories of firing rate coding (London et al., 2010). To test the feasibility of models of cortical coding that rely on spike timing (Gütig and Sompolinsky, 2006; Luczak et al., 2015; Thorpe et al., 2001), we need to understand internal variability and how it arises.

Second, variability could carry information and encode signals itself, for example perceptual uncertainty (Orbán et al., 2016). It is thus essential to understand how to separate intrinsically generated variability that is *bona fide* noise from variability that encodes an additional signal or brain state.

Third, and more generally, optimal coding strategies for neural circuit models depend on where noise enters the circuit (Brinkman et al., 2016). To infer detailed information about the neural code, we need to understand the mechanisms responsible for internally generated variability. Currently, it is impossible to measure all external inputs to a local population of cortical neurons in vivo. As a result, we are still unable to quantify how much of the observed variability is generated internally by the local circuitry, and how much is generated externally.

We therefore sought to address these questions using a detailed, biologically constrained model of a prototypical *neocortical microcircuit* in rat somatosensory cortex (the NMC-model; Markram et al., 2015). The NMC-model, which reproduces a range of in vivo experiments, incorporates several sources of noise, including stochastic synaptic transmission and ion channel noise. Each of these noise sources is constrained to replicate experimentally observed variability. This bottom-up modeling approach provides full control over all noise sources present in the model as well as its external inputs and internal states.

Our approach is fundamentally different from previous theoretical and computational studies on neural variability in cortical networks. While the effects of cellular noise sources on variability have already been studied in isolated biophysical Hodgkin-Huxley type neuron models (Diba et al., 2006; Mendonça et al., 2016; Singh and Levy, 2017; Wang et al., 2010), network models of large populations of cortical neurons tend to make many *ad hoc* assumptions about the relevant level of biological detail (Balaguer-Ballester, 2017). In contrast, the NMC-model is far more detailed than these models and more tightly constrained by the experimental data. In particular, connectivity is established by integrating anatomical data, such as layer-dependent cell type densities, morphologies and bouton densities, to generate a wiring diagram (Reimann et al.,

2015) with highly heterogeneous connectivity (Gal et al., 2017; Reimann et al., 2017a, 2017b). Other features of the model include morphologically-detailed Hodgkin-Huxley type neuron models fitted to electrophysiological data (Van Geit et al., 2016), and probabilistic synapses featuring short-term depression and facilitation. Synaptic parameters are constrained by experimentally determined inter-trial variability, synaptic failure and spontaneous release rates, and corrected to reflect in vivo levels of calcium. Additionally, some of the irregular spiking neurons include stochastic ion-channel models, making it possible to account for the known stochastic firing types in cortex (Petilla Interneuron Nomenclature Group et al., 2008). Finally, the model exhibits a naturally emerging structural and functional balance between excitation and inhibition (Gal et al., 2017), without relying on assumptions about the exact level of coupling between excitatory and inhibitory currents.

To quantify internal variability in the NMC-model, we compared the probabilistic evolution of membrane potentials between simulation trials with identical initial conditions. We observed that membrane potentials for each neuron diverged with a time constant on the order of 10-20 milliseconds, which we found to be robust across dynamical states, and nearly saturated at the scale of the microcircuit. In fully innervated neurons, input from surrounding tissue made little difference to their observed variability.

To characterize how this variability arises, we designed a series of simulation experiments in which we selectively switched noise sources and recurrent network dynamics on and off. These experiments confirmed that the recurrent network dynamics were chaotic, in the sense that any perturbation caused divergence of membrane potentials, as predicted by simplified network models (van Vreeswijk and Sompolinsky, 1996, 1998) and in vivo observations (London et al., 2010). However, the rate by which membrane potentials diverged was determined by the interplay of synaptic noise and recurrent network dynamics. Hence, synaptic noise does not average out, but is a critical component of cortical network dynamics that drives internal variability.

Surprisingly, we found that weak thalamocortical inputs can overcome chaotic network dynamics to produce reliable spike timing. We showed that these reliable responses depend on the recurrent cortical circuitry, with feed-forward thalamocortical input having a much weaker effect at the single cell level. We conclude that recurrent cortical architecture both supports chaotic network dynamics, and allows relatively weak inputs to be transformed into reliable patterns of activity, with millisecond spike timing precision amid the noise and chaos.

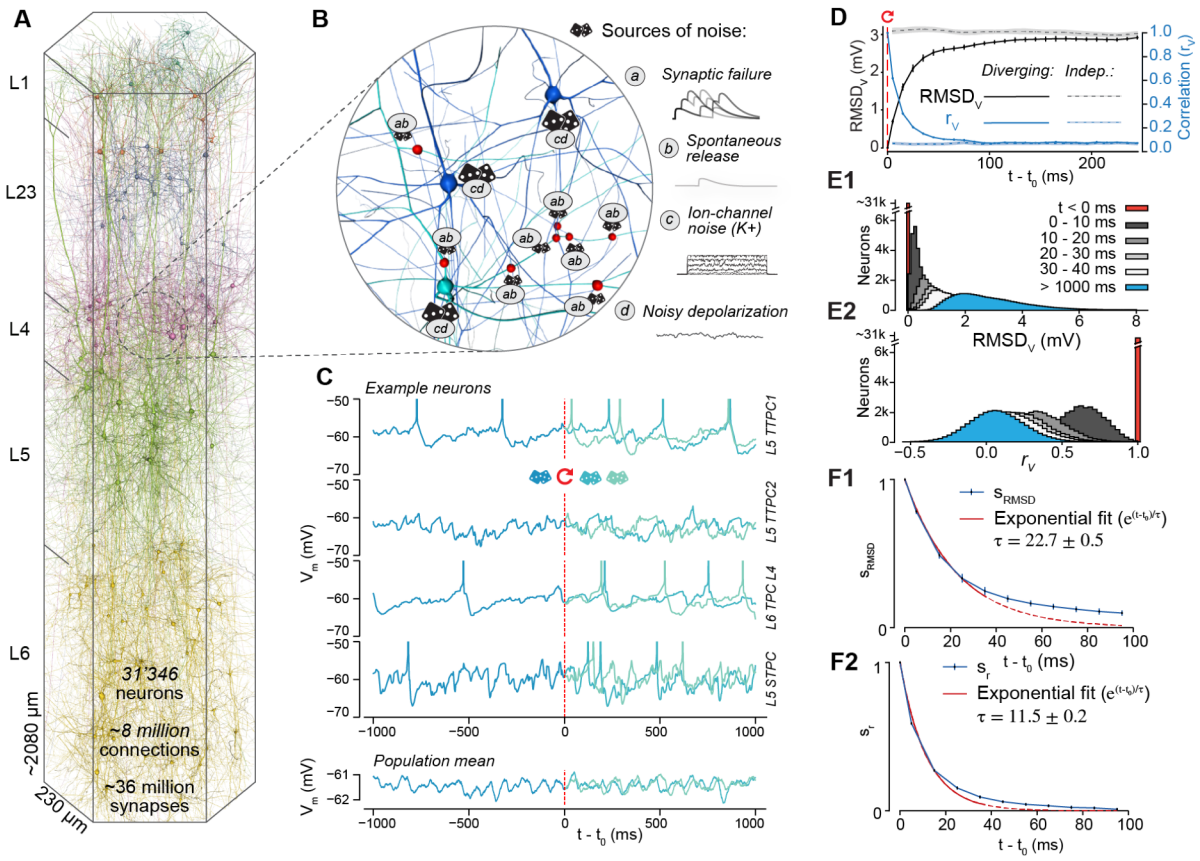

Figure 1: Rapid Divergence of Spontaneous Activity

(A) Morphologically-detailed model of a neocortical microcircuit (NMC); depicted are 100 randomly selected neurons, out of 31'346 in total (~0.3%). Neurons are colored according to their layer. (B) Examples of simulated noise sources in the NMC-model: stochastic synaptic transmission, including (*a*) vesicle release failure and (*b*) spontaneous vesicle release ('miniature PSPs') at all 36 million synapses; (*c*) probabilistic opening and closing of voltage-gated potassium channels in irregularly spiking inhibitory neurons (1'137 out of 31'346 neurons); (*d*) a constant depolarizing current with a weak white noise component ($\sigma_s^2 \ll \mu_s$) injected into the somata of all neurons. (C) The membrane potential of four sample neurons (and population mean of all 31'346 neurons) during a network simulation of in vivo-like spontaneous activity. At t_0, the state of the microcircuit is saved, and then resumed twice with identical initial conditions, but with different random seeds for all noise sources. (D) Root-mean square deviation ($RMSD_V(t)$) and correlation ($r_V(t)$) of the somatic membrane potentials between pairs of resumed simulations diverging from identical initial conditions (mean of all neurons and 40 saved base states ± 95% confidence interval). The dashed lines depict the steady-state $RMSD_V$ and r_V between independent simulations (i.e. resumed from different base states). (E) Time evolution of distributions of mean $RMSD_V$ and r_V values for individual neurons. (F) The similarity of the system (s_{RMSD} and s_r) defined as the difference between the diverging and steady-state $RMSD_V$ and r_V, normalized to lie between 1 (identical) and 0 (fully diverged) (mean ± 95% confidence interval). Exponential fit of s_{RMSD} and s_r for *t* < *40 ms* (estimated time constant ± 68% confidence interval of fit).

Results

Rapid Divergence of Spontaneous Activity

Using the NMC-model of rat somatosensory cortex (31'346 neurons, ~8 million connections, and ~36 million synapses), we simulated in vivo-like spontaneous neuronal activity (Figure 1A). The NMC-model contains three types of biological noise sources, all of which are required to accurately replicate neuronal responses to paired recordings and current injections in vitro (Figure 1B). Each of the 36 million synapses in the model incorporates stochastic models of vesicle release, which display both *failure* of vesicle release (*a*) and *spontaneous release* (*b*). In this way, overall synaptic variability is biologically constrained. The irregular firing electrical neuron types (*e-types*) (1'137 neurons) also contain models of *stochastic potassium channels* (*c*), which induce irregular firing in response to constant current injections in vitro. A fourth tunable noise source consisted of a noisy current (*d*) injected into the soma of each of the 31'346 neurons in the model, making it possible to depolarize neurons to in vivo-like levels (see Methods, Markram et al., 2015) (Figure 1B). In our initial experiments, we maintained the magnitude of this generic noise far below the magnitude of experimentally-constrained noise sources, using it later for sensitivity analysis. Realizations of the stochastic processes underlying these noise sources were determined by *sequences of random numbers*. By generating the sequences with different *random seeds*, we were able to obtain different, but equally valid probabilistic simulation trials.

Independent trials of electrical activity were simulated up to a time t_0, at which point we saved the full dynamical state of the simulation (*base state*). We then *resumed* the simulation two times from the base state, i.e. we used identical initial conditions and histories in each case, but with different sequences of random numbers. This allowed us to obtain two equally valid probabilistic network trajectories for $t > t_0$ for each base state. We observed that somatic membrane potentials (V_m) for individual neurons, and the mean potentials for the population both diverged rapidly between the two simulations (Figure 1C).

To quantify the time course of the divergence for each neuron *n,* we calculated the *root-mean-square deviation* of its somatic membrane potential in time bins of size Δt starting from t_0:

$$RMSD_V(n,k;t) = \sqrt{\int_{t-\Delta t/2}^{t+\Delta t/2} [V_{m,1}(n,k;t') - V_{m,2}(n,k;t')]^2 dt'/\Delta t}, \qquad (1)$$

where $V_{m,1}(n,k;t)$ and $V_{m,2}(n,k;t)$ denote the time series of somatic membrane potentials of neuron *n* in two trials resuming from the same base state *k.* We consequently defined the mean root-mean-square deviation of the microcircuit $RMSD_V(t)$ as the mean of $RMSD_V(n,k;t)$ over all base states (*K=40*) and neurons (*N=31'346*). We observed that $RMSD_V(t)$ diverged rapidly from zero and eventually converged towards a *steady-state* value $RMSD_\infty$, equal to the $RMSD_V$ of independent trials that did not share the same base state (Figure 1D, solid black and solid grey lines). The divergence was fast, with $RMSD_V(t)$ reaching more than 50% of its steady-state value within 20 ms.

While the $RMSD_V(t)$ of the circuit allowed us to accurately track the overall divergence of the whole circuit, $RMSD_V(n,k;t)$ for individual neurons and trials were too noisy for in-depth analysis (Figures 1E1 and S1A). We note that while $RMSD_V(t)$ quantifies the absolute distance between membrane potentials, potentials can still be correlated independent of this distance. To this end, we also computed the *linear correlation* for each neuron for each base state, again for time bins of size Δt:

$$r_V(n,k;t) = \frac{cov\left(V_{m,1}(n,k;t'),\ V_{m,2}(n,k;t')\right)}{\sigma\left(V_{m,1}(n,k;t')\right)\cdot\sigma\left(V_{m,2}(n,k;t')\right)},\ t - \frac{\Delta t}{2} < t' \leq t + \frac{\Delta t}{2} \qquad (2)$$

We found that the mean correlation $r_V(t)$ diverged faster than the absolute distance as measured by $RMSD_V(t)$ (Figure 1D, dashed blue line), again with a broad distribution across individual neurons (Figures 1E2, and S1A).

To better evaluate the difference between $r_V(t)$ and $RMSD_V(t)$, we computed the *similarity* $s_{RMSD}(t)$ of the microcircuit activity as the normalized difference between diverging and steady-state $RMSD_V(t)$ (and similarly $s_r(t)$ for $r_V(t)$). When similarity $s_{RMSD}(t) = 1$, membrane potential traces are identical; when $s_{RMSD}(t) = 0$ membrane potentials have reached their steady-state distance $RMSD_\infty$. Similarly, when $s_r(t) = 1$, membrane potentials have a perfect linear relationship; when $s_r(t) = 0$, they reached their steady-state correlation r_∞. Comparing $s_r(t)$ and $s_{RMSD}(t)$, we observed that $r_V(t)$ diverged approximately twice as fast as $RMSD_V(t)$ (Figures 1F1 vs. 1F2). More precisely, an exponential fit to the first 40 milliseconds revealed divergence time constants of $\tau_{RMSD} = 22.7 \pm 0.5$ ms and $\tau_r = 11.5 \pm 0.2$ ms ($\pm$ 68% confidence interval of fit). These were conserved for different bins sizes Δt, with similar values for bin sizes ranging from 1 ms to 50 ms (Figures S1C1 and S1C2). We observe, however, that simple exponential decay does not provide an adequate description of the whole time-course of the similarity, as the time constant changes continuously, especially in the first several milliseconds (Figure S1B). While the initial divergence is rapid, a small, but statistically significant difference ($p < 0.025$) between diverging and independent activity persists for around 400 ms for $RMSD_V$ (Figure S1D1) and around 200 ms for r_V (Figure S1D2).

Such rapid time-scales of divergence in the absence of any external input suggest that noise in the NMC-model does not average out. Instead, activity is inherently probabilistic, with a high internally generated variability. Throughout the remainder of this study, we will continue to quantify internally generated variability by the divergence of activity from identical initial conditions.

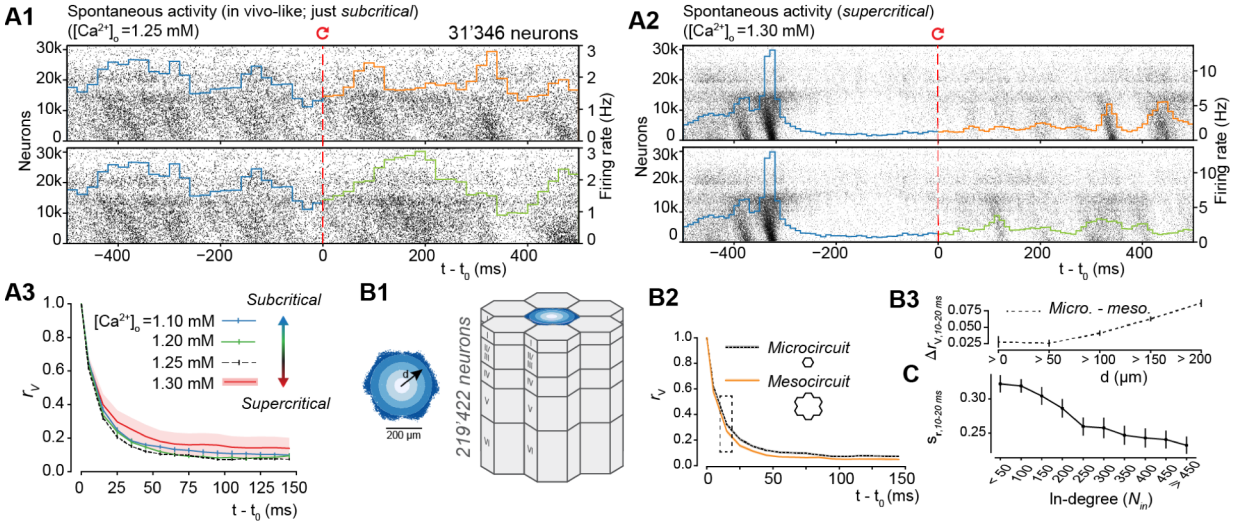

Figure 2: Robust Rapid Divergence Across Dynamical States and Microcircuit Scale

(A1) Population raster plot and population peristimulus time histogram (PSTH) for all 31'346 neurons in the microcircuit, during in vivo-like spontaneous activity. Neurons are ordered according to cortical depth, with deep layers at the bottom and upper layers at the top. Each row represents the spikes of one neuron. For visibility, raster lines extend over dozens of rows for each neuron. For $t < t_0$, the top and bottom raster plots show the same simulation, whereas for $t > t_0$, the raster plots depict two simulations resuming from identical initial conditions at t_0, but using different random number seeds. (A2) Same as A, but for supercritical activity. (A3) $RMSD_V$ and r_V across dynamic regimes (20 saved base states, mean ± 95% confidence interval; same as Figure 1D for in vivo-like regime ($[Ca^{2+}]_o = 1.25\ mM$). (B1) The microcircuit (center, blue), surrounded by 6 other microcircuits (grey), forming a continuous mesocircuit of ~220'000 neurons, with no boundary effects between the circuits. (B2) r_V for the center microcircuit when simulated without surrounding circuits (black), and of the center microcircuit when simulated as a mesocircuit (orange) (microcircuit: 40 saved base states; mesocircuit: 20 saved base states; mean ± 95% confidence interval). (B3) Quantifying edge effects. Difference of r_V between the same neurons in the microcircuit and the mesocircuit at 10-20 ms, plotted according to distance from horizontal center (mean ± 95% confidence interval). (C) Similarity s_r for subsets of neurons grouped by in-degree (bin size: 50; mean ± 95% confidence interval).

Single Cell and Shared Variability is Robust Across Dynamical States

Global fluctuations in spiking activity (Figure 2A1) and the population firing rate (Figure S2) both diverged rapidly for $t > t_0$. These global fluctuations indicate substantial *shared variability* between individual neurons. Thus, internally generated variability is high at both the microscopic and the macroscopic scales. However, cortical network dynamics depend on the balance between excitatory and inhibitory currents (EI-balance) (Brunel, 2000). In the NMC-model, EI-balance is modulated through the effects of extracellular calcium concentration ($[Ca^{2+}]_o$) on vesicle release probabilities (Borst, 2010; Markram et al., 2015). As the dependence on $[Ca^{2+}]_o$ is stronger for excitatory than for inhibitory synapses, increases in the concentration of $[Ca^{2+}]_o$ lead to stronger relative excitation and a sharp transition from asynchronous states (*subcritical*) to more correlated activity (Markram et al., 2015) that is regenerative and synchronous (*supercritical*; Figure 2A2). In the in vivo-like state analyzed here ($[Ca^{2+}]_o$= 1.25 mM), the microcircuit is in a just subcritical (Priesemann et al., 2014) state of asynchronous spontaneous activity, where it reproduces several findings from in vivo experiments (Markram et al., 2015). While this asynchronous state might be important for efficient coding (Beggs and Plenz, 2003; Denève and Machens, 2016), the exact EI-balance in vivo is difficult to determine, and is likely to reconfigure dynamically as a function of the state of arousal and attentiveness of the animal (Constantinople and Bruno, 2011). We therefore investigated the relationship between the time course of divergence and different dynamic regimes. We observed that the rapid divergence of electrical activity was approximately conserved across these different dynamic states (Figure 2A3). While steady-state electrical activity was slightly more de-correlated in the in vivo-like state, the time course of divergence was remarkably similar. We also found that the synchronous state still displayed high shared variability, with unpredictable timing of population bursts (Figure 2A2, $t > t_0$). In our model, therefore, intrinsic variability, as quantified by the time course of divergence, is conserved across a spectrum of dynamic states and does not depend on the exact EI-balance.

Variability is Nearly Saturated at the Scale of the Microcircuit

It is possible that the amount of internally generated variability depends not just on the dynamic state of the model circuit but also on its size. We have previously shown that in models of the size used in the simulations just described, dynamic states stabilize (Markram et al., 2015). At this size, dendritic trees and thus the afferent connections of neurons in the horizontal center of the microcircuit are fully located within the microcircuit. However, a large fraction of their recurrent connections with neurons in the surrounding tissue are with neurons at the periphery of the microcircuit. Since these were not included in the simulations, large portions of synaptic input to peripheral neurons were missing. To quantify the effect of this additional input on variability in the microcircuit, we surrounded the original microcircuit with six additional microcircuits, simulating a much larger *mesocircuit,* which provided missing synaptic input to

the neurons at the periphery of the microcircuit (Figure 2B1, blue and grey). Connectivity in this mesocircuit was homogeneous, both within and between the individual microcircuits.

When we compared the divergence of membrane potentials between micro- and mesocircuit simulations, we found that membrane potentials diverged slightly faster in the mesocircuit, although the time courses of divergence followed similar trends (Figure 2B2). The mean difference in $r_V(t)$ was always below 0.06, and the steady state difference below 0.03. We next focused on the difference at 10-20 ms, which we found to be a good predictor of the relative order of differences at any time. We found that $s_{r,10-20\,ms}$ was directly related to distance from the horizontal center, with the largest differences in neurons at the periphery of the microcircuit (Figure 2B3). At the periphery, the increase in variability between meso- and microcircuit simulations was above 0.08, decreasing toward the center and converging just below 0.03 for neurons within 100 μm of the center. This suggests that direct additional synaptic input onto a neuron increases variability, but that this additional synaptic input has a weak effect on indirectly connected neurons whose inputs are already saturated. Thus, at the scale of the microcircuit, the amount of internally generated variability is nearly saturated, while variability for neurons at the periphery is underestimated.

Highly Connected Neurons Diverge Faster

To directly quantify the dependence of the time course of divergence on the amount of the synaptic input, we examined the relationship between the similarity $s_r(t)$ of a given neuron and the number of connections it receives from within the microcircuit (*in-degree*). Once more, we found that the time course of divergence was faster, the more synaptic inputs a neuron received, as summarized by $s_r(t)$ at 10-20 *ms* (Figure 2C). Thus, it appears that neurons which are more strongly coupled to the local population (Okun et al., 2015) are also more likely to diverge quickly. Repetition of the analysis using $RMSD_V(t)$ instead of $r_V(t)$ gave qualitatively similar results (data not shown). We note that $RMSD_V(t)$ and $r_V(t)$ are generally highly correlated (Figure S3A, *abcd*). In what follows, we hence present the divergence in terms of $r_V(t)$ except when there is a qualitative difference.

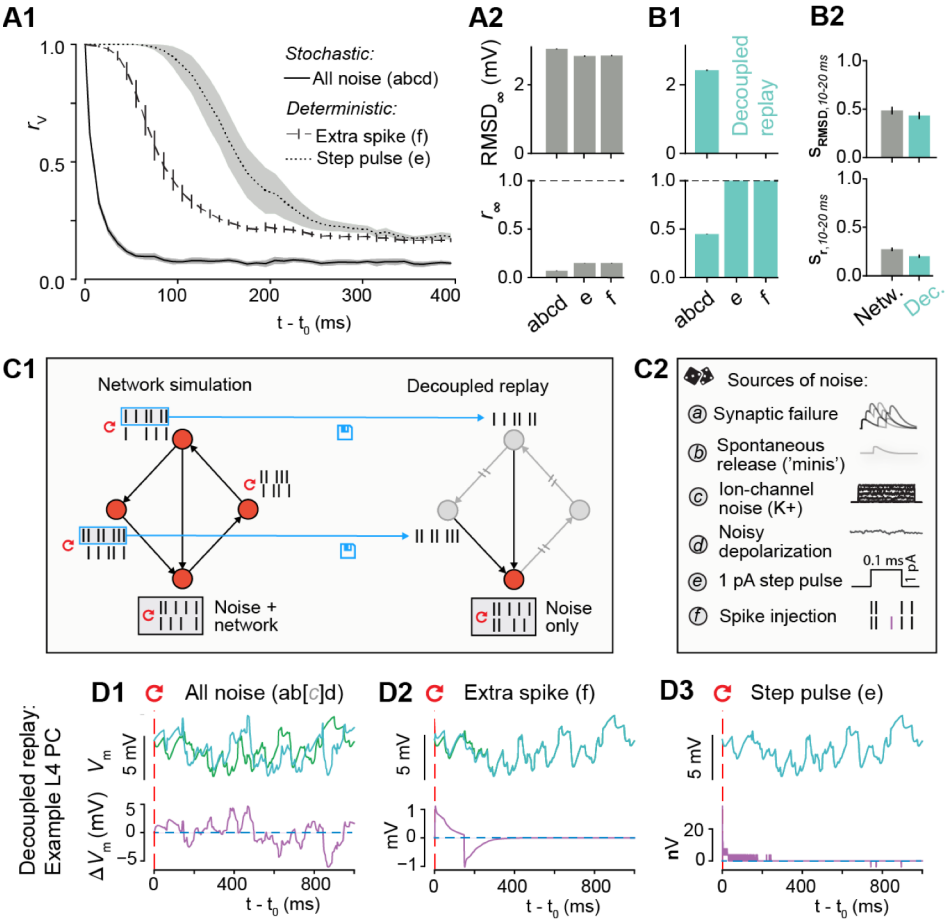

Figure 3: Noise Amplified by Chaos Determines Internally Generated Variability

(A1) Time course of correlation r_V after resuming at t_0 from identical conditions with different forms of perturbation. Full cellular noise as before, solid line (*abcd*); no cellular noise, but perturbing with a single extra spike in one neuron, dashed line (*f*); a miniscule step pulse perturbation in all neurons, dotted line (*e*). (*abcd*: 40 saved base states; *e, f*: 20 saved base states; mean ± 95% confidence interval) (A2) Steady-state root-mean square deviation $RMSD_\infty$ and correlation r_∞ for stochastic (*abcd*) and deterministic simulations (*e*, *f*) as defined in A1 (mean ± 95% confidence interval). (B1) As in A2, but for decoupled, replayed simulations. (B2) Similarity s_{RMSD} and s_r at 10-20 *ms* with all noise sources enabled, for network and decoupled simulations (mean ± 95% confidence interval). (C1) Decoupled replay paradigm. Presynaptic spike trains from a network simulation are saved and then replayed to the synapses of each neuron in a decoupled simulation, thereby removing variability due to feedback network dynamics. (C2) Overview of sources of noise and perturbations. (D) Decoupled replay simulations (see C1) for a representative L4 PC neuron, with somatic membrane potential differences between the two trials only due to cellular noise sources (*ab[c]d*), a single extra presynaptic spike (*f*) or a miniscule step-pulse perturbation (*e*). [*c*] indicates that for some neuron types in the NMC-model, such as L4 PCs, no stochastic ion-channels are present.

Noise Amplified by Chaos Determines Internally Generated Variability

Thus far, we have demonstrated a high level of variability which is robust across dynamical states and nearly saturated at the scale of the microcircuit. We have also shown that divergence is faster for neurons that are more tightly coupled to the local population (Figure 2C). This suggests that the variability of individual neuron activity is driven by the variability of local population activity, or that additional synaptic input simply adds more synaptic noise, or that the noise is determined by some combination of the two effects. In other words, while cellular noise is the only original source of variability in the NMC-model, the question remains to what degree this noise is amplified by recurrent network connectivity.

To address this question, and more generally, to study the interaction of noise sources and recurrent network dynamics, we performed two complementary sets of simulation experiments. In the first set, we sought insights into the role of network dynamics without noise sources, probing the sensitivity of a completely deterministic version of the model to a weak, momentary perturbation. In the second, we studied the opposite case of variability due to stochastic noise sources without amplification by the network.

To implement the first set of simulations, we disabled stochasticity of cellular noise sources, including synaptic transmission, by using a fixed sequence of random numbers, which made the random outcome deterministic (or alternatively by completely replacing the stochastic model with a deterministic one, see below). This enabled us to observe amplification of perturbations through the network without the effect of continuously varying cellular noise sources. As the sole source of perturbation, we injected a single extra spike into one of the neurons in the microcircuit (see Methods). We observed that the network diverged rapidly (Figure 3A1, dashed line), though more slowly than with noise sources enabled (Figure 3A1, solid line). In fact, even a miniscule current injection, which shifted the majority of spike times by less than 0.05 ms (see Methods), eventually led to a divergence of membrane potentials similar to the divergence observed in the full model with noise sources (Figure 3A1, dotted line). The slightly higher steady-state correlation r_∞ in the deterministic simulation was due to identical spontaneous release of neurotransmitter, identical ion-channel opening probabilities, and the small, but identical, noisy component of the depolarizing current injection. However, the relative difference in $RMSD_\infty$ was much smaller than the difference between the deterministic and the stochastic simulations (Figure 3A2, top, vs. Figure 3A2, bottom). That is, any perturbation to the system eventually led to a similarly large steady-state divergence. We conclude that the underlying dynamics of the circuit are chaotic, in the sense that small perturbations, such as one injected spike, lead to completely different, unpredictable activity.

It is important to note that when using a fixed random seed to make the stochastic version of the Tsodyks-Markram synapse model deterministic (Markram et al., 2015; Tsodyks and Markram, 1997), any extra or missing presynaptic spike can change the outcome for the next spike by advancing the sequence of random numbers. To avoid this difficulty, we ran equivalent simulations using the deterministic version of the Tsodyks-Markram synapse model (see

Methods). In these simulations, extra spikes and small perturbations produced qualitatively similar divergence time courses (Figure S4, A vs. B, dark green and pink lines).

We had shown that the network amplifies extra spikes or even small perturbations of membrane potentials. This leads to chaotic divergence of activity with similar steady-state variability, but different time courses. It remained to be seen whether this high level of variability requires network amplification or whether it could be generated directly by the noise sources alone.

To address this question, we implemented a second set of simulations to study the case of ongoing noise sources without network propagation. In these *decoupled replay* simulations, in contrast to regular *network simulations,* synaptic mechanisms were activated by spikes at fixed times, recorded in an earlier simulation experiment (Figure 3C1). In this way, the network was no longer able to amplify neuronal variability and neuronal variability was entirely due either to cellular noise sources or perturbations (Figure 3D). We found with all noise sources turned on, somatic membrane potentials still diverged rapidly, as quantified by $s_{r,10-20\,ms}$ (Figure 3B2) (as mentioned above, we found s_r at 10-20 ms to be a good predictor of the relative order of s_r at any time). However, steady-state r_∞ was higher and $RMSD_\infty$ was lower than in the network simulations (Figure 3B1 vs 3A2). When the decoupled replay paradigm was used with the deterministic version of the model, single extra spikes and brief current injections only evoked small, transient perturbations (Figure 3C2, 3C3). It follows that the high level of variability observed in network simulations was due to chaotic network dynamics which amplified rapid perturbations of activity from cellular noise sources.

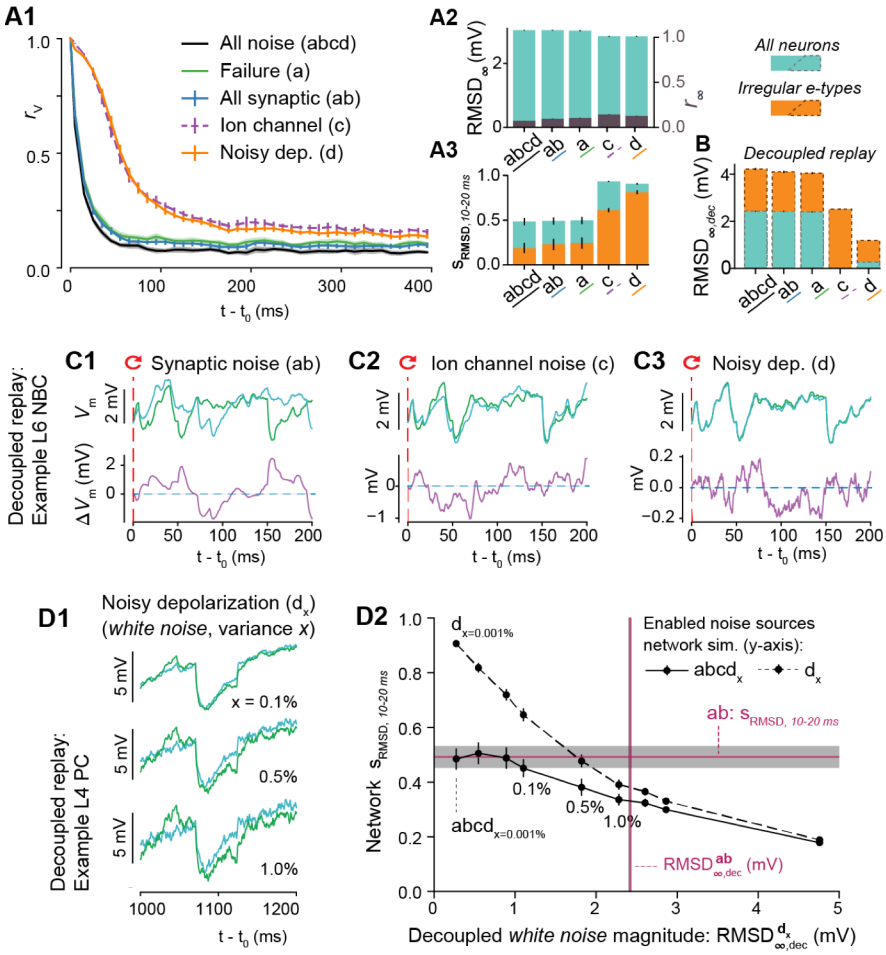

Figure 4: Synaptic Noise Dominates Variability

(A1) Time course of correlation r_V after resuming at t_0 from identical conditions with different noise sources enabled (*abcd*: 40 bases states; *a, ab, c, d*: 20 base states; mean ± 95% confidence interval). (A2) Steady-state root-mean square deviation $RMSD_\infty$ (cyan) and correlation r_∞ (purple) with different noise sources enabled. (A3) Similarity s_{RMSD} at 10-20 ms with different noise sources enabled, for all neurons (cyan) and irregular e-types (orange). (B) Steady-state root-mean square deviation for decoupled simulations, $RMSD_{\infty,dec}$, for all neurons (cyan) and irregular e-types (orange). Only irregular e-types in (c), 1,137 out of 31,346 neurons. (C) Decoupled replay simulations for a representative L6 NBC neuron, with somatic membrane potential differences between the two trials only due to synaptic noise (*ab*), ion-channel noise (*c*) or a noisy current injection (*d*). (D1) The effect of changing random seeds for the noisy depolarization only, for different noise strengths in a decoupled simulation. *x*: white noise variance as percentage of mean injected current (D2) The decoupled steady-state membrane potential fluctuations $RMSD^{d}_{\infty,dec}$evoked by different magnitudes of white noise without network dynamics, versus the similarity s_{RMSD} at 10-20 *ms* during network simulations when either turning on only the white noise depolarization (*d*) or all noise sources (*abcd*). Similarly, in purple, $RMSD^{ab}_{\infty,dec}$ for synaptic noise versus the similarity at 10-20 *ms* when only turning on synaptic noise (*ab*). All error bars and shaded areas indicate 95% confidence intervals. Means for D2 are based on ten base states.

Synaptic Noise Dominates Variability

To understand the contribution of individual noise sources in this interplay of noise and recurrent network dynamics, we designed a series of simulation experiments where we selectively disabled specific subsets of noise sources, instead of all of them as in the deterministic version above. We observed that disabling all noise sources except synaptic failure produced a time course for $r_V(t)$ and steady-state divergence r_∞ which was very similar to observations with all noise sources combined (Figure 4A1, black and green lines). On the other hand, disabling all but ion channel noise or all but the noisy current injection led to much slower divergence (Figure 4A1, orange and purple lines). As before, we quantified the speed of divergence by the similarity s_r at 10-20 ms after t_0 ($s_{r,10-20\,ms}$) (Figure 4A3, cyan). Our results suggest that simulations with synaptic failure give rise to rapid divergence, whereas steady-state r_∞ and $RMSD_\infty$ depend on noise sources only weakly (Figure 4A2). We conclude that in the NMC-model, the time course of divergence depends on synaptic noise, a combination of synaptic failure and spontaneous release, and that other noise sources add little to no additional variability.

Ion-Channel Noise in Irregular Firing Neurons Overshadowed by Synaptic Noise

Synaptic noise in the NMC-model is modeled at every single synapse, while ion-channel noise is limited to irregular firing e-types (see Markram et al., 2015; Petilla Interneuron Nomenclature Group et al., 2008). Irregular e-types are defined by high intrinsic spike-time variability in response to constant current injections in vitro, even in the absence of synaptic noise. In the NMC-model, irregular spiking is modeled with a subset of stochastic ion-channels, in accordance with in vitro findings on the source of the irregular spiking patterns observed in cortical interneurons (Mendonça et al., 2016). In contrast, regular firing e-types do not require noisy ion-channels to replicate in vitro spiking behavior. To better understand the interplay of ion-channel noise and synaptic noise, we focused our next analysis solely on irregular firing e-types. We observed that irregular firing e-types diverged significantly faster than the whole population (Figure 4A3, orange vs. cyan). However, synaptic noise still dominated over ion-channel noise. Enabling ion-channel noise in addition to synaptic noise led to only marginal gains in divergence rate; when ion-channel noise was enabled on its own, divergence was significantly slower (Figure 4A3, orange, *ab* vs. *abcd* and *c*). This suggests that in in vivo conditions, noise from stochastic ion-channels is probably overshadowed by synaptic noise. This contrasts with in vitro conditions, where channel noise is the only major noise source.

Synaptic Noise Acts as Threshold for Other Noise Sources

There are in reality many smaller noise sources that are not included in our model (see Discussion). To understand how additional noise sources of various magnitudes could influence divergence, we analyzed the magnitude of the previously analyzed cellular noise combinations in a decoupled replay, with network propagation removed ($RMSD_{\infty,dec}$) (Figure 4B; see Figure

4C1-3 for representative examples). We found that this magnitude inversely relates to the speed of divergence, $s_{RMSD,10-20\,ms}$ (Figure 4A3). That is, a larger $RMSD_{\infty,dec}$ leads to a faster divergence (as measured by a smaller $s_{RMSD,10-20\,ms}$) (see also Figure S5 for an extensive comparison of noise sources across simulation paradigms). In the NMC-model, synaptic noise has the largest $RMSD_{\infty,dec}$ and determines the rate of divergence. But how strong would any other noise source have to be to generate network variability that is detectable beyond synaptic noise? To answer this question, we studied how the magnitude of an unknown noise source affects the time course of divergence. As a proxy for unknown noise sources, we increased the variance σ_s^2 of the injected white noise depolarizing current. Previously, the variance had been set to 0.001% of the firing threshold for each neuron – a level far lower than other sources of noise. When we increased the variance to values from 0.01% up to 10%, and disabled all other noise sources, we observed that increasing σ_s^2 led to more rapidly diverging network dynamics (Figure S6A). However, when other noise sources were also enabled, the noisy current injection only affected network dynamics beyond a certain threshold (Figure S6B).

To characterize this threshold, we first measured $RMSD_{\infty,dec}^{d}$, that is, the steady-state divergence of membrane potential fluctuations evoked by noisy current injection alone in a decoupled replay, for various levels of σ_s^2 (Figure 4D1). We then compared $RMSD_{\infty,dec}^{d}$ to the time course of divergence in the corresponding network simulations with the same noise conditions (i.e. only noisy depolarization (d); Figure 4D2, dashed line). We found that the rate of divergence as measured by $s_{r,10-20\,ms}$ was strongly dependent on $RMSD_{\infty,dec}^{d}$, with larger values leading to faster divergence. In contrast, when we repeated the analysis with *all* noise sources enabled (Figure 4D2, solid line), the dependence on $RMSD_{\infty,dec}^{d}$ was weaker, indicating a smaller impact of σ_s^2 on $s_{r,10-20\,ms}$. Indeed, σ_s^2 only had a meaningful influence when it was beyond a threshold in the range 0.1% -0.5%. At this threshold, the steady-state divergence in decoupled replays ($RMSD_{\infty,dec}^{d}$) evoked by the noisy current alone was just above 1 mV, approximately half of the value for synaptic noise sources ($RMSD_{\infty,dec}^{ab}$, Figure 4D2, vertical purple line "ab"). When σ_s^2 is increased even more, the curves for $s_{r,10-20\,ms}$ with noisy current alone and with all noise sources eventually began to converge. Thus, when $RMSD_{\infty,dec}^{d}$ was larger than $RMSD_{\infty,dec}^{ab}$ the noisy current injection dominated other noise sources. This suggests that the strongest source of cellular noise dominates over other sources, unless they are of a comparable magnitude. Under biological conditions, we predict that synaptic noise dominates. This prediction matches previous findings that cortical neurons respond very reliably to current injections in vitro (no synaptic noise) (Mainen and Sejnowski, 1995). However, it also suggests an entirely different picture of the reliability of neuronal responses to presynaptic inputs in vivo, with synaptic noise contributing to variability.

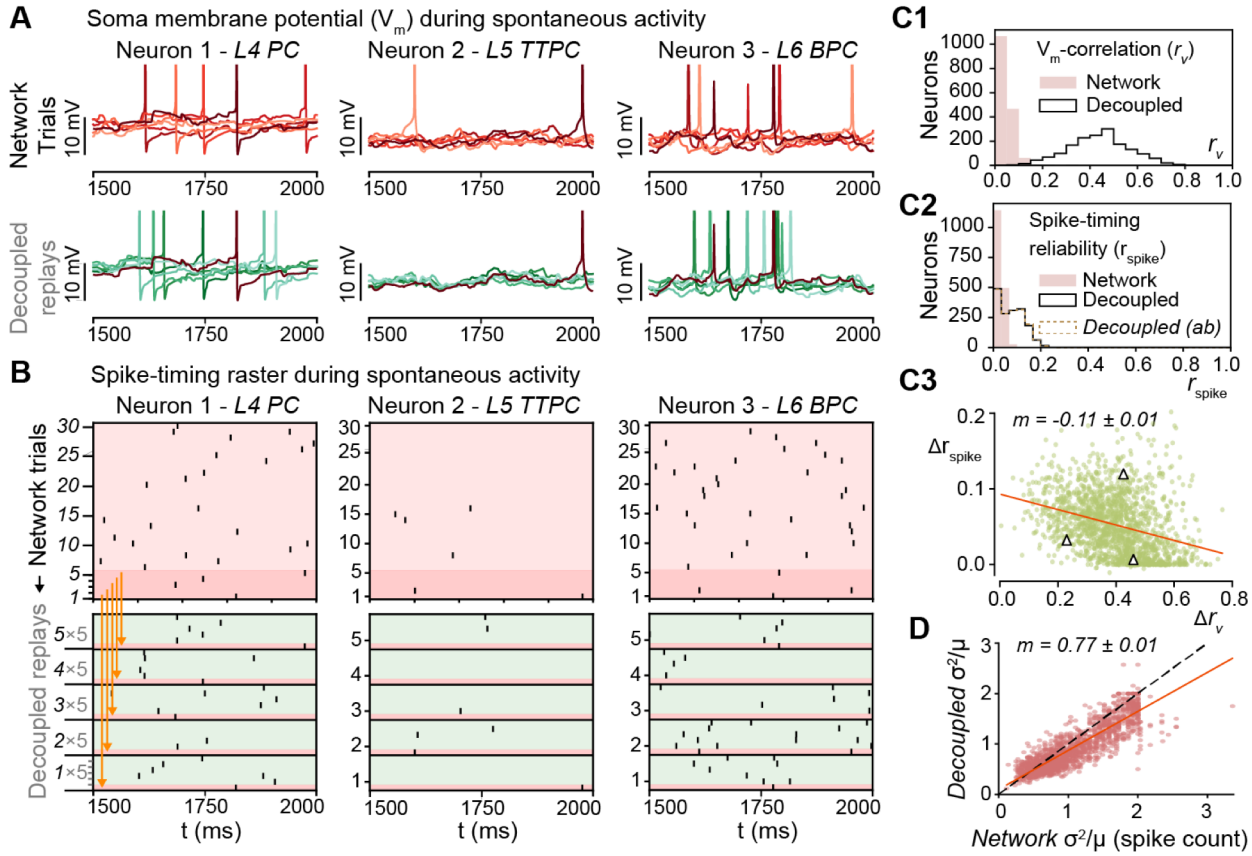

Figure 5: Low Trial-by-Trial Spike-Timing Reliability During Spontaneous Activity

(A) Somatic membrane potentials (V_m) of three representative neurons. Top: during six independent trials of spontaneous activity. Bottom: five decoupled replay trials (green) with the same presynaptic input as during the original network simulation trial (red), but with different random seeds. (B) Top: Raster plot of spike times for the same example neurons as in A, during 30 independent trials of spontaneous activity. Bottom: 5 decoupled replay trials (green) of the same input received during 5 of the 30 original trials (dark red). (C1) Mean somatic membrane potential correlation r_V of the 1666 (*ab*: 1670) most central (and spiking) pyramidal neurons from layers 4, 5, and 6 between independent network simulations, and between decoupled replay simulations with identical presynaptic inputs. (C2) Mean spike-timing reliability r_{spike} of the same neurons. Decoupled and decoupled (ab) are overlapping. (C3) Change in correlation, Δr_V, versus change in spike-timing reliability, Δr_{spike}, for each neuron for decoupled replay simulations relative to network simulations (linear fit with 68% confidence interval on slope *m*, red line). Triangles indicate values of representative neurons in panel B. (D) Comparison of variance of spike count between network and decoupled replay simulations (same neurons as in C-E; linear fit as in C3, red line; identity line, black dashed line).

Low Trial-by-Trial Spike-Timing Reliability During Spontaneous Activity

Neurons transmit signals to postsynaptic partners only in the form of spikes. Their timing is determined by a non-linear transformation of the somatic membrane potential (V_m) we analyzed so far. We therefore next characterized the role of the spike generation mechanism in influencing the reliability of neural responses, and compared both the variability of membrane potentials and spike times between independent network simulations of spontaneous activity over 30 independent trials with different initial conditions. We found that membrane potentials (Figure 5A, top) and the corresponding spike trains (Figure 5B, top) were both highly variable. We then used the spike times recorded from each of these network trials in five decoupled replay simulations per trial. As before, we observed that membrane potentials were less variable and more correlated in decoupled simulations (Figure 5A, bottom). Indeed, the distributions of r_V in decoupled and network simulations were almost completely disjoint (Figure 5C1), with decoupled replay simulations exhibiting much more correlated membrane potentials overall. However, considering just the spike times, we found no clear difference in variability between the network (Figure 5B, top) and the decoupled replay simulations (Figure 5B, bottom). In stark contrast to the results for membrane potentials, quantification of spike time variability using a correlation-based measure, r_{spike} (Schreiber et al., 2003), showed a large overlap in the distributions for decoupled and network simulations (Figure 5C2, red area vs. solid black line; σ_{spike} = 5 ms), in stark contrast to the case for membrane potentials. This suggests that the spike initiation mechanisms cannot transform the increased reliability of V_m into reliable spike trains during spontaneous activity. Indeed, we found that an increase in the magnitude of r_V did not predict a corresponding increase in r_{spike} (Figure 5C3). On the contrary, the two measures displayed a weak inverse correlation. We note that we found no decrease in variability when other sources of noise besides synaptic noise were disabled (Figure 5C2, dashed brown line), as expected in light of our previous result that synaptic noise accounts for a large proportion of variability.

It is possible of course that spike time reliability within tens of milliseconds could be too restrictive a measure of spiking reliability. Therefore, we also compared the variability of spike counts across the entire 0.5 s window analyzed. Use of identical presynaptic inputs produced only a marginal reduction in the variance of spike counts (Figure 5D). In brief, the reliability of spike generation across time-scales is directly, and severely constrained by synaptic noise, even without amplification through network dynamics.

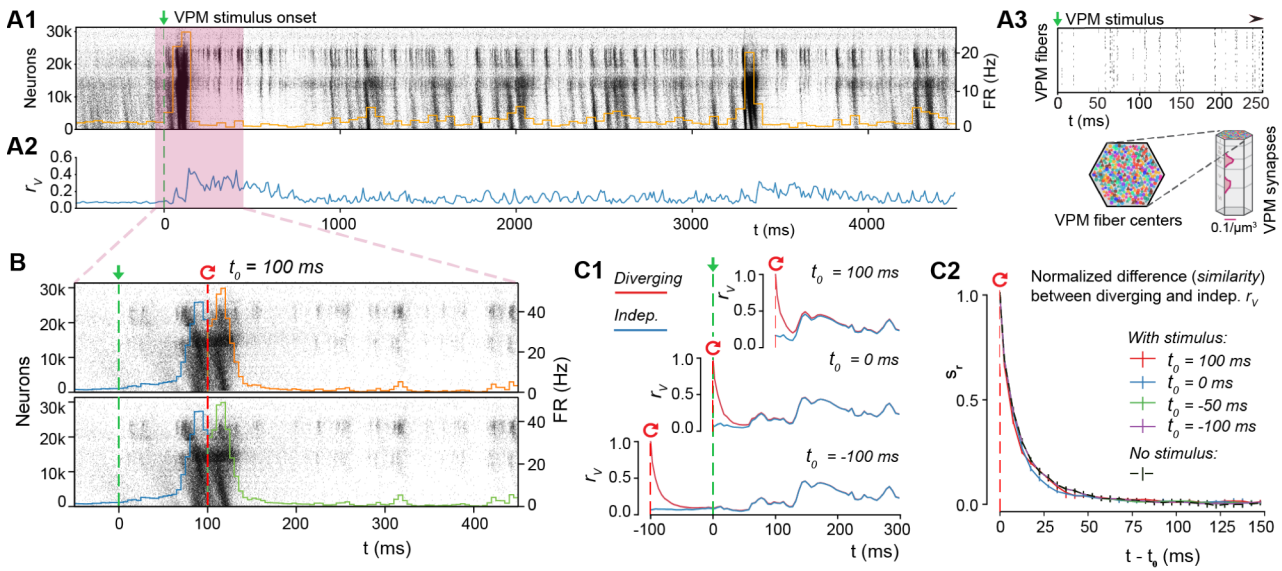

Figure 6: Rapid Divergence of Evoked, Reliable Activity

(A1) Population raster plot and population peristimulus time histogram (PSTH) for all 31'346 neurons in the microcircuit, during evoked activity with a thalamic (VPM) stimulus. Neurons are ordered according to cortical depth, with deep layers at the bottom and upper layers at the top, and each row representing the spikes of one neuron. For visibility, raster lines extend over dozens of rows for each neuron. (A2) Mean somatic membrane potential correlation r_V between independent simulations of the same VPM stimulus (mean ± 95% confidence interval). (A3) Schematic of the VPM stimulus. Top: Raster plot spike times for the first 250 ms of the thalamic stimulus. Bottom: 310 VPM fiber centers are assigned 30 colors, and those with identical colors are provided with duplicate spike trains. The synapse density profile across layers for each fiber is shown to the right. (B) For $t < 100$, the top and bottom raster plots show the same simulation, whereas for $t > 100$, the raster plots depict two resumed simulations starting from the same saved state at $t_0 = 100$, using different random number seeds. (C1) Resuming from identical initial conditions at different times: during (top), at onset (middle), or before the stimulus (bottom). Mean r_V between independent simulations (blue, as in A2), and mean r_V between simulations starting from the same base state (red; mean ± 95% confidence interval). (C2) The similarity, s_r,defined as the difference between the r_V of diverging and independent trials, normalized to lie between 1 (identical) and 0 (fully diverged) (mean ± 95% confidence interval). Means are based on 20 base states, *no stimulus* (spontaneous activity) on 40 as before.

Rapid Divergence of Evoked, Reliable Activity

In the NMC-model, thalamic inputs can evoke responses with varying degrees of reliability among trials (Markram et al., 2015; Reimann et al., 2017a). What then are the roles of synaptic noise and chaotic network dynamics during these evoked responses? To answer this question, we simulated electrical activity in response to a naturalistic thalamocortical stimulus (Figure 6A1), consisting of spike trains recorded in the ventral posteromedial nucleus (VPM) during sandpaper-induced whisker deflection in vivo (Bale et al., 2015). These spike trains were then applied to different feed-forward VPM fibers in the model to achieve a biologically-inspired, time-varying synchronicity among inputs (Figure 6A3; see Methods; Reimann et al., 2017b). To avoid introducing external variability on top of the intrinsically generated microcircuit variability, presynaptic inputs were kept identical across trials, but with thalamocortical synapses subject to the same synaptic noise as cortical synapses. The thalamocortical presynaptic inputs were not subject to recurrent network dynamics. Since this condition excludes variability in the system up to and including the thalamus, it can be considered an intermediate stage between the decoupled replay and regular network simulations. The simulations allowed us to identify an upper bound on the reliability of thalamocortical responses. Mean $r_V(t)$ during evoked activity was stronger than during spontaneous activity, moving between ~0.1 and ~0.4 (Figure 6A2), confirming that that the responses of neuron membrane potentials to the stimulus were relatively more reliable across trials.

To characterize the nature of chaotic network dynamics during this evoked, reliable activity, we again resumed from identical initial conditions, with t_0 at various times relative to the stimulus onset at t = 0 ms (Figure 6B, for t_0 = 100 ms). The population spiking activity across pairs of trials after resuming appeared almost identical, even for time intervals much larger than the divergence time characterized above (Figure 6B). At first glance, it would appear that the input had fully overcome the chaotic divergence. However, quantification of variability by time course of divergence of membrane potentials, $r_V(t)$, showed that it dropped rapidly towards the independent trial average (Figure 6C1, top). When we resumed from identical initial conditions at different times, for example at the onset of evoked activity (Figure 6C1, middle) or before onset (Figure 6C1, bottom), $r_V(t)$ dropped in the same way, subsequently converging to the average for independent trials. Indeed, $s_r(t)$, the normalized difference between the resumed and independent $r_V(t)$ showed a pattern of divergence remarkably similar to the divergence observed in simulations of spontaneous activity (Figure 6C2). Resuming from a base state at the peak of evoked activity, $s_{RMSD}(t)$ drops even faster (Figure S7A). A simpler stimulus, designed to imitate a whisker flick-type experiment (Markram et al., 2015), yielded comparable results (Figure S7B, S7C). Hence, any neuronal activity, whether spontaneous and unpredictable, or evoked and reliable, is ultimately constrained by similar chaotic network dynamics.

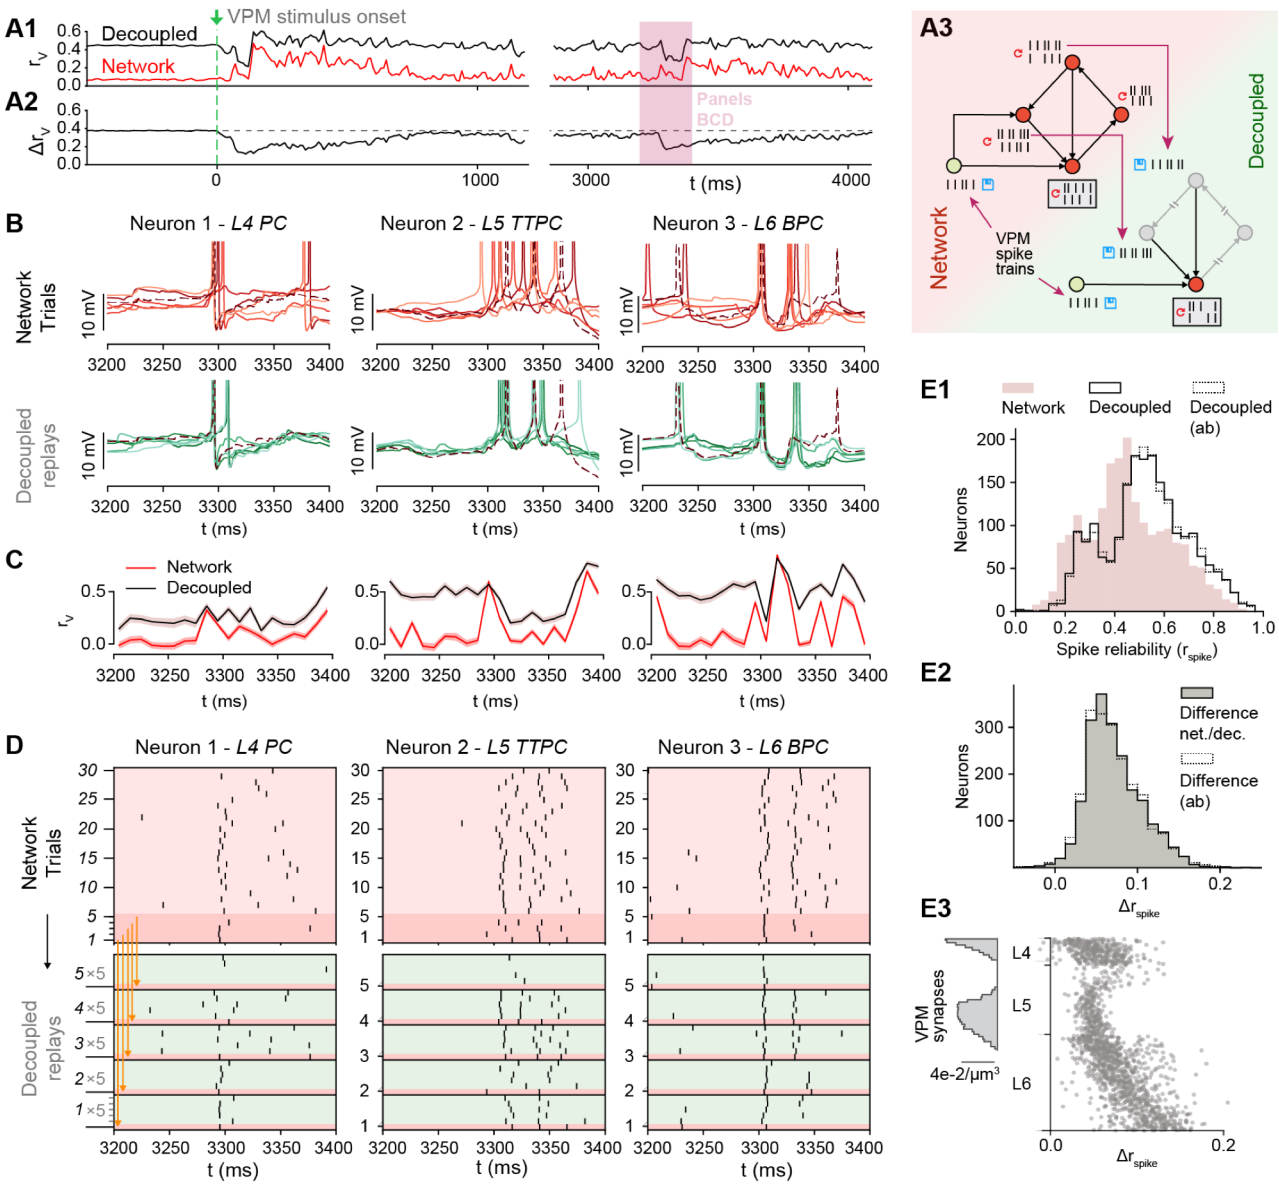

Figure 7: Spike-Timing Reliability Amid Noise and Chaos

(A1) Mean somatic membrane potential correlation, r_V, between independent simulations, and between decoupled replays of those simulations (network simulation identical to Figure 7A2). (A2) Difference in r_V for decoupled and network simulations. (A3) Schematic of network and decoupled replay simulation paradigms, including thalamic input. (B) Somatic membrane potentials (V_m) of three representative neurons for the time interval highlighted by the red box in A. Top: during six independent trials. Bottom: five decoupled replay trials (green) with the same presynaptic input as during the original network simulation trial (red), but with different random seeds. (C) Network and decoupled r_V as in A, but only for the three sample neurons in B. (D) Top: Raster plot of spike times of the same three example neurons as in B, during 30 independent trials of evoked activity. Bottom: Decoupled replay trials (green) of the same input received during 5 of the 30 original trials (dark red). (E1) Mean spike-timing reliability r_{spike} of 2024 pyramidal neurons from layers 4, 5, and 6 between independent network simulations, and between decoupled replay simulations with identical presynaptic inputs. (E2) Difference between r_{Spike} of decoupled and replayed simulations. (E3) Difference between r_{spike} of decoupled and replayed simulations versus position of somata across layers 4,5 and 6 of microcircuit (1675 neurons).

Spike-Timing Reliability Amid Noise and Chaos

At first glance, our observations of reliable population spike responses and chaotic divergence of membrane potentials seem to be mutually contradictory. Could it be that membrane potential reliability is simply not correlated with spike-timing reliability, as we observed for the case of spontaneous activity? To answer this question, we again compared network simulations with decoupled replay simulations with no network dynamics (Figure A3). As before, $r_V(t)$ was much larger in the decoupled simulations (Figure 7A1, black) than in the network simulations (Figure 7A1, red; same as Figure 6A2). However, the difference between the two was always smaller during evoked activity (Figure 7A2, after 0 ms) than during spontaneous activity (Figure 7A2, before 0 ms). This suggests that network dynamics play a reduced role in generating variability during evoked activity. When we focus on individual neurons (Figure 7B), we can see that that the difference between network and decoupled $r_V(t)$ at times collapses to zero (Figure 7C). In other words, variability due to network dynamics can intermittently be completely overcome for a sub-population of neurons in the network. Looking at the corresponding membrane potential traces, we observe that these moments occur during periods of reliable spiking (Figure 7B). This shows that in evoked activity, in contrast to spontaneous activity, moments of reliable membrane potentials can translate into reliable spiking at least for some neurons.

To get an idea of this effect at the population level, we compared spike time reliability r_{spike} with and without network dynamics (Figure 7D). We observed that removing network dynamics only moderately increased spike-timing reliability (Figure 7E1, red vs solid black line). In fact, increases in reliability were small for all neurons (Figure 7E2, solid black line). In stark contrast to the spontaneous case, a small population of neurons in the network simulations achieved spike reliabilities near unity (Figure 7E1). As expected, most of the noise effects could be explained by synaptic noise alone (Figure 7E1, 7E2, dotted black line). We conclude that external stimuli can sparsely and transiently overcome chaotic network dynamics for sub-populations of neurons, though with a substantial residual variability caused by synaptic noise (albeit much smaller than during spontaneous activity).

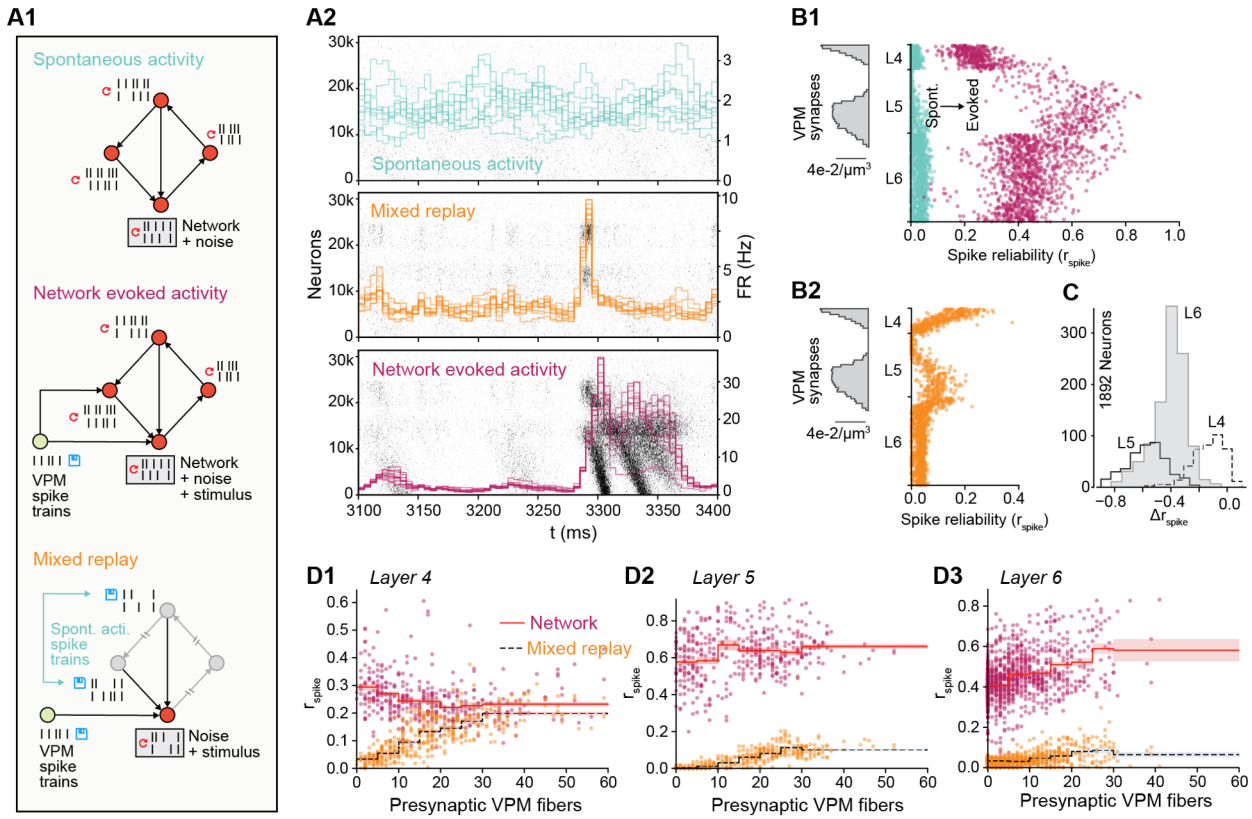

Figure 8: High Reliability Requires Recurrent Cortical Connectivity

(A1) Overview of three simulation paradigms: *spontaneous activity*, *network evoked activity* (with network propagation intact and VPM input), and *mixed replay* (with network propagation replaced by replays of spontaneous activity spike trains, and VPM input) (A2) Examples of population spiking activity during the three simulation paradigms. (B1) Spike-timing reliability, r_{spike}, during spontaneous (blue) and evoked (purple) activity for 1675 excitatory neurons in the center of layers 4, 5 and 6. (B2) Spike-timing reliability, r_{spike}, during a mixed replay with VPM input but with network propagation disabled for the same neurons as in B1. (C) Difference in r_{spike} between evoked activity with and without network propagation for 1892 excitatory neurons in the center of layers 4, 5 and 6 (same for D1-3). (D1) The number of presynaptic VPM fibers from which each neuron receives input versus r_{spike} in evoked simulations with (*network*) and without (*mixed replay*) network propagation.

High Reliability Requires Recurrent Cortical Connectivity

It is conceivable that the spike-timing reliability we observed could simply be a result of direct and feed-forward input from VPN (Wang et al., 2010). Indeed, when we look at changes in reliability without network dynamics, the strongest increase in reliability is in neurons at the bottom of layer six that receive comparatively little direct VPM input (Figure 7E3). To test whether the intermittent suppression of chaotic dynamics is simply an effect of the feed-forward input, we designed a new simulation paradigm similar to our previous decoupled replay, where each neuron received a combination of replayed presynaptic inputs from a simulation of spontaneous activity and from the direct feed-forward VPM input it received in the evoked network simulations (Figure 8A1). That is, each neuron receives input as in a spontaneous activity trial through its recurrent synaptic contacts, and input as in an evoked trial through its feed-forward synaptic contacts.

In this *mixed replay* paradigm, the population response was much weaker (Figure 8A2). While in simulations of evoked activity all neurons showed higher reliability than in simulations of spontaneous activity (Figure 8B1), in the mixed replay, the only cells that showed increased reliability were those close to the VPM synapses (Figure 8B2). Furthermore, the only neurons to display similar reliability, with and without recurrent network propagation, were a small group in layer 4 (Figure 8C). Taken together, these findings suggest that feed-forward VPM input alone is not enough to make the majority of neurons spike reliably.

To test this hypothesis, we compared the reliability between the two simulation paradigms to the number of presynaptic VPM fibers innervating each neuron (Figure 8D1-3). We can see that neurons in layer 4 that receive little direct VPM input responded more reliably with the network than neurons that receive a lot of VPM input with no network effect (Figure 8D1). Similarly, neurons in layers 5 and 6 were more reliable in mixed replay when they had more presynaptic VPM connections. However, this reliability increases drastically when network dynamics are enabled (Figure 8D2-3). We conclude that the reliable spiking observed in response to VPM inputs is enabled and propagated by recurrent cortical connectivity, and that this is true both for neurons that receive large direct VPM input, and for neurons that receive little or no such input. In brief, in spontaneous activity, recurrent connectivity amplifies variability; in the evoked state it amplifies reliability.

Discussion

This study provides a first estimate of the internally generated variability that arises directly from local neocortical circuitry, and the constraints this places on the reliability of signal encoding. Our simulation-based approach was enabled by the NMC-model, a biophysical model of neocortical microcircuitry in the somatosensory cortex of the two-week old rat, with biological noise sources calibrated to experimental data. Internally generated variability during spontaneous activity was quantified by the rate at which somatic membrane potentials probabilistically diverged from identical initial conditions, showing rapidly diverging activity, with a time constant on the order of 10-20 milliseconds (Figure 1). The time constant was approximately conserved across a range of dynamical states, from asynchronous to synchronous and regenerative (Figure 2). This suggests that all dynamical states have similar chaotic properties. Furthermore, internally generated variability was nearly saturated at the scale of the microcircuit, whose radial size just accommodates the full dendritic trees of the neurons at its center (Figure 2). Taken together these findings show that the cortex is in a constant state of perturbation due to intrinsic noise sources, and thus that cortical coding schemes need to be robust to such perturbations (London et al., 2010).

Our simulations show that the observed variability arises predominantly due to a combination of synaptic noise sources that are amplified by chaotic network dynamics (Figure 4). Thus, synaptic noise does not average out. On the contrary, it is a defining component of cortical network dynamics. Any network model studying variability, response reliability and coding must therefore take into account the role of synaptic noise. We predict that synaptic noise likely dominates over other cellular noise sources (Figure 4). This result is consistent with in vitro recordings of cortical neurons showing reliable responses to somatic current injections (Mainen and Sejnowski, 1995). However, it also shows that assessments of response reliability need to consider the impact of synaptic noise. Our results show, not only that synaptic noise drives chaotic network dynamics, but also that it causes high variability in spike times, even when network dynamics are disabled (Figure 5). However, we also found that a relatively weak thalamic stimulus can evoke reliable responses amid the noise and chaos (Figure 6) and that chaotic network dynamics are fully overcome during intermittent periods of reliable spiking (Figure 7). This last finding confirms predictions from previous computational and theoretical studies with simplified network models (Lajoie et al., 2013, 2016; Rajan et al., 2010). We observe, however, that synaptic noise always causes residual spike-timing variability. Interestingly, under external drive, a small fraction of neurons approach near perfect average reliability (Figure 7E1). This suggests that thalamocortical input can produce sustained reliability in a few select neurons.

There is strong experimental and theoretical evidence that cortical circuitry is both chaotic and sensitive to perturbations (London et al., 2010; van Vreeswijk and Sompolinsky, 1996). On the other hand, there are also many observations of cortical neurons responding with millisecond precision to repeated sensory inputs (Hires et al., 2015; Kayser et al., 2010). Some authors have

conjectured that this is simply an effect of direct input from the sensory periphery, for example the thalamus (Wang et al., 2010). We show, however, that these reliable events are not a simple consequence of feed-forward thalamic input suppressing recurrent activity, and that they require recurrent cortical connections (Figure 8). In fact, previous estimates of how many synchronous thalamic inputs are needed to achieve reliable responses in neurons in cat visual cortex (Wang et al., 2010) may well overestimate synaptic reliability: synaptic release probability is lower in vivo than in vitro, both in general (Borst, 2010) and in this specific pathway (Sedigh-Sarvestani et al., 2017). Thus, the same recurrent cortical architecture that produces chaotic dynamics also provides mechanisms to overcome chaos in response to the right stimuli. This result matches previous findings that patterns of activity generated by cortical circuitry in response to sensory stimuli have millisecond spike-timing precision (Luczak et al., 2015; Sato et al., 2012). We conclude that chaos and reliability are not mutually exclusive, but may, in fact be generated by the same underlying cortical architecture.

The Biological Details Matter

A wide range of cortical network models have been used to study variability, often with substantial differences in terms of modeling detail (Balaguer-Ballester, 2017). While these models can predict that a specific kind of network architecture leads to activity that resembles in vivo cortical activity, it is impossible to know whether the architecture is implemented in real neuronal circuitry. For example, cortical network models with dense, homogeneous connectivity exhibit asynchronous activity and negligible spike correlations due to their tightly coupled EI-balance (Renart et al., 2010). A more biologically accurate, heterogeneous pattern of connectivity breaks these dynamics. Including additional biological details such as spike-frequency adaptation can recover the original dynamics (Landau et al., 2016). Similarly, taking account of other biological details, such as distance dependent connectivity (Rosenbaum et al., 2017), strength of feedback inhibition (Stringer et al., 2016) or differences in synaptic time scales (Huang et al., 2017), can produce emergent, internally generated variability that is shared across neurons.

These findings suggest mechanisms that could generate variability. However, they also suggest that adding more detail or changing parameters would fundamentally alter the dynamics. In contrast, the details included in the NMC-model are constrained by experimental data. Thus, the NMC-model contains cellular noise sources, which are typically ignored by studies with simplified models. While it has already been suggested that stochastic vesicle release generates high neural variability in postsynaptic neurons (Moreno-Bote, 2014; Reich and Rosenbaum, 2013), and that synaptic noise can reduce neural correlations (Doiron et al., 2016; Rosenbaum et al., 2013), this study provides the first direct evidence that synaptic noise is a crucial component in shaping internally generated neural variability in cortex, both at the single neuron and the population level. We conclude that mechanistically understanding cortical variability requires a

holistic approach: leaving out biological details can lead to activity that matches in vivo activity, but which could, in reality, be generated by other mechanisms.

Missing Noise Sources Could Moderately Increase Internally Generated Variability

While the NMC-model is one of the most detailed and rigorously data-constrained models of neocortical circuitry to date, there are at least several potentially important features missing. First, while synaptic noise is modeled at every synapse, other noise sources intrinsic to the local circuitry are incomplete or absent. The most important neglected noise source is ion-channel noise, with other electrical noise sources, such as thermal noise, likely orders of magnitude smaller (Faisal et al., 2008). We showed that ion-channel noise in irregular firing neurons, which facilitates irregular spike-timing in response to constant current injections in vitro (Mendonça et al., 2016), is still overshadowed by synaptic noise under vivo-like conditions. Moreover, regular firing pyramidal cells respond very regularly to current injections in vitro (Mainen and Sejnowski, 1995), with irregular behavior only emerging in response to variable presynaptic inputs in vivo. We therefore conclude that ion-channel noise is overshadowed by synaptic noise in relation to the timing of action potential (AP) initiation.

What is the impact of ion-channel noise in axons and dendrites? In dendrites, ion-channel noise is thought to evoke little to no variability in isolated back propagating action potentials (Diba et al., 2006). Thus, mean ion-channel models are likely sufficient to accurately represent their electrical behavior. In axons, there is evidence that APs generated by neocortical pyramidal neurons reliably invade axonal arbors without failures (Cox et al., 2000). As they propagate along axons, however, the timing of APs becomes increasingly variable. In fact, simulations predict that ion-channel noise affects AP timing in all axons with a diameter below 0.05 μm, with the standard deviation of AP variability predicted to increase by 0.6 ms per 2 mm in 0.02 μm axons (Faisal and Laughlin, 2007). In the NMC-model, axons have a mean axonal diameter of around 0.03 μm and are modeled deterministically. Potentially, therefore, some longer axons could cause increases in variability of up to several milliseconds.

Another unknown noise source is local neuromodulation. Some cholinergic neurons in cortex directly modulate the firing rates of local neurons by increasing spontaneous vesicle release (Engelhardt et al., 2007). Little is known about whether this leads to an increase or a decrease in internally generated variability. However, the time scale for variability generated in this way is probably much slower than for synaptic noise. In Figure 4, we showed that any additional noise sources would have to increase variability beyond 1 mV to be relevant. We therefore predict that, if at all, missing intrinsic cellular noise sources could only moderately increase the divergence rate beyond synaptic noise.

More Reliable Synapses Could Moderately Decrease Internally Generated Variability

What about synaptic noise? The reliability of synaptic transmission increases with the number of readily releasable vesicles (Rudolph et al., 2015). While some cortical connection types are probably univesicular (Silver et al., 2003), others could have up to ten releasable vesicles per synapse (Loebel et al., 2009). The current version of the NMC-model assumes one readily releasable vesicle per synapse, thus potentially underestimating synaptic reliability. To estimate the potential impact, we repeated the simulation experiments with an increasing number of readily releasable vesicles (n_{rrp}) at all synapses (Figure S8). As expected, the time course of divergence slowed down with increasing n_{rrp}. Nonetheless, for realistic mean n_{rrp} values ($n_{rrp} = 2 - 3$; data not shown), it was still dominated by synaptic noise. More importantly, n_{rrp} would likely vary across synapses, with small (n_{rrp} = 1) numbers of vesicles for some synapses and larger numbers for others. However, the exact effect of biologically-realistic n_{rrp} distributions remains to be studied. The robustness of divergence to synapse reliability is also evidenced by the robustness to extracellular calcium concentrations modulating release probabilities across the different dynamical states.

Gap junctions, another type of intercellular connection that is not included in the NMC-model, could also change the electrical behavior of inhibitory cells, for example by causing them to synchronize their irregular firing (Amsalem et al., 2016; Mendonça et al., 2016). However, gap junctions are predominantly formed between certain types of inhibitory neurons, and their effects would likely tend to turn individual variability into shared variability, rather than remove it.

Reliable Coding Amid Synaptically Driven Chaos

We have demonstrated that chaotic network dynamics, and reliable encoding and propagation of signals with millisecond spike-timing precision can both be supported by the same cortical architecture and are not mutually exclusive. We have also shown that packets of reliable activity (Luczak et al., 2015) rely on amplification by the same cortical circuitry that also induces chaotic network dynamics. In the mouse visual cortex, neurons that receive input from the same thalamocortical axons are also more likely to be connected to each other (Morgenstern et al., 2016). This provides a possible mechanism for generating packets, which the intrinsic properties of synapses, neurons and connectivity allow to propagate, changing amplification of noise into amplification of reliability. The exact mechanism for this reversal, and the means by which signals are reliably propagated through the circuitry remains a subject for future investigation. One possible explanation is that certain connectivity motifs could amplify reliability through redundant connectivity. Candidate motifs have already been identified in the NMC-model, such as common neighbor motifs (Perin et al., 2011) and high-dimensional cliques (Reimann et al., 2017a). Dendritic nonlinearities, such as N-Methyl-D-aspartate (NMDA)-mediated plateau potentials evoked by clustered synaptic inputs onto the dendritic tree could further play an important role (Wilson et al., 2016).

This study provides, for the first time, a conceptual framework to quantify these transitions between variable and reliable responses, and thereby a solid foundation on which biological theories of cortical coding can be evaluated.

Methods

Simulation

Model of Neocortical Microcircuitry (NMC)

Simulations of electrical activity were performed on a previously published model of a neocortical microcircuit in two-week old rat. Reconstruction and simulation methods are described extensively in Markram et al. (2015). In our study, we used a microcircuit consisting of 31,346 biophysical Hodgkin-Huxley NEURON models and around 7.8 million connections forming roughly 36.4 million synapses. Synaptic connectivity between 55 distinct morphological types of neurons (*m-types*) was predicted algorithmically and constrained by experimental data (Reimann et al., 2015). The densities of ion-channels on morphologically-detailed neuron models were optimized to reproduce the behavior of different electrical neuron types (*e-types*) as recorded in vitro (Van Geit et al., 2016). We also used a larger mesocircuit comprising seven microcircuits (mean of 36.5 million synapses per circuit), with no boundaries between the peripheral circuits and the original microcircuit in the center (only shown in Figure 2B). Simulations were run on a BlueGene/Q supercomputer (BlueBrain IV). NEURON models and the connectome are available online at bbp.epfl.ch/nmc-portal (Ramaswamy et al., 2015).

Simulation of In Vivo-Like Spontaneous Activity

In the in vivo-like state, release probabilities for all synapses were modulated according to the extracellular calcium concentration found in vivo, leading to substantially lower reliability than in vitro (Borst, 2010). As described in Markram et al. (2015), the u_{SE} parameter for synaptic transmission was modulated differentially as a function of extracellular calcium concentration ($[Ca^{2+}]_O$), allowing transitions from in vitro to in vivo-like dynamics. Neurons were depolarized with a somatic current injection, with currents expressed as a percent of first spike threshold for each neuron, to mimic, for example, the effect of depolarization due to missing neuromodulators. Apart from a small white-noise component (with a variance of 0.001% of the mean injected current per neuron, unless stated otherwise), the current injection was constant. With mean injected currents at around 100% of first spike threshold and $[Ca^{2+}]_O$ at 1.25 mM, the microcircuit exhibits in vivo-like spontaneous activity (Markram et al., 2015).

Simulation of Evoked Activity

The microcircuit is innervated by 310 (virtual) thalamic fibers (Markram et al., 2015). In vivo spike train recordings from 30 VPM neurons were randomly assigned to the 310 fibers, to

achieve varying degrees of naturalistic synchronous thalamic inputs. Spike trains were recorded during replayed texture-induced whisker motion in anesthetized rats (Bale et al., 2015). Full methods are described in (Reimann et al., 2017a). The second stimulus consisted of synchronous spikes at the 60 central thalamic fibers, with all 60 virtual thalamic neurons firing simultaneously, to approximate a whisker 'flick' (see Markram et al., 2015).

Save-Resume

After running a simulation for some amount of biological time, the final states of all variables in the system were written to disk using NEURON's *SaveState* class. For large-scale simulations, this required the various processes to coordinate how much data each needed to write, so that each rank could then seek the appropriate file offset and together write in parallel without interfering with the others. After restoring a simulation, the user could specify new random seeds (see below).

Random Numbers

In our simulations, we used random number generators (RNGs) to model all stochastic processes: noisy current injection, stochastic ion channels, probabilistic release of neurotransmitters and generation of spontaneous release events. Each synapse had two RNGs. One was used to determine vesicle release on the arrival of an action potential. The other determined the spontaneous release signal. Similarly, each stochastic K^+-channel model had a RNG determining voltage-dependent opening and closing times. Finally, the white noise process underlying the noisy depolarization was determined by one RNG per neuron. By using different random seeds to initialize the RNGs, we obtained different sequences of random numbers, and consequently different but equally valid simulation outcomes. In earlier versions of the NEURON microcircuit simulation software, the user was given only a single random seed parameter with which to alter the random number streams generated by all RNGs. We added the option to separately change random seeds for RNGs for a specific type of stochastic component. For example, "IonChannelSeed <value>" allows the specification of a seed which is only given to the RNGs used by ion channel instances.

Stochastic Ion-Channels

In some interneuron models, a potassium channel type with a stochastic implementation was added using methods described in (Diba et al., 2006) and (Mendonça et al., 2016). This made it possible to model ion channel noise. Instead of a mean field model, the equations used explicitly track the number of channels in a certain state and allow these numbers to evolve stochastically. When the seed of the random process changes, the small fluctuation caused by the channel noise change with it, but the mean behavior of the ion channel remains the same.

Stochastic Synapses

The synapse models are described in full detail by Markram et al. (2015). Each synapse has one RNG to determine vesicle release upon action potential arrival at the synapse. A second RNG is used to determine the signal for spontaneous miniature post synaptic potentials. When the synapse receives the signal for a spontaneous release event, it is treated as a presynaptic action potential. Therefore, changing the random seed for the minis will eventually change the random number stream for vesicle release for presynaptic spikes, and therefore lead to a "pseudo-deterministic" synapse.

Multivesicular Release

The synapse model used in this study (see Markram et al., 2015) supports multivesicular release (MVR): each release event activates a fraction of the maximal postsynaptic conductance (g_{max}) proportional to the size of the readily-releasable pool of vesicles (n_{rrp}). In the univesicular case ($n_{rrp} = 1$), the release of one vesicle is sufficient to completely activate the postsynaptic conductance. However, when $n_{rrp} > 1$, full activation requires the release of all available presynaptic resources. This allowed us to independently control the *mean* postsynaptic response to synapse activation (which depends on g_{max}, but not n_{rrp}) and its *instantaneous* profile (where n_{rrp} matters).

Deterministic Synapse Model

In the deterministic synapse model, the u_{SE} variable is interpreted as the fraction of consumed resources, rather than a release probability. That is, each release event activates a fraction of postsynaptic conductance proportional to u_{SE}. For this reason, *DetAMPANMDA* and *DetGABAAB* are identical to their stochastic (multivesicular) counterparts in the limit as $n_{rrp} \rightarrow \infty$.

Single Spike Injection

We injected single spikes in twenty different layer 4 pyramidal neurons (and twenty random neurons across the circuit, data not shown) by *replaying* (see below) an additional spike event in one neuron per simulation. Thus, there were no shifted or missing spikes, as may occur when injecting a spike in vivo. The spike was injected *0.1 ms* after resuming the simulation from identical initial conditions.

Step-Pulse Perturbation

We applied a microscopic current step-pulse to all neurons at their soma *0.1 ms* after resuming the simulation (duration: 0.1 ms, amplitude: *1* pA,). The current was chosen to have an almost negligible effect on individual neurons, and was near the limit of the NEURON integrator. On average, 108 ± 8 neurons out of 31,346 neurons had any changes in their spike times (mean of 19 trials ± STD). The majority of the shifted spikes were shifted by less than 0.05 ms (59.1%: < 0.05; 33.1%: < 1 ms.; 5.5%: < 20 ms; 1.8%: < 100 s; 0.5%: < 1 s). Finally, 3 ± 2 neurons had extra or missing spikes. The median first occurrence of an extra or missing spike was at 257 ms (min: 11 ms, max: 946 ms after resuming).

Decoupled Replay

When resuming a simulation at t_0, we decoupled all connections by setting the connection weights to zero, ensuring that APs would be delivered to the synapses of postsynaptic neurons. At the same time, we started *replaying* AP times from a previous resumed simulation, activating the synapses of postsynaptic neurons as if the presynaptic neuron had fired an AP, but actually replaying presynaptic APs from the previous simulation. For computational reasons, spikes that had not been delivered at the save time t_0 were not delivered in the decoupled replay (meaning that a couple of presynaptic spikes per neuron may have been lost).

Analysis

RMSD and Correlation

All analysis was performed using custom scripts written in Python 2.7 using the *NumPy, matplolib* and *SciPy* libraries. Scripts were executed on a Linux cluster connected to the same IBM GPFS file system that the simulation output was written to. Root-mean-square deviation $RMSD_V$ and correlation r_V as defined in Equations 1 and 2 were implemented with *NumPy*.

Similarity

The similarity measure $s(t)$ was defined as the normalized difference between diverging $r_V(t)$ (or $RMSD_V(t)$), and steady-state $r_V(t)$ (or $RMSD_V(t)$). The steady-state value was defined, as the continuous $r_{V,shuffle}(t)$ computed by shuffling the soma voltages between simulation trials, so that instead of 40 deviating pairs of trajectories, we compared 40 independent pairs of trajectories. As an alternative, we defined it as the mean stationary, fully deviated r_∞ for $t > 1000$ *ms* after resuming from identical initial conditions.

Firing Rate

Firing rate was defined as the average number of spikes in a time interval of size Δt, divided by Δt ($\Delta t = 10\ ms$, unless stated otherwise).

Neuron Selection

We selected all excitatory neurons in layers 4, 5 and 6 that belonged to the 30 *minicolumns* (out 310 in total) in the center of microcircuit (n = 2024). The analysis was restricted to neurons that spiked at least once in each of the compared simulation paradigms.

Spike-Timing Reliability

Spike-timing reliability was measured using a correlation-based measure first proposed by Schreiber et al. (2003). Briefly, the spike times of each neuron in each trial were convolved with a Gaussian kernel of width $\sigma_s = 5\ ms$ to yield filtered signals $s(n, k; t)$ for each neuron n and each trial k ($\Delta t_s = 1\ ms$). The spike-timing reliability for each neuron was then defined as the mean inner product between pairs of signals divided by their magnitude: $r_{spike}(n) = \frac{2}{K(K-1)} \sum_{k \neq l} \frac{s(n,k;t) \cdot s(n,l;t)}{|s(n,k;t)| \cdot |s(n,l;t)|}$, (K = 30; independent trials). Decoupled replay: there are M=5 replays of each of the K=30 trials, and thus $r_{spike}(n) = \frac{2}{KM(M-1)} \sum_m \sum_{k \neq l} \frac{s_m(n,k;t) \cdot s_m(n,l;t)}{|s_m(n,k;t)| \cdot |s_m(n,l;t)|}$.

Errors and Statistical Tests

Error bars and shared areas indicate 95%-confidence intervals (CI), unless stated otherwise. t-based CIs (n = 20; or n = 40 if stated) were computed using *scipy.stats.sem* and *scipy.stats.t.ppf* to compute P-values from the CIs. Errors for fit parameters, obtained with *scipy.optimize.curve_fit*, are given as the square-root of the variance of the parameter estimate.

Author Contributions

Conceptualization, M.N., M.R., H.M., E.M.; Methodology, M.N., M.R., H.M., E.M.; Software, M.N., J.K.; Validation, M.N., J.K.; Investigation, M.N; Visualization, M.N.; Writing – Original Draft, M.N., M.R.; Writing – Review & Editing, M.N., M.R., H.M., E.M.; Supervision, H.M., E.M.; Project Administration, E.M.; Funding Acquisition, H.M.

Acknowledgements

We thank Giuseppe Chindemi and Srikanth Ramaswamy for help with synapse models, Werner Van Geit for help with ion channels, and the rest of the Blue Brain team for developing and

maintaining the microcircuit model and computational infrastructure. We thank Madineh Sedigh-Sarvestani and Taylor Newton for discussions and critical comments on the manuscript. We thank Oren Amsalem, Mickey London and Idan Segev for helpful discussions. We thank Richard Walker for feedback on the manuscript. This work was supported by funding from the ETH Domain for the Blue Brain Project. The Blue Brain Project's IBM BlueGene/Q system, BlueBrain IV, was funded by the ETH Board and hosted at the Swiss National Supercomputing Center (CSCS).

References

Amsalem, O., Van Geit, W., Muller, E., Markram, H., and Segev, I. (2016). From Neuron Biophysics to Orientation Selectivity in Electrically Coupled Networks of Neocortical L2/3 Large Basket Cells. Cereb. Cortex *26*, 3655–3668.

Balaguer-Ballester, E. (2017). Cortical Variability and Challenges for Modeling Approaches. Front. Syst. Neurosci. *11*.

Bale, M.R., Ince, R.A.A., Santagata, G., and Petersen, R.S. (2015). Efficient population coding of naturalistic whisker motion in the ventro-posterior medial thalamus based on precise spike timing. Front. Neural Circuits 50.

Beggs, J.M., and Plenz, D. (2003). Neuronal Avalanches in Neocortical Circuits. J. Neurosci. *23*, 11167–11177.

Borst, J.G.G. (2010). The low synaptic release probability in vivo. Trends Neurosci. *33*, 259–266.

Brinkman, B.A.W., Weber, A.I., Rieke, F., and Shea-Brown, E. (2016). How Do Efficient Coding Strategies Depend on Origins of Noise in Neural Circuits? PLOS Comput. Biol. *12*, e1005150.

Brunel, N. (2000). Dynamics of Sparsely Connected Networks of Excitatory and Inhibitory Spiking Neurons. J. Comput. Neurosci. *8*, 183–208.

Constantinople, C.M., and Bruno, R.M. (2011). Effects and Mechanisms of Wakefulness on Local Cortical Networks. Neuron *69*, 1061–1068.

Cox, C.L., Denk, W., Tank, D.W., and Svoboda, K. (2000). Action potentials reliably invade axonal arbors of rat neocortical neurons. Proc. Natl. Acad. Sci. *97*, 9724–9728.

Denève, S., and Machens, C.K. (2016). Efficient codes and balanced networks. Nat. Neurosci. *19*, 375–382.

Diba, K., Koch, C., and Segev, I. (2006). Spike propagation in dendrites with stochastic ion channels. J. Comput. Neurosci. *20*, 77–84.

Doiron, B., Litwin-Kumar, A., Rosenbaum, R., Ocker, G.K., and Josić, K. (2016). The mechanics of state-dependent neural correlations. Nat. Neurosci. *19*, 383–393.

Engelhardt, J. von, Eliava, M., Meyer, A.H., Rozov, A., and Monyer, H. (2007). Functional Characterization of Intrinsic Cholinergic Interneurons in the Cortex. J. Neurosci. *27*, 5633–5642.

Faisal, A.A., and Laughlin, S.B. (2007). Stochastic Simulations on the Reliability of Action Potential Propagation in Thin Axons. PLOS Comput. Biol. *3*, e79.

Faisal, A.A., Selen, L.P.J., and Wolpert, D.M. (2008). Noise in the nervous system. Nat. Rev. Neurosci. *9*, 292–303.

Gal, E., London, M., Globerson, A., Ramaswamy, S., Reimann, M.W., Muller, E., Markram, H., and Segev, I. (2017). Rich cell-type-specific network topology in neocortical microcircuitry. Nat. Neurosci. *advance online publication*.

Gütig, R., and Sompolinsky, H. (2006). The tempotron: a neuron that learns spike timing–based decisions. Nat. Neurosci. *9*, 420–428.

Hires, S.A., Gutnisky, D.A., Yu, J., O'Connor, D.H., and Svoboda, K. (2015). Low-noise encoding of active touch by layer 4 in the somatosensory cortex. ELife *4*, e06619.

Huang, C., Ruff, D., Pyle, R., Rosenbaum, R., Cohen, M., and Doiron, B.D. (2017). Circuit-based models of shared variability in cortical networks. BioRxiv 217976.

Kayser, C., Logothetis, N.K., and Panzeri, S. (2010). Millisecond encoding precision of auditory cortex neurons. Proc. Natl. Acad. Sci. *107*, 16976–16981.

Lajoie, G., Lin, K.K., and Shea-Brown, E. (2013). Chaos and reliability in balanced spiking networks with temporal drive. Phys. Rev. E *87*, 052901.

Lajoie, G., Lin, K.K., Thivierge, J.-P., and Shea-Brown, E. (2016). Encoding in Balanced Networks: Revisiting Spike Patterns and Chaos in Stimulus-Driven Systems. PLOS Comput. Biol. *12*, e1005258.

Landau, I.D., Egger, R., Dercksen, V.J., Oberlaender, M., and Sompolinsky, H. (2016). The Impact of Structural Heterogeneity on Excitation-Inhibition Balance in Cortical Networks. Neuron *92*, 1106–1121.

Loebel, A., Silberberg, G., Helbig, D., Markram, H., Tsodyks, M., and Richardson, M. (2009). Multiquantal release underlies the distribution of synaptic efficacies in the neocortex. Front. Comput. Neurosci. *3*.

London, M., Roth, A., Beeren, L., Häusser, M., and Latham, P.E. (2010). Sensitivity to perturbations in vivo implies high noise and suggests rate coding in cortex. Nature *466*, 123–127.

Luczak, A., McNaughton, B.L., and Harris, K.D. (2015). Packet-based communication in the cortex. Nat. Rev. Neurosci. *16*, 745–755.

Mainen, Z.F., and Sejnowski, T.J. (1995). Reliability of spike timing in neocortical neurons. Science *268*, 1503–1506.

Markram, H., Muller, E., Ramaswamy, S., Reimann, M.W., Sanchez, C., Ailamaki, A., Alonso-Nanclares, L., Antille, N., Arsever, S., and Kahou, G. (2015). Reconstruction and Simulation of Neocortical Microcircuitry. Cell.

Markram H, Lübke J, Frotscher M, Roth A, and Sakmann B (1997). Physiology and anatomy of synaptic connections between thick tufted pyramidal neurones in the developing rat neocortex. J. Physiol. *500*, 409–440.

Masquelier, T. (2013). Neural variability, or lack thereof. Front. Comput. Neurosci. *7*.

Mendonça, P.R., Vargas-Caballero, M., Erdélyi, F., Szabó, G., Paulsen, O., and Robinson, H.P. (2016). Stochastic and deterministic dynamics of intrinsically irregular firing in cortical inhibitory interneurons. ELife *5*, e16475.

Mohajerani, M.H., Chan, A.W., Mohsenvand, M., LeDue, J., Liu, R., McVea, D.A., Boyd, J.D., Wang, Y.T., Reimers, M., and Murphy, T.H. (2013). Spontaneous cortical activity alternates between motifs defined by regional axonal projections. Nat. Neurosci. *16*, 1426–1435.

Moreno-Bote, R. (2014). Poisson-Like Spiking in Circuits with Probabilistic Synapses. PLOS Comput. Biol. *10*, e1003522.

Morgenstern, N.A., Bourg, J., and Petreanu, L. (2016). Multilaminar networks of cortical neurons integrate common inputs from sensory thalamus. Nat. Neurosci. *19*, 1034–1040.

Niell, C.M., and Stryker, M.P. (2010). Modulation of Visual Responses by Behavioral State in Mouse Visual Cortex. Neuron *65*, 472–479.

Okun, M., Steinmetz, N.A., Cossell, L., Iacaruso, M.F., Ko, H., Barthó, P., Moore, T., Hofer, S.B., Mrsic-Flogel, T.D., Carandini, M., et al. (2015). Diverse coupling of neurons to populations in sensory cortex. Nature *521*, 511–515.

Orbán, G., Berkes, P., Fiser, J., and Lengyel, M. (2016). Neural Variability and Sampling-Based Probabilistic Representations in the Visual Cortex. Neuron *92*, 530–543.

Pachitariu, M., Lyamzin, D.R., Sahani, M., and Lesica, N.A. (2015). State-Dependent Population Coding in Primary Auditory Cortex. J. Neurosci. *35*, 2058–2073.

Perin, R., Berger, T.K., and Markram, H. (2011). A synaptic organizing principle for cortical neuronal groups. Proc. Natl. Acad. Sci. *108*, 5419–5424.

Petersen, R.S., Panzeri, S., and Diamond, M.E. (2001). Population Coding of Stimulus Location in Rat Somatosensory Cortex. Neuron *32*, 503–514.

Petilla Interneuron Nomenclature Group, Ascoli, G.A., Alonso-Nanclares, L., Anderson, S.A., Barrionuevo, G., Benavides-Piccione, R., Burkhalter, A., Buzsáki, G., Cauli, B., Defelipe, J., et al. (2008). Petilla terminology: nomenclature of features of GABAergic interneurons of the cerebral cortex. Nat. Rev. Neurosci. *9*, 557–568.

Priesemann, V., Wibral, M., Valderrama, M., Pröpper, R., Le Van Quyen, M., Geisel, T., Triesch, J., Nikolić, D., and Munk, M.H.J. (2014). Spike avalanches in vivo suggest a driven, slightly subcritical brain state. Front. Syst. Neurosci. *8*.

Rajan, K., Abbott, L.F., and Sompolinsky, H. (2010). Stimulus-dependent suppression of chaos in recurrent neural networks. Phys. Rev. E *82*, 011903.

Ramaswamy, S., Courcol, J.-D., Abdellah, M., Adaszewski, S.R., Antille, N., Arsever, S., Atenekeng, G., Bilgili, A., Brukau, Y., Chalimourda, A., et al. (2015). The neocortical microcircuit collaboration portal: a resource for rat somatosensory cortex. Front. Neural Circuits 44.

Reich, S., and Rosenbaum, R. (2013). The impact of short term synaptic depression and stochastic vesicle dynamics on neuronal variability. J. Comput. Neurosci. *35*, 39–53.

Reimann, M.W., King, J.G., Muller, E.B., Ramaswamy, S., and Markram, H. (2015). An algorithm to predict the connectome of neural microcircuits. Front. Comput. Neurosci. 120.

Reimann, M.W., Nolte, M., Scolamiero, M., Turner, K., Perin, R., Chindemi, G., Dłotko, P., Levi, R., Hess, K., and Markram, H. (2017a). Cliques of Neurons Bound into Cavities Provide a Missing Link between Structure and Function. Front. Comput. Neurosci. *11*.

Reimann, M.W., Horlemann, A.-L., Ramaswamy, S., Muller, E.B., and Markram, H. (2017b). Morphological Diversity Strongly Constrains Synaptic Connectivity and Plasticity. Cereb. Cortex *27*, 4570–4585.

Renart, A., and Machens, C.K. (2014). Variability in neural activity and behavior. Curr. Opin. Neurobiol. *25*, 211–220.

Renart, A., Rocha, J. de la, Bartho, P., Hollender, L., Parga, N., Reyes, A., and Harris, K.D. (2010). The Asynchronous State in Cortical Circuits. Science *327*, 587–590.

Ribrault, C., Sekimoto, K., and Triller, A. (2011). From the stochasticity of molecular processes to the variability of synaptic transmission. Nat. Rev. Neurosci. *12*, 375–387.

Rosenbaum, R., Rubin, J.E., and Doiron, B. (2013). Short-term synaptic depression and stochastic vesicle dynamics reduce and shape neuronal correlations. J. Neurophysiol. *109*, 475–484.

Rosenbaum, R., Smith, M.A., Kohn, A., Rubin, J.E., and Doiron, B. (2017). The spatial structure of correlated neuronal variability. Nat. Neurosci. *20*, 107–114.

Rudolph, S., Tsai, M.-C., von Gersdorff, H., and Wadiche, J.I. (2015). The ubiquitous nature of multivesicular release. Trends Neurosci. *38*, 428–438.

Sato, T.K., Nauhaus, I., and Carandini, M. (2012). Traveling Waves in Visual Cortex. Neuron *75*, 218–229.

Schreiber, S., Fellous, J.M., Whitmer, D., Tiesinga, P., and Sejnowski, T.J. (2003). A new correlation-based measure of spike timing reliability. Neurocomputing *52–54*, 925–931.

Sedigh-Sarvestani, M., Vigeland, L., Fernandez-Lamo, I., Taylor, M.M., Palmer, L.A., and Contreras, D. (2017). Intracellular, In Vivo, Dynamics of Thalamocortical Synapses in Visual Cortex. J. Neurosci. *37*, 5250–5262.

Shadlen, M.N., and Newsome, W.T. (1998). The Variable Discharge of Cortical Neurons: Implications for Connectivity, Computation, and Information Coding. J. Neurosci. *18*, 3870–3896.

Silver, R.A., Lübke, J., Sakmann, B., and Feldmeyer, D. (2003). High-Probability Uniquantal Transmission at Excitatory Synapses in Barrel Cortex. Science *302*, 1981–1984.

Singh, C., and Levy, W.B. (2017). A consensus layer V pyramidal neuron can sustain interpulse-interval coding. PLOS ONE *12*, e0180839.

Stern, E.A., Kincaid, A.E., and Wilson, C.J. (1997). Spontaneous Subthreshold Membrane Potential Fluctuations and Action Potential Variability of Rat Corticostriatal and Striatal Neurons In Vivo. J. Neurophysiol. *77*, 1697–1715.

Stringer, C., Pachitariu, M., Steinmetz, N.A., Okun, M., Bartho, P., Harris, K.D., Sahani, M., and Lesica, N.A. (2016). Inhibitory control of correlated intrinsic variability in cortical networks. ELife *5*, e19695.

Thorpe, S., Delorme, A., and Van Rullen, R. (2001). Spike-based strategies for rapid processing. Neural Netw. *14*, 715–725.

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

Tsodyks, M.V., and Markram, H. (1997). The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. Proc. Natl. Acad. Sci. *94*, 719–723.

Van Geit, W., Gevaert, M., Chindemi, G., Rössert, C., Courcol, J.-D., Muller, E.B., Schürmann, F., Segev, I., and Markram, H. (2016). BluePyOpt: Leveraging Open Source Software and Cloud Infrastructure to Optimise Model Parameters in Neuroscience. Front. Neuroinformatics *10*.

van Vreeswijk, C., and Sompolinsky, H. (1996). Chaos in Neuronal Networks with Balanced Excitatory and Inhibitory Activity. Science *274*, 1724–1726.

van Vreeswijk, C., and Sompolinsky, H. (1998). Chaotic Balanced State in a Model of Cortical Circuits. Neural Comput. *10*, 1321–1371.

Wang, H.-P., Spencer, D., Fellous, J.-M., and Sejnowski, T.J. (2010). Synchrony of Thalamocortical Inputs Maximizes Cortical Reliability. Science *328*, 106–109.

Wilson, D.E., Whitney, D.E., Scholl, B., and Fitzpatrick, D. (2016). Orientation selectivity and the functional clustering of synaptic inputs in primary visual cortex. Nat. Neurosci. *19*, 1003–1009.

Supplementary Figures (S1 –S8)

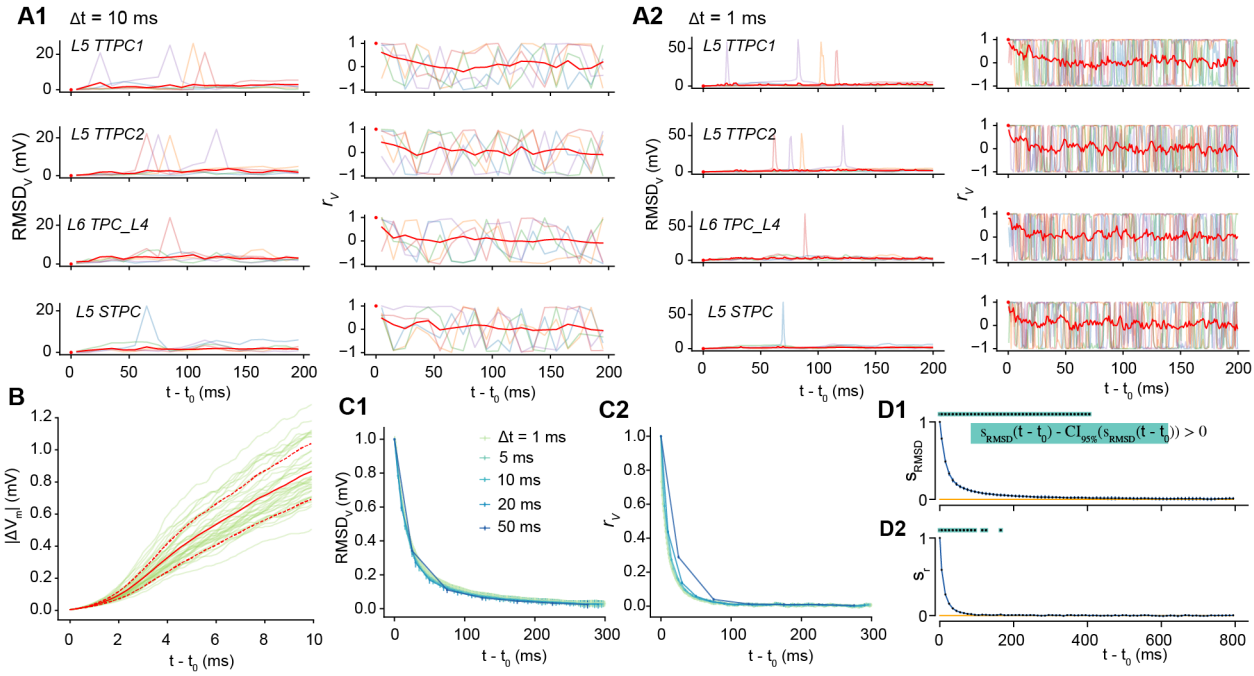

Figure S1: Quantifying the Rapid Divergence of Electrical Activity

(A1) Root-mean square deviation ($RMSD_V$) and correlation (r_V) of the somatic membrane potentials between pairs of resumed simulations diverging from identical conditions, for five different base states (faded colors) and the mean of 40 saved base states (red), with $\Delta t = 10\ ms$. Same neurons as in Figure 1C. (A2) Same as A1, but with $\Delta t = 1ms$. (B) Mean divergence in the first 10 ms, with $\Delta t_V = 0.1\ ms$ (mean of all neurons and 40 saved base states ± standard deviation). (C) $RMSD_V$ and r_V for different analysis bin sizes Δt. The time step for the soma voltage is $\Delta t_V = 0.1\ ms$. (D) The similarity (s_{RMSD} and s_r) (mean ± 95% confidence interval). Dots signal where s_{RMSD} and s_r are larger than 0, by a 95% confidence interval ($p < 0.025$).

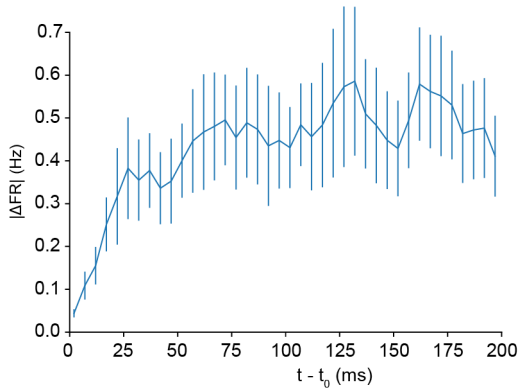

Figure S2: Rapid Divergence of Population Firing Rate

Mean population firing rate difference ($\Delta t = 5\ ms$) between pairs of simulations diverging from identical initial conditions (mean of all neurons and of 40 saved base states ± 95% confidence interval).

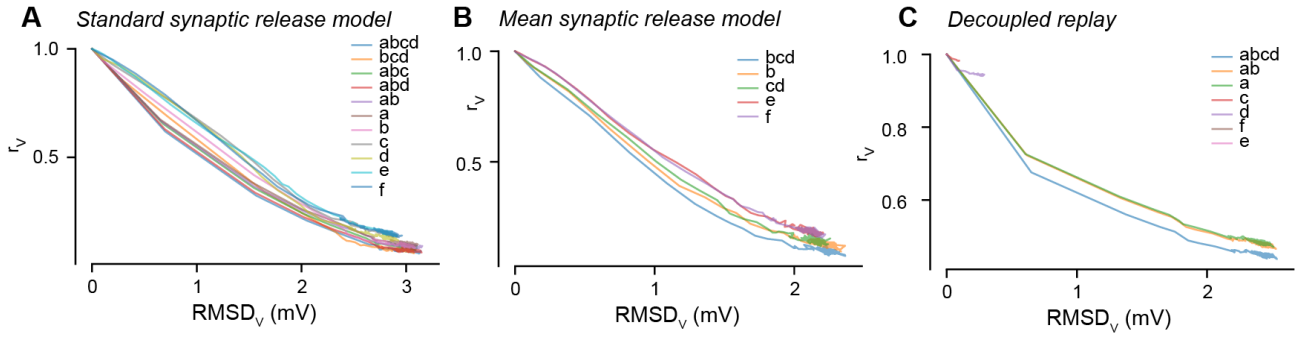

Figure S3: Linear Relationship between $RMSD_V$ and r_V

Root-mean square deviation ($RMSD_V$) and correlation (r_V) of the somatic membrane potentials between pairs of simulations diverging from identical initial conditions (mean of all neurons and saved base states). (A) Changing random seeds for subsets of noise sources with the standard stochastic release model. (B) Changing random seeds for subsets of noise sources with a mean release model. (C) Standard stochastic release model for decoupled, replayed simulations.

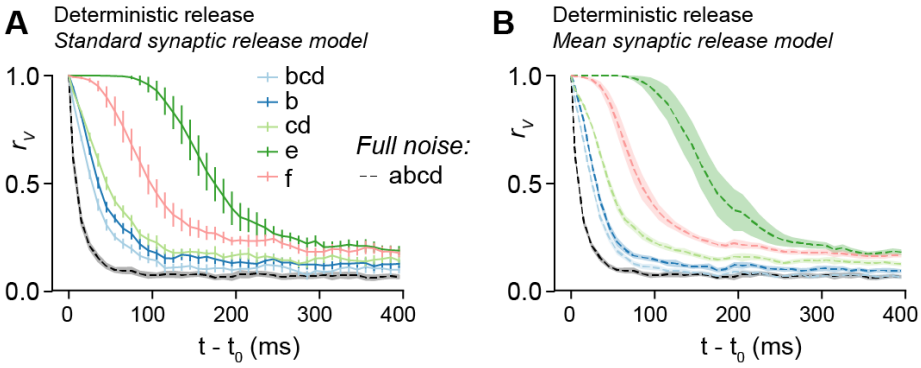

Figure S4: Mean Synaptic Release Model

(A) Correlation r_V (as in Figure 4 and Figure S5), with pseudo-deterministic synaptic release by not changing the random seeds for vesicle release (but with a change in '*mini*' signals for *b*). (B) As in A, but with deterministic synaptic release (mean release model), apart from *abcd* which has the fully stochastic model.

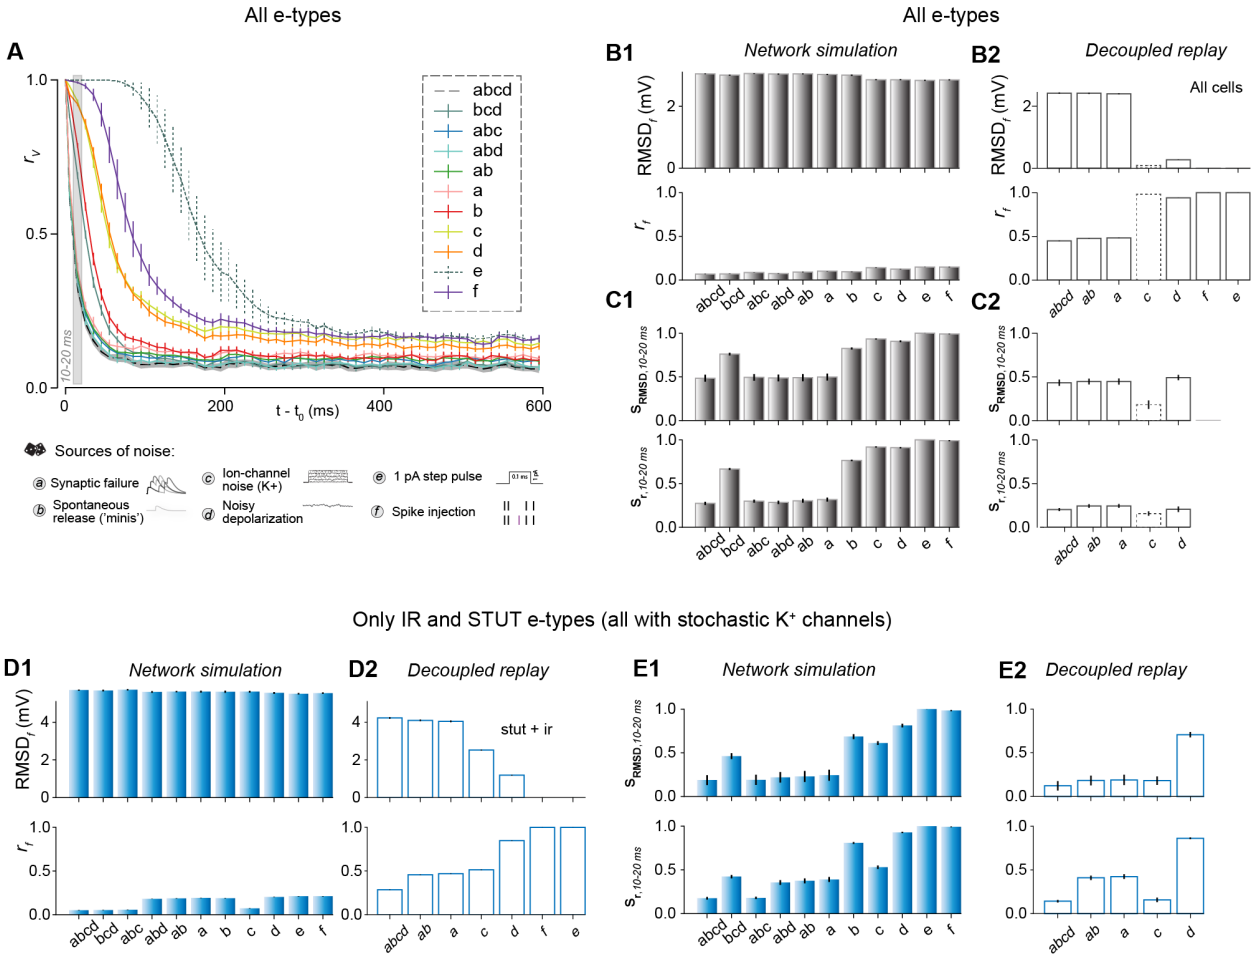

Figure S5: Unravelling Noise Sources

(A) Correlation r_V from identical initial conditions with different cellular noise sources turned on, and when turning of cellular noise, but perturbing the system by a single extra spike (in one neuron) or a miniscule perturbation in all neurons. (B) Steady-state membrane potential fluctuations ($RMSD_\infty$) and correlations (r_∞) for network simulations (B1) and decoupled, replayed simulations (B2) for different noise sources. (C) Similarity $s_{r/RMSD}$ at 10-20 *ms* for network simulations (C1) and decoupled, replayed simulations (C2) for different noise sources. (D-E) Same as B-C, but only for the subset of neurons that have stochastic ion-channels (irregularly firing e-types, 1'137 out of 31'346 neurons). All error bars indicate 95% confidence intervals, based on 20 pairs of simulations (40 for abcd).

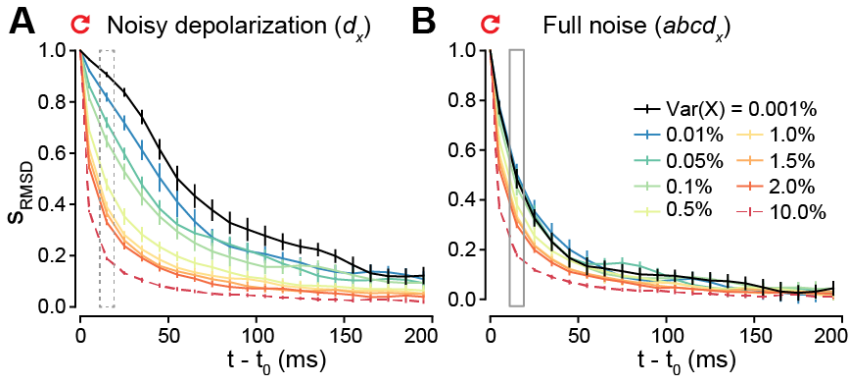

Figure S6: Predicting Impact of Other Noise Sources

(A) Correlation r_V when only changing random seeds for noisy depolarization, but with different magnitudes of noise. (B) As in A, but with all noise sources enabled by changing random seeds.

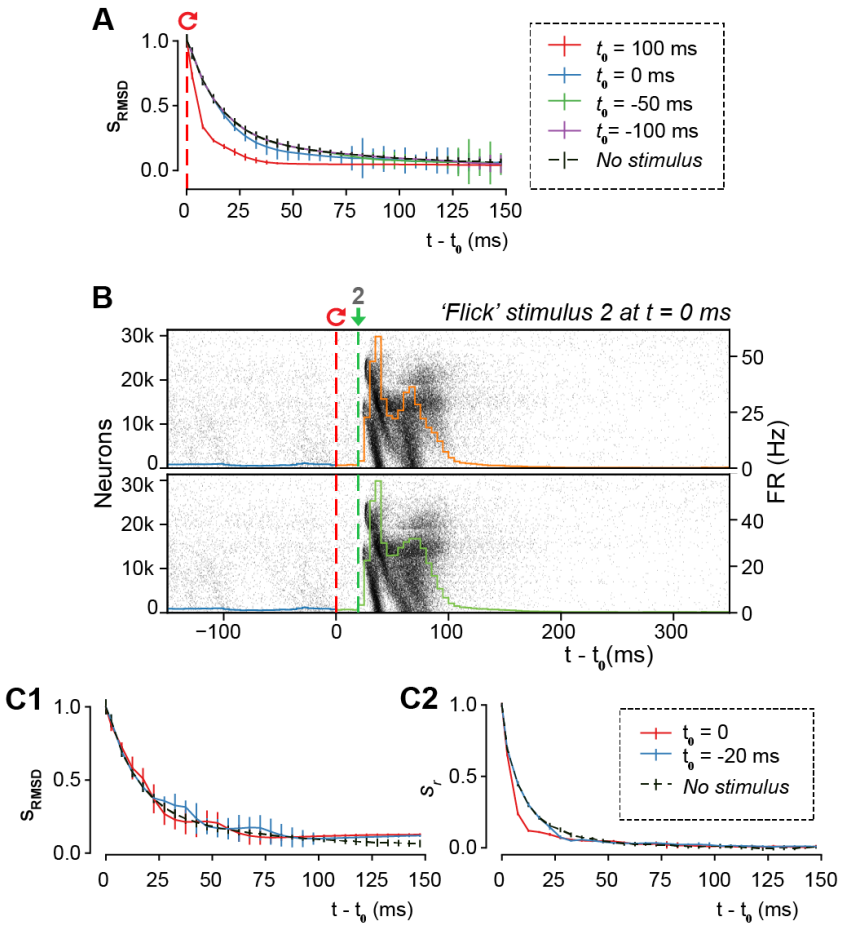

Figure S7: Divergence of Evoked Activity

(A) The similarity s_{RMSD} defined as the difference between the $RMSD_V$ of diverging and independent trials, normalized to lie between 1 (identical) and 0 (fully diverged) (mean ± 95% confidence interval), for the thalamic stimulus. (B) Population raster plot and population peristimulus time histogram (PSTH) of all 31'346 neurons in the microcircuit, during evoked activity with a simplified "whisker flick" stimulus (60 VPM neurons are firing at the same time, one spike). (C1) As A, but for the "whisker flick" stimulus. (C2) As C1, but for s_r instead of s_{RMSD}.

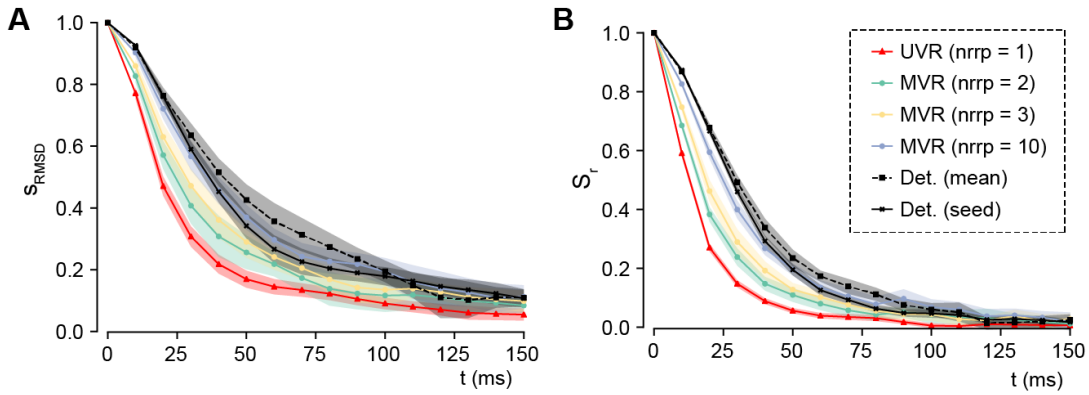

Figure S8: Multivesicular Release

Change in divergence time course depending on the size of the pool of readily releasable vesicles (*nrrp*). Quantified by similarity of the somatic membrane potentials diverging from identical initial conditions: (A) s_{RMSD} and (B) s_r. (mean of all neurons and *n* base states ± 95% confidence interval). (UVR: *n* = 40; all others: *n* = 20).