An Oscillatory Mechanism for Multi-level Storage in Short-term Memory

Kathleen P. Champion[1#*], Olivia Gozel[2,3*], Benjamin S. Lankow[4*], G. Bard Ermentrout[6],

Mark S. Goldman[4,5]

1. Department of Applied Mathematics, University of Washington, Seattle, WA 98195, USA

2. Departments of Neurobiology and Statistics, University of Chicago, Chicago, IL 60637, USA

3. Grossman Center for Quantitative Biology and Human Behavior, University of Chicago, Chicago, IL 60637, USA

4. Center for Neuroscience, University of California, Davis, Davis, CA 95618, USA

5. Department of Neurobiology, Physiology, and Behavior, and Department of Ophthalmology and Visual Science, University of California, Davis, Davis, CA 95618, USA

6. Department of Mathematics, University of Pittsburgh, Pittsburgh, PA, 15213, USA

* These authors contributed equally to this work
# This work was completed outside current employment at Amazon.com, Inc.

Correspondence:

Mark Goldman
Center for Neuroscience
UC Davis
1544 Newton Court Davis, CA 95616
(530) 757-8749

msgoldman@ucdavis.edu

Bard Ermentrout
Department of Mathematics
University of Pittsburgh
Pittsburgh, PA, 15213

bard@pitt.edu

Abstract

Oscillatory activity is commonly observed during the maintenance of information in short-term memory, but its role remains unclear. Non-oscillatory models of short-term memory storage are able to encode stimulus identity through their spatial patterns of activity, but are typically limited to either an all-or-none representation of stimulus amplitude or exhibit a biologically implausible exact-tuning condition. Here, we demonstrate a simple phase-locking mechanism by which oscillatory input enables a circuit to generate persistent or sequential activity patterns that encode information not only in their location but also in a discrete set of amplitudes.

The maintenance of information in short-term memory is a key component of a wide array of cognitive (1, 2) and non-cognitive (3, 4) functions. However, the biophysical mechanisms that enable memory storage over the seconds-long time scale remain unclear. Single-unit studies have demonstrated a neural correlate of memory maintenance in the persistent activation of neurons whose population activity spans the memory period (reviewed in (2, 5, 6)). Theoretical studies have shown how such persistent activity can be generated by recurrent network feedback (7–9), but simple instantiations of this idea are either implausibly sensitive to mistuning or can only maintain a single elevated firing rate that is unrealistically high (the 'low firing rate problem', reviewed in (4)), limiting storage about a given item to a single bit ('on' or 'off') of information.

Separately, previous studies have identified distinct bands of oscillatory activity in field potential recordings and EEG during the maintenance of working memory (reviewed in (10) . Such activity can be generated through cell-intrinsic mechanisms, local circuitry, or long-range interactions (11–13). However, it remains an open question whether oscillatory activity is necessary, sufficient, or even beneficial for working memory storage. Previous work has proposed how oscillations can contribute to a variety of memory functions such as the generation or maintenance of persistent activity (14, 15); the structuring of spatial codes through frequency coupling (16); and the coordination, control, and gating of memory-related activity (17–26). By contrast, other studies have suggested that oscillations could be an epiphenomenon of other computational or network mechanisms (27–29). Here, we demonstrate a potential mechanistic role for oscillations, regardless of source or frequency, by showing how the addition of oscillatory inputs to simple recurrent feedback circuits can enable both low firing rate persistent activity and a discrete set of elevated firing rates that increases the information capacity of a memory network.

To illustrate the core challenges that arise when generating biologically plausible models of persistent activity, consider an idealized circuit consisting of a memory neuron (or lumped population) connected to itself through positive feedback (Fig. 1A); this basic motif of recurrent excitation is the key component of most circuit models of persistent neural activity (reviewed in (4)). This simple circuit receives a brief stimulus (Figs. 1A, C, E external input) and needs to store

it through persistent activity. Stable persistent activity ($\frac{dr}{dt} = 0, r > 0$) is achieved only when the intrinsic decay of the neuron (represented by the term$-r$) and the recurrent drive to the neuron ($f(r)$) are equal in magnitude and cancel each other. This condition imposes two separate, but related, problems that depend on whether the rate function $f(r)$is linear or nonlinear. In the typical nonlinear case (Fig. 1B,C), if the stimulus is too weak, the memory neuron's low initial firing rate provides insufficient recurrent feedback to overcome the post-stimulus intrinsic decay of activity (Fig. 1B, left of open circle). As a result, the firing rate of the network returns to a low (or zero) baseline firing rate (Fig. 1C, orange, purple, green traces). By contrast, if the stimulus is stronger, the memory neuron's initial firing provides recurrent feedback that exceeds the rate of intrinsic decay (Fig. 1B, right of open circle), leading to a reverberatory amplification of activity in which the rate rises until some saturation process brings the rate to rest at an elevated persistent level (Fig. 1C, blue and red traces). Thus, the only possibilities are that activity decays to its baseline level or that activity runs away to saturation at a high level of activity that, for typical neuronal nonlinearities, is unrealistically large. A different problem emerges in the case where the rate function $f(r)$ is linear (Figs. 1D,E). The linearity of the rate function in this case allows a continuum of persistent rates, corresponding to the continuous set of points at which the feedback and decay lines overlap, to be stored (Figure 1D, blue line), unlike the nonlinear case . This comes at the cost of a 'fine-tuning' condition: the strength of the recurrent synapse(s) must be exactly tuned to counterbalance the strength of the rate decay; an arbitrarily small violation of this condition causes the rate to explode (Figure 1D,E red trace) or decay to a low baseline (Figure 1D, E orange trace). Although presented here for a very simple example, these problems are also commonly observed in larger neural networks (30).

We next illustrate what happens when a network with the same positive feedback architecture is provided with a subthreshold oscillatory drive (Fig. 2). We demonstrate this in the more biologically realistic case of a conductance-based spiking neuron model (31) that facilitates the phase-locking phenomenon that we will describe. Without an oscillatory input, the model exhibits the 'low firing-rate problem' (Figure 2A,B) and can only maintain persistent activity at a high spiking rate or not spike at all. When a subthreshold oscillation is added to the model (Figure 2 C-F), the oscillatory drive has two effects. First, it provides extra input that allows small initial

inputs to trigger low-rate spiking. Second, spiking of the memory neuron does not lead to runaway feedback because, before the feedback can run away, the oscillatory drive returns towards its trough, causing a cessation of spiking. The net result is that the spike-driven feedback becomes discretized, forming a staircase whose step heights correspond to the number of spikes emitted by the neuron per oscillation cycle (Figure 2D-F). The phase locking of the spiking to the subthreshold oscillatory drive constrains these spike numbers to be integer multiples of the oscillation.

The key requirements for this mechanism to enable multi-level persistent activity are the following: First, the oscillation must be strong enough to reset the activity at its troughs. Second, there must be some process that enables the activity from one cycle of the oscillation to carry through to the start of the next cycle and consequently enable renewed spiking as the oscillatory input heads towards its peak. For the simple case illustrated here, where all neurons receive oscillatory inputs that are perfectly aligned in phase, the mechanism enabling inter-cycle memory is a slow NMDA-like (or local dendritic) synaptic time constant (3, 32, 33). Alternatively, we show in the supplemental material that, if there is heterogeneity in the phases of the oscillations received by the individual neurons in the network, the time between cycles may be bridged by the firing of other neurons in the network (Figure S2). Some degree of tuning of the feedback is required to have multiple levels of response -- such a tuning requirement is generic of models of analog or finely discretized persistent activity. In the present case, the width of the steps of the staircase provides a moderate level of robustness to mistuning, especially for higher oscillation frequencies (Figure 2D-F). Mechanistically, this robustness occurs because errors in the tuning of feedback that are insufficient to systematically add or subtract an extra spike per cycle do not persist from cycle to cycle, unlike in models that have no oscillatory trough to reset (error-correct) the spiking activity. We illustrate this robustness to weight changes in Figure 3, where we compare the oscillatory autapse memory model (Fig. 3A,B) to an approximately linear autapse model (34, 35) that can produce (nearly) graded persistent activity (Fig. 3C,D). Each model receives an arbitrary sequence of positive and negative input pulses, and must temporally accumulate and store the pulses in persistent activity. The linear spiking autapse model requires fine tuning to maintain persistent activity: very small deviations from the tuned autapse weight

lead to activity that grows to a saturating level or decays to zero activity (Figure 3C,D). In contrast, the same synaptic weight deviations have negligible effect on the accumulation and multi-level storage capability of the nonlinear spiking neuron with oscillatory drive (Figure 3A,B).

The above examples demonstrate the basic mechanism by which oscillatory input may permit multiple levels of firing rate to be robustly stored in a recurrent excitatory network model of persistent activity. We next explored applications of this basic principle in the case of three different network architectures: a spatially uniform (all-to-all) network that temporally integrates its inputs (Fig. 4); a 'ring-like' architecture whose activity can store both a spatial location and discrete levels at that location (Fig. 5); and a chain-like architecture that can generate sequences of activity with multiple discrete-valued amplitudes (Fig. 6).

We first extended the demonstration of temporal integration, shown in Figure 3, to a spatially homogeneous (all-to-all) network composed of 1000 neurons (Fig. 4A,B). This permitted us to not only examine the systematic mistuning of weights shown in the autapse network, which produces identical results in the averaged activity of the 1000 neuron network (Methods), but also to examine the robustness to four different sources of noise and variability: input noise, in which each neuron in the network received independent exponentially-filtered noise added to the subthreshold oscillatory drive (Fig 4C); noise in the connection weights, in which each synapse in the network was initialized with added random noise (Fig 4D); randomly shuffled phases of the subthreshold oscillatory drive, in which each neuron received an oscillatory signal whose phase was randomly picked from the uniform distribution on $[0, 2\pi)$ at initialization (Fig S1A); and noisy oscillation frequency and amplitude, in which the parameters of the subthreshold oscillatory drive underwent noisy drift (Ornstein-Uhlenbeck process) during the simulations (Fig S1B). In all of these cases, the network was able to accurately maintain multi-level persistent activity despite moderate perturbations. Figures 4C-D and Figure S1B illustrate the conditions where the magnitude of the perturbations began to adversely affect network performance -- for noise less than this amount, persistent activity was accurately maintained over a timescale of seconds, whereas larger noise levels led to progressively larger drifts of activity.

Next we demonstrate that a similar temporal integration of inputs can also occur in spatially structured networks. We consider a classic "ring model" architecture commonly used to model spatial working memory tasks in which stimuli can be presented at any of various locations arranged in a circular (ring-like) layout. The model consists of a ring of neurons with local excitatory connectivity and functionally wider inhibitory connectivity (Fig. 5A, Methods). Such models can generate persistent activity at any spatial location along the ring, but typically have only a binary "on-off" representation at a given spatial location (Fig. 5B). When we added an oscillatory input stimulus to such a ring model, the network could store multiple, discrete levels of activity at any spatial location (Fig. 5C) and could temporally integrate location-specific inputs into discretized levels (Fig. 5D). While the spatial memory (bump attractor) networks proposed in (36, 37) are capable of generating graded persistent activity, the network presented here represents, to our knowledge, the first spatial memory network to encode multi-level activity without requiring an exact tuning condition.

Recent studies have shown that memory activity during a delay period also may take the form of a sequence of activity that spans the delay (38–40). Models of such activity typically generate chain-like patterns of activity that attain only a single, stereotyped level of firing rate. Consistent with this, when we constructed a network with a chain-like architecture (Fig. 6), we found that, in the absence of oscillatory input, the sequential network activity either quickly decayed when the initial stimulus amplitude was too small or converged to a single saturated level of activity for larger stimuli (Fig. 6B). By contrast, in the presence of a subthreshold oscillatory input, the network could exhibit multiple levels of sequential activity for the same pattern of input (Fig. 6C). Thus, as in the persistently active networks, the oscillatory sequential memory network could encode multiple, discretized stimulus levels.

In summary, this work demonstrates a simple mechanism by which oscillatory input to a memory network can transform it from storing only binary amplitudes to maintaining multiple levels of persistent activity. Memory networks using this mechanism require a cellular, synaptic, or network process that can span the period of the oscillation, suggesting a possible tradeoff in memory storage: higher frequency oscillations do not require long timescale processes to span the oscillation cycle, but due to their short period may only store one or a few values (Figure 2F);

lower frequency oscillations could store more items, but require a process with longer timescale to bridge the troughs occurring in each cycle. Our work complements traditional attractor models of working memory that typically fall into two classes: bistable models that robustly maintain two levels of activity (Figures 1B and 2B) and continuous attractor models can maintain nearly analog storage of memory but require very precise tuning of connection weights (Figures 1D and 3C). Our model represents an intermediate possibility with relatively moderate tuning requirements (Fig. 3B) and a discrete set of response levels. Previous work (41, 42) has suggested how multiple, spatially distinct bistable processes in a cell can be coupled together to form multiple stable levels of firing activity; here we demonstrate a complementary mechanism for forming multi-stable representations that relies on temporal, rather than spatial, patterning of inputs. Altogether, this work suggests a potential mechanism by which oscillatory activity, which is commonly observed during working memory tasks, may expand short-term memory capacity.

Methods

The Wang-Buzsaki model neuron used for most spiking neuron simulations is based upon the original model described in (31). Below, we show the equations for the dynamical variables most relevant to the maintenance of multi-level persistent activity. The full model equations are included in the Supplemental Material. The membrane potential of the Wang-Buzsaki neuron is given by:

$$C_m \frac{dV_i}{dt} = -I_{Na}(V_i, h_i) - I_K(V_i, n_i) - I_L(V_i) + I_{syn,i}(s_1, \dots, s_N) + I_0 + I(t) + I_{ext,i}(t) \quad (1)$$

$$I_{syn,i}(s_1, \dots, s_N) = g_{syn} \sum_{j=1}^{N} w_{ij} s_j \quad (2)$$

$$\tau_{syn} \frac{ds_i}{dt} = -s_i + \sum_{t_i^{spike}}^{n} \delta\left(t - t_i^{spike}\right) \quad (3)$$

$$I(t) = \psi \cos(\omega t) \quad (4)$$

where h and n are time-varying channel variables (Supplemental Material). The parameter values used are specified in Tables S1 and S2. The Wang-Buzsaki neuron receives several sources of inputs: (1) $I_{syn,i}(s1,\dots,sN)$ represents recurrent feedback to neuron i, the strength of which is determined by a weight matrix w_{ij} defining the strength of the connection from neuron j to neuron i, (2) I_0 is a constant current that shifts the resting potential, and could represent tonic background input or intrinsic currents not explicitly modeled, (3) $I(t)$ is the external oscillatory input (I(t)=0 for models with no oscillatory input), and (4) $I_{ext,i}(t)$ represents the external inputs to be accumulated and stored by the memory network. To calculate spike times in equation 3, we used the time of the peak of the action potential, with only action potentials exceeding a voltage

of 0 mV included. Integration was performed numerically using the fourth order Runge-Kutta method with a time step $\Delta t = 10^{-2}$ ms.

In the single neuron case, there is a single recurrent synaptic weight, w $[\mu A/cm^2]$. Values for all simulation parameters are included in Table S2. In Figures 4-6, we study three different network architectures composed of Wang-Buzsaki neurons: an all-to-all connectivity (Figure 4), a ring structure (Figure 5), and a directed structure (Figure 6).

The all-to-all connected networks of Figure 4 are composed of 1000 Wang-Buzsaki neurons. Figures 4B,C implement a network with uniform connection strengths $w_{ij} = c\frac{w}{N}$ $[\mu A/m^2]$. Figure 4D implements a network in which these uniform connection strengths have been perturbed by adding static Gaussian noise of mean zero independently to each connection. Exponentially filtered temporally white noise (Ornstein-Uhlenbeck process) input was implemented in the network illustrated in Figure 4C; for each neuron i, the additive noise was given independently by:

$$n_i(t) = n_{i,t-\Delta t} - n_{i,t-\Delta t}\frac{\Delta t}{\tau_n} + \sigma_n \eta_{i,t}\sqrt{\frac{\Delta t}{\tau_n}} \tag{5}$$

$$\eta_{i,t} \sim N(0,1) \tag{6}$$

where σ_n is the standard deviation of the noise.

For the ring connectivity structure in Figure 5, the connection strength from neuron j to neuron i is described by:

$$w_{ij} = A + B\,cos\left(\frac{2\pi(i-j)}{N}\right)\ [\mu A/cm^2] \tag{7}$$

The directed structure illustrated in Figure 6 resembles the ring structure, but results in a drift of the 'activity bump' in one direction. The connection strength from neuron *j* to neuron *i* in this case is defined by:

$$w_{ij} = A + B\ cos\left(\frac{4\pi(i-j)}{N} + 0.1\right) H(C - |i-j|)\ [\mu A/cm^2] \tag{8}$$

where H is the Heaviside step function and *C* controls the spatial extent of the connectivity.

Comparison to linear spiking autapse model

In Figure 3, we compare the robustness of multi-level persistent activity of the phase-locking nonlinear spiking model described above, to that of a spiking autapse model in which analogue persistent activity is enabled by excitatory feedback that is tuned to cancel the intrinsic decay of activity. The autapse model is described in detail in (35); equations describing the dynamics of the model are included in the supplemental methods.

Simple rate model

The equation for the simple rate model implemented for Figure 1 is given in Figure 1A. The nonlinear term used for Figure 1B is:

$$f(x) = \sqrt{w[x - x_{thr}]_+} \tag{9}$$

with $w = 75$, $x_{thr} = 10\ \ [Hz]$.

References

1. J. Jonides, *et al.*, The Mind and Brain of Short-Term Memory. *Annual Review of Psychology* **59**, 193–224 (2008).
2. G. Major, D. Tank, Persistent neural activity: prevalence and mechanisms. *Current Opinion in Neurobiology* **14**, 675–684 (2004).
3. H. S. Seung, How the brain keeps the eyes still. *PNAS* **93**, 13339–13344 (1996).
4. X. J. Wang, Synaptic reverberation underlying mnemonic persistent activity. *Trends Neurosci* **24**, 455–463 (2001).
5. J. Zylberberg, B. W. Strowbridge, Mechanisms of Persistent Activity in Cortical Circuits: Possible Neural Substrates for Working Memory. *Annu. Rev. Neurosci.* **40**, 603–627 (2017).
6. C. D. Brody, R. Romo, A. Kepecs, Basic mechanisms for graded persistent activity: discrete attractors, continuous attractors, and dynamic representations. *Current Opinion in Neurobiology* **13**, 204–211 (2003).
7. O. Barak, M. Tsodyks, Working models of working memory. *Current Opinion in Neurobiology* **25**, 20–24 (2014).
8. M. S. Goldman, A. Compte, X.-J. Wang, "Neural Integrator Models" in *Encyclopedia of Neuroscience*, L. R. Squire, Ed. (Academic Press, 2009), pp. 165–178.
9. R. Chaudhuri, I. Fiete, Computational principles of memory. *Nat Neurosci* **19**, 394–403 (2016).
10. F. Roux, P. J. Uhlhaas, Working memory and neural oscillations: alpha–gamma versus theta–gamma codes for distinct WM information? *Trends in Cognitive Sciences* **18**, 16–25 (2014).
11. X.-J. Wang, Neurophysiological and Computational Principles of Cortical Rhythms in Cognition. *Physiological Reviews* **90**, 1195–1268 (2010).
12. N. Kopell, M. A. Kramer, P. Malerba, M. A. Whittington, Are Different Rhythms Good for Different Functions? *Front. Hum. Neurosci.* **0** (2010).
13. A. Bollimunta, J. Mo, C. E. Schroeder, M. Ding, Neuronal Mechanisms and Attentional Modulation of Corticothalamic Alpha Oscillations. *J. Neurosci.* **31**, 4935–4943 (2011).
14. J. E. Lisman, M. A. Idiart, Storage of 7 +/- 2 short-term memories in oscillatory subcycles. *Science* **267**, 1512–1515 (1995).
15. N. Kopell, M. A. Whittington, M. A. Kramer, Neuronal assembly dynamics in the beta1 frequency range permits short-term memory. *PNAS* **108**, 3779–3784 (2011).
16. J. E. Lisman, O. Jensen, The Theta-Gamma Neural Code. *Neuron* **77**, 1002–1016 (2013).

17. A. K. Engel, P. Fries, W. Singer, Dynamic predictions: Oscillations and synchrony in top–down processing. *Nature Reviews Neuroscience* **2**, 704–716 (2001).
18. A. Palmigiano, T. Geisel, F. Wolf, D. Battaglia, Flexible information routing by transient synchrony. *Nat Neurosci* **20**, 1014–1022 (2017).
19. J. Riddle, J. M. Scimeca, D. Cellier, S. Dhanani, M. D'Esposito, Causal Evidence for a Role of Theta and Alpha Oscillations in the Control of Working Memory. *Current Biology* **30**, 1748-1754.e4 (2020).
20. M. J. Jutras, P. Fries, E. A. Buffalo, Oscillatory activity in the monkey hippocampus during visual exploration and memory formation. *Proceedings of the National Academy of Sciences* **110**, 13144–13149 (2013).
21. K. L. Hoffman, *et al.*, Saccades during visual exploration align hippocampal 3–8 Hz rhythms in human and non-human primates. *Front. Syst. Neurosci.* **7** (2013).
22. J. E. Pina, M. Bodner, B. Ermentrout, Oscillations in working memory and neural binding: A mechanism for multiple memories and their interactions. *PLOS Computational Biology* **14**, e1006517 (2018).
23. C. Börgers, S. Epstein, N. J. Kopell, Gamma oscillations mediate stimulus competition and attentional selection in a cortical network model. *PNAS* **105**, 18023–18028 (2008).
24. M. Dipoppa, B. S. Gutkin, Flexible frequency control of cortical oscillations enables computations required for working memory. *Proceedings of the National Academy of Sciences* **110**, 12828–12833 (2013).
25. S. Kornblith, T. J. Buschman, E. K. Miller, Stimulus Load and Oscillatory Activity in Higher Cortex. *Cerebral Cortex* **26**, 3772–3784 (2016).
26. E. K. Miller, M. Lundqvist, A. M. Bastos, Working Memory 2.0. *Neuron* **100**, 463–475 (2018).
27. M. Chalk, B. Gutkin, S. Denève, Neural oscillations as a signature of efficient coding in the presence of synaptic delays. *eLife* **5**, e13824 (2016).
28. J. F. Mejias, J. D. Murray, H. Kennedy, X.-J. Wang, Feedforward and feedback frequency-dependent interactions in a large-scale laminar network of the primate cortex. *Science Advances* **2**, e1601335 (2016).
29. M. Lundqvist, *et al.*, Gamma and Beta Bursts Underlie Working Memory. *Neuron* **90**, 152–164 (2016).
30. A. Compte, N. Brunel, P. S. Goldman-Rakic, X.-J. Wang, Synaptic Mechanisms and Network Dynamics Underlying Spatial Working Memory in a Cortical Network Model. *Cerebral Cortex* **10**,

910–923 (2000).

31. X.-J. Wang, G. Buzsáki, Gamma Oscillation by Synaptic Inhibition in a Hippocampal Interneuronal Network Model. *J. Neurosci.* **16**, 6402–6413 (1996).
32. X.-J. Wang, Synaptic Basis of Cortical Persistent Activity: the Importance of NMDA Receptors to Working Memory. *J Neurosci* **19**, 9587–9603 (1999).
33. G. Major, A. Polsky, W. Denk, J. Schiller, D. W. Tank, Spatiotemporally Graded NMDA Spike/Plateau Potentials in Basal Dendrites of Neocortical Pyramidal Neurons. *Journal of Neurophysiology* **99**, 2584–2601 (2008).
34. O. Shriki, D. Hansel, H. Sompolinsky, Rate Models for Conductance-Based Cortical Neuronal Networks. *Neural Computation* **15**, 1809–1841 (2003).
35. H. S. Seung, D. D. Lee, B. Y. Reis, D. W. Tank, The Autapse: A Simple Illustration of Short-Term Analog Memory Storage by Tuned Synaptic Feedback. *J Comput Neurosci* **9**, 171–185 (2000).
36. S. Carroll, K. Josić, Z. P. Kilpatrick, Encoding certainty in bump attractors. *J Comput Neurosci* **37**, 29–48 (2014).
37. S. Lim, M. S. Goldman, Balanced Cortical Microcircuitry for Spatial Working Memory Based on Corrective Feedback Control. *Journal of Neuroscience* **34**, 6790–6806 (2014).
38. E. Pastalkova, V. Itskov, A. Amarasingham, G. Buzsaki, Internally Generated Cell Assembly Sequences in the Rat Hippocampus. *Science* **321**, 1322–1327 (2008).
39. C. J. MacDonald, K. Q. Lepage, U. T. Eden, H. Eichenbaum, Hippocampal “Time Cells” Bridge the Gap in Memory for Discontiguous Events. *Neuron* **71**, 737–749 (2011).
40. C. D. Harvey, P. Coen, D. W. Tank, Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* **484**, 62–68 (2012).
41. A. A. Koulakov, S. Raghavachari, A. Kepecs, J. E. Lisman, Model for a robust neural integrator. *Nature Neuroscience* **5**, 775–782 (2002).
42. M. S. Goldman, Robust Persistent Neural Activity in a Model Integrator with Multiple Hysteretic Dendrites per Neuron. *Cerebral Cortex* **13**, 1185–1195 (2003).

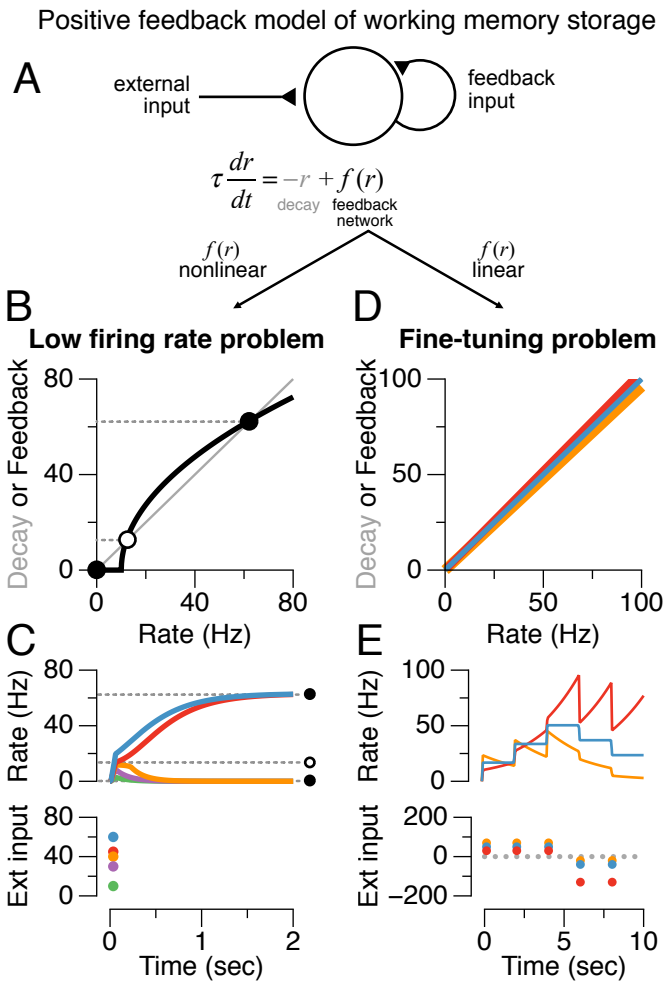

Figure 1. Failures of the positive feedback model of working memory storage. (A) Description of positive feedback models; the location and stability of fixed points depend on whether the model is implemented with linear or nonlinear feedback functions. (B) Nonlinear implementations exhibit the 'low firing rate problem': There is not an attracting fixed point (solid circles) within the range of firing rates typically observed during persistent neural activity (top). Firing rates below the unstable fixed point (open circle) decay to zero, while firing rates above the unstable fixed point run off to uncharacteristically high rates (bottom). (C) Linear implementations exhibit the 'fine tuning problem': minute changes in the strength of feedback (red: +5%, orange: -5%) relative to the tuned value (blue) result in instability, and the inability to maintain stable persistent activity.

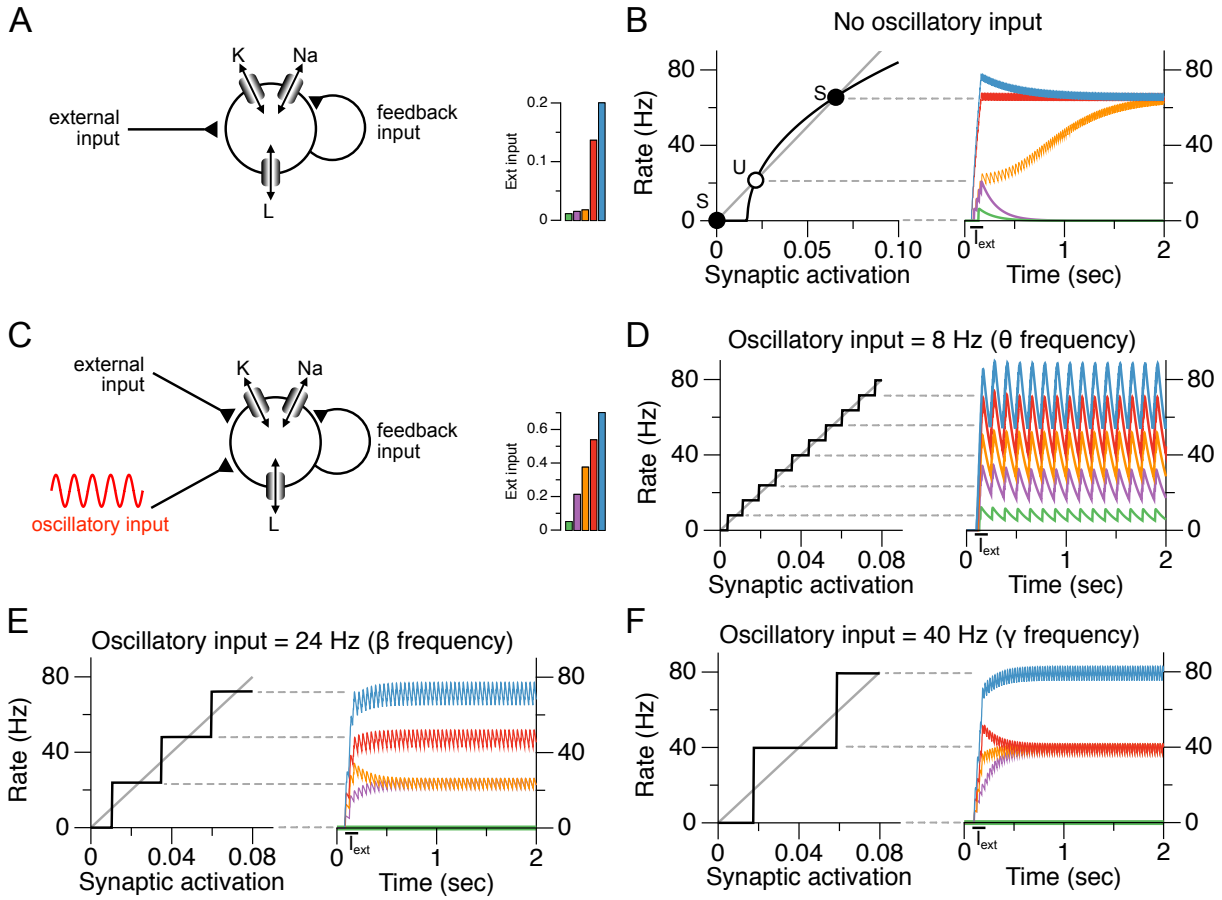

Figure 2. Baseline oscillatory input allows robust maintenance of discretely graded persistent activity levels in a conductance-based (spiking) model neuron. (A) Schematic of conductance-based autapse model. The model described by (31) is composed of potassium, sodium, and leak conductances, and receives current input from an external source; the model is expanded here to include a recurrent feedback input. (B) Manifestation of 'low-firing-rate' problem in the conductance-based model without oscillatory input. Similar to the nonlinear firing rate model depicted in Figure 1B, the conductance-based spiking model exhibits stable fixed points only at zero and high rates (filled circles). Spiking rates between these two fixed points will decay to zero from below the unstable fixed point (open circle) or will run off to high rates from above the unstable fixed point. (C) Schematic of conductance-based neuron with the addition of an oscillatory baseline input. (D-F) Maintenance of discretely-graded persistent activity levels enabled by baseline oscillatory input. Phase-locking to oscillatory input creates stable fixed points at integer multiples of the baseline frequency. (D) Lower-frequency oscillations allow a higher number of discrete rates to be stored than higher frequencies (E,F), but these panels illustrate the trade-off between the robustness of the fixed points created by high-frequency oscillations (F), which are related to the size of the 'stairs' shown, and the storage capacity of more closely-spaced fixed points created by lower-frequency oscillations (D).

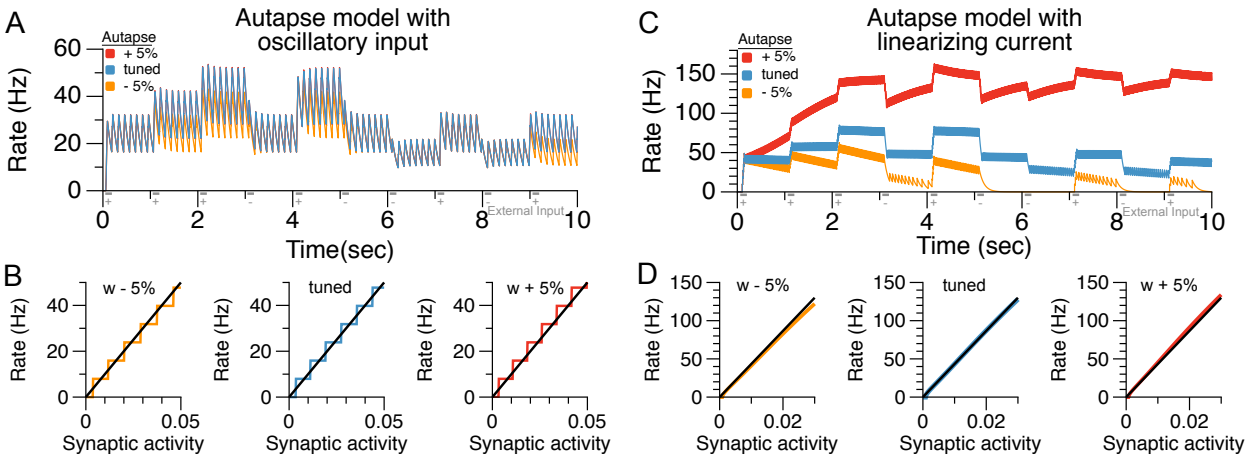

Figure 3. Robustness of oscillator-integrator model to detuning of recurrent feedback. (A) Responses of model neuron to a sequence of positive and negative input pulses. Red and yellow traces show the conditions in which the recurrent feedback strength has been detuned by plus or minus 5%, respectively. There was little change to the red trace, which is overlapped by the output of the tuned (blue) model. (B) Steady-state firing rates as a function of synaptic activity that is held at steady values; mistuning the autapse strength by 5% has no effect on the existence and location of the stable fixed points (intersections of black lines and horizontal stairs). (C -D) Comparison to another conductance-based spiking model of persistent neural activity. The model by (35) demonstrates the reliance on exact tuning of feedback strength that is overcome by the oscillator-integrator model. Here, detuning the recurrent feedback strength by 5% (orange and red traces) causes spiking activity to decay to 0 (orange, decreasing feedback strength) or run off to high rates (red, increasing feedback strength). These qualitatively large differences are driven by very small deviations from the tuned case (D).

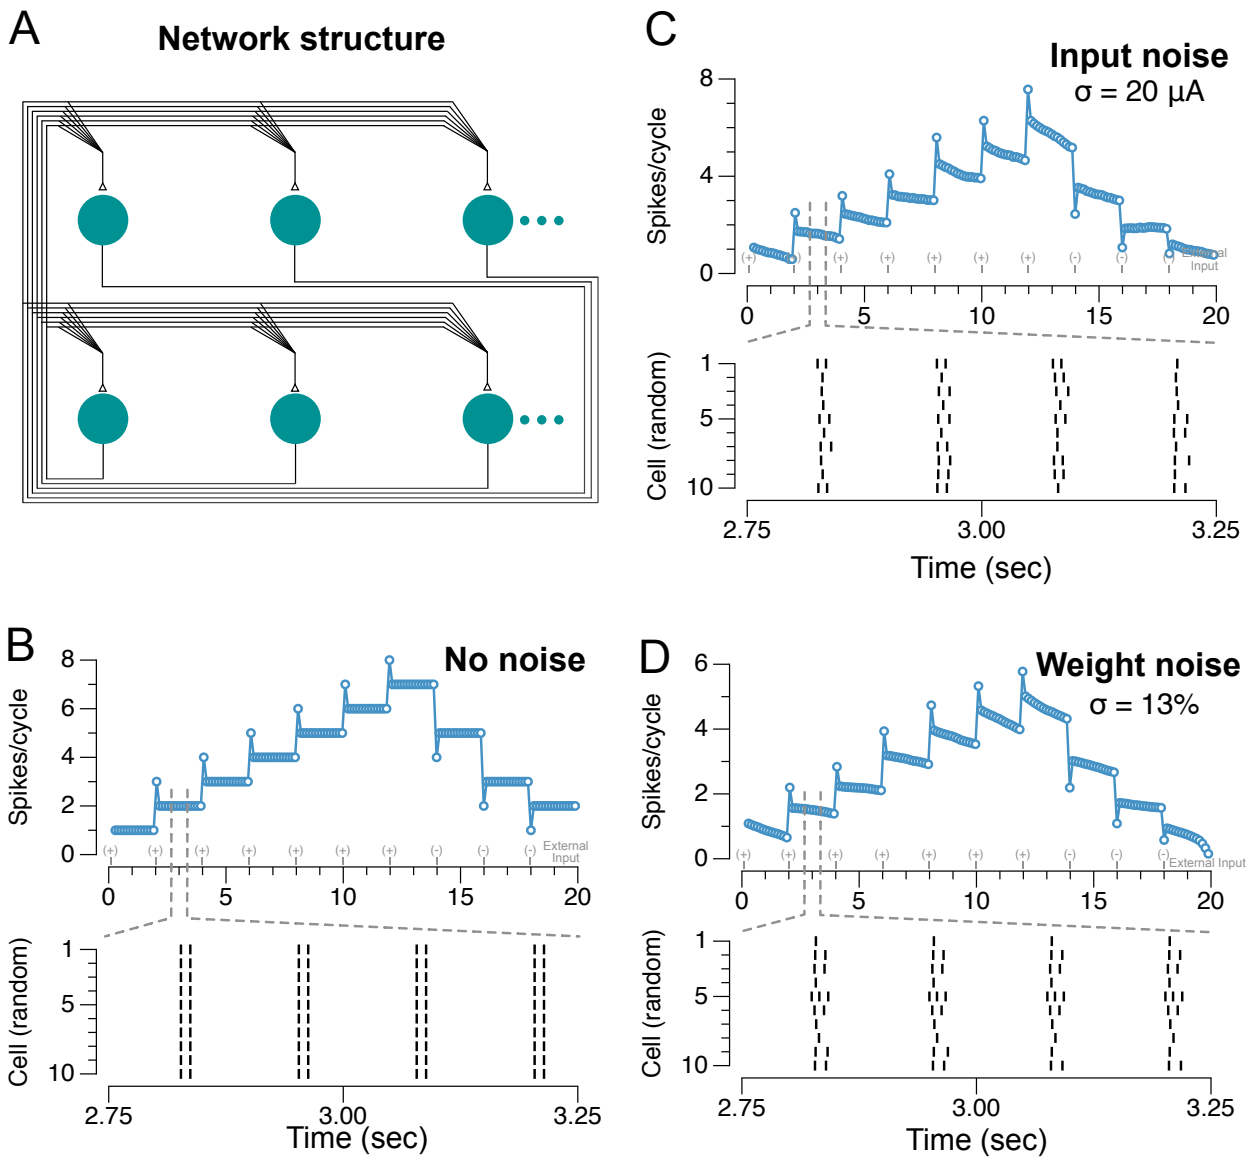

Figure 4. Maintenance of persistent activity and robustness to noise in a fully-connected network of 1000 neurons. (A) Schematic of network. All nodes make synapses on all other nodes, with mean synapse weight 1/N. (B) Spiking responses of network neurons to a sequence of positive and negative input pulses, with spike raters plotted for the time window shown by the dashed grey line. (C) Spiking responses of network to the same input sequence in the presence of ongoing input noise. The noise has a mean of zero and the standard deviation is roughly 30% of the magnitude of the individual input pulses. The network is clearly able to maintain persistent activity despite the noisy input. (D) Spiking responses of the network initialized with random noise in the connection weights. Noise with mean zero and standard deviation of 15% of the mean connectivity strength is added to the individual connection weights between neurons. While individual neurons in the network respond with different rates, the network is able to maintain persistent activity at many levels.

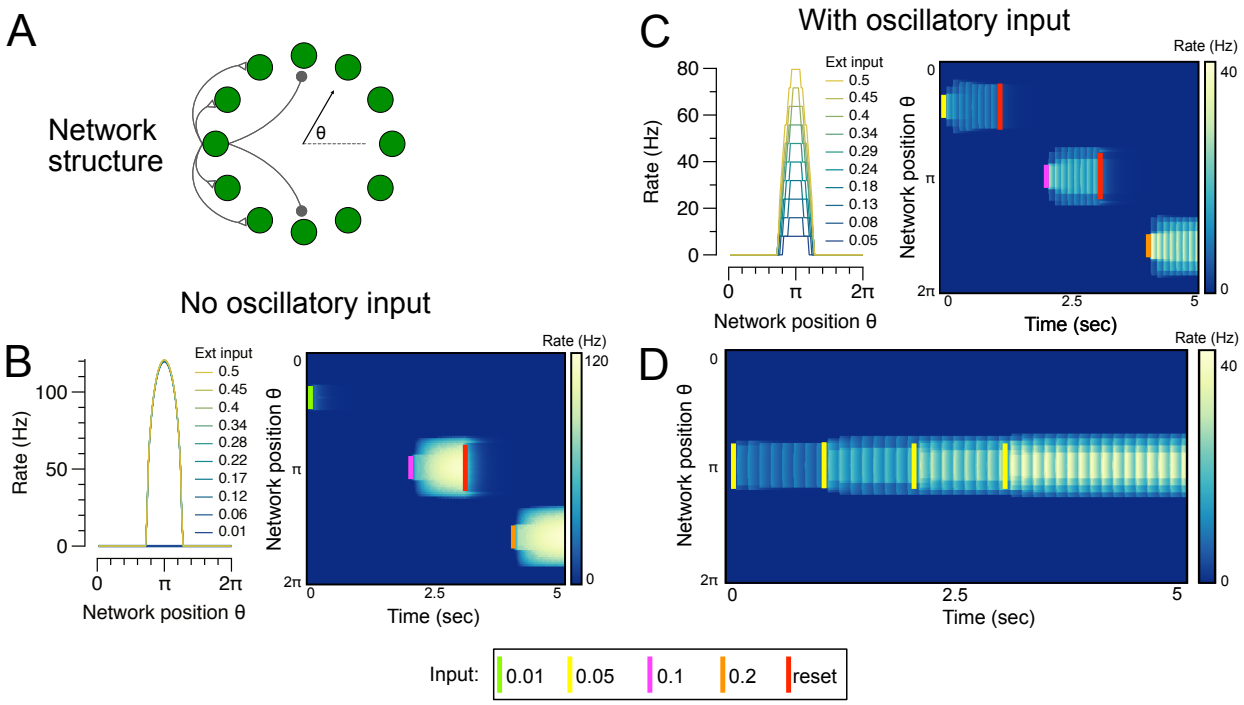

Figure 5. Maintenance of discretely-graded bumps of persistent activity in a ring network. (A) Schematic of network structure. The spatial position of network neurons in (B-C) are indexed by the angle theta from an arbitrary reference neuron. (B) Illustration of low-rate problem in a ring network of conductance-based spiking neurons without oscillatory input. Left column plots the steady-state firing rate response of neurons in the network to input pulses of graded levels at locations centered around network position=pi. Heatmap plots at right illustrate network's temporal firing-rate responses to short pulses of inputs at network locations labeled by colored bars. The network is unable to maintain bump activity levels between the low and high fixed points. (C) Ring network of oscillator-integrator neurons is able to maintain discretely-graded bumps of persistent activity. (D) Temporal integration in the ring network. Short input pulses to the network (labeled by colored bars) are integrated and stored in persistent activity.

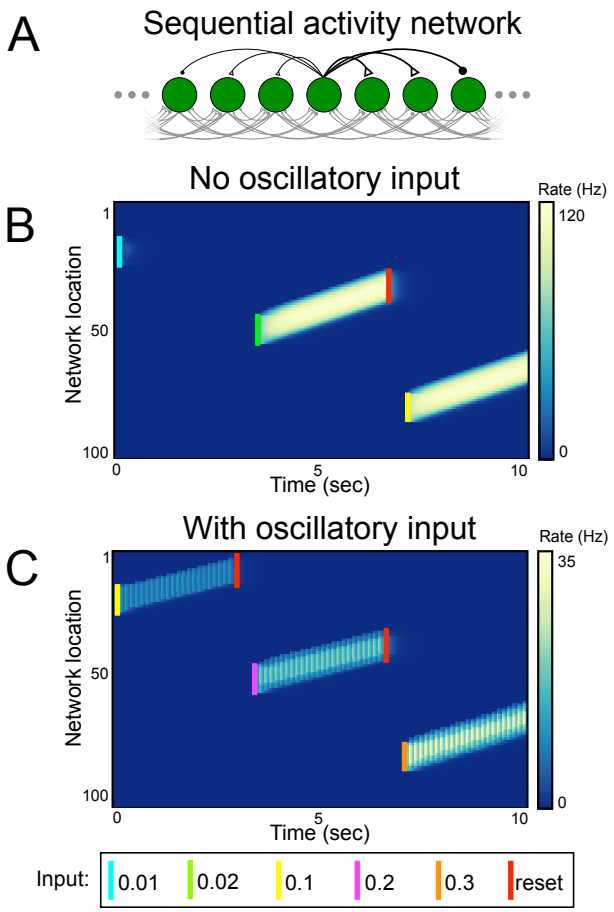

Figure 6. Generation of sequences of discretely graded activity in network with asymmetric connectivity. (A) Schematic of network structure. Asymmetric connectivity underlies slow drift of activity bumps. (B) Illustration of low-rate problem in a sequential-activity network of conductance-based spiking neurons without oscillatory input. Drifting bumps of activity in the network initiated by short pulses (labeled by colored bars) exhibit only a single level of activity. (C) A sequential-activity network of oscillator-integrator neurons is able to maintain drifting bumps at discretely graded levels of activity.

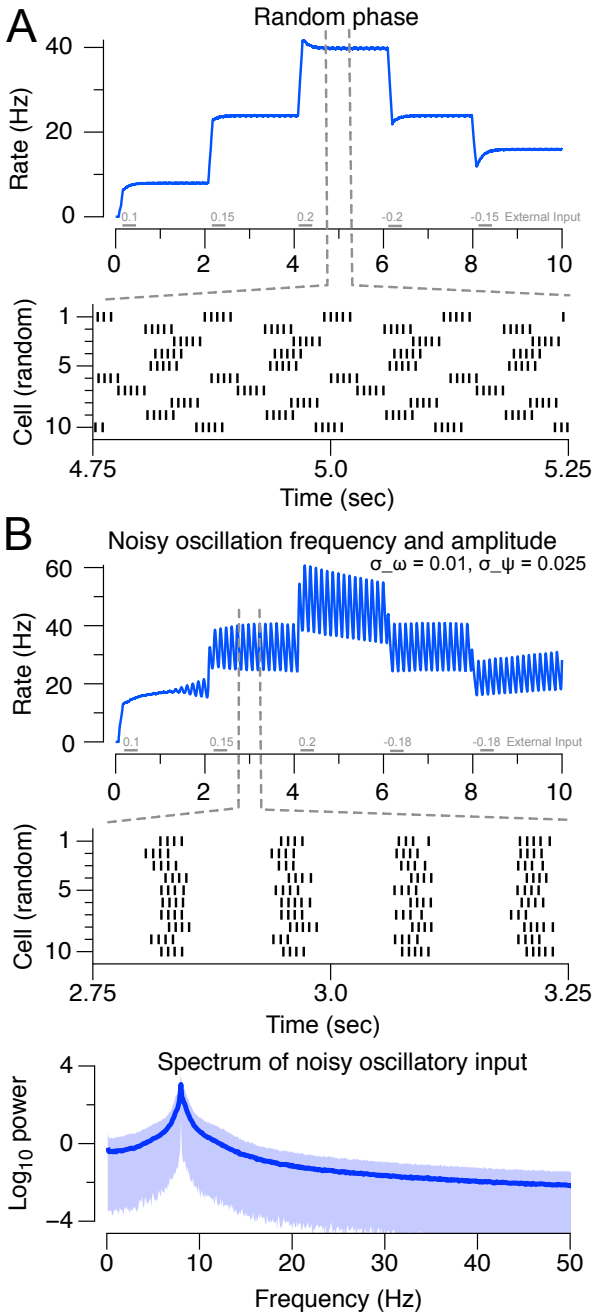

Figure S1. Robustness of persistent activity in a fully-connected network of 1000 oscillator-integrator neurons receiving noisy oscillatory input or oscillatory input with randomized phase across neurons. Noisy oscillatory activity is modeled multiplicatively (see methods). Noise is independent (uncorrelated) between neurons of the network. A, mean firing rate and rasters depicting responses of neurons in the network to a sequence of brief (100-ms) input pulses while oscillatory phase of the input has been randomly initialized. B, Firing rate and rasters of network neuron responses to sequence of brief input pulses with noisy oscillatory input. Bottom, power spectrum of oscillatory input.

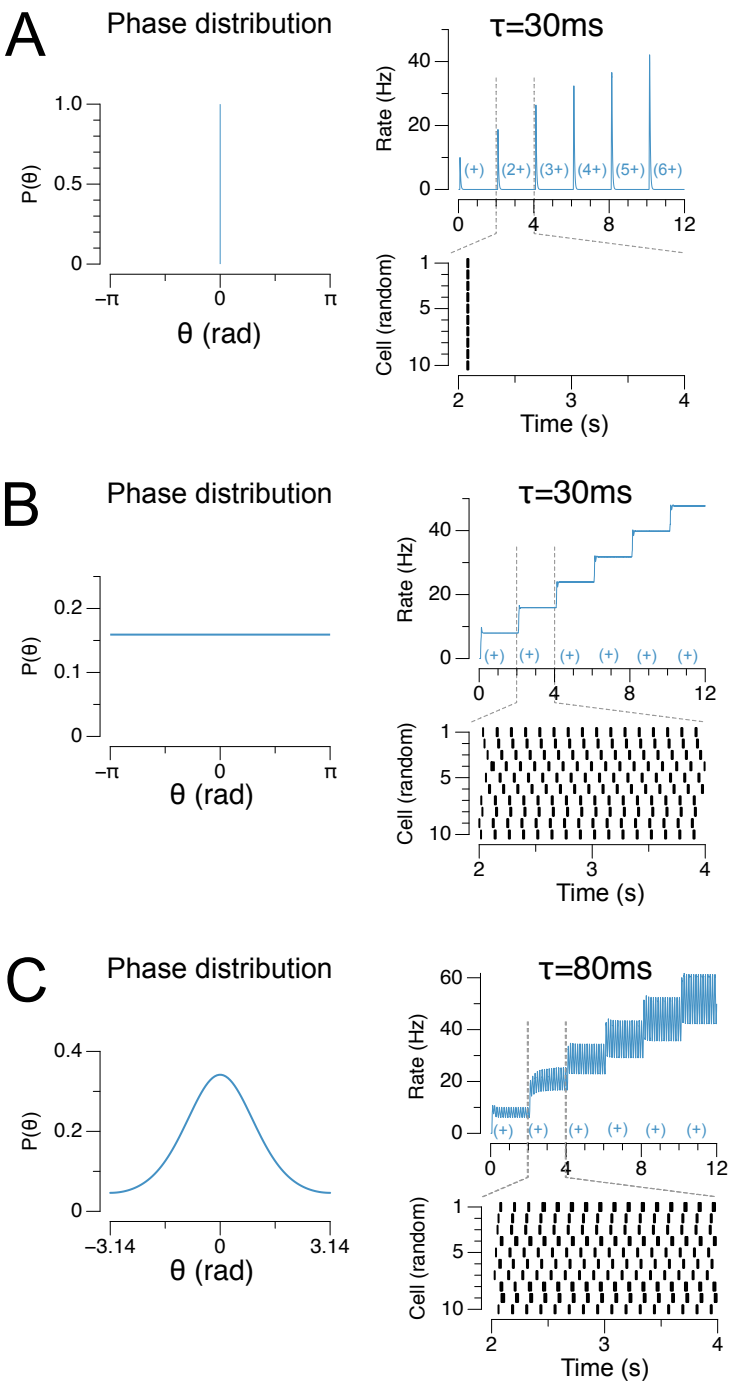

Figure S2. Heterogeneity in the phases of oscillations in the network enables multi-level activity storage using synaptic time constants that are shorter than the oscillation cycle. (A) Failure of network with perfectly aligned phases and short synaptic time constants to maintain multi-level persistent activity. When all oscillatory phases are perfectly aligned (left), fast synapses (tau=30 ms) cannot bridge the oscillation cycles, and the network fails to maintain persistently elevated firing (right). Times and sizes of external pulses of increasing magnitude are depicted along the x-axis with (6+) denoting a pulse of six times the magnitude of (+), etc. (B) Distributed phases of oscillatory inputs enables the network to store multi-level activity even with short (30 ms) synaptic time constants. When phases of the oscillatory inputs to the network are independently drawn from a uniform distribution (left), The relative spread of spiking responses in the network can bridge the oscillation cycles and enable multi-level persistent activity. Responses to a sequence of equally-sized short input pulses shown. (C) Intermediate between the phase distributions illustrated in (A,B), a more biologically-plausible distribution of oscillatory inputs to the network (left), given by sampling the phases from a Von Mises distribution (concentration parameter k=1.0), the network can maintain multi-level persistent activity with synaptic time constants that are consistent with commonly observed values (shown here, tau = 80 ms).