Supplemental Material Inventory:

Supplemental Figure 1: Divergence values are sensitive to higher-order correlations. Related to Figures 2 and 3.

Supplemental Figure 2: Analysis of data from most well-isolated neurons. Related to Figures 6 and 7.

Supplemental Figure 3: Properties of V1-V2 CCGs. Related to Figures 4 and 5.

Supplemental Figure 4: Comparing divergence analysis for middle-layer V2 neurons with and without a significant V1-V2 CCG. Related to Figures 6, 7, and 8.

Supplemental Figure 5: Divergence in the middle-layers of V2 neurons, conditioned on spiking in other V2 layers. Related to Figures 1, 6, and 7.

Supplemental Figure 6: Generating synthetic population responses with physiologically-realistic correlation timescales. Related to Figure 2.

Supplemental Experimental Procedures: Equations and parameters for the simulations. Related to Figures 1 and 8.

Supplemental References

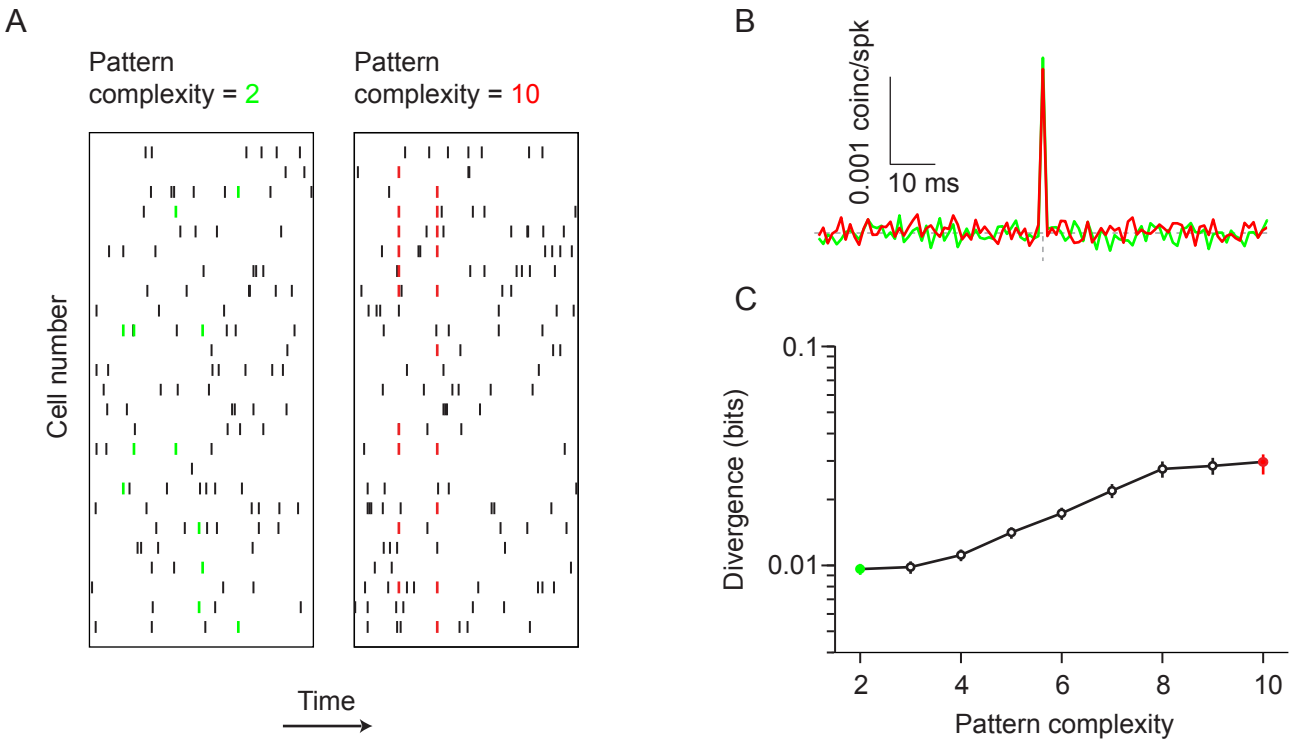

Supplemental Figure 1: **Divergence values are sensitive to higher-order correlations.** Related to Figures 2 and 3.

In Figure 3 we showed that our divergence metric is sensitive to the magnitude of pairwise correlations, but not to the firing rates of individual neurons. However, cortical population responses are not fully specified by first- and second-order statistics; cortical activity may contain higher-order dependencies (Ohiorhenuan et al., 2010) and these can strongly influence the drive provided to downstream networks (Kuhn et al., 2003).

Because our divergence metric relies on quantifying the difference between empirical distributions and those based on jittered responses, it should be sensitive to variations in higher-order structure. To test this intuition, we generated synthetic population responses in which we varied the strength of higher-order correlations. This required using a distinct algorithm from that used in the main text, which only specifies first- and second-order statistics. We generated population responses using 100 independent Poisson processes; correlations were introduced by adding coordinated events consisting of 2-10 cells firing together (termed pattern complexity, two examples plotted in Figure S1A). In all cases, the firing rate of each individual cell was fixed at 20 Hz and the number of coordinated events added was chosen to produce pairwise CCGs with a peak of 0.002 coinc/spk. As shown in Figure S1B, the CCGs were nearly identical for populations with the lowest (green) and highest (red) pattern complexity.

We then calculated divergence values for the resultant populations, for a jitter window of 10 ms (larger jitter windows gave identical results since the introduced patterns involve perfect synchrony). Divergence values rose monotonically with pattern complexity (Figure S1C). Thus, divergence values reflect not only pairwise (Figure 3) but also higher-order response statistics. This behavior follows from the fact that our metric reflects the divergence between network patterns and those expected given the same surrogate data, in which pairwise and higher-order response statistics have been removed.

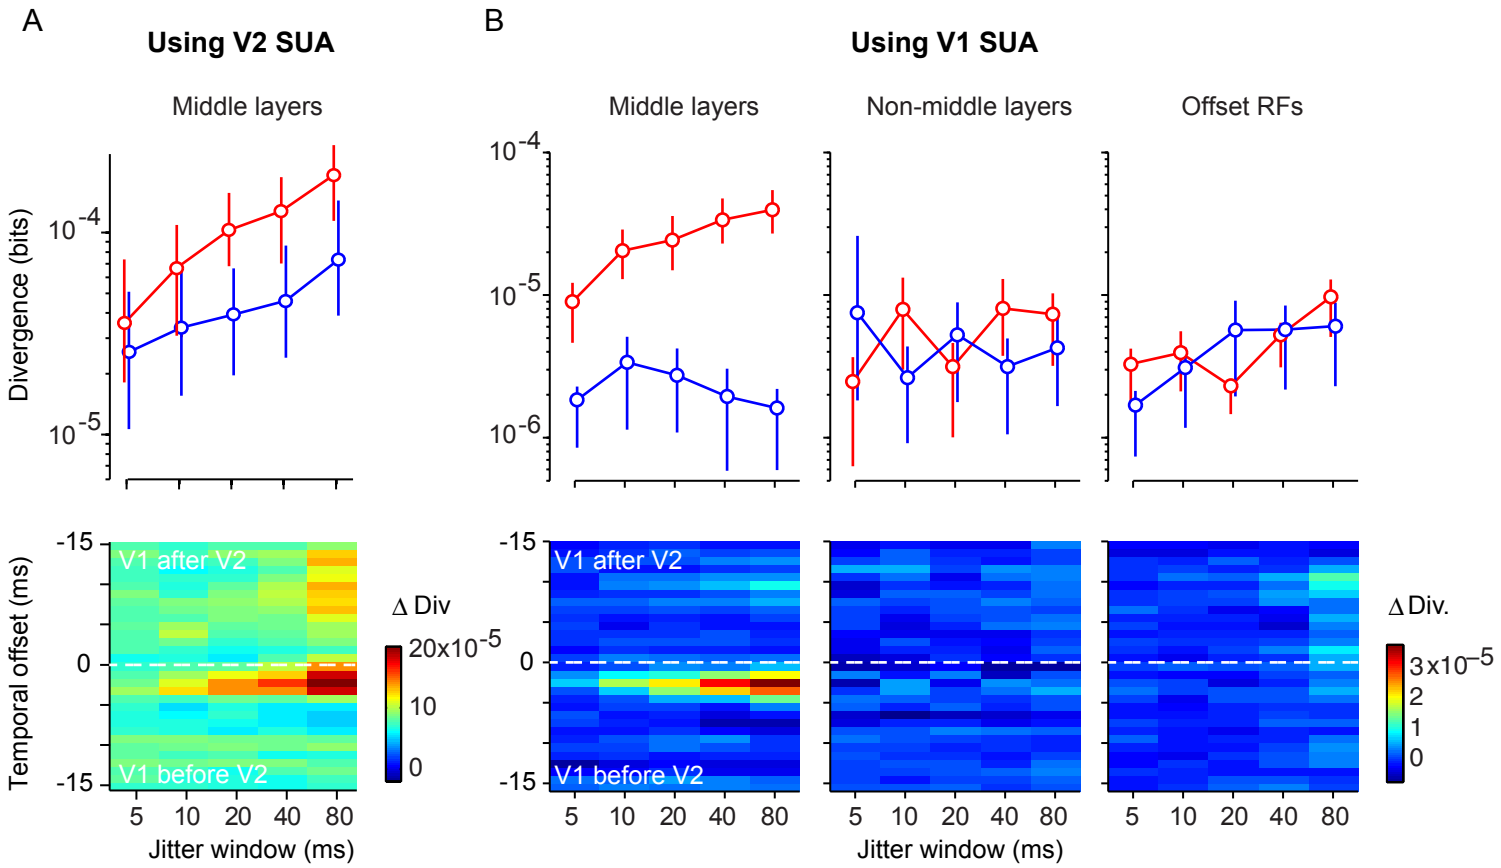

Supplemental Figure 2: Analysis of data from most well-isolated neurons. Related to Figures 6 and 7.

We computed the signal-to-noise ratio (SNR) of each recorded unit as the ratio of the average waveform amplitude to the SD of the waveform noise (Kelly et al., 2007). In the main text we included units with an SNR of 2 or greater. This threshold provides data consisting of both single units and small clusters of units (MUA). Importantly, the MUA consists of clearly defined waveforms, but because of the low amplitude or diversity of shapes, we cannot be sure they come from a single unit. In previous work (Wissig and Kohn, 2012) we have compared the peak firing rate and orientation selectivity, for our MUA and SUA recordings. The peak firing rate of these combined signals (15.2 sp/s) was only slightly higher than that of well-isolated single units (12.1 sp/s), and the orientation selectivity was similar as well. Thus, the MUA activity in our recordings likely consists of the responses of only a few distinct neurons.

Here we use a stringent SNR threshold of 3, to select well-isolated V2 single units. We first apply this selection criterion to the middle layers of V2, yielding 43 neurons. Figure S2A (top) shows the results of our divergence analysis were similar for this subpopulation of V2 cells, compared to the full population shown in Figure 7A of the main text: divergence was significantly elevated in epochs preceding V2 spiking, relative to control epochs, for most jitter windows ($p=0.3$ for 5 ms jitter window; $p=0.08$ for 10 ms; $p<0.01$ for all other comparisons). Figure S2A (bottom) shows the difference in divergence between epochs around V2 spiking and control epochs, for a range of jitter windows (abscissa) and temporal offsets (ordinate). This figure can be compared to Figure 7D of the main text. Divergence is elevated several milliseconds before the spiking of single units in the middle layers of V2. Applying this selection criterion to neurons in the non-middle layers of V2 or to V2 neurons with offset RFs revealed

no difference in V1 divergence in epochs associated with V2 spiking compared to control epochs (not shown), as observed for the full data set.

We then applied our SNR selection criteria to the V1 population (yielding populations of 19.7±1.1 neurons), retaining all V2 neurons. Figure S2B shows results in an identical format to those shown in Figure 7 of the main text. Divergence is elevated in a population of V1 single units before spiking in the middle layers of V2 ($p<0.001$ for all jitter windows). V1 divergence is not elevated before spiking in the superficial or deep layers of V2, or before spiking at offset V2 sites.

Note that the divergence values for V1 SUA are different from those reported in the main text. This may indicate a different level of coordination between single-units than multi-units, but it also reflects the fact that the divergence measurement here is based on substantially smaller populations. (This is also responsible for the difficulty of detecting brief time scale coordination in the control epochs.) Our divergence metric is most appropriate for comparing the relative level of coordination in a given population at different times, as we have done in the main text.

In summary, the effects illustrated in Figure S2 are similar those presented in the main text. Thus, our primary finding is also evident when the analysis is based on the activity of single units —spiking in the middle layers of V2, but not other layers, is preceded by elevated V1 coordination.

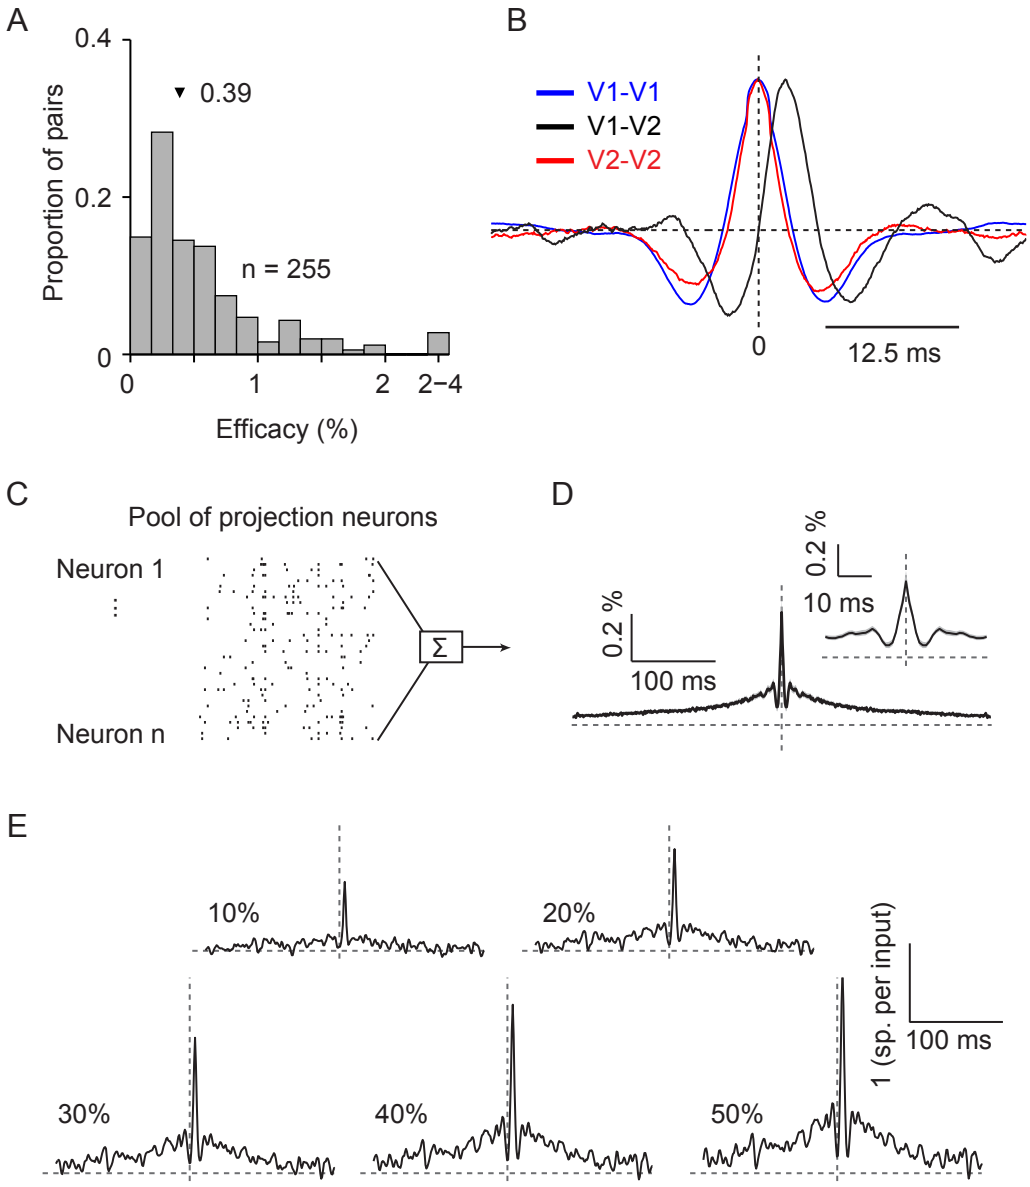

Supplemental Figure 3: Properties of V1-V2 CCGs. Related to Figures 4 and 5.

Figures 4 and 5 of the main text described the key properties of V1-V2 CCGs, which we used to identify V2 sites functionally connected with the monitored V1 network. Here we describe additional properties of the CCGs with significant narrow peaks, focusing on their amplitude, shape, and the temporal offset.

Amplitude: We quantified the strength of V1-V2 coupling by calculating the integral of the CCG in a window ±2.5 ms around the peak. CCGs were normalized by the V1 firing rate, so that their amplitude indicates the probability that each V1 spike will trigger a response in V2 (Levick et al., 1972).

To quantify efficacy across the population, we focused on CCGs whose peak offset was larger than 1 ms (n = 255), as expected for delays between V1 and V2. We did so to minimize the contamination of our estimates by false positives that happened to cross our significance criterion. In the analyzed pairs, the median efficacy was 0.39% (Figure S3A, mean 0.69 ± 0.06%), meaning that after a V1 spike there was an increased probability of the V2 cell firing by this amount.

This estimated efficacy of V1-V2 functional connections is 5-10 fold lower than that reported for thalamocortical projections in visual (Reid and Alonso 1995) and somatosensory systems (Roy and Alloway 2001). There are two factors which may contribute to this difference. First, weaker connections may be explained by the fact that, compared to LGN-V1 connections, there are far more cells in V1 projecting to V2. Because there are fewer thalamocortical projections, each synapse in that system may need to be stronger to ensure a postsynaptic response. *In vitro* evidence suggests that thalamocortical synapses are more powerful than intra-areal corticocortical synapses (Stratford et al., 1996; Gil et al., 1999). To our knowledge, there is no equivalent set of measurements for the efficacy of corticocortical connections. However, anatomical studies have shown that the size of V1 boutons in V2 are smaller those of geniculocortical afferents, consistent with a weaker efficacy (Freund et al., 1989; Anderson and Martin, 2009; Marion et al., 2013).

Second, the weaker V1-V2 efficacy compared to previous thalamocortical measurements could also be due to issues of statistical power. At each site, our recordings lasted for 2-3 hrs, providing us with a large number of spikes and thus the ability to detect weak connections. Previous studies may have provided larger mean estimates of efficacy because weak connections went undetected.

V1-V2 shape: Previous measurements of V1-V2 CCGs (Frien et al., 1994; Nowak et al., 1999; Roe et al., 1999) revealed robust pairwise spiking coordination consistent with the broad peaks in our pairwise measurements (Figure 4 of the main text). However, short-latency, narrow CCG peaks were either absent (Frien et al., 1994; Nowak et al., 1999) or exceedingly rare (Roe et al., 1999) in those studies, presumably because these peaks require sampling from neurons with precisely aligned RFs in specific layers of cortex.

Short-latency, narrow CCG peaks are indicative of a functional interaction between neurons, and have been interpreted as evidence for direct synaptic connectivity (Reid and Alonso, 1995). When a neuron provides direct input to another, the shape of the resultant CCG is expected to show no structure near 0 ms time lag, and to have a sharp rise time and slower decay, related to the shape of the postsynaptic potential (Kirkwood 1979). Our V1-V2 CCGs with narrow peaks did not have these features. This is illustrated in Figure S3B, which shows the average V1-V1 (blue), V2-V2 (red) and V1-V2 (black) jitter-corrected CCGs, using all significant cases (as in Figure 5B of the main text). The CCGs are normalized to the same height to allow direct comparison of their shape. The shape of the V1-V2 CCG is similar to the intra-areal cases, but its peak is clearly shifted rightward.

If V1-V2 CCGs reflect input from V1 to V2, why are the expected features not evident? The answer lies in considering that each V2 neuron receives convergent input from a population of correlated V1 cells. Each spike in a recorded V1 neuron will be accompanied by some number of near synchronous spikes in other, unobserved V1 projection neurons. As a result, the timescale of V1-V1 correlations will affect those calculated for V1-V2 pairs. This blurs the shape of the V1-V2 CCG.

We conducted simple simulations to confirm this basic intuition. We defined a target cell, which fired whenever it received drive from any cell in the input population (Figure S3C). The input population consisted of 100 cells, and different proportions of this population were made to be correlated. When the proportion of correlated cells in the input population was low (10%), the average CCG between the target cell and the input population had a sharp peak clearly offset from zero (Figure S3E). As the

proportion of correlated cells in the input population grew, the input-target cell CCGs included a broad peak and negative side lobes (e.g. 50%; Figure S3E), features present in the input population CCGs (Figure S3D).

These simulations show that V1-V2 CCGs are likely to reflect features of V1-V1 correlations, in addition to the delay between V1 and V2 spiking.

Latency offset: We note that we found few cases in which V1-V2 CCG peaks occurred at negative latencies (i.e. V2 firing before V1; Figure 5B), as one might expect for feedback connections from V2 to V1. Inspection of the few CCGs with negative latencies that exceeded our statistical criterion revealed little convincing structure. A portion of our recordings were performed in the deep and superficial layers of V2, where feedback to the superficial layers of V1 arises (Felleman and Van Essen, 1991). Thus, the paucity in our sample is not due solely to the laminar location of our recordings. A more likely explanation is that feedback synapses are weak and modulatory (Sherman and Guillery, 1998), and thus do not generate the type of measurable functional interactions evident in the feedforward direction.

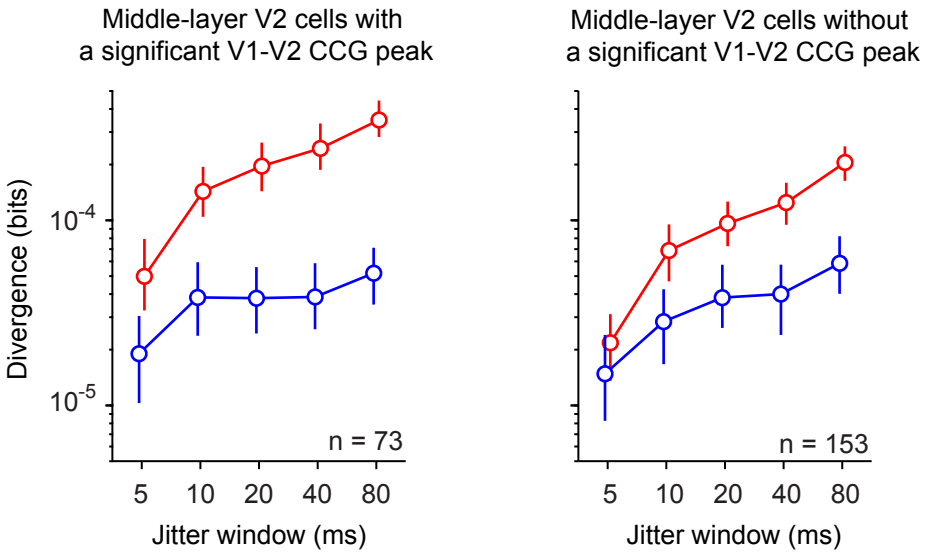

Supplemental Figure 4: Comparing divergence analysis for middle-layer V2 neurons with and without a significant V1-V2 CCG. Related to Figures 6, 7, and 8.

The analysis of Figures 6 and 7A,D was performed for all V2 neurons recorded at sites inferred to be in layer 4 (on the basis of a significant site-average V1-V2 CCG peak). The selection of inferred middle-layer sites was done separately for each electrode, to account for the possibility that the electrodes were not precisely aligned in depth. We adopted this approach because we reasoned that if the electrode was in layer 4, we could not exclude individual cells at that site just because we did not happen to record a V1 partner, or because the CCG between the V1 population and any individual V2 neuron at that site was not significant.

To test whether the results of our divergence analysis were different when based on V2 cells with a V1 partner, compared to middle-layer V2 neurons without a V1 partner, we reanalyzed our middle-layer recordings. We split the data into two subsets: one consisting of middle-layer V2 neurons with a V1 partner, and the other of middle-layer V2 neurons that did not have a recorded V1 partner.

Figure S4 shows that coordination in V1 was elevated in epochs preceding either subset of V2 cells, evident as higher divergence values associated with V2 spiking than in the corresponding control epochs. The increase in divergence was slightly stronger, when based on V2 neurons with a V1 partner (left compared to right).

These results are consistent with the intuition provided by the model simulations of Figure 8 (specifically, Figure 8G-I)): elevated divergence does not require that the downstream neuron has a recorded partner within the source network. Rather, the recorded neurons in the source network must be correlated with the projecting population.

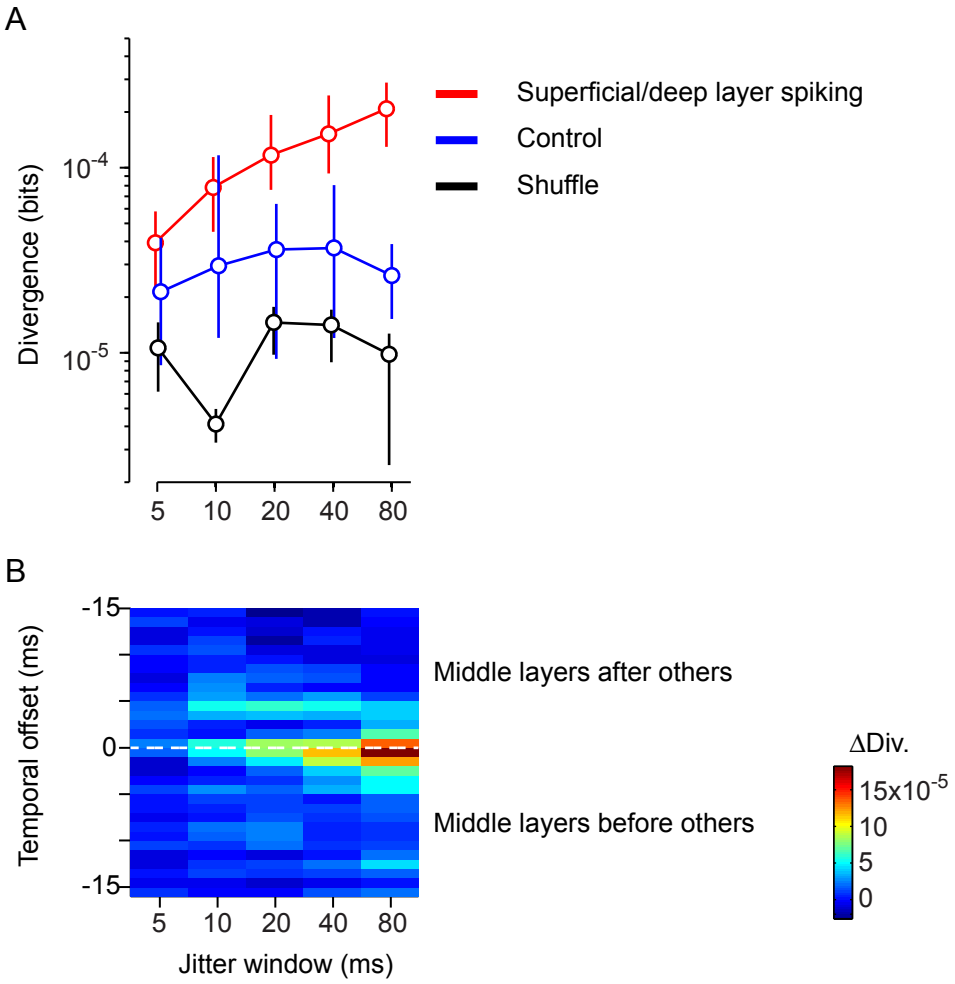

Supplemental Figure 5: Divergence in the middle-layers of V2 neurons, conditioned on spiking in other V2 layers. Related to Figures 1, 6, and 7.

Our analysis focused on understanding the role of coordinated network activity in relaying signals to downstream networks, through corticocortical connections. We found that the spiking of retinotopically-aligned neurons in the middle layers of V2 was preceded by elevated V1 network coordination, but this was not the case for spiking in other layers of V2 or at V2 sites with offset RFs. Here we ask whether spiking in the superficial and deep layers of V2 is related to the coordination of spiking activity within the V2 middle layers. We have shown in previous work that there is robust pairwise correlation of spiking activity in the middle layers of V2 (Smith et al., 2013).

Figure S5A shows the divergence for V2 middle-layer populations (11.5±0.6 cells on average), in epochs 0.5±1 ms before spiking in the superficial and deep layers. We combined epochs conditioned on superficial and deep layers because it provided greater statistical power. Middle-layer divergence was stronger in epochs preceding spiking in superficial and deep layers, than in control epochs. The elevation in divergence was significant for all jitter windows greater than 10 ms ($p=0.02$ for jitter window of 20 ms; $p=0.009$ for 40 ms; $p<0.001$ for 80 ms). Figure S5B shows the difference in divergence (epochs preceding 'downstream' spiking compared to control) as a function of jitter window (abscissa) and temporal offset (ordinate). The difference in divergence is largest in the bin just preceding downstream spiking, but divergence is also elevated after downstream spiking ($p=0.03$ for 5 ms jitter

window; $p=0.007$ for 10 ms; $p<0.001$ for all other windows). Thus, at least some of the elevated divergence likely reflects near-synchronous activity across layers, with the smaller asymmetry of the effect perhaps indicating some role for middle-layer coordination in relaying signals through intra-areal circuitry.

We note that the divergence values in this analysis cannot be compared to those in the main text. There are several notable differences between the two data sets. First, while the V1-V2 analysis involved computing divergence for a large V1 population, the analysis here involved much smaller V2 populations. This is because our V2 recordings used a smaller numbers of electrodes, and this smaller number was further divided between those placed in the middle layers and those in the superficial/deep layers. Second, the smaller number of cases—not all superficial/deep layer recordings were paired with middle layer recordings—resulted in noisier distributions of network states. This is evident in the greater divergence for shuffled data in this analysis (black line in Figure S5A; compared to shuffled data of Figure 7A of the main text). As discussed in the main text, the divergence for shuffled data reflects the noisiness of the recorded distributions, since with infinite data there should be no difference between shuffled data and the corresponding jittered surrogate data (i.e. there is no coordination in either case). Finally, the recorded V2 neurons tended to be closer together since they often involved a few nearby tetrodes, as opposed to the larger V1 array. As a result the correlations between neurons could be higher than those observed in V1 on average.

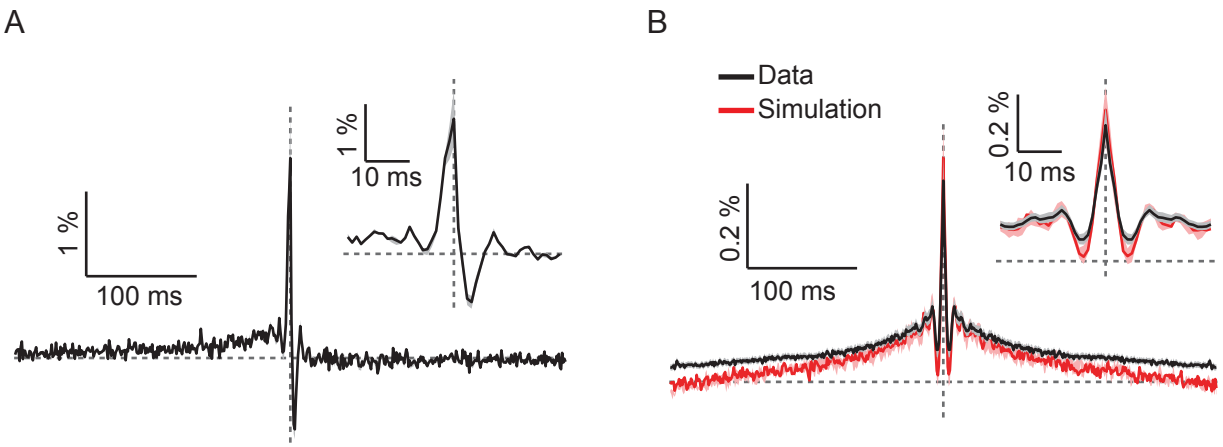

Supplemental Figure 6: Generating synthetic population responses with physiologically-realistic correlation timescales. Related to Figure 2.

We simulated V1 population responses using an algorithm related to that developed by Macke et al. (2009), as explained in the main text. Here we describe the modifications we made to that algorithm, to allow us to simulate populations with correlations arising over a range of time scales, as typically observed in cortex (Kohn and Smith, 2005; Smith and Kohn, 2008).

The Macke algorithm generates the spike train of each unit by drawing samples from an n-dimensional (where n is the number of cells) Gaussian distribution U and then thresholding:

$$X_i = 1 \text{ iff } U_i > 0 \text{ where } U \sim N(\gamma, \Lambda), \quad \text{(S1)}$$

where γ is the mean and Λ is the covariance of the Gaussian distribution. The resultant spike trains are constrained by user-defined first- (rate) and second-order (correlations) statistics, but otherwise have nearly maximal entropy.

The algorithm in its original form does not account for correlations occurring across time. To introduce correlations involving non-synchronous spikes, we used a temporally-correlated Gaussian time series. Generating time series from a multidimensional temporally-correlated distribution is computationally complex. To make the problem tractable, we fixed the rate and correlations among all simulated cells (i.e. we considered homogeneous populations). A population consisting of n cells responding over t time bins was simulated by combining random variables from two uncorrelated Gaussian distributions: an independent variable unique for each of the n cells and a common variable which was shared among all of them. The temporal correlation was implemented by convolving the common variable with a temporal filter (shown in Figure S6A, and at high temporal resolution in the inset). The temporal filter was selected so that its autocorrelation function would be equal to the average V1-V1 CCG. According to the Wiener–Khinchin theorem, the Fourier transform of the autocorrelation of a random process is the power density of that process. We thus estimated the temporal filter by the square root of Fourier transform of the average V1-V1 CCG.

The outcome of this process is a temporally-correlated multidimensional Gaussian time series. The strength of correlation is controlled by the variances of the independent and common variables. We

chose these to match the correlations and firing rate of the recorded data. Spike trains were then generated by thresholding the Gaussian time series (similar to equation S1). We confirmed that this produced CCGs with the appropriate shape (Figure S6B) for time lags of roughly ±20 ms, precisely the time scale on which our analysis is focused.

Note that these simulations were only used to provide intuition for our analysis, and to aid in its interpretation. The simulated data were not used in the analysis itself.

Supplemental Experimental Procedures: Equations and parameters for the simulations.

Figure 1: Each neuron in the network was modeled as a leaky, current-based, integrate-and-fire neuron.

$$C\frac{dV}{dt} \quad g_{passive}(V \quad V_{rest}) \quad g_{Na}(V \quad V_{Na}) \quad g_{K\ fast}(V \quad V_K) \quad g_{K\ slow}(V \quad V_K) \quad I_{syn}(t)$$

The $g_{passive}$ represents the passive membrane conductance and is equal to C/τ with τ representing the membrane time constant. We used C = 250 pF, τ = 10 ms, V_{rest} = -70 mV, V_{Na}=45, and V_K=-75. We solved the differential equation numerically, using the Runge-Kutta method (0.1 ms step size).

The conductances g_{Na}, $g_{K\text{-}fast}$ and $g_{K\text{-}slow}$ were triggered only when the membrane potential reached the threshold of -55 mV. They represent the active membrane dynamics responsible for generation of spikes and post-spike refractory period. Each of these conductances were calculated using:

$$g(t) \quad g_0(e^{\ t/\ 1} \quad e^{\ t/\ 2})$$

where t represents time after reaching threshold, g_0 was 13 μS for Na, 5.2 μS for K-Fast and 0.21 μS for K-Slow. τ_1 was 0.3 ms for Na, 3 ms for K-fast and 20 ms for K-slow. τ_2 was 0.1 ms for Na, 1 ms for K-fast and K-slow. Action potentials were followed by a resetting of the membrane potential at V_{rest} for a 1 ms absolute refractory period.

Each presynaptic spike generated a postsynaptic current, I_{syn}, which was modeled as an alpha function with a peak of 45.6 pA and time constant of 0.32 ms. Inhibitory inputs produced currents identical to the excitatory currents, but with negative sign. All of these parameters were identical to those of Diesmann et al. (1999), as further detailed in Gewaltig (2000).

Output spike times were perturbed by a random delay of 0.25-2.25 ms, as in Troyer et al. (1998). This perturbation does not affect the network's ability to propagate spike packets, but counteracts its tendency to cluster spikes in increasingly temporally precise synchronous events.

Each layer of the network contained 100 neurons. The neurons in layers 2-10 received excitatory input from cells in the preceding layer (in addition to Poisson inputs), whose number we varied to implement different convergence regimes (40, 70, 100, with only the latter used in Diesmann et al. (1999)). Coordinated epochs were modeled by near synchronous spiking (standard deviation 2.5 ms, so the packet width—3 SDs—was ~15 ms) in all the cells in the first layer. These epochs occurred randomly at a mean rate of 2 Hz.

Figure 8: In Figure 8 we modeled a downstream V2 neuron as a leaky, integrate-and-fire cell (identical to the simulated neurons of Figure 1; described above). The cell received input from 1,500 excitatory and 1,140 inhibitory cells, each firing at 20 sp/s. The majority of these inputs were Poisson, but a subset was chosen from a pool of 500 correlated "V1" neurons (mean pairwise correlations of 0.1 and rate of 20 sp/s). Each of the scenarios involved a distinct subset of the correlated V1 population. In Figure 8A, the V1 inputs consisted of 75 excitatory cells. We calculated divergence values based on sampling the activity of 100 cells from the V1 population, including 10 of the 75 cells that projected to V2; the V1-V2 CCG was the average calculated from 10 V1 neurons (of the 100 sampled) and the V2 target. In Figure 8B, the V2 cell received excitation from 75 V1 neurons, and inhibition from 150 cells (chosen from the original pool of 500, and altering the sign of their input). Divergence values and CCGs were again calculated based on a sample of 100 V1 cells (none inhibitory). The simulations of Figure 8C were

identical to those of 8A except the V2 cell received excitatory input from 150 V1 cells. The sampled V1 population did not include any cells providing input to the V2 neuron. For all cases, we simulated 400 trials of activity, each lasting 1280 ms. This was repeated for 100 V2 cells for each scenario.

Supplemental References

Anderson JC, Martin KA (2009) The synaptic connections between cortical areas V1 and V2 in macaque monkey. J Neurosci 29:11283–11293.

Freund TF, Martin KA, Soltesz I, Somogyi P, Whitteridge D (1989) Arborisation pattern and postsynaptic targets of physiologically identified thalamocortical afferents in striate cortex of the macaque monkey. J Comp Neurol 289: 315–336.

Gewaltig M-O (2000) Evolution of synchronous spike volleys in cortical networks: network simulations and continuous probabilistic models. Shaker Verlag.

Gil Z, Connors BW, Amitai Y (1999) Efficacy of thalamocortical and intracortical synaptic connections: quanta, innervation, and reliability. Neuron 23: 385-397.

Kirkwood PA (1979) On the use and interpretation of cross-correlations measurements in the mammalian central nervous system. J Neurosci Methods. 1: 107-132.

Kohn A, Smith MA (2005) Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. J Neurosci. 25: 3661-3673.

Levick WR, Cleland BG, Dubin MW (1972) Lateral geniculate neurons of cat: retinal inputs and physiology. Invest Ophthalmol 5:302–311.

Stratford KJ, Tarczy-Hornoch K, Martin KA, Bannister NJ, Jack JJ (1996) Excitatory synaptic inputs to spiny stellate cells in cat visual cortex. Nature 382: 258-261.

Marion R, Li K, Purushothaman G, Jiang Y, Casagrande VA (2013) Morphological and neurochemical comparisons between pulvinar and V1 projections to V2. J Comp Neurol. 521: 813-832.

Ohiorhenuan IE, Mechler F, Purpura KP, Schmid AM, Hu Q, Victor JD (2010) Sparse coding and high-order correlations in fine-scale cortical networks. Nature 466: 617-621.

Sherman SM, Guillery RW (1998) On the actions that one nerve cell can have on another: distinguishing "drivers" from "modulators". Proc Natl Acad Sci USA 95: 7121-7126.

Stratford KJ, Tarczy-Hornoch K, Martin KA, Bannister NJ, Jack JJ (1996) Excitatory synaptic inputs to spiny stellate cells in cat visual cortex. Nature 382: 258-261.

Troyer TW, Krukowski AE, Priebe NJ, Miller KD (1998) Contrast-invariant orientation tuning in cat visual cortex: thalamocortical input tuning and correlation-based intracortical connectivity. J Neurosci. 18: 5908-5927.

Wissig SC, Kohn A (2012) The influence of surround suppression on adaptation effects in primary visual cortex. J Neurophysiol. 107: 3370-3384.