

Correlated neural variability in persistent state networks

Amber Polk^{a,b,c}, Ashok Litwin-Kumar^{a,c}, and Brent Doiron^{b,c,1}

^aProgram for Neural Computation, Carnegie Mellon University and University of Pittsburgh, Pittsburgh, PA 15213; ^bDepartment of Mathematics, University of Pittsburgh, Pittsburgh, PA 15260; and ^cCenter for the Neural Basis of Cognition, Pittsburgh, PA 15213

Edited by Ranulfo Romo, Universidad Nacional Autonoma de Mexico, Mexico City D.F., Mexico, and approved February 28, 2012 (received for review December 26, 2011)

Neural activity that persists long after stimulus presentation is a biological correlate of short-term memory. Variability in spiking activity causes persistent states to drift over time, ultimately degrading memory. Models of short-term memory often assume that the input fluctuations to neural populations are independent across cells, a feature that attenuates population-level variability and stabilizes persistent activity. However, this assumption is at odds with experimental recordings from pairs of cortical neurons showing that both the input currents and output spike trains are correlated. It remains unclear how correlated variability affects the stability of persistent activity and the performance of cognitive tasks that it supports. We consider the stochastic long-timescale attractor dynamics of pairs of mutually inhibitory populations of spiking neurons. In these networks, persistent activity was less variable when correlated variability was globally distributed across both populations compared with the case when correlations were locally distributed only within each population. Using a reduced firing rate model with a continuum of persistent states, we show that, when input fluctuations are correlated across both populations, they drive firing rate fluctuations orthogonal to the persistent state attractor, thereby causing minimal stochastic drift. Using these insights, we establish that distributing correlated fluctuations globally as opposed to locally improves network's performance on a two-interval, delayed response discrimination task. Our work shows that the correlation structure of input fluctuations to a network is an important factor when determining long-timescale, persistent population spiking activity.

decision-making | neural correlations | persistent neural activity | spiking model

The dynamics of the nervous system are inherently stochastic at all levels from ion channel (1) to single cell (2), neural population (3, 4), and whole-brain activity (5). A central challenge in neuroscience is to relate the neural variability expressed at the cellular and circuit levels to the variability of cognitive behavior. Short-term memory (6–8) and decision-making (9, 10) are tasks where this link can be established, because there is trial to trial variability of both subject response and associated neural correlates. Furthermore, neural representations based on probabilistic codes (11) as opposed to trial averaged responses have been successful at capturing key aspects of these behaviors (12), and spiking variability is inherent in these cognitive theories. However, relating the statistical structure of the inputs that cause spiking variability to the variability of population-wide neural activity is often difficult (13–16).

Spiking activity with timescales much longer than the timescales of typical synapses and cellular membrane potentials is an established neural correlate of working memory and decision-making. Most famously, persistent activity lasting for many seconds has been reported in prefrontal cortices of primates (6, 17, 18) and rodents (19) during delayed response tasks and premotor areas driving oculomotor gaze fixation (20, 21). In these cases, an initial transient stimulus recruits persistent firing activity, providing a neural basis for short-term working memory. These

observations have led to models, both of large-scale networks of spiking neurons (22–26) and reduced attractor-based frameworks (25–34), aimed at reproducing a continuum of persistent states. A prevailing assumption in all spiking network models of persistent activity (to our knowledge) is that neurons receive statistically independent input fluctuations from sources external to the network. This assumption results in population activity with limited dynamic variability because of population averaging (16, 35), a distinct advantage when building networks with the primary task of maintaining fixed firing rates over long timescales. However, the dense architecture of cortex often provides pairs of neurons common inputs from overlapping presynaptic targets (36–38). External sensory input or intrinsic cortical dynamics may provide correlated input fluctuations to pairs of cortical cells through this presynaptic overlap (39–42), leading to trial to trial covariability of output spike trains (3). Although reported correlations are sometimes small (3, 15), even a minuscule pairwise correlation can lead to very large population variability when responses are pooled over large groups of neurons (16, 35). Hence, neuronal circuits that promote robust persistent activity despite such correlated fluctuations are critical for proper cortical function.

In a two-interval, delayed response discrimination task, the work by Romo et al. (18) combined short-term memory and decision tasks and identified populations of neurons in primate prefrontal cortex that participate in both these computations. Machens et al. (26) presented a model with mutual inhibition between two distinct populations of neurons that captured both the persistent activity and competitive decision dynamics seen in experiment. The model used independent fluctuations to all neurons; however, simultaneous recordings from neuron pairs in the same (43) and related (44) experiments show positive trial-to-trial covariability, even for neurons with opposite stimulus tuning. It is well known that mutually inhibitory networks with deterministic dynamics suppress the influence of input similarities, while amplifying input differences (27, 45). We build on these results and show that the stability of the persistent state in a trial variable stochastic network is enhanced when weakly correlated input fluctuations are distributed globally across neurons in both model populations compared with the case when correlations are only local within each population. A reduced random walk attractor model shows that the increased stability is caused by an opposing interaction between the anticorrelating mutual inhibitory circuitry and the globally correlating fluctuations. Finally, this enhanced stability of the memory state in the presence of global correlations improves task performance by a large degree. Thus, although past studies have identified constructive roles for correlations in static population coding frameworks (43, 46, 47), we

Author contributions: A.P., A.L.-K., and B.D. designed research; A.P. performed research; A.P., A.L.-K., and B.D. analyzed data; and B.D. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

¹To whom correspondence should be addressed. E-mail: bdoiron@pitt.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1121274109/-DCSupplemental.

analyze how correlated variability affects the dynamic stability of persistent neural activity.

Results

Persistent State Spiking Model. Recurrent inhibition between populations of neurons produces a competitive, winner take all dynamic, and it is a common architecture in decision-making (9, 28, 29) and persistent activity circuits (21, 26, 27). We considered the spiking responses of a network of leaky integrate and fire model neurons grouped into two functional populations (labeled A and B) having strong and symmetric interpopulation inhibition (Fig. 1A and *SI Materials and Methods*). The populations were presented with a static excitatory stimulus for an initial period of 500 ms, after which the stimulus was removed; the resulting long-timescale population dynamics were analyzed (Fig. 1B). The population response (ensemble-averaged firing rate) (*Materials and Methods*) was persistent, with rates remaining at near constant levels for seconds after stimulus removal. Furthermore, there was a continuum of such persistent states, each determined by different initial stimulus configurations (three example cases are shown in Fig. 1C–E). In this way, the poststimulus activity was an accurate reflection of the initial stimulus, providing a model of short-term memory (26).

The network retained memory through a balance of competitive forces. The population with a higher stimulus-evoked firing rate suppressed the other population, effectively disinhibiting itself; therefore, it maintained high firing rates after the stimulus

was removed. This balance was dynamically maintained and hence, sensitive to firing rate fluctuations. In our spiking model, the membrane potential fluctuations were substantial, consistent with *in vivo* neural activity (Fig. 1C). However, the fluctuations were also independently distributed to all of the neurons, producing a relatively asynchronous population dynamic (Fig. 1D). In this case, population rate variability scaled roughly inversely to population size. For reasonably sized networks, the population activity had little variability, allowing rates to remain stable for long periods of time (Fig. 1E). Thus, despite the significant single neuron variability, population averaging produced a stable persistent network state. In the next section, we relax the assumption of independent variability and consider the consequences for the stability of persistent network states.

Global Correlated Variability Preserves Persistent Network States.

The connectivity within cortical networks is dense, with groups of neurons receiving a significant amount of common synaptic input (36–38). This overlap of synaptic architecture correlates the membrane potential fluctuations of distinct neurons, and in several cortical areas, the magnitude of correlation can be significant (39–42). In complex cortical circuits, it is difficult to determine the precise structure of the joint fluctuations to different neuron populations. This structure may depend on, for example, physical proximity (48) or functional similarity (49) between populations. We studied two simple schemes in which correlated stochastic inputs were incorporated into our two-population persistent activity model. The first scheme was a local framework, in which only neurons within the same population received correlated inputs (Fig. 2A, *Upper*). The second scheme was a global framework, in which all neurons, both within and between populations, received a common source of input fluctuations (Fig. 2A, *Lower*). The local model assumes that shared input fluctuations respect the network's functional segregation (population A vs. B), whereas the global model considers correlated fluctuations that ignore any network divisions. In both cases, the magnitude of input correlation was small (0.05), meaning that individual neurons still received a significant amount of independent variability. Despite this small degree of pairwise input correlation, its impact was dramatic at the population level (16, 35), because common fluctuations could not be attenuated through population averaging. In our model, integrated population variance increased by a factor of 25 compared with the uncorrelated case.

Although both local and global correlation models received significant population-wide input variability, they nevertheless supported a continuum of persistent states during the initial poststimulus period (Fig. S1). To ease presentation, we hereafter study the symmetric case, where the persistent state has equal rates for both populations (Fig. 1C–E, blue); however, our results are by no means limited to the case. For uncorrelated inputs, the persistent state remained stable for long periods (>10 s) after stimulus removal (example trials are in Fig. 2B, *Top*). In contrast, local correlations yielded a persistent state that was stable for a short period directly after stimulus removal, but the trial-to-trial population firing rates began to drift after a few seconds (example trials are in Fig. 2B, *Middle*). Finally, the population rates for the globally correlated model were stable long after (~10 s) stimulus removal (example trials are in Fig. 2B, *Bottom*). Although the stability for the global model persistent state was similar to the uncorrelated model, the overall variability of firing rates was significantly larger, which is expected from simple pooling arguments (16, 35).

To better quantify the drift in the population rates, we performed a large number of trials (10^4) and estimated the probability density of the joint population rates as it evolved over time. When fluctuations were local, there was a significant drift of the trial-to-trial population activity along the antidiagonal firing rate axis (Fig. 2C, *Upper* and *Movie S1*). This drift was much reduced when fluctuations were global (Fig. 2C, *Lower* and *Movie S1*).

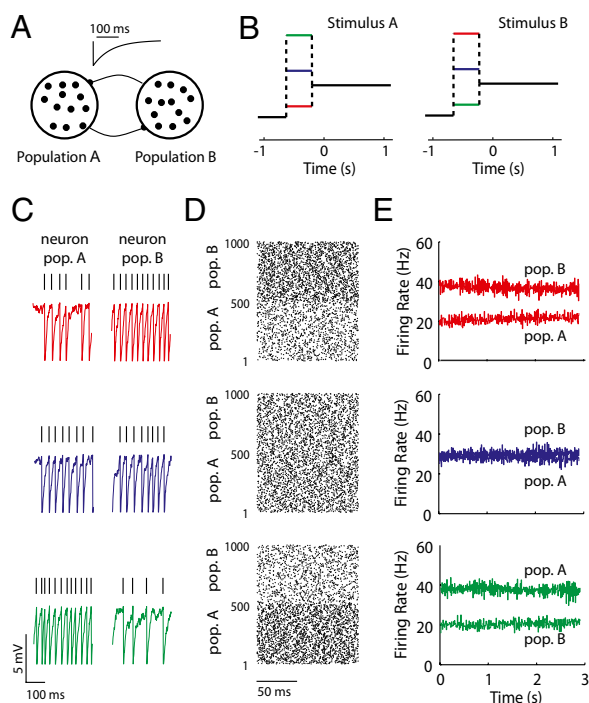


Fig. 1. Recurrent inhibition produces a continuum of persistent states in a spiking network model. (A) Schematic of two populations (A and B) each having $N (= 500)$ leaky integrate and fire neuron models. Neurons between populations inhibited one another; the time course of a single spike-induced inhibitory current is shown (*Inset*). (B) Stimulation protocol showing an initial stimulus period (< 0 s) when populations A and B received distinct stimuli followed by a long period (> 0 s) when the populations received an equal bias. The colors indicate three example stimulation sequences. (C) Representative membrane potential traces from a neuron in either population. (D) Spike time rasters for the entire network. Population membership is indicated on the vertical axis. (E) The evolution of populations A and B firing rates (*Materials and Methods*). Colors in C–E match the stimulation protocols in B.

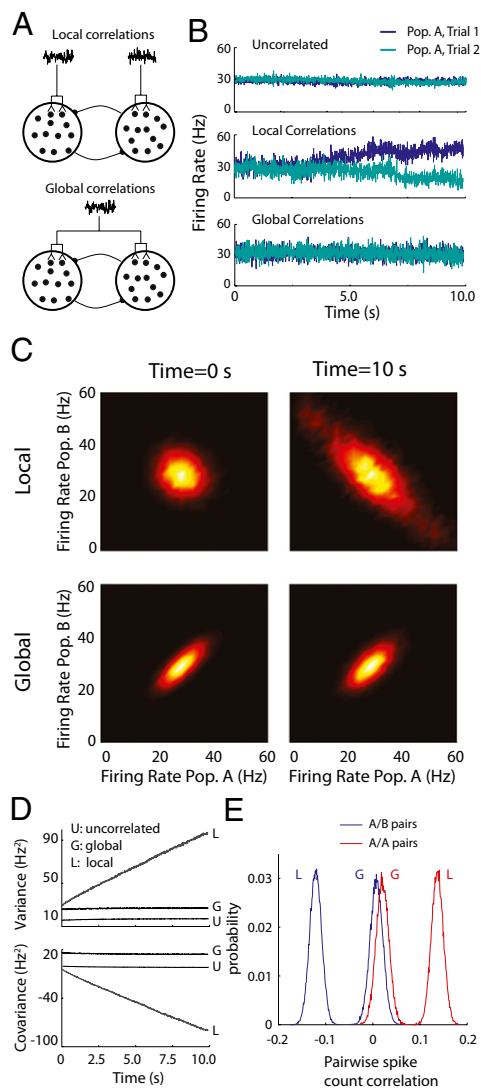


Fig. 2. Influence of local vs. global correlated input variability on the stability of persistent activity. (A) Schematic of the distribution of common fluctuations in the local (*Upper*) and global (*Lower*) correlation frameworks. (B) Example pair of trials showing the temporal evolution of population firing rates with uncorrelated (*Top*), locally correlated (*Middle*), or globally correlated (*Bottom*) inputs. (C) The probability density of population rates $P(r_A, r_B, t)$ for $t = 0$ s (*Left*) and $t = 10$ s (*Right*) estimated from repeated trials of the spiking network model. This estimation was done for both the local (*Upper*) and global (*Lower*) correlation schemes. The (r_A, r_B) grid was sampled at 1 Hz, and $P(r_A, r_B, t)$ was smoothed for presentation. (D) Population rate variance and covariance as a function of time for the uncorrelated (U), local (L), and global (G) correlation schemes. (E) Pairwise spike count correlation ($T = 100$ ms) for all neuron pairs both within and between populations during the local and global correlation schemes.

The trial-to-trial population rate variance and covariance between populations A and B computed over trials mirrored these differences (Fig. 2D). Furthermore, although the globally correlated model showed time stationary variability, its overall variability was nevertheless larger than the uncorrelated model (Fig. 2D), again because of population-wide pooling of activity. Thus, when the correlations of input fluctuations were global to the entire network rather than local to subpopulations within the network, the stability of persistent neural activity was increased.

The correlations between the spike trains of individual neuron pairs were substantially different for networks with local or global

input correlation. Local input correlations produced a significant positive spike count correlation for neuron pairs that were within the same population but an equally strong negative correlation for neurons that were members of distinct populations (Fig. 2E). In contrast, global input correlations produced weak pairwise correlations (Fig. 2E). This attenuation was caused by cancellation mechanisms similar to those operative in balanced cortical circuits (15). Furthermore, for the model with global input correlations, neuron pairs that were in different populations had a mean correlation that was positive (0.01), albeit smaller than the correlation for neurons within the same population (0.02) (compare blue G with red G curves in Fig. 2E). This finding was in contrast to the large negative correlation (-0.13) observed for neurons in different populations when input correlations were local (compare blue L with blue G curves in Fig. 2E). Paired recordings in prefrontal cortex of primates (43) participating in a delayed response task showed a net positive correlation for all neuron pairings (within-population pairs had a mean correlation coefficient of 0.15, and across-population pairs had a mean correlation coefficient of 0.05) (table 1 in ref. 43). Although we do not expect a quantitative match between our model and the data, the qualitative match that all correlations were positive gives strong experimental justification for the global input model.

In total, although globally correlated fluctuations represent a larger source of overall input covariability, they do not destroy long-term persistent activity and result in weaker pairwise correlations than a model with only local correlations. We next study a reduced framework to uncover the core mechanism for these findings.

Random Walk Firing Rate Model. Phenomenological models often consider the aggregated firing rate of a population of cells rather than individual spiking activity (50). In particular, firing rate models of mutually inhibitory populations have captured key aspects of decision-making (28, 29) and short-term memory (26) circuits. We considered a firing rate model of our persistent state network (Fig. 3A) where the population firing rates $r_A(t)$ and $r_B(t)$ obeyed the stochastic differential equations (Eq. 1)

$$\tau \frac{dr_A}{dt} = \mu - (r_A + r_B) + \sigma \left(\sqrt{1-c} \xi_A(t) + \sqrt{c} \xi_c(t) \right) \quad [1]$$

and (Eq. 2)

$$\tau \frac{dr_B}{dt} = \mu - (r_B + r_A) + \sigma \left(\sqrt{1-c} \xi_B(t) + \sqrt{c} \xi_c(t) \right), \quad [2]$$

where μ is a symmetric mean input and τ is the population timescale. The coupling between the populations caused degeneracy in the dynamics, with a continuum of neutrally stable states (i.e., a line attractor) along the anti-diagonal $r_B = \mu - r_A$ (Fig. 3B, green line). The population activity and coupling was linear to facilitate analysis; for a general class of nonlinear models, our model is an appropriate limiting case (*SI Materials and Methods* and Fig. S2).

Population-wide input fluctuations were modeled by the combination of three Gaussian white noise processes: $\xi_A(t)$ and $\xi_B(t)$ were private fluctuations to populations A and B, respectively, and $\xi_c(t)$ was a global fluctuation to the two populations. The total fluctuations to a population had intensity σ , with the input correlations between populations set by c (Fig. 3A). In our reduced model, $c = 0$ corresponds to the local correlation spiking network, whereas $c > 0$ corresponds to a globally correlated network. The model is a 2D Ornstein–Uhlenbeck process (51), and the associated Fokker–Planck equation gave an analytic solution for the time-dependent probability density $P(r_A, r_B, t)$ (*SI Materials and Methods* shows a full derivation).

The density $P(r_A, r_B, t)$ is Gaussian, with a mean (μ, μ) and a time-dependent variance and covariance. Importantly, the

in firing rate coordinates (26). Globally correlated fluctuations were orthogonal to the direction of instability and thus, best mitigated drift of the persistent state caused by input fluctuations.

Global Correlations Improve Discrimination in Delayed Response Task. The motivating experiments for the mutual inhibition model of persistent activity were from primate prefrontal cortex recordings during a delayed response, two-interval discrimination task (18, 26, 43). Briefly, a vibrating input with a fixed frequency was presented to the fingertips, and neurons sensitive to this stimulus could be grouped into two classes (our populations A and B) based on their response: neurons either increased or decreased their firing rate to increasing input frequency (Fig. 4A, *Left*). This initial stimulus loaded the network in a specific location on the line attractor (Fig. 4B, *Left*). The stimulus was removed (Fig. 4A, *Center*), and the loaded state was maintained in working memory during a delay period (Fig. 4B, *Center*). Finally, a second stimulus was presented, and the monkey was directed to indicate if the first stimulus vibrated at a frequency that was faster than the second stimulus (Fig. 4A, *Right*). This decision task had the neural correlate that one population's firing rate reached a high value, whereas the other decreased its firing rate, with the specific choice depending on which stimulus had the higher frequency. The initial model by Machens et al. (26) replicated this memory and decision task. However, when input fluctuations were correlated in our model, the sizable trial to trial variability sometimes caused a misclassification error by the spiking network (Fig. 4C). We investigated the effects of local and global correlations on the performance of network discrimination in the delayed, two-interval discrimination task.

Without loss of generality, we assumed that the first stimulus (s_1) loaded the network to the symmetric $(r_A, r_B) = (\mu, \mu)$ state. We computed the density $P(r_A, r_B, t_2)$ from the firing rate model, where $t_2 = 3s$ was the presentation time of the second stimulus. The value of the second stimulus ($s_2 > s_1$) defined a line in (r_A, r_B) space, and a reasonable approximation to the decision process is that realizations with joint rates that were below (above) the line led to correct (incorrect) discriminations (Fig. 4D). For discriminations between first and second stimuli with large differences ($\Delta s = s_2 - s_1$), the distance between the loaded firing rate (μ, μ) and the decision line was large. In this case, the decision task was robust to moderate variability, and both local and global models performed adequately (Fig. 4D, *Left*). However, for finer discriminations, the reduction of spread along the line attractor for global correlations greatly enhanced the probability of correct decision compared with the model with local fluctuations (Fig. 4D, *Right*). To explore this result, we integrated $P(r_A, r_B, t_2)$ below the decision line to find the probability of correct decision over a large range of $(\Delta s, c)$ parameter space (Fig. 4E and *SI Materials and Methods*). When $|\Delta s|$ was small, the probability of correct decision was dramatically improved with increasing c , whereas for large $|\Delta s|$, the decision process was less affected by c . Finally, we verified this prediction with simulations of the spiking network during the decision process for both local and global input correlations (Fig. 4F). For small $\Delta s (= 0.5)$, the presence of global fluctuations increased task performance from near 65% correct decision for local to almost 95% for global input correlations.

In summary, global as opposed to local input correlations greatly enhanced performance for fine discriminations in a delayed response task. This enhancement was because of the increased stability of the persistent state for global correlations, enhancing the fidelity of the final stimulus comparison.

Discussion

Inhibitory interaction between distinct populations of neurons is an often proposed mechanism for persistent neural activity. It is known that mutual inhibition between populations amplifies input differences to populations while suppressing input similarities (27,

45). However, how these principles operate in a persistent activity model that exhibits trial to trial variability has not been addressed. Previous spiking models of persistent activity have assumed that external input fluctuations are independent for all neurons (22–26, 33), and thus, population pooling attenuates the variability at the network level (16). Heuristic firing rate models with stochastic inputs are not limited to small variability, and past work with them has given insight about how fluctuations promote stochastic drift along a continuum attractor (33, 34). However, these studies have also assumed that the fluctuations to distinct populations are uncorrelated, being equivalent to the local correlation model in our study. Our primary finding was that robust stability of the persistent state required common fluctuating inputs to be globally distributed over the network as opposed to only locally within each population. In our spiking model, the recurrent inhibition and global correlations balance to give overall weak positive pairwise spike correlations, consistent with the experimental recordings that motivated the model (43). Finally, this enhanced stability with global fluctuations improved network performance in a two-interval, delayed response discrimination task.

Our network modeled functionally segregated populations that were distinguished by their stimulus preference. Global correlations to such a network require trial-variable inputs that are not related to this segregation and rather indiscriminately correlate the activity of all neurons. Although this assumption is unreasonable for populations with a clear spatial lateralization (21, 52), it is appropriate for persistent activity networks in prefrontal cortex that are involved in working memory (18, 43, 44). In these areas, simultaneous recordings from neuron pairs show fast-timescale cross-correlograms of spiking activity peaked at a zero lag (43, 44), consistent with a source of common input. The magnitude of output spike train correlation is larger for neuron pairs that have similar tuning than those pairs with dissimilar tuning, and therefore, output correlations are related to population segregation. However, this finding does not necessarily imply that external input correlations are local, because it is also consistent with neurons in the same population receiving common inhibition from within the network (this finding is the case for our network) (Fig. 2E). Recently, mean field techniques have been developed that can combine internal and external sources of correlation (15, 53–55). Using these techniques to analyze the dynamic behavior in our network is an interesting potential avenue of study.

It is known that the combination of an anticorrelated stimulus tuning and positive trial to trial correlations improve population-based codes compared with codes in which neurons have independent variable activity (46, 47). Furthermore, when decisions involve a comparison between two population rates, a subtractive computation can benefit from correlated variability (43). In both of these cases, there is an anticorrelating deterministic structure (trial averaged), and thus, positive correlations in stochastic activity (trial variable) best mitigate the impact of fluctuations on computation. Our model extends this general principle of improving network function by matching the structure of network dynamics with the joint structure of stochastic inputs, to the case of persistent state networks. Other models of persistent activity involve either recurrent excitation (22, 23) or feedforward architectures (30). Correlated activity in these networks will no doubt also affect persistent neural activity, but how the structure of input fluctuations will interact with circuit-based dynamics is unclear and open to future investigation.

Materials and Methods

A detailed description of the spiking network model can be found in *SI Materials and Methods*. We let the spike train from neuron i of population X be $y_i^X(t) = \sum_k \delta(t - t_{ik}^X)$, where $\{t_{ik}^X\}$ is a set of spike times ($X = A, B$). The firing rate of population X over time window $(t - T, t)$ is the random variable $r^X(t) = \frac{1}{NT} \int_{t-T}^t \sum_{i=1}^N y_i^X(t') dt'$. Population rate variance was measured as

$\langle r_X(t)^2 \rangle - \langle r_X(t) \rangle^2$, and covariance was measured as $\langle r_X(t)r_Y(t) \rangle - \langle r_X(t) \rangle \langle r_Y(t) \rangle$; here, $\langle \cdot \rangle$ is an expectation over trials. We take $T = 10$ ms, and we use 10^4 trials to compute population statistics. The reduced firing rate model and associated analysis are presented in *SI Materials and Methods*.

ACKNOWLEDGMENTS. We thank Eric Shea-Brown and Kresimir Josic for useful conversations and Christian Machens for providing the spiking model code. Funding was provided by National Science Foundation Grants NSF-DMS-1121784 and EMSW21-RTG0739261. B.D. is a Sloan Research Fellow.

- White JA, Rubinstein JT, Kay AR (2000) Channel noise in neurons. *Trends Neurosci* 23:131–137.
- Britten KH, Shadlen MN, Newsome WT, Movshon JA (1993) Responses of neurons in macaque MT to stochastic motion signals. *Vis Neurosci* 10:1157–1169.
- Cohen MR, Kohn A (2011) Measuring and interpreting neuronal correlations. *Nat Neurosci* 14:811–819.
- Arieli A, Sterkin A, Grinvald A, Aertsen A (1996) Dynamics of ongoing activity: Explanation of the large variability in evoked cortical responses. *Science* 273:1868–1871.
- Fox MD, Raichle ME (2007) Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nat Rev Neurosci* 8:700–711.
- Goldman-Rakic PS (1995) Cellular basis of working memory. *Neuron* 14:477–485.
- Wang XJ (2001) Synaptic reverberation underlying mnemonic persistent activity. *Trends Neurosci* 24:455–463.
- Curtis CE, D'Esposito M (2003) Persistent activity in the prefrontal cortex during working memory. *Trends Cogn Sci* 7:415–423.
- Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD (2006) The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol Rev* 113:700–765.
- Gold JI, Shadlen MN (2007) The neural basis of decision making. *Annu Rev Neurosci* 30:535–574.
- Ma WJ, Beck JM, Latham PE, Pouget A (2006) Bayesian inference with probabilistic population codes. *Nat Neurosci* 9:1432–1438.
- Churchland AK, et al. (2011) Variance as a signature of neural computations during decision making. *Neuron* 69:818–831.
- van Vreeswijk C, Sompolinsky H (1996) Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* 274:1724–1726.
- Amit DJ, Brunel N (1997) Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cereb Cortex* 7:237–252.
- Renart A, et al. (2010) The asynchronous state in cortical circuits. *Science* 327:587–590.
- Rosenbaum R, Trousdale J, Josic K (2011) The effects of pooling on spike train correlations. *Front Neurosci*, 10.3389/fnins.2011.00058.
- Niki H, Watanabe M (1979) Prefrontal and cingulate unit activity during timing behavior in the monkey. *Brain Res* 171:213–224.
- Romo R, Brody CD, Hernández A, Lemus L (1999) Neuronal correlates of parametric working memory in the prefrontal cortex. *Nature* 399:470–473.
- Narayanan NS, Laubach M (2006) Top-down control of motor cortex ensembles by dorsomedial prefrontal cortex. *Neuron* 52:921–931.
- McFarland JL, Fuchs AF (1992) Discharge patterns in nucleus prepositus hypoglossi and adjacent medial vestibular nucleus during horizontal eye movement in behaving macaques. *J Neurophysiol* 68:319–332.
- Aksay E, et al. (2007) Functional dissection of circuitry in a neural integrator. *Nat Neurosci* 10:494–504.
- Seung HS, Lee DD, Reis BY, Tank DW (2000) Stability of the memory of eye position in a recurrent network of conductance-based model neurons. *Neuron* 26:259–271.
- Koulakov AA, Raghavachari S, Kepecs A, Lisman JE (2002) Model for a robust neural integrator. *Nat Neurosci* 5:775–782.
- Deco G, Rolls ET (2003) Attention and working memory: A dynamical model of neuronal activity in the prefrontal cortex. *Eur J Neurosci* 18:2374–2390.
- Wong KF, Wang XJ (2006) A recurrent network mechanism of time integration in perceptual decisions. *J Neurosci* 26:1314–1328.
- Machens CK, Romo R, Brody CD (2005) Flexible control of mutual inhibition: A neural model of two-interval discrimination. *Science* 307:1121–1124.
- Cannon SC, Robinson DA, Shamma S (1983) A proposed neural network for the integrator of the oculomotor system. *Biol Cybern* 49:127–136.
- Usher M, McClelland JL (2001) The time course of perceptual choice: The leaky, competing accumulator model. *Psychol Rev* 108:550–592.
- Brown E, et al. (2005) Simple networks that optimize decisions. *Int J Bifurcations Chaos* 15:803–826.
- Goldman MS (2009) Memory without feedback in a neural network. *Neuron* 61:621–634.
- Rolls ET, Deco G (2010) *The Noisy Brain: Stochastic Dynamics as a Principle of Brain Function* (Oxford University Press, New York).
- Machens CK, Brody CD (2008) Design of continuous attractor networks with monotonic tuning using a symmetry principle. *Neural Comput* 20:452–485.
- Wu S, Hamaguchi K, Amari S (2008) Dynamics and computation of continuous attractors. *Neural Comput* 20:994–1025.
- Lim S, Goldman MS (2012) Noise tolerance of attractor and feedforward memory models. *Neural Comput* 24:332–390.
- Zohary E, Shadlen MN, Newsome WT (1994) Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* 370:140–143.
- Yoshimura Y, Dantzker JLM, Callaway EM (2005) Excitatory cortical neurons form fine-scale functional networks. *Nature* 433:868–873.
- Oswald AMM, Doiron B, Rinzel J, Reyes AD (2009) Spatial profile and differential recruitment of GABAB modulate oscillatory activity in auditory cortex. *J Neurosci* 29:10321–10334.
- Fino E, Yuste R (2011) Dense inhibitory connectivity in neocortex. *Neuron* 69:1188–1203.
- Lamp I, Reichova I, Ferster D (1999) Synchronous membrane potential fluctuations in neurons of the cat visual cortex. *Neuron* 22:361–374.
- Poulet JFA, Petersen CCH (2008) Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. *Nature* 454:881–885.
- Okun M, Lampl I (2008) Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities. *Nat Neurosci* 11:535–537.
- Yu J, Ferster D (2010) Membrane potential synchrony in primary visual cortex during sensory stimulation. *Neuron* 68:1187–1201.
- Romo R, Hernández A, Zainos A, Salinas E (2003) Correlated neuronal discharges that increase coding efficiency during perceptual discrimination. *Neuron* 38:649–657.
- Constantinidis C, Franowicz MN, Goldman-Rakic PS (2001) Coding specificity in cortical microcircuits: A multiple-electrode analysis of primate prefrontal cortex. *J Neurosci* 21:3646–3655.
- Seung HS (2003) *Amplification, Attenuation, and Integration. The Handbook of Brain Theory and Neural Networks*, ed Arbib MA (MIT Press, Cambridge, MA), 2nd Ed, pp 94–97.
- Abbott LF, Dayan P (1999) The effect of correlated variability on the accuracy of a population code. *Neural Comput* 11:91–101.
- Averbeck BB, Latham PE, Pouget A (2006) Neural correlations, population coding and computation. *Nat Rev Neurosci* 7:358–366.
- Oihorhenuan IE, et al. (2010) Sparse coding and high-order correlations in fine-scale cortical networks. *Nature* 466:617–621.
- Kohn A, Smith MA (2005) Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J Neurosci* 25:3661–3673.
- Ermentrout GB, Terman D (2010) *Mathematical Foundations of Neuroscience* (Springer, New York).
- Risken H (1996) *The Fokker–Planck Equation: Methods of Solution and Applications* (Springer, New York), 2nd Ed.
- Hanks TD, Ditterich J, Shadlen MN (2006) Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. *Nat Neurosci* 9:682–689.
- Buice MA, Cowan JD, Chow CC (2010) Systematic fluctuation expansion for neural network activity equations. *Neural Comput* 22:377–426.
- Bressloff PC (2009) Stochastic neural field theory and the system-size expansion. *SIAM J Appl Math* 70:1488–1521.
- Touboul JD, Ermentrout GB (2011) Finite-size and correlation-induced effects in mean-field dynamics. *J Comput Neurosci* 31:453–484.