When do correlations increase with firing rates in recurrent networks?

Andrea K. Barreiro[1]*, **Cheng Ly**[2]

1 Department of Mathematics, Southern Methodist University, Dallas, Texas, United States of America,
2 Department of Statistical Sciences and Operations Research, Virginia Commonwealth University, Richmond, Virginia, United States of America

* abarreiro@smu.edu

Data Availability Statement: The authors confirm that all data underlying the findings are fully available without restriction. Software implementing the computational methods used in this paper can be found in a git repository at: https://github.com/andreakbarreiro/LR_CondBased. All other relevant data are within the paper and its Supporting Information files.

Abstract

A central question in neuroscience is to understand how noisy firing patterns are used to transmit information. Because neural spiking is noisy, spiking patterns are often quantified via pairwise correlations, or the probability that two cells will spike coincidentally, above and beyond their baseline firing rate. One observation frequently made in experiments, is that correlations can increase systematically with firing rate. Theoretical studies have determined that stimulus-dependent correlations that increase with firing rate can have beneficial effects on information coding; however, we still have an incomplete understanding of what circuit mechanisms do, or do not, produce this correlation-firing rate relationship. Here, we studied the relationship between pairwise correlations and firing rates in recurrently coupled excitatory-inhibitory spiking networks with conductance-based synapses. We found that with stronger excitatory coupling, a positive relationship emerged between pairwise correlations and firing rates. To explain these findings, we used linear response theory to predict the full correlation matrix and to decompose correlations in terms of graph motifs. We then used this decomposition to explain why covariation of correlations with firing rate—a relationship previously explained in feedforward networks driven by correlated input—emerges in some recurrent networks but not in others. Furthermore, when correlations covary with firing rate, this relationship is reflected in low-rank structure in the correlation matrix.

Author summary

A central question in neuroscience is to understand how noisy firing patterns are used to transmit information. We quantify spiking patterns by using pairwise *correlations*, or the probability that two cells will spike coincidentally, above and beyond their baseline firing rate. One observation frequently made in experiments is that correlations can increase systematically with firing rate. Recent studies of a type of output cell in mouse retina found this relationship; furthermore, they determined that the increase of correlation with firing rate helped the cells encode information, provided the correlations were stimulus-dependent. Several theoretical studies have explored this basic structure, and found that it is generally beneficial to modulate correlations in this way. However—aside from mouse retinal cells referenced here—we do not yet have many examples of real neural circuits

that show this correlation-firing rate pattern, so we do not know what common features (or mechanisms) might occur between them.

In this study, we address this question via a computational model. We set up a computational model with features representative of a generic cortical network, to see whether correlations would increase with firing rate. To produce different firing patterns, we varied excitatory coupling. We found that with stronger excitatory coupling, there was a positive relationship between pairwise correlations and firing rates. We used a network linear response theory to show why correlations could increase with firing rates in some networks, but not in others; this could be explained by how cells responded to fluctuations in inhibitory conductances.

Introduction

One prominent goal of modern theoretical neuroscience is to understand how the features of cortical neural networks lead to modulation of spiking statistics [1–3]. This understanding is essential to the larger question of how sensory information is encoded and transmitted, because such statistics are known to impact population coding [4–8]. Both experimental and theoretical inquiries are complicated by the fact that neurons are widely known to have heterogeneous attributes [9–14].

One family of statistics that is implicated in nearly all population coding studies is trial-to-trial variability (and co-variability) in spike counts; there is now a rich history of studying how these statistics arise, and how they effect coding of stimuli [15–19]. Recent work by numerous authors has demonstrated that the information content of spiking neural activity depends on spike count correlations and its relationship (if any) with stimulus tuning [15, 17, 19–21]. Since a population of sensory neurons might change their firing rates in different ways to stimuli, uncovering the general mechanisms for when spiking correlations increases with firing rate (or when they do not) is important in the context of neural coding. Thus, we study this question in a general recurrent neural network model.

One observation that has been made in some, but not all, experimental studies is that pairwise correlations increase with firing rates. This relationship has been observed *in vitro* [22] and in several visual areas: area MT [23], V4 [24], V1 [25, 26], and notably, in ON-OFF directionally sensitive retinal ganglion cells [21, 27]. The retinal studies involved cells with a clearly identified function, and therefore allowed study of the coding consequences of this correlation/firing rate relationship. Both studies found that the *stimulus-dependent* correlation structure observed compared favorably to a structure in which *stimulus-independent* correlations were matched to their (stimulus-)averaged levels. This finding reflects a general principle articulated in other studies [17, 19], that stimulus-dependent correlations are beneficial when they serve to spread the neural response in a direction *orthogonal* to the signal space.

While many studies have illustrated the connection between stimulus-dependent correlation structure and coding, these have (until recently: see [21, 25, 27]) largely taken the correlation structure as given, leaving open the question of how exactly a network might produce the hypothesized correlation structure [6, 7] (see also the theoretical calculations in [21, 27]). Theoretical studies of the *mechanisms* that contribute to correlation distributions have largely analyzed homogeneous networks (i.e. cells are identical, aside from E/I identity) [2, 3, 28, 29], which does not allow an exploration of a correlation/firing rate relationship. Thus, how correlation coefficients can vary across a population of heterogeneously-tuned neurons is not yet well understood despite its possible implications for coding.

In this paper we investigated the relationship between correlations and firing rates in conductance-based leaky integrate-and-fire (**LIF**) neural network models, consisting of excitatory (E) and inhibitory (I) cells that are recurrently and randomly coupled. We introduced neural heterogeneity by allowing thresholds to vary across the population, which induced a wide range of firing rates, and explored different firing regimes by varying the strength of recurrent excitation. We found that with relatively strong excitation, pairwise correlations increased with firing rate.

In theoretical studies, this correlation-firing rate trend has been explained in feed-forward networks driven by common input [22, 30, 31]. Here we investigated whether the correlation/firing relationship in recurrent networks can be explained by this theory, but where the source of input correlations is internally generated; i.e., from overlapping projections within the recurrent network. We first adapted a network linear response theory, to decompose predicted correlations into contributions from different graph motifs, which are subgraphs which form the building blocks of complex networks [28, 32, 33].

We found that in all networks studied here, second-order motifs—and specifically *inhibitory common input*—were the dominant contributor to overall pairwise correlations. This allowed us to generalize theory from [22], and describe pairwise correlations in terms of a single-cell susceptibility function. Surprisingly, we found that correlations from inhibitory common input could either increase *or* decrease with firing rate, depending on how cells responded to fluctuations in inhibitory conductances.

We further show that a correlation-firing rate relationship has an important consequence for heterogeneous networks; it can shape low-dimensional structure in the correlation matrix. Low-dimensional structure—often modeled with a low-rank approximation to the correlation matrix—is important because it can be used to improve estimation [34] and even to reconstruct full correlation matrices from incomplete data [35–37]; such structure has been observed in experimental data [25, 38–41] but its origin is not always known. We demonstrate in our networks that when correlation co-varies with firing rate, the (E-E) correlation matrix could be accurately modeled with a low-rank approximation, and the low-rank projection in this approximation was strongly associated with firing rate. Thus we demonstrate that low-rank structure can result from recurrent activity modulated by single-cell characteristics, as well as from a global input or a top-down signal [38].

Results

We studied asynchronous recurrent networks of leaky integrate-and-fire model neurons, and varied the strength of excitation to get different firing behaviors. We found that the covariation of correlations with firing rates—a phenomenon observed in feed-forward networks—occurs here in one firing regime, but not the other. We then found that this could be explained in terms of how single cells responded to fluctuations in inhibitory conductance. Finally, we show that when correlations covary with firing rates, the correlation matrix admits a low-rank approximation.

Asynchronous firing in heterogeneous networks

We performed Monte Carlo simulations of recurrent, randomly connected E/I networks, as described in **Methods: Neuron model and network setup**. To connect to previous literature on asynchronous spiking, we compared networks with and without single-cell variability—referred to as *heterogeneous* and *homogeneous* respectively. Heterogeneity was introduced by allowing cell threshold to vary, which induced a corresponding range of firing rates (see **Methods: Neuron model and network setup** for details). We first chose parameters so that the
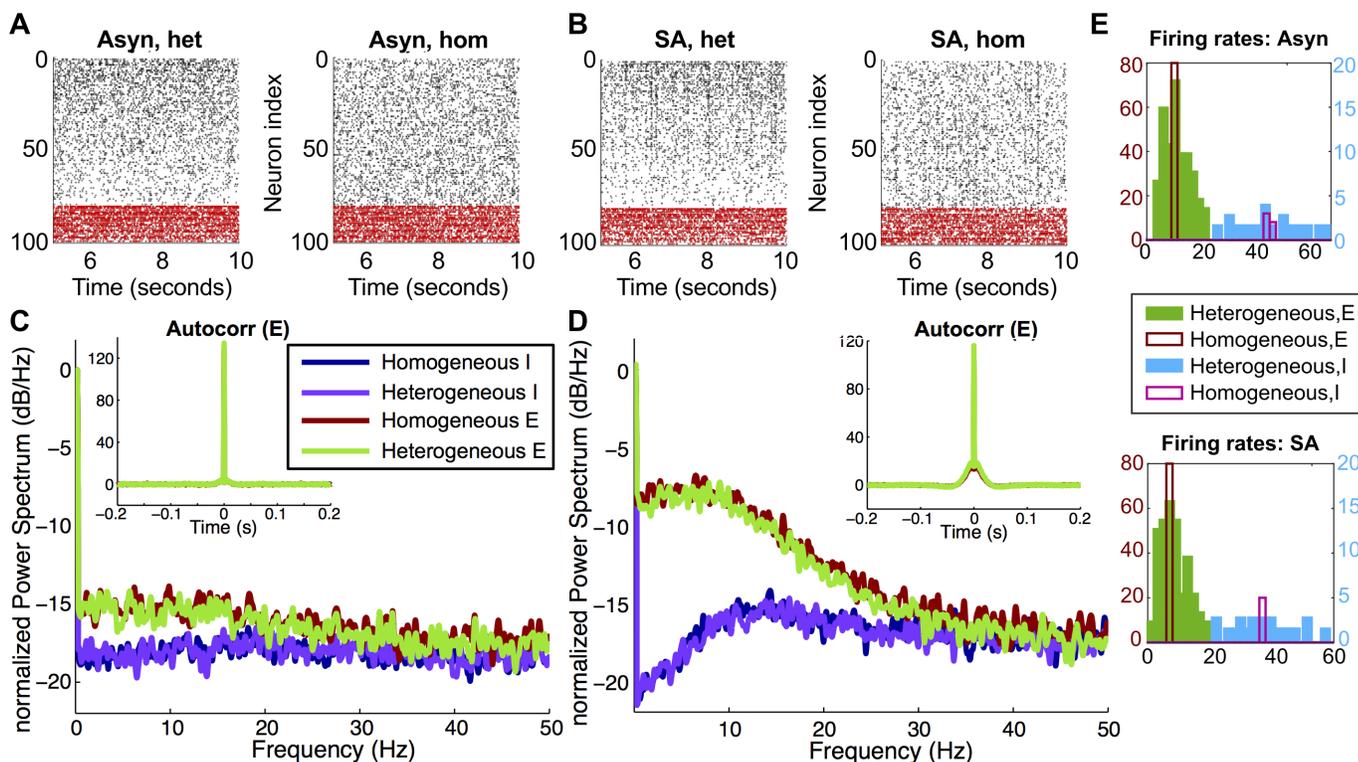
Fig 1. Two firing regimes in heterogeneous networks. Monte Carlo simulations illustrating two firing regimes we consider in this paper. (A) Raster plots from the asynchronous (**Asyn**) regime. (B) Raster plots from the *strong asynchronous* (**SA**) regime, showing occasional bursts of activity. (C) Power spectra in the asynchronous regime. (D) Power spectra in the strong asynchronous regime. (E) Firing rates in the asynchronous (top panel) and SA (bottom panel) regimes. In (A-B), cells are ordered by increasing threshold value. Power spectra (C-D) are normalized to their maximum value and expressed in decibels/Hz.

https://doi.org/10.1371/journal.pcbi.1005506.g001

networks exhibited the classical *asynchronous irregular* (**Asyn**) regime, in which each neuron has irregular Poisson-like spiking, correlations are low, and the population power spectra are flat [42]. In Fig 1A we show raster plots from both the heterogeneous and homogeneous networks, in this regime. The heterogeneous network shows a gradient in its raster plot, because cells are ordered by decreasing firing rate. The population power spectra were flat, for both E and I cells and in both homogeneous and heterogeneous networks (Fig 1C).

When we increased excitation (by increasing both W_{EE} and W_{IE}, where W_{XY} is the conductance strength from type Y to X; see Table 1 for parameter values), we observed occasional bursts of activity. However, the bursts do not occur at regular intervals and do not involve the entire population (we found excitatory bursts involved at most 25% of the population). The

Table 1. Excitatory connection strengths mediate between different firing regimes.

Parameter	$W_{EI}(I \rightarrow E)$	$W_{IE}(E \rightarrow I)$	W_{EE}	W_{II}	$\sigma_i(i \in E)$	$\sigma_i(i \in I)$
Asynchronous	10	5	0.5	5	$2/\sqrt{2}$	$3/\sqrt{2}$
Str. Asynch.	10	8	9	5	$1.5/\sqrt{2}$	$2.5/\sqrt{2}$
% connectivity	35%	20%	40%	40%		

Here W^{YX} denotes $X \rightarrow Y$ connections; e.g. W^{IE} denotes the strength of excitatory connections onto inhibitory neurons. The parameter σ_i denotes the strength of background noise in units of (scaled) voltage, and depends only on cell type (E or I).

https://doi.org/10.1371/journal.pcbi.1005506.t001

network is still moderately inhibition-dominated and neurons are spiking irregularly; example raster plots are shown in Fig 1B. The population power spectra (Fig 1D) are no longer flat (compare to the asynchronous regime, Fig 1C); they show local maxima around 8 Hz, but it is not a pronounced peak. We will refer to this as the *strong asynchronous* (**SA**) regime [43].

In both Fig 1C and 1D, we note that—despite the apparent differences in the distribution of spikes across the network, evident in the raster plots—both the autocorrelation functions (Fig 1C and 1D, insets) and the power spectra from the heterogeneous and homogeneous networks are very similar. Thus, we have a fair comparison to examine the role of heterogeneity, independent of other characteristics of the network.

The distribution of both excitatory and inhibitory firing rates are extremely narrow in the homogeneous network, but broad in the heterogeneous network (Fig 1E). This is expected, as each excitatory (inhibitory) cell in the homogenous network has the same uncoupled firing rate; because the number of synaptic inputs is likewise fixed, population variability in synaptic input is limited. The heterogeneous networks have a range of firing rates, which allows us to investigate the possibility of a relationship between (variable) firing rate and pairwise correlations. Population-averaged firing rates were very similar between the heterogeneous and homogeneous networks: in the asynchronous regime $\langle\langle v_E \rangle\rangle$ = 10.6 Hz (heterogeneous) and $\langle\langle v_E \rangle\rangle$ = 10.1 Hz (homogeneous), while $\langle\langle v_I \rangle\rangle$ = 44.3 Hz (heterogeneous) and $\langle\langle v_I \rangle\rangle$ = 43.5 Hz (homogeneous). In both regimes Fano factors ranged between 0.9 and 1.1, consistent with Poisson-like spiking (more statistics are given in S1 and S3 Tables).

Correlation increases with firing rate in the strong asynchronous regime

We next sought a possible relationship between pairwise correlations—quantified via the Pearson's correlation coefficient for spike counts, $\rho_{ij} \equiv \mathrm{Cov}_T(n_i, n_j) / \sqrt{\mathrm{Var}_T(n_i)\,\mathrm{Var}_T(n_j)}$—and single-cell firing rates. Such relationships have been found in feed-forward networks [22, 30, 31], and impact information transfer when considered in concert with stimulus selectivity (i.e. signal correlations) [7, 8, 15, 19]. In heterogeneous networks, the large range of firing rates—equivalently the large range of operating points—admits the possibility that cells at different operating points may differ in their ability to transfer correlations.

To investigate this we plotted pairwise correlations for each distinct excitatory pair ρ_{ij}, versus the geometric mean of the firing rates $\sqrt{v_i v_j}$, in both regimes (asynchronous and strong asynchronous), for a range of time scales (blue stars in Fig 2). We focus here on excitatory-excitatory (E-E) pairs, because excitatory synaptic connections provide the predominant means of propagating cortical sensory information to higher layers. Our results show a striking difference between the two spiking regimes; while there is no clear relationship with firing rate in the asynchronous regime (Fig 2, top row), the strong asynchronous regime shows a distinct positive trend with firing rate (Fig 2, bottom row). We can quantify a hypothesized relationship between v and ρ with linear regression, and indeed find that geometric mean firing rate explains a substantial part of the variability of correlations in the strong asynchronous regime obtained from the Monte Carlo simulations, with R^2 values (i.e. percentage of variability explained) of 0.41, 0.37, and 0.34 for time windows of T = 5, 50, and 100 ms respectively (in contrast, R^2 values for the asynchronous network are below 0.005).

In recurrent networks, the response of each cell is shaped by both direct and indirect connections through the network. We used the linear response theory described in **Methods: Linear Response Theory** and **Methods: Computing statistics from linear response theory** to predict the full correlation matrix $\mathbf{C}_T$ at various time scales, including the limit of long time scales: $\tilde{\mathbf{C}}(0) = \lim_{T\to\infty} \frac{1}{T}\mathbf{C}_T$. We found that this theory successfully captured E-E correlations,
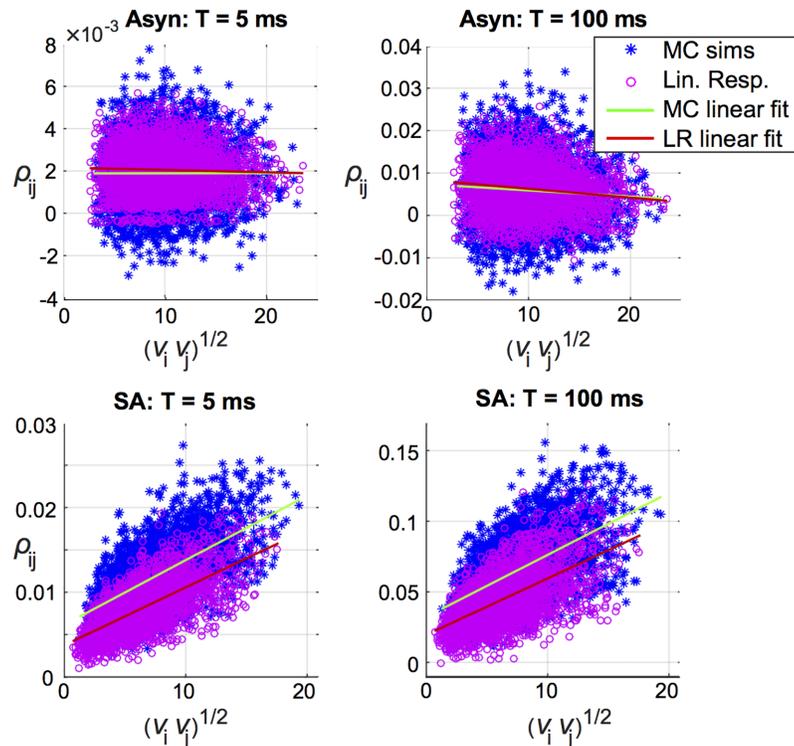
Fig 2. Correlation increases with firing rate in the strong asynchronous regime. E-E correlation ρ_{ij} vs. geometric mean firing rate $\sqrt{\nu_i \nu_j}$, cell-by-cell comparison of Monte Carlo simulations (blue stars) and linear response (magenta circles), in a heterogeneous network. Left to right: time window $T = 5$ ms and 100 ms. Top row: asynchronous regime. Bottom row: strong asynchrony.

both the full distribution of values and coefficients of individual cell pairs (details, including figures, can be found in: **Supporting Information: S1 Text**).

We then plotted the predicted correlation, $\tilde{\mathbf{C}}_{ij}/\sqrt{\tilde{\mathbf{C}}_{ii}\tilde{\mathbf{C}}_{jj}}$, vs. geometric mean firing rate $\sqrt{\nu_i \nu_j}$ (magenta circles in Fig 2). The predicted correlations captured the same positive relationship observed in Monte Carlo results, with R^2 values of 0.47, 0.4, and 0.36.

Decomposition of correlation by graph motifs shows strong role for second-order motifs

Why does a correlation/firing rate relationship emerge in one spiking regime, but not the other? In feed-forward networks, a positive correlation/firing rate relationship results from transferring common input through fluctuation-driven, asynchronously-firing cells [22, 30]. In contrast, the amount of shared input into two cells in a recurrent network is determined by both direct and indirect connections through the network. To separate the impact of different network pathways, we decomposed the *linear response-predicted* correlations at long time scales (i.e. $\tilde{\mathbf{C}}(0) = \lim_{T\to\infty} \frac{1}{T}\mathbf{C}_T$) into normalized contributions from n-th order motifs, as described in **Methods: Quantifying the role of motifs in networks**. Common input from a divergent connection, for example, results from the 2nd-order motif $\mathbf{K}^*\mathbf{C}^0\mathbf{K}$. In Fig 3, we plot the summed contributions up to sixth order—i.e. $\tilde{\mathbf{R}}_{ij}^k$, for $k = 1, 2, \ldots 6$—versus geometric
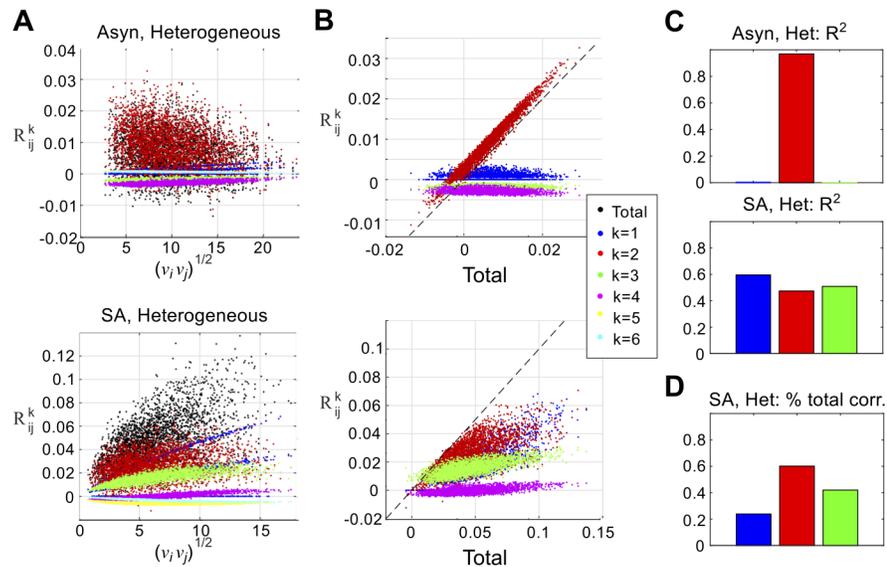
Fig 3. Pairwise correlations are built from graph motifs. Contributions of different orders to prediction of E-E correlations with linear response theory. (A) Normalized contributions to pairwise correlation ($\tilde{\mathbf{R}}_{ij}^k$) vs. geometric mean firing rate ($\sqrt{v_i v_j}$) for heterogeneous networks in the asynchronous (top panel) and strong asynchronous (bottom panel) regimes. (B) As in (A), but plotted vs. total predicted normalized correlation ($\tilde{\mathbf{C}}_{ij}/\sqrt{\tilde{\mathbf{C}}_{ii}\tilde{\mathbf{C}}_{jj}}$). See main text for further discussion. (C) To quantify the relative importance of different motifs, we report the fraction of variance explained (R^2) from linear regressions, in which we regressed total correlation ($\tilde{\mathbf{C}}_{ij}/\sqrt{\tilde{\mathbf{C}}_{ii}\tilde{\mathbf{C}}_{jj}}$) against the contributions at each specific order ($\tilde{\mathbf{R}}_{ij}^k$). As suggested by (B), second-order contributions (red) overwhelmingly determine total correlations in the asynchronous network (R^2 values for first- and third-order terms are shown, but barely visually distinguishable; R^2 values for higher orders are also small, within 0.08 up to $k = 6$). (D) Fraction of total correlation from each order, strong asynchronous regime.

mean firing rate, $\sqrt{v_i v_j}$. The total normalized correlation, $\tilde{\mathbf{C}}_{ij}/\sqrt{\tilde{\mathbf{C}}_{ii}\tilde{\mathbf{C}}_{jj}}$, is shown as well. In all cases, we plot long time scale correlations $\omega = 0$; each distinct E-E pair is represented.

In the asynchronous regime (top panel of Fig 3A), first-order contributions ($\tilde{\mathbf{R}}^1$) separate into three distinct "curves", reflecting a 1-1 relationship with firing rate conditioned on first-order connectivity (no connection between i and j; one connection between i and j; bidirectional connection between i and j). Second-order contributions are overall positive while third-order contributions are overall negative (consistent with [28]); neither appear to have a relationship with firing rate. Second-order contributions are conspicuously dominant; fifth and sixth order terms are near zero.

This qualitative picture changes when we consider the strong asynchronous regime (bottom panel of Fig 3A). First-order contributions follow a similar pattern as in the asynchronous regime, and second-order contributions are likewise positive. However, third-order contributions are positive, and in the heterogeneous network they have a distinctly positive relationship with firing rate (top panel). Thus, in the asynchronous regime, negative third-order contributions partially cancel with positive second-order contributions; in the strong asynchronous regime, first, second, and third-order motifs reinforce each other, contributing to an overall positive relationship with firing rate (black dots).

Despite these differences, second-order contributions are the major determinant of total correlation in both regimes. In Fig 3B we plot the same data ($\tilde{\mathbf{R}}_{ij}^k$) vs. total correlation, rather

than geometric mean firing rate. In the asynchronous regime, second-order contributions cluster near the unity line, suggesting they are strongly predictive of total correlation. To quantify this intuition we computed the fraction of variance explained (R^2) by performing a linear regression of total normalized correlation ($\tilde{\mathbf{C}}_{ij}/\sqrt{\tilde{\mathbf{C}}_{ii}\tilde{\mathbf{C}}_{jj}}$) against contributions of each order (Fig 3C); in the asynchronous regime R^2 values for $\tilde{\mathbf{R}}^1_{ij}$, $\tilde{\mathbf{R}}^2_{ij}$, and $\tilde{\mathbf{R}}^3_{ij}$ were 0.004, 0.969, and 0.0002, respectively. R^2 values for higher orders were likewise small: for $\tilde{\mathbf{R}}^4_{ij}$, $\tilde{\mathbf{R}}^5_{ij}$, and $\tilde{\mathbf{R}}^6_{ij}$ they were 0.047, 0.034, and 0.074.

This statistic was more ambiguous in the strong asynchronous regime, where R^2 values for $\tilde{\mathbf{R}}^1_{ij}$, $\tilde{\mathbf{R}}^2_{ij}$, and $\tilde{\mathbf{R}}^3_{ij}$ were 0.595, 0.474, and 0.509 respectively. However, note that $\tilde{\mathbf{R}}^1_{ij}$ and $\tilde{\mathbf{R}}^3_{ij}$ were positive for all cell pairs; $\tilde{\mathbf{R}}^2_{ij}$ and total correlation were negative for less than 0.3% of cell pairs. Thus, we considered how each motif contributed to the total correlation by taking the ratio of each contribution to the total, averaged over all cell pairs (Fig 3D). By this measure, second-order contributions were largest; fraction explained for $\tilde{\mathbf{R}}^1_{ij}$, $\tilde{\mathbf{R}}^2_{ij}$, and $\tilde{\mathbf{R}}^3_{ij}$ were 0.239, 0.601, and 0.420, respectively. Note that this measure cannot be used for the asynchronous (**Asyn**) regime because of the negative values of $\tilde{\mathbf{R}}^k_{ij}$.

Taken together, this evidence points to a distinguished role for second-order motifs ($\tilde{\mathbf{R}}^2_{ij}$) in determining total correlation. In the asynchronous regime in particular, $\tilde{\mathbf{R}}^2_{ij}$ is a near-perfect predictor of total correlation.

Inhibitory common input is the most important second-order motif

We next analyzed contributions from specific **second-order motifs** in Fig 4. There are four distinct second-order motifs that can correlate two E cells. There are two types of chains, from
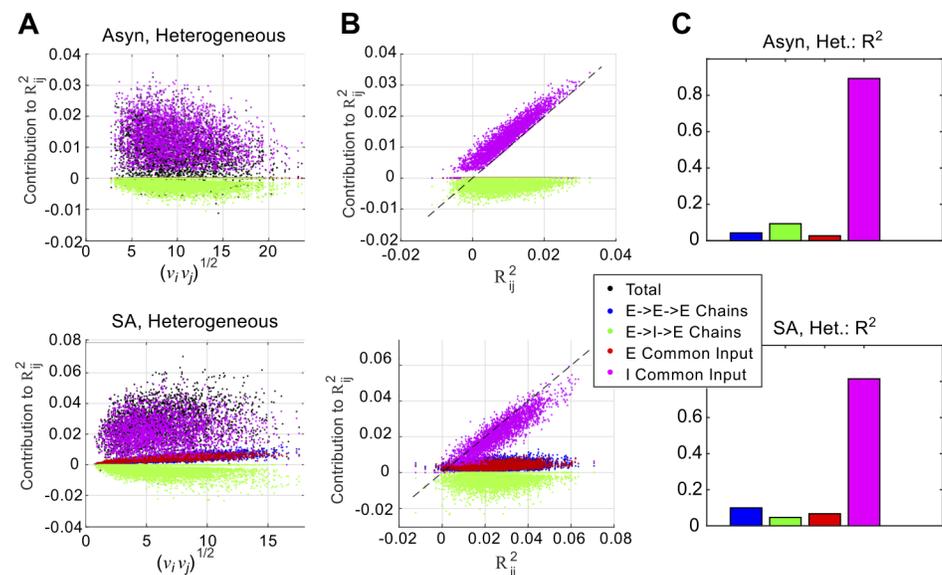
Fig 4. Inhibitory common input is the dominant second-order motif in both asynchronous and strong asynchronous networks. (A) Contributions of different 2nd-order motifs to prediction of E-E correlations in a heterogeneous network, in the asynchronous (top) and strong asynchronous (bottom) regimes. (B) As in (A), but plotted vs. total contribution from second-order motifs $\tilde{\mathbf{R}}^2$. In both panels, inhibitory common input (magenta) clusters near the unity line. (C) To quantify the relative importance of different motifs, we report the fraction of variance explained (R^2) from linear regressions, in which we regressed the total contribution from second-order motifs ($\tilde{\mathbf{R}}^k_{ij}$) against the contribution from specific motifs types.

https://doi.org/10.1371/journal.pcbi.1005506.g004

$\mathbf{K}^2\mathbf{C}^0$ and $\mathbf{C}^0 (\mathbf{K}^*)^2$. An $E \to E \to E$ chain tends to positively correlate, while an $E \to I \to E$ chain will negatively correlate; these are shown as blue and green respectively. There are two types of common input, from $\mathbf{K}\mathbf{C}^0(\mathbf{K}^*)$; they correspond to common input from E and I cells, i.e. $E \leftarrow E \to E$ and $E \leftarrow I \to E$. They *both* lead to positive correlations and are shown as red and magenta respectively.

In the asynchronous regime (left panel of Fig 4A), the dominant contributions are I common input (magenta) and negative ($E \to I \to E$) chains (green); correlating chains (blue) and excitatory common input (red) are barely visible, as they are clustered near zero. In the strong asynchronous case (right panel), blue and red dots are now visible and show a clear 1-1 trend with firing rate. In both regimes, *inhibitory common input* appears to be the dominant second-order motif. In Fig 4B we plot the contribution from different second-order motifs vs. the total contribution from second-order motifs, $\tilde{\mathbf{R}}_{ij}^2$ (rather than geometric mean firing rate, $\sqrt{\nu_i\nu_j}$). In both panels, the inhibitory common input (magenta) clusters around the unity line, showing it is the best predictor of the total second-order contribution. In Fig 4C we quantify this observation by reporting fraction of variance explained (R^2) from linear regressions: the R^2 value for inhibitory common input exceeds 0.8 in both networks, while the R^2 values for all other motifs types are less than 0.1.

In conclusion, decomposition of pairwise correlations into graph motifs has shown us two important things: first, while third-order motifs probably contribute to the positive correlation/firing rate relationship observed in the SA regime, second-order motifs still dominate in both regimes. Second, **inhibitory common input is the most important second-order motif in both regimes** (Fig 4B).

Susceptibility to inhibition can either increase or decrease with firing rate

In feedforward networks—i.e., in the absence of a path between two cells—correlations in *outputs* (i.e. spike trains) must arise from correlations in *inputs*; for example, through shared or common inputs. We have found that inhibitory common input is the dominant contributor to pairwise correlations in both the asynchronous and strong asynchronous regimes; we now turn our attention to modeling this term (inhibitory common input) specifically.

Previous work that analyzed the relationship between the long-time correlation and firing rate in feedforward networks [22, 30] quantified a *susceptibility* function that measures the ratio between output and input correlations:

$$S \approx \frac{\rho}{c}. \tag{1}$$

If both cells receive a large (but equal) number of uncorrelated inputs, c would be the fraction of inputs that are common to both i and j.

In the networks examined here, each cell had a fixed in-degree for both excitatory and inhibitory cells; however, for any given *pair* of cells i and j, the number of E and I inputs that synapsed onto both cells will vary from pair to pair. Thus, we next considered the possibility that our (negative) finding in the asynchronous network could be explained by accounting for variable c_{ij}.

We focus on inhibitory common input, which is the dominant second-order contribution in the asynchronous network (Fig 4). We segregated pairs by whether they had 0, 1, 2, etc.. common inhibitory inputs; we then use this number as a proxy for c (recall that each excitatory cell had exactly 7 inhibitory inputs, so that this number divided by 7 approximates the common input fraction; two common inputs imply $c \approx 0.28$ for example). We plot the results for the asynchronous network in Fig 5A, top panel (data for each distinct value of c is presented by color). As we might expect, correlation increases as c increases. However, for a fixed c, there is
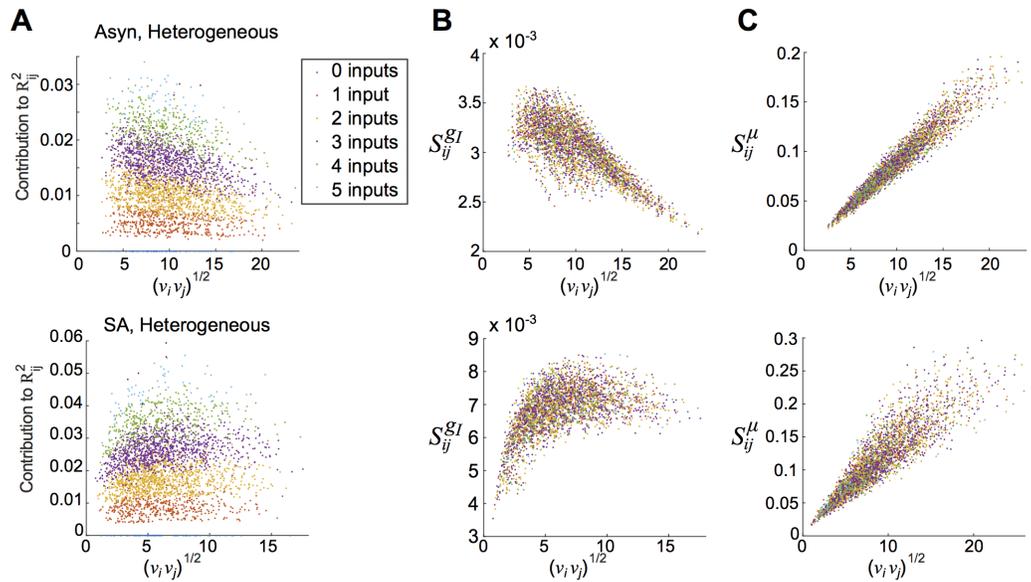
Fig 5. Susceptibility to conductance fluctuations can explain correlation-firing rate relationships. In (A-C): heterogeneous asynchronous (top) and heterogeneous strong asynchronous (bottom). (A) Correlation (ρ) from I common inputs vs. firing rate, segregated based on the number of common inhibitory inputs. (B) Estimated correlation susceptibility to fluctuations in inhibitory *conductances* vs. firing rate ($S_{ij}^{g_I}$). (C) Correlation susceptibility to fluctuations in inhibitory *currents* vs. firing rate (S_{ij}^{μ}).

not an apparent relationship between firing rate and correlation; if anything, there appears to be a slight decrease. Correlation also increases with c in the strong asynchronous network (Fig 5A, bottom panel); however, here we also see a modest increase with geometric mean firing rate $\sqrt{v_i v_j}$.

Previous theoretical work [22, 30] identified an increase in susceptibility with firing rates in *current-driven* neurons; we next considered the possibility that this fails to hold for conductance-driven neurons. As described in **Methods: Quantifying correlation susceptibility**, we estimated correlation susceptibility for each pair of neurons, by using the susceptibility function for each neuron to conductance fluctuations (computed as part of the linear response theory), divided by a measure of the long-timescale spike count variance:

$$S_{ij}^{\langle g_I \rangle} = \frac{\tilde{A}_{\langle g_I \rangle, i}(0) \tilde{A}_{\langle g_I \rangle, j}(0)}{\sqrt{\tilde{\mathbf{C}}_{ii}(0) \tilde{\mathbf{C}}_{jj}(0)}} \tag{2}$$

We plotted the results for both networks in Fig 5B; while susceptibility increases with firing rate in the strong asynchronous network (except for the largest firing rates), it actually decreases with firing rate in the asynchronous network.

We can contrast with the estimated susceptibility to *current* fluctuations (i.e. $A_{\mu,i}$, with μ_i, $\tau_{\text{eff},i}$, and $\sigma_{\text{eff},i}$ as in Eq 29) which we also computed for the same set of cell pairs, shown in Fig 5C.

$$S_{ij}^{\mu} = \frac{\tilde{A}_{\mu,i}(0) \tilde{A}_{\mu,j}(0)}{\sqrt{\tilde{\mathbf{C}}_{ii}(0) \tilde{\mathbf{C}}_{jj}(0)}} \tag{3}$$

Here, we see that S_{ij}^{μ} increases with firing rates, in both networks.

Understanding the susceptibility function

We next sought to understand how susceptibility depends on neural parameters; that is, we define the single-cell susceptibility

$$S_i^{\langle g_I \rangle} \equiv \frac{\tilde{A}_{\langle g_I \rangle, i}(0)}{\sqrt{\nu_i}} \tag{4}$$

where

$$\nu_i = f\left(\langle g_{I,i} \rangle, \sigma_{I,i}, \langle g_{E,i} \rangle, \sigma_{E,i}, \sigma_i, \theta_i\right) \tag{5}$$

$$\tilde{A}_{\langle g_I \rangle, i}(0) = \frac{\partial f}{\partial x_1}\left(\langle g_{I,i} \rangle, \sigma_{I,i}, \langle g_{E,i} \rangle, \sigma_{E,i}, \sigma_i, \theta_i\right). \tag{6}$$

("$\frac{\partial}{\partial x_1}$" indicates that derivative is taken with respect to the first argument, $\langle g_{I,i} \rangle$). We have also used the asynchronous spiking assumption, that $\tilde{\mathbf{C}}_{ii}(0) \approx \nu_i$ (compare with Eq 3).

This quantity is shown in Fig 6, where it is plotted vs. firing rate ν_i (blue stars). Note that this is a negative quantity; since the susceptibility for a neuron pair $S_{ij}^{\langle g_I \rangle} = S_i^{\langle g_I \rangle} S_j^{\langle g_I \rangle}$ is the product (and therefore positive), an *increase* in $S_i^{\langle g_I \rangle}$ will result in a *decrease* in $S_{ij}^{\langle g_I \rangle}$ and vice versa.

In principle, the firing rate function (Eq 5)—and therefore susceptibility—can depend on all *six* parameters defining the cell: our next step was to reduce the dimensionality of the problem. We first looked for any possible relationship between single-cell firing rates and cell parameters (see S1 Text: **Approximating single-cell susceptibility in a heterogeneous network**, S6 and S7 Figs): in both networks, only threshold θ_i had an obvious relationship with firing rate. Among the remaining parameters, the mean inhibitory conductance $\langle g_I \rangle$ had the greatest relative range of values in the asynchronous network (S6A Fig). Therefore, we hypothesized that we could accurately capture $S_i^{\langle g_I \rangle}$, by approximating it as a function of the two parameters θ_i and $\langle g_I \rangle$.
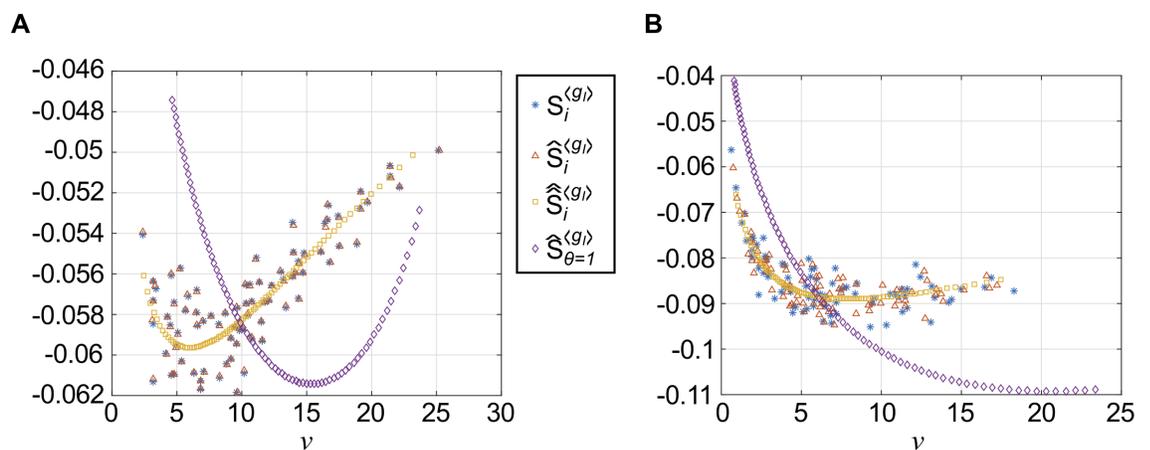
Fig 6. How firing rate diversity is achieved in a heterogeneous network will affect susceptibility. Single-cell susceptibility function(s) for a conductance-based LIF neuron, as a function of firing rate ν. Successive approximations shown are: original single-cell susceptibility, $S_i^{\langle g_I \rangle}$ (Eq 4, blue stars); most parameters set to average value, $\hat{S}_i^{\langle g_I \rangle}$ (Eq 7, red triangles); all parameters but θ_i set to average value, $\hat{\hat{S}}_i^{\langle g_I \rangle}$ (Eq 9, gold squares); and θ fixed, $\hat{S}_{\theta=1}^{\langle g_I \rangle}$ (Eq 10, purple diamonds). (A) Asynchronous regime. (B) Strong asynchronous regime.

https://doi.org/10.1371/journal.pcbi.1005506.g006

We reevaluated the firing rate function, where $\sigma_{I,i}$, $\langle g_{E,i} \rangle$, $\sigma_{E,i}$ and σ_i have been replaced by their average values: i.e.

$$\hat{S}_i^{\langle g_I \rangle} \equiv \frac{1}{\sqrt{F(\langle g_{I,i} \rangle, \theta_i)}} \frac{\partial F}{\partial x_1} (\langle g_{I,i} \rangle, \theta_i) \tag{7}$$

where

$$F(\langle g_{I,i} \rangle, \theta_i) \equiv f \left(\langle g_{I,i} \rangle, \langle \sigma_{I,i} \rangle_p, \langle \langle g_{E,i} \rangle \rangle_p, \langle \sigma_{E,i} \rangle_p, \langle \sigma_i \rangle_p, \theta_i \right) \tag{8}$$

and $\langle \cdot \rangle_p$ denotes the population average. The results are also illustrated in Fig 6 (red triangles). In the asynchronous regime (Fig 6A), the results are remarkably close to the original quantities, indicating that using average parameter values has little effect; in the strong asynchronous regime (Fig 6B) the difference is larger, but the points appear to occupy the same "cloud". However, we can now visualize the susceptibility as a function of only *two* parameters, and we do so in Fig 7 by evaluating $\hat{S}_i^{\langle g_I \rangle}$ on a $(\theta, \langle g_I \rangle)$ grid; the points corresponding to the actual excitatory cells in our network are illustrated in red. In both the asynchronous and strong asynchronous regimes, the red stars form a scattered cloud around the average value $\langle \langle g_{I,i} \rangle \rangle_p$, with no obvious relationship with θ_i.

This fact motivated a further simplification,

$$\hat{\hat{S}}_i^{\langle g_I \rangle} \equiv \frac{1}{\sqrt{F(\langle \langle g_{I,i} \rangle \rangle_p, \theta_i)}} \frac{\partial F}{\partial x_1} \left(\langle \langle g_{I,i} \rangle \rangle_p, \theta_i \right) \tag{9}$$

i.e., we replaced $\langle g_{I,i} \rangle$ with its population average, $\langle \langle g_{I,i} \rangle \rangle_p$, in essence approximating a one-dimensional "path" that the cells take through parameter space. The results are shown in Fig 6 (gold squares) and, as we should expect, allow us to discern a functional relationship with firing rate ν_i; importantly, it appears to capture the average behavior of the actual susceptibility values $S_i^{\langle g_I \rangle}$. Here, we can see clearly that in the asynchronous regime, correlations should
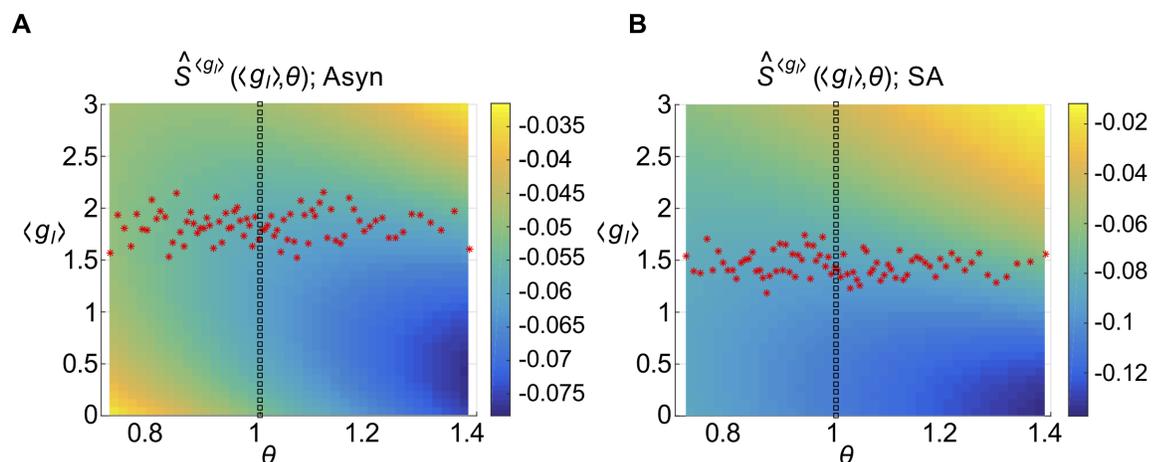
Fig 7. Susceptibility as a function of inhibitory conductance and threshold. Single-cell susceptibility function for a conductance-based LIF neuron, as a function of mean inhibitory conductance $\langle g_I \rangle$ and threshold θ: $\hat{S}^{\langle g_I \rangle}(\langle g_I \rangle, \theta)$ (defined in Eq 7). Other parameters are set to the population average. Overlays show $(\langle g_{I,i} \rangle, \theta_i)$ values of the actual cells in the network (red stars) and an alternative curve through the plane, $(\langle g_I \rangle, 1)$, along which comparable firing rate diversity can be observed (black squares). (A) Asynchronous regime. (B) Strong asynchronous regime.

actually *decrease* with firing rate, for $v_i > 5$ Hz. In the strong asynchronous regime, correlations will *increase* with firing rate, saturating around 10–15 Hz.

Susceptibility depends on the mechanism underlying firing rate diversity

Finally, recall that our actual network sampled a relatively small part of the $(\theta, \langle g_I \rangle)$ plane. This may be attributed to the fact that we generated firing rate diversity (and therefore heterogeneity), by modulating cell excitability through the cell threshold θ_i. How might our results have changed, if we had generated firing rate diversity through some other mechanism? In both regimes, we can increase firing rates by either decreasing $\langle g_{I,i} \rangle$, or by decreasing θ (see S7 Fig). To explore this, we computed susceptibility values along another curve in the $(\theta, \langle g_I \rangle)$ plane; specifically, we held θ fixed and instead varied $\langle g_I \rangle$ (illustrated with black squares on Fig 7); i.e.

$$\hat{S}_{\theta=1}^{\langle g_I \rangle}(\langle g_I \rangle) \equiv \frac{1}{\sqrt{G(\langle g_I \rangle, \theta)}} \frac{\partial G}{\partial x_1}(\langle g_I \rangle, \theta)\Bigg|_{\theta=1} \tag{10}$$

where

$$G(\langle g_I \rangle, \theta) = f\left(\langle g_I \rangle, \langle \sigma_{I,i} \rangle_p, \langle \langle g_{E,i} \rangle \rangle_p, \langle \sigma_{E,i} \rangle_p, \langle \sigma_i \rangle_p, \theta\right)$$

Results are shown in Fig 6 (purple diamonds) and show a strikingly different relationship with firing rate; in the asynchronous regime, correlations should increase with firing rate for $v < 15$ Hz; in the strong asynchronous regime correlations will increase with firing rate, saturating near 20 Hz.

To summarize the previous two subsections, we first defined a single-cell susceptibility function (Eq 4), which captures a linear approximation to the cell's response to input. This quantity relies on an underlying firing rate, which is a function of all parameters that define single-cell dynamics; in this case, six. Each cell occupies a point in this six-dimensional parameter space. We found that in each network studied here, the occupied points approximately lie along a one-dimensional path through this parameter space, along which we could visualize the susceptibility. Finally, we considered the consequences of taking other paths through this parameter space: these paths can be interpreted as generating firing rate heterogeneity using other network mechanisms.

Low-rank structure in neural correlations is mediated by correlation-firing rate relationship

Previous work has identified low-dimensional structure in neural correlation matrices [25, 38–41]; its origin is not always known [3]. We next hypothesized that the positive correlation-firing rate relationship we observed in the strong asynchronous regime, might be reflected in low-dimensional structure in the correlation matrix. For simplicity, suppose that correlations were really represented by a function of firing rate (as in [22]): i.e. $\rho_{ij} = cS(v_i)S(v_j)$. Then we could represent the *off-diagonal* part of the correlation matrix as $\mathbf{C}_T = c\mathbf{S}\mathbf{S}^T$, where $\mathbf{S}$ is a length N vector such that $\mathbf{S}_i = S(v_i)$; that is, $\mathbf{C}_T$ would be a rank-one matrix.

We followed the procedure outlined in **Methods: Low-rank approximation to the correlation matrix** to approximate each correlation matrix, $\mathbf{C}_T$, as the sum of a diagonal matrix and low-rank matrix:

$$\mathbf{C}_T \approx \mathbf{C}_T^{\text{diag+R1}} = \lambda \mathbf{I} + (\sigma_1 - \lambda)\mathbf{u}_1\mathbf{u}_1^T \tag{11}$$
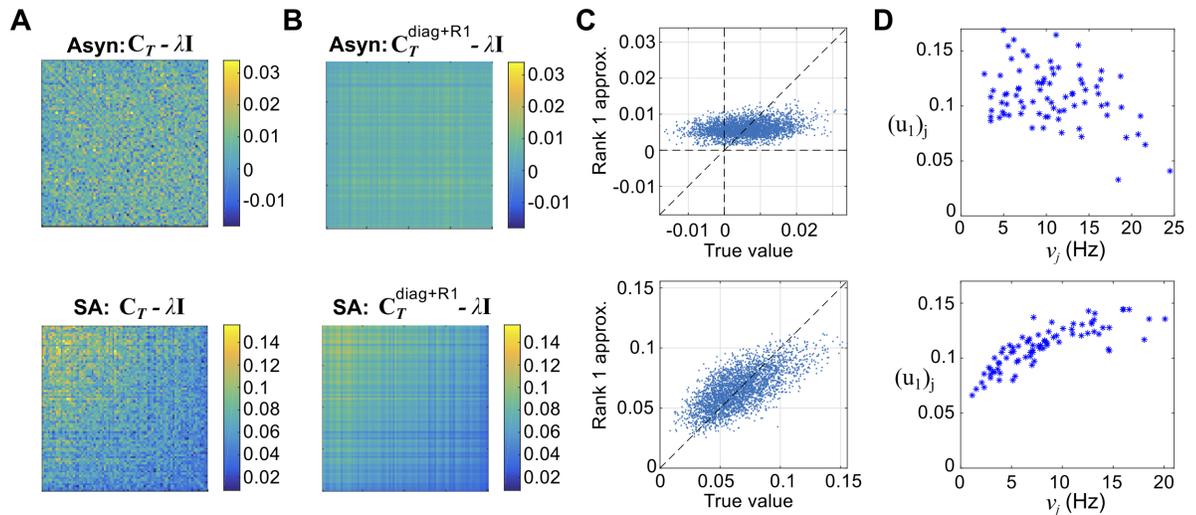
Fig 8. Low-rank structure in correlation matrices. Approximating correlation matrices for the heterogeneous networks as a diagonal plus rank-one. Neurons are ordered by firing rate (highest to lowest). In each column of (A-D), the asynchronous (top) and strong asynchronous (bottom) regimes are shown; T = 100 ms. (A) The shifted E-E correlation matrix, $\mathbf{C}_T - \lambda\mathbf{I}$, for an appropriately chosen λ. (B) A rank-one approximation to $\mathbf{C}_T - \lambda\mathbf{I}$. (C) True correlation coefficients vs. rank-one approximation, cell-by-cell. (D) Weight in the first singular vector, $\mathbf{u}_1$ vs. geometric mean firing rate $\sqrt{v_i v_j}$.

https://doi.org/10.1371/journal.pcbi.1005506.g008

where λ is given in closed form by the eigenvalues of $\mathbf{C}_T$:

$$\lambda = \lambda_1 - \frac{\sum_{j>1}(\lambda_1 - \lambda_j)^2}{\sum_{j>1}\lambda_1 - \lambda_j} \tag{12}$$

and σ_1, $\mathbf{u}_1$ are the first singular value and singular vector of $\mathbf{C}_T$.

In Fig 8, we show the results from heterogeneous networks in both the asynchronous (top panel in each subfigure) and strong asynchronous (bottom panel in each subfigure) regimes. We first show $\mathbf{C}_T - \lambda\mathbf{I}$, where λ is given by Eq 12, in Fig 8A. Cells are ordered by (decreasing) firing rate. While no pattern is visible in the asynchronous state (top panel), the strong asynchronous state (bottom panel) shows larger values in the upper left corner, suggesting that correlation increases with firing rate. This is even more visible in the rank one approximation, $(\sigma_1 - \lambda)\mathbf{u}_1\mathbf{u}_1^T$, shown in Fig 8B.

We now use $\mathbf{C}_T^{\text{diag+R1}}$ to approximate $\mathbf{C}_T$, and compare the results, cell pair-by-cell pair (Fig 8C). In the asynchronous network, the approximated correlations take on a narrow range (between 0 and 0.01, compared to between −0.015 and 0.03 for the measured coefficients) and do not show an obvious positive relationship. In the strong asynchronous regime, the range is more accurate (between 0.02 and 0.1, vs. 0.01 and 0.15 for the measured coefficients) and the points cluster around the unity line.

In Fig 8D, we plot the weight of each cell in the first singular vector, $(\mathbf{u}_1)_j$ vs. the firing rate v_j. We can clearly see a positive relationship in the strong asynchronous regime (bottom panel), suggesting that the positive relationship between correlation and firing rate is related to the success of the low-rank approximation.

Discussion

We simulated heterogeneous, asynchronous networks of leaky integrate-and-fire model neurons in order to investigate a possible relationship between firing rates and pairwise

correlations in recurrent networks. We found that correlations can either increase *or* decrease with firing rates; this could be attributed to differences in how cells responded to fluctuations in inhibitory conductances. When correlations *did* increase with firing rates, this relationship was reflected in low-dimensional structure in the correlation matrix.

This study offers an example of a practical consequence of the difference between treating synaptic inputs as *conductances* rather than *currents*; while most synaptic currents are more accurately modeled as conductances, current-based formulations are often used for analytical and computational simplicity. Although it is known that neural models responding to currents vs. conductances differ in their response dynamics [44–46], this approach is supported by findings that steady-state firing rates are qualitatively similar in both settings (e.g. [47]). Here, we found that refined features of the steady-state firing rate surface will govern susceptibility to common input in asynchronous networks; two "cuts" through this surface may yield divergent behavior with respect to correlation susceptibility, despite yielding similar firing rates.

In other words, the relationship between pairwise correlations and firing rate will depend on the means through which firing rate diversity is achieved. In our study, we created firing rate diversity by regulating cell excitability; if we had instead varied mean inhibitory input (by varying the number of inhibitory connections) or a background excitatory current (which would model diversity in stimulus tuning from feed-forward inputs), we would likely have seen a different pattern. Finally, the recurrent network will also shape the path the cells follow through the "firing rate surface"; to generalize [22] to recurrent networks, we need to identify both how firing rates are produced and how they are shaped by the recurrent network.

Thus far, we have not directly connected the presence or absence of a positive correlation-firing rate relationship to other firing statistics (such as being in the asynchronous vs. strong asynchronous state, for example). We believe this will be a challenging question to answer; indeed, we showed earlier (in Results**: Susceptibility depends on the mechanism underlying firing rate diversity**) that we can construct a network with asynchronous firing, but where correlations *do* increase with firing rate. Therefore, our goal with this paper is to present a detailed procedure for analyzing how correlations will vary with firing rates in recurrent networks, along with a few illustrative (and nonintuitive) examples.

Finally, while the networks considered in this paper had fixed in-degrees (rather than Erdős-Rényi), this is not necessary. We have reproduced these results in Erdős-Rényi networks, in which parameters are identical to those chosen here, except that each network connection was chosen independently with a probability that depended only on E/I identity. In S9 Fig, we show correlations vs. geometric firing rates, for all excitatory pairs (as in Fig 2), confirming that correlations increased with firing rates in the strong asynchronous network, but not in the asynchronous network. Furthermore, we hypothesize that any difference between fixed in-degree and Erdős-Rényi networks will become less rather than more important with increasing network size, as the variance in synaptic inputs decreases.

Low-dimensional structure has been a common finding in many large-scale neural recordings [25, 38–41]; while the origin is not always known, it is often interpreted as arising from a global input or top-down signal. This is an interpretation that arises naturally from the technique of *factor analysis*, in which one seeks to explain a data vector as the sum of a random vector and the linear combination of some number of latent factors [48] (for Gaussian random variables, each latent factor can literally be interpreted as a global input with a distinct pattern of projection onto the observed variables). In our network, we found that a single latent factor was effective at capturing correlations in the strong asynchronous regime; however, this latent factor did not reflect common input (there was no global external input into the network) but rather modulation from single-cell characteristics. Thus, we identify a novel mechanism that can contribute to low-dimensional structure in neural recordings.

Impact on stimulus coding

The networks studied here were not encoding a stimulus; correlations were generated by recurrent activity, given that each neuron had a baseline firing rate in the absence of recurrent input. However, we can readily connect this network to a stimulus coding task, in order to understand how the correlation-firing rate relationship can impact coding.

Consider a population of cells that is responsible for encoding a single scalar stimulus θ, such as movement direction or orientation of a visual stimulus, and that each cell has roughly a bell-shaped tuning curve. Furthermore, we model an incoming stimulus by modulating a stimulus-dependent background current $I_{i,}(\theta)$; i.e., cells which prefer the current stimulus have a higher level of current, and thus a higher firing rate, than cells which prefer an orthogonal or opposite stimulus. The network we studied here would model the response to a single stimulus θ_0; that is, the firing rate diversity we observe is present because some cells are strongly tuned to the current stimulus, while others are not.

We could extend this model, by resetting background currents to model a complete set of stimuli $\{\theta_1, \theta_2, \ldots \theta_{n-1}\}$. For each stimulus θ_j, correlations would show the rough firing rate dependence displayed in the strong asynchronous network, resulting in a *stimulus-dependent* correlation structure in which pairwise correlations vary like geometric mean firing rate. This is the structure analyzed in [21, 27]: the authors found that such a stimulus-dependent correlation code enhances information, when compared to a stimulus-independent code with the same average correlation level. Intuitively, the mean population response lives on the surface of a (hyper-)sphere in neural response space; the population encodes location on this surface. Positive correlations between similarly tuned cells produce response distributions that are stretched in the radial direction, "orthogonal" to this sphere, and thus have a minimal impact on the encoded variable.

Moreover, the mechanism that produced stimulus-dependent correlations in [21, 27] was similar to that shown here (see also [25]); common input modulated by stimulus-dependent gain factors. Here, we demonstrated how these stimulus-dependent gain factors might arise (or not) in a recurrent network. If excitation is tuned to put the network in the strong asynchronous regime, then the (stimulus-dependent) correlation structure that results will be favorable to coding. If excitation is tuned to put the network in the asynchronous regime, then correlations are overall low and not stimulus-dependent (although, given that average correlations are not matched, we do not here compare information contained within the two networks).

Future work

This work has, necessarily, focused only on a subset of network attributes that might affect firing statistics. One important feature is the frequency of higher-order graph motifs; experiments have shown that specific motifs will occur more frequently, than would be expected in an Erdős-Rényi network with fixed single-cell connection probability [49]. Theoretical work has found that in networks of integrate-and-fire neurons, an overabundance of divergent and chain motifs will lead to enhanced correlation [33] (this finding does depend on the dynamical regime; different motifs impact correlations in networks of coupled oscillators [32]). In [33], the authors use the assumption of homogeneous single-cell characteristics to find parsimonious and instructive formulae for the average correlation, and give a roadmap for how this might be generalized to heterogeneous networks. We look forward to considering the combined effect of single-cell *and* network heterogeneity in future work.

Another source of cell-to-cell heterogeneity is how cells respond to stimuli, as emphasized in the previous discussion [17, 20, 21, 27, 50] (see [19] for a review). Here, we did not consider

a specific sensory system with tuning but rather focus on the general question of how the distribution of correlation values arise in recurrent networks. Given the previous discussion, one next step will be to investigate how correlations covary with firing rates, when cell-to-cell heterogeneity is produced by stimulus tuning in a structured network responding to a single variable (such as direction or orientation).

Finally, for numerical tractability our simulations here were performed in relatively small networks. While high average correlations have been measured in experiments [51], theoretical models of asynchronous networks have found that correlations must go to zero as the system becomes large ($N \to \infty$) [2]. However, recent work has found that this does not have to be true, as long as spatial structure is introduced into the network [52]. We anticipate that this may carry over to other forms of heterogeneity, such as single-cell variability, and that therefore the effect we observe here persists for larger networks. We look forward to reporting on this in future work.

Methods

Neuron model and network setup

We considered randomly connected networks of excitatory and inhibitory neurons. Each cell was a linear integrate-and-fire model with second-order alpha-conductances, i.e. membrane voltage v_i was modeled with a stochastic differential equation, as long as it remained beneath a threshold θ_i:

$$\tau_m \frac{dv_i}{dt} = -v_i - g_{E,i}(t)(v_i - \mathcal{E}_E) - g_{I,i}(t)(v_i - \mathcal{E}_I) + \sigma_i \sqrt{\tau_m} \xi_i(t), \tag{13}$$

When v_i reaches θ_i, it is reset to 0 following a refractory period:

$$v_i(t + \tau_{\text{ref}}) \to 0, \qquad v_i(t) \geq \theta_i \tag{14}$$

Each neuron was driven by a Gaussian, white background noise, with magnitude σ_i depending only on the cell type; that is, $\langle \xi_i(t) \rangle = 0$ and $\langle \xi_i(t)\xi_i(t + s) \rangle = \delta(s)$. The membrane time constant, τ_m, and excitatory and inhibitory synaptic reversal potentials, $\mathcal{E}_E$ and $\mathcal{E}_I$, are the same for every cell in the network.

Each cell responded to synaptic input through conductance terms, $g_{E,i}$ and $g_{I,i}$, which are each governed by a pair of differential equations:

$$\tau_{d,X} \frac{dg_{X,i}}{dt} = -g_{X,i} + g_{X,i}^{(1)} \tag{15}$$

$$\tau_{r,X} \frac{dg_{X,i}^{(1)}}{dt} = -g_{X,i}^{(1)} + \tau_{r,X} \alpha_X \left(\frac{W_{YX}}{N_{YX}} \right) \sum_{j \in X, j \to i} \sum_k \delta(t - t_{j,k}) \tag{16}$$

where $Y = \{E, I\}$ denotes the type of cell i and $X = \{E, I\}$ denotes the type of the source neuron j. Each spike is modeled as a delta-function that impacts the auxiliary variable $g_{X,i}^{(1)}$; here $t_{j,k}$ is the k-th spike of cell j. The rise and decay time constants $\tau_{r,X}$ and $\tau_{d,X}$ and pulse amplitude α_X depend only on the type of the source neuron; i.e. they are otherwise the same across the population. The parameter W_{YX} denotes the strength of $X \to Y$ synaptic connections, which are (once given the type of source and target neurons) identical across the population. The "raw" synaptic weight (listed in Table 1) is divided by N_{YX}, the total number of $X \to Y$ connections received by each Y-type cell.

We chose connections to be homogeneous and relatively dense, consistent with the local architecture of cortex. Connection probabilities ranged from 20%–40%, consistent with

Table 2. Other parameters used in network simulations.

Parameter	Definition	X = E	X = I
$\tau_{r,X}$	Synaptic rise time	1 ms	2 ms
$\tau_{d,X}$	Synaptic decay time	5 ms	10 ms
τ_m	Membrane time constant	20 ms	20 ms
τ_{ref}	Refractory time	2 ms	2 ms
α_X	Pulse amplitude	1	2
$\mathcal{E}_X$	Synaptic reversal potential	6.5	-0.5

experimentally measured values [53–55]. For our baseline network state, we then chose synaptic weights so the network is moderately inhibition-dominated ($\alpha_E W_{IE} < \alpha_I W_{II}$ and $\alpha_E W_{EE} < \alpha_I W_{EI}$); that is both *E* and *I* cells receive more inhibition than excitation) and shows noisy spiking consistent with the classical asynchronous state. Each neuron receives a fixed number of incoming connections, the identities of which are chosen randomly. (The specific cell ID numbers differ in the different simulations shown below.) For most of the networks we discuss here $N = 100$ with the 80/20 ratio typical of cortex (i.e $n_E = 80$, $n_I = 20$). Each excitatory cell receives $N_{EE} = 32$ (40%) excitatory and $N_{EI} = 7$ (35%) inhibitory connections; each inhibitory cell receives $N_{IE} = 16$ (20%) and $N_{II} = 8$ (40%) inhibitory projections.

In heterogeneous networks, the threshold θ_i varied across the population. For both excitatory and inhibitory neurons, the thresholds θ_i were chosen from a log-normal distribution between 0.7 and 1.4 (where the rest potential, $V_r = 0$). To be precise, $\log \theta_i$ was chosen from a (truncated) normal distribution with mean $-s_\theta^2/2$ and standard deviation s_θ. With this choice, θ_i has mean 1 and variance: $e^{s_\theta^2} - 1$. Thus we can view s_θ as a measure of the level of threshold heterogeneity.

Throughout this paper, we set $s_\theta = 0.2$, which results in a wide range of firing rates compared to the homogeneous case. This was the only source of cell-to-cell heterogeneity; all other parameters were identical across the population, conditioned on neuron type (values listed in Table 2). In homogeneous networks, the threshold was the same across the population: $\theta_i = 1$.

Monte Carlo simulations were performed using the stochastic forward- Euler method (Euler-Maruyama), with a time step much smaller than any time scale in the system ($\Delta t = 0.01$ ms). Each network was simulated for one second of simulation time, after an equilibration time. Then, a large number of realizations of this interval ($n_R = 10^5$) were simulated. Spike counts were retained in each 1 ms window (for a total of 1000 windows) within a realization. With this large number of realizations/trials, the error bars on the resulting time-dependent firing rates were small. Therefore we emphasize that the firing rate pattern is largely driven by network connectivity; while firing is driven by random fluctuations in the background noise, any cell-to-cell variability in the *trial-averaged* firing rates are not an artifact of the finite number of trials.

Linear response theory

In general, computing the response of even a single neuron to an input requires solving a complicated, nonlinear stochastic process. However, it often happens that the presence of background noise linearizes the response of the neuron, so that we can describe this response as a perturbation from a background state. This response is furthermore linear in the perturbing input and thus referred to as *linear response* theory [56]. The approach can be generalized to yield the dominant terms in the coupled network response, as well; we will use the theory to predict the covariance matrix of activity.

We first consider the case of a single cell: an LIF neuron responding to a mean zero current $\epsilon X_i(t)$

$$\tau_m \frac{dv_i}{dt} = -(v_i - E_L) + E_i + \sigma_i \sqrt{\tau_m} \xi_i(t) + \epsilon X_i(t).$$

(otherwise, the mean of X_i can simply be absorbed into E_i).

For a fixed input $\epsilon X_i(t)$, the output spike train $y_i(t)$ will be slightly different for each realization of the noise $\xi_i(t)$ and initial condition $v_i(0)$. Therefore we try to work with the time-dependent firing *rate*, $v_i(t) \equiv \langle y_i(t) \rangle$, which is obtained by averaging over all realizations and initial conditions. Linear response theory proposes the ansatz that the firing rate can be described as a perturbation from a baseline rate proportional to the input ϵX_i:

$$v_i(t) = v_{i,0} + (A_i * \epsilon X_i)(t); \tag{17}$$

$v_{i,0}$ is the baseline rate (when $X = 0$) and $A_i(t)$ is a *susceptibility function* that characterizes this firing rate response up to order ϵ [22, 29, 57].

We now consider the theory for networks; here cell i responds to the spike train of cell j, $y_j(t)$, via the synaptic weight matrix $\mathbf{W}$, after convolution with a synaptic filter $F_j(t)$:

$$\tau_m \frac{dv_i}{dt} = -(v_i - E_L) + E_i + \sigma_i \sqrt{\tau_m} \xi_i(t) + \sum_j \mathbf{W}_{ij} F_j * y_j(t)$$

In order to consider joint statistics, we need the trial-by-trial response of the cell. We first propose to approximate the response of each neuron as:

$$y_i(t) \approx y_i^0(t) + \left(A_i * \sum_j (\mathbf{J}_{ij} * y_j) \right)(t); \tag{18}$$

that is, each input X_i has been replaced by the synaptic input, and $\mathbf{J}_{ij} = \mathbf{W}_{ij} F_j(t)$ includes both the $i \leftarrow j$ synaptic weight W_{ij} and synaptic kernel F_j (normalized to have area 1); $A_i(t)$ is the susceptibility function from Eq 17. In the frequency domain this becomes

$$\tilde{y}_i(\omega) = \tilde{y}_i^0 + \tilde{A}_i(\omega) \left(\sum_j \tilde{\mathbf{J}}_{ij}(\omega) \tilde{y}_j(\omega) \right) \tag{19}$$

where $\tilde{y}_i = \mathcal{F}[y_i - v_i]$ is the Fourier transform of the mean-shifted process (v_i is the average firing rate of cell i) and $\tilde{f} = \mathcal{F}[f]$ for all other quantities. In matrix form, this yields a self-consistent equation for $\tilde{y}$ in terms of $\tilde{y}^0$:

$$\left(\mathbf{I} - \tilde{\mathbf{K}}(\omega) \right) \tilde{y} = \tilde{y}^0 \Rightarrow \tilde{y} = \left(\mathbf{I} - \tilde{\mathbf{K}}(\omega) \right)^{-1} \tilde{y}^0 \tag{20}$$

where $\tilde{\mathbf{K}}_{ij}(\omega) = \tilde{A}_i(\omega) \tilde{\mathbf{J}}_{ij}(\omega)$ is the interaction matrix, in the frequency domain. The cross-spectrum is then computed

$$\langle \tilde{y}(\omega) \tilde{y}^*(\omega) \rangle = \left(\mathbf{I} - \tilde{\mathbf{K}}(\omega) \right)^{-1} \langle \tilde{y}^0(\omega) \tilde{y}^{0*}(\omega) \rangle \left(\mathbf{I} - \tilde{\mathbf{K}}^*(\omega) \right)^{-1} \tag{21}$$

To implement this calculation, we first solve for a self-consistent set of firing rates: that is, v_i is the average firing rate of

$$\tau_m \frac{dv_i}{dt} = -(v_i - E_L) + (E_i + \mathbb{E}[f_i]) + \sigma_i \sqrt{\tau_m} \xi_i(t) \tag{22}$$

where $\mathbb{E}[f_i] = \Sigma_j \mathbf{W}_{ij} v_j$.

We must then compute the power spectrum $\langle \tilde{y}^0(\omega) \tilde{y}^{0*}(\omega) \rangle$ and the susceptibility $A_i(\omega)$, which is the (first order in ϵ) response in the firing rate $r_i(t) = r_i^0 + \epsilon A_i(\omega) \exp(\imath \omega t)$ in response to an input current perturbation $X(t) = \epsilon \exp(\imath \omega t)$ (here $\imath$ is used for $\sqrt{-1}$, while i denotes an index). Both can be expressed as the solution to (different) first-order boundary value problems and solved via Richardson's threshold integration method [47, 58].

In our simulations, we used conductance-based neurons; this requires modification, compared with the simpler current-based models. We first approximate each conductance-based neuron as an effective current-based neuron with reduced time constant, following the discussion in [59]. First, separate each conductance into mean and fluctuating parts; e.g. $g_{E,i} \rightarrow \langle g_{E,i} \rangle + (g_{E,i} - \langle g_{E,i} \rangle)$. Then we identify an effective conductance $g_{0,i}$ and potential μ_i, and treat the fluctuating part of the conductances as noise, i.e. $g_{E,i} - \langle g_{E,i} \rangle \rightarrow \sigma_{E,i} \xi_{E,i}(t)$:

$$\tau_m \frac{dv_i}{dt} = -g_{0,i}(v_i - \mu_i) + \sigma_{E,i} \xi_{E,i}(t)(v_i - \mathcal{E}_E) + \sigma_{I,i} \xi_{I,i}(t)(v_i - \mathcal{E}_I) + \sqrt{\sigma_i^2 \tau_m} \xi_i(t) \tag{23}$$

where

$$g_{0,i} = 1 + \langle g_{E,i} \rangle + \langle g_{I,i} \rangle \tag{24}$$

$$\mu_i = \frac{E_L + E_i + \langle g_{E,i} \rangle \mathcal{E}_E + \langle g_{I,i} \rangle \mathcal{E}_I}{g_{0,i}} \tag{25}$$

$$\sigma_{E,i}^2 = \text{Var}\left[g_{E,i}(t)\right] = \text{E}\left[\left(g_{E,i}(t) - \langle g_{E,i} \rangle\right)^2\right] \tag{26}$$

$$\sigma_{I,i}^2 = \text{Var}\left[g_{I,i}(t)\right] = \text{E}\left[\left(g_{I,i}(t) - \langle g_{I,i} \rangle\right)^2\right] \tag{27}$$

We next simplify the noise terms by writing

$$v_i - \mathcal{E}_E = v_i - \mu_i + \mu_i - \mathcal{E}_E \tag{28}$$

and assume that the fluctuating part of the voltage, $v_i - \mu_i$, is mean-zero and uncorrelated with the noise terms $\xi_{E,i}(t)$ [59]. That allows us to define an effective equation

$$\tau_{\text{eff},i} \frac{dv_i}{dt} = -(v_i - \mu_i) + \sqrt{\sigma_{\text{eff},i}^2 \tau_{\text{eff},i}} \eta_{\text{eff},i}(t) \tag{29}$$

where

$$\tau_{\text{eff},i} = \frac{\tau_m}{g_{0,i}} \tag{30}$$

$$\sigma_{\text{eff},i}^2 = \frac{\sigma_{E,i}^2(\mu_i - \mathcal{E}_E)^2 + \sigma_{I,i}^2(\mu_i - \mathcal{E}_I)^2 + \sigma_i^2 \tau_m}{g_{0,i} \tau_m} \tag{31}$$

and the fluctuating voltage, $v_i(t) - \mu_i$, now makes no contribution to the effective noise variance.

Finally, we consider how to model the conductance mean and variance, e.g. $\langle g_{E,i} \rangle$ and $\sigma_{E,i}^2$. In our simulations, we used second order α-functions: each conductance $g_{X,i}$ is modeled by

two equations that take the form

$$\tau_{r,X} \frac{dg_{X,i}^{(1)}}{dt} = -g_{X,i}^{(1)} + \tau_{r,X} \hat{\alpha}_{X,i} \sum_k \delta(t - t_k) \tag{32}$$

$$\tau_{d,X} \frac{dg_{X,i}}{dt} = -g_{X,i} + g_{X,i}^{(1)} \tag{33}$$

where $X = E, I$ and the summation is over all type-X spikes incoming to cell i. (For notation purposes, $\hat{\alpha}_{X,i}$ includes all factors that contribute to the pulse size in Eq 16, including synapse strength and pulse amplitude.) The time constants $\tau_{r,X}, \tau_{d,X}$ may depend on synapse type; the spike jumps $\hat{\alpha}_{X,i}$ may depend on synapse type and target cell identity. We assume that each spike train is Poisson, with a constant firing rate: i.e. each spike train is modeled as a stochastic process $S(t)$ with

$$\langle S(t) \rangle = v$$
$$\langle S(t)S(t + \tau) \rangle - v^2 = v\delta(\tau)$$

Then a straightforward but lengthy calculation shows that

$$\langle g_{X,i}(t) \rangle = \hat{\alpha}_{X,i} v_{X,i} \tau_{r,X} \tag{34}$$

$$\text{Var}\left[g_{X,i}(t)\right] = \left(\frac{1}{2} \hat{\alpha}_{X,i}^2 v_{X,i} \tau_{r,X}\right) \left(\frac{\tau_{r,X}}{\tau_{r,X} + \tau_{d,X}}\right) \tag{35}$$

where $v_{X,i}$ is the total rate of type-X spikes incoming to cell i.

We now describe how these considerations modify the linear response calculation. First, for the self-consistent firing rate calculation, Eq 22 is replaced by an equation with a modified time constant, conductance, and noise (Eq 29).

We next compute the susceptibility in response to parameters associated with the conductance, i.e. $\langle g_{E,i} \rangle$ and $\sigma_{E,i}^2$. This differs from the current-based case in two ways: first, there is voltage-dependence in the diffusion terms, which results in a different Fokker-Planck equation (and thus a different boundary value problem to be solved for the power spectrum $\langle \tilde{y}^0(\omega) \tilde{y}^{0*}(\omega) \rangle$). Second, modulating the rate of an incoming spike train will impact *both* the mean and variance of the input to the effective equation, Eq 23 (via μ_i and $\sigma_{X,i}$). Furthermore, this impact may differ for excitatory and inhibitory neurons, giving us a total of *four* parameters that can be varied in the effective equation. However, neither consideration presents any essential difficulty [47].

Therefore we apply Richardson's threshold integration method directly to Eq 23:

$$\tau_m \frac{dv_i}{dt} = -g_{0,i}(v_i - \mu_i) + \sigma_{E,i} \xi_{E,i}(t)(v_i - \mathcal{E}_E) + \sigma_{I,i} \xi_{I,i}(t)(v_i - \mathcal{E}_I) + \sqrt{\sigma_i^2 \tau_m} \xi_i(t) \tag{36}$$

When we compute susceptibilities, the parameter to be varied is either a mean conductance—$\langle g_{E,i} \rangle \rightarrow \langle g_{E,i} \rangle_0 + \langle g_{E,i} \rangle_1 \exp(\iota\omega t)$ or $\langle g_{I,i} \rangle \rightarrow \langle g_{I,i} \rangle_0 + \langle g_{I,i} \rangle_1 \exp(\iota\omega t)$—or a variance—$\sigma_{E,i}^2 \rightarrow (\sigma_{E,i}^2)_0 + (\sigma_{E,i}^2)_1 \exp(\iota\omega t)$ or $\sigma_{I,i}^2 \rightarrow (\sigma_{I,i}^2)_0 + (\sigma_{I,i}^2)_1 \exp(\iota\omega t)$. Thus we have a total of four susceptibility functions $\tilde{A}_{\langle g_E \rangle, i}(\omega)$, $\tilde{A}_{\langle g_I \rangle, i}(\omega)$, $\tilde{A}_{\sigma_E^2, i}(\omega)$, and $\tilde{A}_{\sigma_I^2, i}(\omega)$. Since the Fokker-Planck equation to be solved is linear, we can compute both susceptibilities separately and then add their

effects. We now have the interaction matrix:

$$\tilde{\mathbf{K}}_{ij}(\omega) \quad = \quad \begin{cases} \tilde{A}_{\langle g_E \rangle, i}(\omega)\tilde{\mathbf{J}}_{ij}(\omega) + \tilde{A}_{\sigma_E^2, i}(\omega)\tilde{\mathbf{L}}_{ij}(\omega), & j \text{ excitatory} \\[2mm] \tilde{A}_{\langle g_I \rangle, i}(\omega)\tilde{\mathbf{J}}_{ij}(\omega) + \tilde{A}_{\sigma_I^2, i}(\omega)\tilde{\mathbf{L}}_{ij}(\omega), & j \text{ inhibitory} \end{cases} \tag{37}$$

where $\tilde{\mathbf{L}}(\omega)$ plays a similar role as $\tilde{\mathbf{J}}$, but for the effect of incoming spikes on the *variance* of conductance. Its relationship to $\tilde{\mathbf{J}}$ (either in the frequency or time domain) is given by the same simple scaling shown in Eq 35: i.e., for j excitatory,

$$\tilde{\mathbf{L}}_{ij}(\omega) \quad = \quad \tilde{\mathbf{J}}_{ij}(\omega) \times \left(\frac{\hat{\alpha}_{E,i}}{2}\right) \times \left(\frac{\tau_{r,E}}{\tau_{r,E} + \tau_{d,E}}\right) \tag{38}$$

where the first factor comes from the effective spike amplitude $\hat{\alpha}_{E,i}$ (and is the scale factor proposed in [47], Eq (64)), and the second arises from using second-order (vs. first-order) alpha-functions.

We use a modified version of the implementation given by [29] for Richardson's threshold integration algorithm [47, 58] to compute rate v_i, power $\langle \tilde{y}_i^0(\omega)\tilde{y}_i^{0*}(\omega)\rangle$, and the various susceptibilities ($\tilde{A}_{\langle g_E \rangle, i}(\omega)$, $\tilde{A}_{\langle g_I \rangle, i}(\omega)$, $\tilde{A}_{\sigma_E^2, i}(\omega)$, and $\tilde{A}_{\sigma_I^2, i}(\omega)$) for an LIF neuron. We validated our code using exact formulas known for the LIF [60], and qualitative results from the literature [61].

Computing statistics from linear response theory

Linear response theory yields the cross spectrum of the spike train, $\langle \tilde{y}_i(\omega)\tilde{y}_j^*(\omega)\rangle$, for each distinct pair of neurons i and j (see Eq 21). To recover a representative set of statistics, we rely on several standard formulae relating this function to other statistical quantities.

The cross correlation function, $\mathbf{C}_{ij}(\tau)$, measures the similarity between two processes at time lag τ, while the cross spectrum measures the similarity between two processes at frequency ω:

$$\mathbf{C}_{ij}(\tau) \equiv \langle (y_i(t) - v_i)(y_j(t + \tau) - v_j)\rangle \tag{39}$$

$$\tilde{\mathbf{C}}_{ij}(\omega) \equiv \langle \tilde{y}_i(\omega)\tilde{y}_j(\omega)\rangle \tag{40}$$

The Weiner-Khinchin theorem [56] implies that $\{\mathbf{C}_{ij}, \tilde{\mathbf{C}}_{ij}\}$ are a Fourier transform pair: that is,

$$\tilde{\mathbf{C}}_{ij}(\omega) = \int_{-\infty}^{\infty} \mathbf{C}_{ij}(t)e^{-2\pi \iota \omega t}\, dt \tag{41}$$

In principle, the crosscorrelation $\mathbf{C}(t)$ and cross-spectrum $\tilde{\mathbf{C}}(\omega)$ matrices are functions on the real line, reflecting the fact that correlation can be measured at different time scales. In particular, for a stationary point process the covariance of spike counts over a window of length T, n_i and n_j, can be related to the crosscorrelation function $\mathbf{C}_{ij}$ by the following formula [4]:

$$\text{Cov}_T(n_i, n_j) = \int_{-T}^{T} \mathbf{C}_{ij}(\tau)(T - |\tau|)\, d\tau \tag{42}$$

The variance of spike counts over a time window of length T, n_i, is likewise given by integrating the autocorrelation function $\mathbf{C}_{ii}$:

$$\text{Var}_T(n_i) = \int_{-T}^{T} \mathbf{C}_{ii}(\tau)(T - |\tau|)\, d\tau \tag{43}$$

It can be helpful to normalize by the time window, i.e.

$$\frac{\text{Cov}_T(n_i, n_j)}{T} = \int_{-T}^{T} \mathbf{C}_{ij}(\tau)\left(1 - \frac{|\tau|}{T}\right) d\tau; \tag{44}$$

we can now see that for an integrable cross correlation function (and bearing in mind that the cross-spectrum is the Fourier transform of the cross correlation), that

$$\lim_{T \to \infty} \frac{\text{Cov}_T(n_i, n_j)}{T} = \int_{-\infty}^{\infty} \mathbf{C}_{ij}(\tau)d\tau = \tilde{\mathbf{C}}_{ij}(0) \tag{45}$$

while

$$\lim_{T \to 0} \frac{\text{Cov}_T(n_i, n_j)}{T^2} = \frac{1}{T}\int_{-T}^{T} \mathbf{C}_{ij}(\tau)\left(1 - \frac{|\tau|}{T}\right) d\tau \approx \mathbf{C}_{ij}(0) \tag{46}$$

Thus, we can use $\tilde{\mathbf{C}}_{ij}(0)$ and $\mathbf{C}_{ij}(0)$ as measures of long and short time correlations respectively.

Finally, the Pearson's correlation coefficient of the spike count defined as:

$$\rho_{T,ij} = \frac{\text{Cov}_T(n_i, n_j)}{\sqrt{\text{Var}_T(n_i)\,\text{Var}_T(n_j)}} \tag{47}$$

is a common normalized measure of noise correlation, with $\rho \in [-1, 1]$. While Cov_T and Var_T grow linearly with T (for a Poisson process, for example), $\rho_{T,ij}$ in general will not (although it may increase with T). In general, $\rho_{T,ij}$ depends on the time window T; however for readability we will often suppress the T-dependence in the notation (and use ρ_{ij} instead).

Quantifying the role of motifs in networks

We next explain how we can use the results of linear response theory to give insight into the role of different paths in the network. We begin with our predicted cross-spectrum (Eqs 21 and 40) and apply a standard series expansion for the matrix inverse:

$$\tilde{\mathbf{C}}(\omega) = \left(\mathbf{I} - \tilde{\mathbf{K}}(\omega)\right)^{-1}\tilde{\mathbf{C}}^0(\omega)\left(\mathbf{I} - \tilde{\mathbf{K}}^*(\omega)\right)^{-1} \tag{48}$$

$$= \left[\sum_{k=0}^{\infty} \left(\tilde{\mathbf{K}}(\omega)\right)^k\right]\tilde{\mathbf{C}}^0(\omega)\left[\sum_{l=0}^{\infty} \left(\tilde{\mathbf{K}}(\omega)\right)^l\right] \tag{49}$$

$$= \sum_{k=0}^{\infty}\sum_{l=0}^{\infty} \left(\tilde{\mathbf{K}}(\omega)\right)^k\tilde{\mathbf{C}}^0(\omega)\left(\tilde{\mathbf{K}}(\omega)\right)^l \tag{50}$$

where $\tilde{\mathbf{C}}^0(\omega)$ is a diagonal matrix containing the power spectra of the unperturbed processes; i.e. $\tilde{\mathbf{C}}^0_{ii} \equiv \langle \tilde{y}_i(\omega)\tilde{y}_i(\omega)\rangle$. This double sum will converge as long as the spectral radius of $\tilde{\mathbf{K}}$ is less than 1 [29].

By truncating this double sum to contain terms such that $k + l \leq n$, we define the nth approximation to the cross-spectrum:

$$\tilde{\mathbf{C}}(\omega) \approx \tilde{\mathbf{C}}^n(\omega) \tag{51}$$

$$= \tilde{\mathbf{C}}^0(\omega) + \sum_{k=1}^{n}\left[\sum_{l=0}^{k} \left(\tilde{\mathbf{K}}(\omega)\right)^{k-l}\tilde{\mathbf{C}}^0(\omega)\left(\tilde{\mathbf{K}}^*(\omega)\right)^l\right] \tag{52}$$

Each distinct term in the inner sum can be attributed to a particular undirected path of length k. Terms of the form $\tilde{\mathbf{K}}^k\tilde{\mathbf{C}}^0$ and $\tilde{\mathbf{C}}^0(\tilde{\mathbf{K}}^*)^k$ account for unidirectional paths from $j \to i$ and $i \to j$ respectively; the term $(\tilde{\mathbf{K}}(\omega))^{k-l}\tilde{\mathbf{C}}^0(\omega)(\tilde{\mathbf{K}}^*(\omega))^l$ captures the contribution from a cell that has a length l path onto cell j and a length $k - l$ path onto cell i. Thus, we can use Eq 52 to decompose the correlation into contributions from different motifs ([28], see also [31, 62]).

We can also consider the contribution from all length-n paths; that is,

$$\tilde{\mathbf{P}}^n = \tilde{\mathbf{C}}^n(\omega) - \tilde{\mathbf{C}}^{n-1}(\omega) = \sum_{l=0}^{n} (\tilde{\mathbf{K}}(\omega))^{n-l}\tilde{\mathbf{C}}^0(\omega)(\tilde{\mathbf{K}}^*(\omega))^l$$

If the sum in Eq 50 converges, we should expect the magnitude of contributions to decrease as n increases.

We will also show the *normalized* contribution from length-n paths, which we define as follows: let $\mathbf{\Lambda}(\omega)$ be the diagonal matrix with $\mathbf{\Lambda}_{ii}(\omega) = \tilde{\mathbf{C}}_{ii}(\omega)$. Then we define the matrix of contributions from length-n paths $\tilde{\mathbf{R}}^n$ as follows:

$$\tilde{\mathbf{R}}^n(\omega) = \mathbf{\Lambda}^{-1/2}(\omega)\tilde{\mathbf{P}}^n(\omega)\mathbf{\Lambda}^{-1/2}(\omega) \tag{53}$$

Equivalently, $\tilde{\mathbf{R}}_{ij}^n(\omega) = \tilde{\mathbf{P}}_{ij}^n(\omega)/\sqrt{\tilde{\mathbf{C}}_{ii}(\omega)\tilde{\mathbf{C}}_{jj}(\omega)}$. This effectively normalizes the cross correlation by the autocorrelation; in particular, we can use this to decompose the correlation coefficient (Eq 47) for long time windows, because $\lim_{n\to\infty} \sum_{k=0}^{n} \tilde{\mathbf{R}}^k(0) = \lim_{T\to\infty} \rho_{T,ij}$.

In general, we will show long-timescale correlation (e.g. $\tilde{\mathbf{C}}(0)$ or $\tilde{\mathbf{R}}^n(0)$) (Eq 45); results were qualitatively similar for other timescales.

Quantifying correlation susceptibility

We next consider how to quantify the (linear) susceptibility of correlation to a change in parameter. Returning to Eq 17, but written in terms of the single-cell response:

$$y_i(t) = y_{i,0} + (A_{\mu,i} * X_\mu)(t) \qquad \Rightarrow \tag{54}$$

$$\tilde{y}_i(\omega) = \tilde{y}_{i,0}(\omega) + \tilde{A}_{\mu,i}(\omega)\tilde{X}_\mu(\omega) \tag{55}$$

Here, $X_\mu(t)$ is a (possibly) time-dependent change in a parameter, such as input current or mean inhibitory conductance; $y_{i,0}$ is the baseline spike train (when $X = 0$). $A_{\mu,i}(t)$ is a *susceptibility function* that characterizes the cell's response (to the parameter variation) as long as $X_\mu(t)$ is small [22, 29, 57]. Following [22], the cross-spectrum of y can now be approximated as:

$$\tilde{C}_{ij}(\omega) \equiv \langle \tilde{y}_i^* \tilde{y}_j \rangle \approx \langle \tilde{y}_{i,0}^* \tilde{y}_{j,0} \rangle + \langle \tilde{A}_{\mu,i}^* \tilde{X}_\mu^* \tilde{y}_{j,0} \rangle + \langle \tilde{A}_{\mu,j} \tilde{X}_\mu \tilde{y}_{i,0}^* \rangle + \tilde{A}_{\mu,i}^* \tilde{A}_{\mu,j} \langle \tilde{X}_\mu^* \tilde{X}_\mu \rangle \tag{56}$$

$$= \tilde{A}_{\mu,i}^*(\omega)\tilde{A}_{\mu,j}(\omega)\tilde{C}_\mu(\omega) \tag{57}$$

where $\tilde{C}_\mu(\omega)$ is the spectrum of the parameter variation. The susceptibility has an appealing interpretation in the limit $\omega \to 0$, as the derivative of the classical *f-I curve*:

$$\lim_{\omega\to 0} \tilde{A}_{\mu,i}(\omega) = \frac{dv_i}{d\mu} \tag{58}$$

where v_i is the steady-state firing rate of cell i, assuming we can measure it for specific values of

the parameter μ.

$$\lim_{T\to\infty}\rho_{T,ij} \quad = \quad \lim_{T\to\infty}\frac{\mathrm{Cov}_T(n_i,n_j)}{\sqrt{\mathrm{Var}_T(n_i)\,\mathrm{Var}_T(n_j)}} = \frac{\tilde{\mathbf{C}}_{ij}(0)}{\sqrt{\tilde{\mathbf{C}}_{ii}(0)\tilde{\mathbf{C}}_{jj}(0)}} \tag{59}$$

$$\approx \quad \frac{\tilde{A}_{\mu,i}(0)\tilde{A}_{\mu,j}(0)}{\sqrt{\tilde{\mathbf{C}}_{ii}(0)\tilde{\mathbf{C}}_{jj}(0)}}\,\tilde{C}_{\mu}(0) \tag{60}$$

This motivates the definition of a *correlation susceptibility*, which approximates the change in pairwise correlation induced by a parameter change experienced by both cells i and j:

$$S_{ij}^{\mu} \quad = \quad \frac{\tilde{A}_{\mu,i}(0)\tilde{A}_{\mu,j}(0)}{\sqrt{\tilde{\mathbf{C}}_{ii}(0)\tilde{\mathbf{C}}_{jj}(0)}} \tag{61}$$

If this increases with firing rate—that is, if $\frac{dS_{ij}^{\mu}}{dv} > 0$—then correlations will also increase with firing rate.

We can further analyze this quantity by making an assumption for asynchronous spiking, that spike count variance is equal to spike count mean; i.e. $\mathrm{Var}_T(n_i) = Tv_i \Rightarrow \tilde{\mathbf{C}}_{ii} = v_i$. Then

$$S_{ij}^{\mu} \quad \approx \quad \frac{1}{\sqrt{v_i v_j}}\tilde{A}_{\mu,i}(0)\tilde{A}_{\mu,j}(0) = \frac{\tilde{A}_{\mu,i}(0)}{\sqrt{v_i}}\frac{\tilde{A}_{\mu,j}(0)}{\sqrt{v_j}} \tag{62}$$

which motivates the definition of the single-cell quantity

$$S_i^{\langle g_I\rangle} \equiv \frac{\tilde{A}_{\langle g_I\rangle,i}(0)}{\sqrt{v_i}}$$

In general, the firing rate depends on *all* single cell parameters included in Eqn.; i.e. there exists some function f such that

$$v_i = f\left(\langle g_{I,i}\rangle, \sigma_{I,i}, \langle g_{E,i}\rangle, \sigma_{E,i}, \sigma_i, \theta_i\right) \tag{63}$$

$$\tilde{A}_{\langle g_I\rangle,i}(0) \quad = \quad \frac{\partial f}{\partial x_1}\left(\langle g_{I,i}\rangle, \sigma_{I,i}, \langle g_{E,i}\rangle, \sigma_{E,i}, \sigma_i, \theta_i\right) \tag{64}$$

(recall that the susceptibility for $\omega = 0$ is the derivative of the firing rate with respect to the appropriate parameter (here, mean inhibitory conductance $\langle g_I\rangle$).

Low-rank approximation to the correlation matrix

We consider the correlation matrix of spike counts, as measured from Monte Carlo simulations; while these are in principle related to the cross-correlation functions $\mathbf{C}(t)$ defined in **Methods: Computing statistics from linear response theory** we will use $\mathbf{C}_T$ to denote the matrix of correlation coefficients measured for time window T; i.e.

$$(\mathbf{C}_T)_{ij} = \rho_{T,ij} \tag{65}$$

Furthermore, we will restrict to the E-E correlations; i.e. $\mathbf{C}_T$ will be a $n_E \times n_E$ matrix, with ones on the diagonal (as $\rho_{T,ii} = 1$).

When we examined the singular values of the E-E correlation matrices obtained from Monte Carlo simulations, we noticed a consistent trend: there was usually one large cluster

with one positive outlier. This motivates the following simple idea: by subtracting off a multiple of the identity matrix, $\lambda \mathbf{I}$, we shift the cluster towards zero; consequently $\mathbf{C}_T - \lambda \mathbf{I}$ is close to a rank-1 matrix. We then propose to use the sum of the two as an approximation to $\mathbf{C}_T$:

$$\mathbf{C}_T \approx \lambda \mathbf{I} + (\sigma_1 - \lambda)\mathbf{u}_1 \mathbf{u}_1^T. \tag{66}$$

We seek the value λ which maximizes the fraction of the Frobenius norm explained by the first singular vector: i.e. in terms of the singular values,

$$\lambda = \max_{\lambda} \frac{\tilde{\sigma}_1^2}{\sum_{j=1}^{r} \tilde{\sigma}_j^2} \tag{67}$$

$$= \max_{\lambda} \frac{(\sigma_1 - \lambda)^2}{\sum_{j=1}^{r}(\sigma_j - \lambda)^2} \tag{68}$$

Since $\mathbf{C}_T$ is symmetric semi-positive definite, the singular values σ_j are equal to the eigenvalues λ_j: here $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_r \geq 0$ and r is the rank of $\mathbf{C}_T$. This has an exact solution:

$$\lambda = \lambda_1 - \frac{\sum_{j>1}(\lambda_1 - \lambda_j)^2}{\sum_{j>1}\lambda_1 - \lambda_j} \tag{69}$$

Because we have subtracted a multiple of the identity matrix, none of the singular vectors will have changed. We then have

$$\mathbf{C}_T \equiv \lambda \mathbf{I} + (\mathbf{C}_T - \lambda \mathbf{I}) \tag{70}$$

$$= \lambda \mathbf{I} + \sum_{i=1}^{r}(\sigma_i - \lambda)\mathbf{u}_i \mathbf{u}_i^T \tag{71}$$

By truncating this sum, we approximate C with a shifted low-rank matrix:

$$\mathbf{C}_T \approx \mathbf{C}_T^{\text{diag+R1}} \equiv \lambda \mathbf{I} + (\sigma_1 - \lambda)\mathbf{u}_1 \mathbf{u}_1^T \tag{72}$$

This procedure is similar to *factor analysis*, in which one seeks to explain a data vector as the sum of a random vector ($\mathbf{u}$) and the linear combination of some number of latent factors ($\mathbf{z}$) [48]:

$$\mathbf{x} = \Lambda \mathbf{z} + \mathbf{u};$$

the entries of $\mathbf{x}$ would then have the correlation matrix $\Psi + \Lambda\Lambda^T$, where Ψ is a diagonal matrix containing the variances of $\mathbf{u}$.

Supporting information

S1 Text. Includes supplementary analysis of statistics and numerical methods, including discussion of supplementary figures.
(PDF)

S1 Fig. Theory predicts population statistics in the asynchronous regime.
(TIF)

S2 Fig. Theory predicts population statistics in the strong asynchronous regime.
(TIF)

S3 Fig. Theory predicts cell-by-cell statistics in the asynchronous regime.
(TIF)

S4 Fig. Theory predicts cell-by-cell statistics in the strong asynchronous regime.
(TIF)

S5 Fig. Theory captures low-rank structure in correlation matrices.
(TIF)

S6 Fig. Effective parameters in the heterogeneous network: Asynchronous regime.
(TIF)

S7 Fig. Effective parameters in the heterogeneous network: Strong asynchronous regime.
(TIF)

S8 Fig. Firing rate as a function of inhibitory conductance and threshold.
(TIF)

S9 Fig. Correlation increases with firing rate in the strong asynchronous regime: Erdős-Rényi networks.
(TIF)

S1 Table. Statistics from heterogeneous vs. homogeneous networks: Asynchronous regime.
(PDF)

S2 Table. Statistics in recurrent networks: Monte Carlo vs. linear response theory, asynchronous regime.
(PDF)

S3 Table. Statistics from heterogeneous vs. homogeneous networks: Strong asynchronous regime.
(PDF)

S4 Table. Statistics in recurrent networks: Monte Carlo vs. linear response theory, strong asynchronous regime.
(PDF)

Acknowledgments

Author Contributions

Conceptualization: AKB CL.

Formal analysis: AKB CL.

Methodology: AKB CL.

Software: AKB CL.

Writing – original draft: AKB CL.

Writing – review & editing: AKB CL.

References

1. van Vreeswijk C, Sompolinsky H. Chaos in neuronal networks with balanced excitatory and inhibitory activity. Science. 1996; 274:1724–1726. https://doi.org/10.1126/science.274.5293.1724 PMID: 8939866

2. Renart A, de la Rocha J, Bartho P, Hollender L, Parga N, Reyes A, et al. The asynchronous state in cortical circuits. Science. 2010; 327:587–590. https://doi.org/10.1126/science.1179850 PMID: 20110507

3. Doiron B, Litwin-Kumar A, Rosenbaum R, Ocker GK, Josić K. The mechanics of state-dependent neural correlations. Nature neuroscience. 2016; 19(3):383–393. https://doi.org/10.1038/nn.4242 PMID: 26906505

4. Kay SM. Fundamentals of Statistical Signal Processing, Volume 1: Estimation Theory. Prentice Hall PTR; 1993.

5. Dayan P, Abbott LF. Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems. Taylor & Francis; 2001.

6. Shamir M, Sompolinsky H. Implications of neuronal diversity on population coding. Neural Computation. 2006; 18:1951–1986. https://doi.org/10.1162/neco.2006.18.8.1951 PMID: 16771659

7. Josić K, Shea-Brown E, Doiron B, de la Rocha J. Stimulus-dependent correlations and population codes. Neural Computation. 2009; 21:2774–2804. https://doi.org/10.1162/neco.2009.10-08-879 PMID: 19635014

8. Hu Y, Zylberberg J, Shea-Brown E. The Sign Rule and Beyond: Boundary Effects, Flexibility, and Noise Correlations in Neural Population Codes. PLoS Computational Biology. 2014; 10(2):e1003469. https://doi.org/10.1371/journal.pcbi.1003469 PMID: 24586128

9. Azouz R, Gray CM. Cellular mechanisms contributing to response variability of cortical neurons *in vivo*. The Journal of Neuroscience. 1999; 19:2209–2223. PMID: 10066274

10. Padmanabhan K, Urban NN. Intrinsic biophysical diversity decorrelates neuronal firing while increasing information content. Nature Neuroscience. 2010; 13:1276–1282. https://doi.org/10.1038/nn.2630 PMID: 20802489

11. Marsat G, Maler L. Neural heterogeneity and efficient population codes for communication signals. Journal of Neurophysiology. 2010; 104:2543–2555. https://doi.org/10.1152/jn.00256.2010 PMID: 20631220

12. Marder E, Goaillard JM. Variability, compensation and homeostasis in neuron and network function. Nature Reviews Neuroscience. 2006; 7:563–574. https://doi.org/10.1038/nrn1949 PMID: 16791145

13. Marder E. Variability, compensation, and modulation in neurons and circuits. Proceedings of the National Academy of Sciences. 2011; 108:15542–15548. https://doi.org/10.1073/pnas.1010674108 PMID: 21383190

14. Harrison PM, Badel L, Wall MJ, Richardson MJE. Experimentally Verified Parameter Sets for Modelling Heterogeneous Neocortical Pyramidal-Cell Populations. PLOS Computational Biology. 2015; 11. https://doi.org/10.1371/journal.pcbi.1004165 PMID: 26291316

15. Averbeck BB, Latham PE, Pouget A. Neural correlations, population coding and computation. Nature Reviews Neuroscience. 2006; 7:358–366. https://doi.org/10.1038/nrn1888 PMID: 16760916

16. Ecker AS, Berens P, Tolias AS, Bethge M. The effect of noise correlations in populations of diversely tuned neurons. The Journal of Neuroscience. 2011; 31(40):14272–14283. https://doi.org/10.1523/JNEUROSCI.2539-11.2011 PMID: 21976512

17. Moreno-Bote R, Beck J, Kanitscheider I, Pitkow X, Latham P, Pouget A. Information-limiting correlations. Nature Neuroscience. 2014; 17:1410–1417. https://doi.org/10.1038/nn.3807 PMID: 25195105

18. Ruff DA, Cohen MR. Attention can either increase or decrease spike count correlations in visual cortex. Nature Neuroscience. 2014; 17(11):1591–1597. https://doi.org/10.1038/nn.3835 PMID: 25306550

19. Kohn A, Coen-Cagli R, Kanitscheider I, Pouget A. Correlations and Neuronal Population Information. Annual review of neuroscience. 2016; 39(0). https://doi.org/10.1146/annurev-neuro-070815-013851 PMID: 27145916

20. da Silveira RA, Berry MJ. High-Fidelity Coding with Correlated Neurons. PLoS Computational Biology. 2014; 10(11). https://doi.org/10.1371/journal.pcbi.1003970 PMID: 25412463

21. Zylberberg J, Cafaro J, Turner MH, Shea-Brown E, Rieke F. Direction-selective circuits shape noise to ensure a precise population code. Neuron. 2016; 89(2):369–383. https://doi.org/10.1016/j.neuron.2015.11.019 PMID: 26796691

22. de la Rocha J, Doiron B, Shea-Brown E, Josić K, Reyes A. Correlation between neural spike trains increases with firing rate. Nature. 2007; 448:802–806. https://doi.org/10.1038/nature06028 PMID: 17700699

23. Bair W, Zohary E, Newsome WT. Correlated firing in macaque visual area MT: time scales and relationship to behavior. The Journal of Neuroscience. 2001; 21(5):1676–1697. PMID: 11222658

24. Cohen MR, Maunsell JH. Attention improves performance primarily by reducing interneuronal correlations. Nature Neuroscience. 2009; 12:1594–1600. https://doi.org/10.1038/nn.2439 PMID: 19915566

25. Lin IC, Okun M, Carandini M, Harris KD. The nature of shared cortical variability. Neuron. 2015; 87:644–656. https://doi.org/10.1016/j.neuron.2015.06.035 PMID: 26212710

26. Schulz DPA, Sahani M, Carandini M. Five key factors determining pairwise correlations in visual cortex. Journal of Neurophysiology. 2015; 114:1022–1033. https://doi.org/10.1152/jn.00094.2015 PMID: 26019310

27. Franke F, Fiscella M, Sevelev M, Roska B, Hierlemann A, da Silveira RA. Structures of Neural Correlation and How They Favor Coding. Neuron. 2016; 89:409–422. https://doi.org/10.1016/j.neuron.2015.12.037 PMID: 26796692

28. Pernice V, Staude B, Cardanobile S, Rotter S. How structure determines correlations in neuronal networks. PLoS Computational Biology. 2011; 7(5):e1002059. https://doi.org/10.1371/journal.pcbi.1002059 PMID: 21625580

29. Trousdale J, Hu Y, Shea-Brown E, Josic K. Impact of network structure and cellular response on spike time correlations. PLoS Computational Biology. 2012; 8(3):e1002408–e1002408. https://doi.org/10.1371/journal.pcbi.1002408 PMID: 22457608

30. Shea-Brown E, Josic K, de la Rocha J, Doiron B. Correlation and synchrony transfer in integrate-and-fire neurons: basic properties and consequences for coding. Physical Review Letters. 2008; 100:108102. https://doi.org/10.1103/PhysRevLett.100.108102 PMID: 18352234

31. Ostojic S, Brunel N, Hakim V. How connectivity, background activity, and synaptic properties shape the cross-correlation between spike trains. The Journal of Neuroscience. 2009; 29:10234–10253. https://doi.org/10.1523/JNEUROSCI.1275-09.2009 PMID: 19692598

32. Zhao L, Beverlin B II, Netoff T, Nykamp DQ. Synchronization from second order network connectivity statistics. Frontiers in Computational Neuroscience. 2011; 5:1–16. https://doi.org/10.3389/fncom.2011.00028 PMID: 21779239

33. Hu Y, Trousdale J, Josic K, Shea-Brown E. Motif statistics and spike correlations in neuronal networks. Journal of Statistical Mechanics: Theory and Experiment. 2013; https://doi.org/10.1088/1742-5468/2013/03/P03012

34. Yatsenko D, Josić K, Ecker AS, Froudarakis E, Cotton RJ, Tolias AS. Improved estimation and interpretation of correlations in neural circuits. PLoS Comput Biol. 2015; 11(3):e1004083. https://doi.org/10.1371/journal.pcbi.1004083 PMID: 25826696

35. Candes EJ, Plan Y. Matrix completion with noise. Proceedings of the IEEE. 2010; 98(6):925–936. https://doi.org/10.1109/JPROC.2009.2035722

36. Saunderson J, Chandrasekaran V, Parrilo PA, Willsky AS. Diagonal and Low-Rank Matrix Decompositions, Correlation Matrices, and Ellipsoid Fitting. SIAM J Matrix Anal Appl. 2012; 33(4):1395–1416. https://doi.org/10.1137/120872516

37. Bishop WE, Yu BM. Deterministic Symmetric Positive Semidefinite Matrix Completion. In: Advances in Neural Information Processing Systems; 2014. p. 2762–2770.

38. Goris RL, Movshon JA, Simoncelli EP. Partitioning neuronal variability. Nature neuroscience. 2014; 17(6):858–865. https://doi.org/10.1038/nn.3711 PMID: 24777419

39. Ecker AS, Berens P, Cotton RJ, Subramaniyan M, Denfield GH, Cadwell CR, et al. State dependence of noise correlations in macaque primary visual cortex. Neuron. 2014; 82:235–248. https://doi.org/10.1016/j.neuron.2014.02.006 PMID: 24698278

40. Kanashiro T, Ocker GK, Cohen MR, Doiron B. Attentional modulation of neuronal variability in circuit models of cortex; 2017.

41. Cunningham J, Yu B. Dimensionality reduction for large-scale neural recordings. Nature Neuroscience. 2014; 17(11):1500–1509. https://doi.org/10.1038/nn.3776 PMID: 25151264

42. Brunel N. Dynamics of Sparsely Connected Networks of Excitatory and Inhibitory Spiking Neurons. Journal of Computational Neuroscience. 2000; 8:183–208. https://doi.org/10.1023/A:1008925309027 PMID: 10809012

43. Ostojic S. Two types of asynchronous activity in networks of excitatory and inhibitory spiking neurons. Nature neuroscience. 2014; 17(4):594–600. https://doi.org/10.1038/nn.3658 PMID: 24561997

44. Lindner B, Schimansky-Geier L. Transmission of noise coded versus additive signals through a neuronal ensemble. Physical Review Letters. 2001; 86(14):2934. https://doi.org/10.1103/PhysRevLett.86.2934 PMID: 11290076

45. Shelley M, McLaughlin D, Shapley R, Wielaard J. States of high conductance in a large-scale model of the visual cortex. Journal of computational neuroscience. 2002; 13(2):93–109. https://doi.org/10.1023/A:1020158106603 PMID: 12215724

46. Apfaltrer F, Ly C, Tranchina D. Population density methods for stochastic neurons with realistic synaptic kinetics: Firing rate dynamics and fast computational methods. Network: Computation in Neural Systems. 2006; 17:373–418. https://doi.org/10.1080/09548980601069787 PMID: 17162461

47. Richardson MJ. Firing-rate response of linear and nonlinear integrate-and-fire neurons to modulated current-based and conductance-based synaptic drive. Physical Review E. 2007; 76:021919. https://doi.org/10.1103/PhysRevE.76.021919 PMID: 17930077

48. Ghahramani Z, Hinton GE. The EM Algorithm for Mixtures of Factor Analyzers. CRG; 1996. CRG-TR-96-1.

49. Song S, Sjöström PJ, Reigl M, Nelson S, Chklovskii DB. Highly nonrandom features of synaptic connectivity in local cortical circuits. PLoS Biology. 2005; 3(3):e68. https://doi.org/10.1371/journal.pbio.0030068 PMID: 15737062

50. Abbott LF, Dayan P. The effect of correlated variability on the accuracy of a population code. Neural Computation. 1999; 11:91–101. https://doi.org/10.1162/089976699300016827 PMID: 9950724

51. Cohen MR, Kohn A. Measuring and interpreting neuronal correlations. Nature Neuroscience. 2011; 14:811–819. https://doi.org/10.1038/nn.2842 PMID: 21709677

52. Rosenbaum R, Smith MA, Kohn A, Rubin JE, Doiron B. The spatial structure of correlated neuronal variability. Nature Neuroscience. 2016;. https://doi.org/10.1038/nn.4433 PMID: 27798630

53. Oswald AM, Doiron B, Rinzel J, Reyes AD. Spatial profile and differential recruitment of GABAB modulate oscillatory activity in auditory cortex. The Journal of Neuroscience. 2009; 29:10321–10334. https://doi.org/10.1523/JNEUROSCI.1703-09.2009 PMID: 19692606

54. Fino E, Yuste R. Dense inhibitory connectivity in neocortex. Neuron. 2011; 69:1188–1203. https://doi.org/10.1016/j.neuron.2011.02.025 PMID: 21435562

55. Levy RB, Reyes AD. Spatial profile of excitatory and inhibitory synaptic connectivity in mouse primary auditory cortex. The Journal of Neuroscience. 2012; 32(16):5609–5619. https://doi.org/10.1523/JNEUROSCI.5158-11.2012 PMID: 22514322

56. Risken H. The Fokker-Planck equation: methods of solutions and applications. New York, NY: Springer-Verlag; 1989. https://doi.org/10.1007/978-3-642-61544-3

57. Lindner B, Doiron B, Longtin A. Theory of oscillatory firing induced by spatially correlated noise and delayed inhibitory feedback. Physical Review E. 2005; 72(6):061919. https://doi.org/10.1103/PhysRevE.72.061919

58. Richardson MJ. Spike-train spectra and network response functions for non-linear integrate-and-fire neurons. Biological Cybernetics. 2008; 99:381–392. https://doi.org/10.1007/s00422-008-0244-y PMID: 19011926

59. Gerstner W, Kistler WM, Naud R, Paninski L. Neuronal Dynamics: From single neurons to networks and models of cognition. Cambridge University Press; 2014. https://doi.org/10.1017/CBO9781107447615

60. Ricciardi LM. Diffusion processes and related topics in biology. vol. 14. Springer Science & Business Media; 1977.

61. Lindner B. Coherence and Stochastic Resonance in Nonlinear Dynamical Systems. Humboldt University of Berlin; 2002.

62. Rangan AV. Diagrammatic expansion of pulse-coupled network dynamics. Phys Rev Lett. 2009; 102:158101. https://doi.org/10.1103/PhysRevLett.102.158101 PMID: 19518674