

Stochastic transitions into silence cause noise correlations in cortical circuits

Gabriela Mocho^{a,b}, Ainhoa Hermoso-Mendizabal^b, Shuzo Sakata^{c,d}, Kenneth D. Harris^{e,f}, and Jaime de la Rocha^{b,1}

^aNencki Institute of Experimental Biology, 02-093 Warsaw, Poland; ^bInstitut d'Investigacions Biomèdiques August Pi i Sunyer (IDIBAPS), 08036 Barcelona, Spain; ^cCentre for Neuroscience and ^dStrathclyde Institute of Pharmacy and Biomedical Sciences, University of Strathclyde, Glasgow G4 0RE, United Kingdom; and ^eUCL Institute of Neurology and ^fUCL Department of Neuroscience, Physiology, and Pharmacology, University College London, London WC1E 6DE, United Kingdom

Edited by Tony Movshon, New York University, New York, NY, and approved January 26, 2015 (received for review June 10, 2014)

The spiking activity of cortical neurons is highly variable. This variability is generally correlated among nearby neurons, an effect commonly interpreted to reflect the coactivation of neurons due to anatomically shared inputs. Recent findings, however, indicate that correlations can be dynamically modulated, suggesting that the underlying mechanisms are not well understood. Here, we investigate the hypothesis that correlations are dominated by neuronal coinactivation: the occurrence of brief silent periods during which all neurons in the local network stop firing. We recorded spiking activity from large populations of neurons in the auditory cortex of anesthetized rats across different brain states. During spontaneous activity, the reduction of correlation accompanying brain state desynchronization was largely explained by a decrease in the density of the silent periods. The presentation of a stimulus caused an initial drop of correlations followed by a rebound, a time course that was mimicked by the instantaneous silence density. We built a rate network model with fluctuation-driven transitions between a silent and an active attractor and assumed that neurons fired Poisson spike trains with a rate following the model dynamics. Variations of the network external input altered the transition rate into the silent attractor and reproduced the relation between correlation and silence density found in the data, both in spontaneous and evoked conditions. This suggests that the observed changes in correlation, occurring gradually with brain state variations or abruptly with sensory stimulation, are due to changes in the likeliness of the microcircuit to transiently cease firing.

neuronal variability | noise correlations | brain state | auditory cortex | stochastic network dynamics

Neuronal noise correlations are defined as common fluctuations in the spiking activity of neurons under conditions of constant sensory input or motor output. Traditionally, they have been thought to arise from the dense connectivity of the cortex, such that neighboring neurons sharing a fraction of their inputs should also share a fraction of their output variability (1). Several observations are consistent with this hypothesis: pairwise correlations in the cortex decrease with cell pair distance (2) or with the difference in stimulus selectivity (3), dependencies that could follow from a variation in shared input given the anatomy of cortical circuits. Recent findings, however, challenge this simple interpretation. Recordings in the primate visual cortex have shown that attention or task context can change correlation structure (4–6) and that the magnitude of averaged correlation can be very low (7). In anesthetized rodents correlations decrease with brain state desynchronization (8, 9) or when animals switch from quiet wakefulness to active whisking during waking (10). Moreover, the commonly observed drop of spiking variability following stimulus onset (11–13) seems to occur jointly with a transient decrease in correlation (2, 14, 15). These observations suggest that correlations reflect the dynamical state of the circuit more than its hardwired connectivity.

Despite substantial progress in understanding the mechanisms giving rise to large individual variability in recurrent networks

(9, 16–18), we still lack a canonical model that can generate correlations with the same magnitude and spatiotemporal structure as those observed in cortical circuits. Balanced networks, for instance, a common model that reproduces the large variability of cortical neurons (9, 18, 19), show near-zero averaged correlations (9). Numerous studies have investigated the generation of synchronous firing (20), but whether short bursts of population activity can quantitatively account for the spike count correlations found in the data is unclear. Recurrent networks can also generate fast oscillations in the population activity, but, in a regime of low rates, typical of cortical circuits, average spike count correlations are negligible (21). Network models producing nonzero average correlations are those exhibiting up and down dynamics (22–29). Most of these studies have focused on investigating the mechanisms underlying the slow oscillatory activity observed in cortical slices (30), under anesthesia (31, 32), or during slow-wave sleep (33). Only recently the impact of up and down switching on trial-to-trial response variability (25) and on the probability distribution of multi-unit activity (29) across brain states has been investigated. Whether the alternation between up and down phases could quantitatively account for the pairwise correlations observed in different brain states and describe their stimulus-evoked dynamics remains an open question.

To investigate the mechanisms producing correlated firing, we recorded the spiking activity of large populations of neurons from the auditory cortex of anesthetized rats. During spontaneous activity, changes in correlation were largely explained by variation of the occurrence rate of periods during which neurons in the circuit stopped firing. Furthermore, the time course of correlation in response to an acoustic stimulus reflected the transient variation of this silence density. A computational rate model

Significance

Neurons in the cerebral cortex emit action potentials in a seemingly random manner. One puzzling aspect of this neuronal “noise” is that it is correlated among neighboring neurons, something thought to reflect the tendency of neurons to fire together. Here, we recorded the activity from populations of cortical neurons in rats and found that correlations could be largely explained by the tendency of cortical neurons to stop firing together. A computational network model whose activity alternated between periods of activity and silence was able to reproduce the pattern of correlations found in the experiments. Our findings shed light on the mechanisms causing neuronal variability and may contribute to elucidate its role in a neural code.

Author contributions: G.M., K.D.H., and J.R. designed research; G.M., A.H.-M., S.S., and J.R. performed research; G.M. analyzed data; and G.M. and J.R. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

¹To whom correspondence should be addressed. Email: jrochav@clinic.ub.es.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1410509112/-DCSupplemental.

with fluctuation-driven transitions between silent and active attractors could explain the experimentally observed time course of correlation and its relation to silence density. Our findings suggest that the dynamics of these transitions play a fundamental role in generating noise correlations among cortical neurons.

Results

We recorded spontaneous and stimulus-evoked population activity from the primary auditory cortex (A1) of urethane-anesthetized rats ($n = 6$) using multisite silicon microelectrodes. We sorted spikes off-line and obtained multiple spike trains from isolated single units (range, 44–147) as well as some spike trains from multiunit activity (range, 3–103). During the experiments, the brain activity underwent spontaneous transitions across a continuum of brain states varying between a synchronized state exhibiting alternations between active and silent periods (Fig. 1A, Right), called up and down phases, respectively (30–33), and a desynchronized state with no apparent up and down alternations (9, 25) (Fig. 1A, Left). To quantitatively characterize the full spectrum of brain states we used silence density (S) computed from the pooled population activity of merged single units and multiunits during spontaneous conditions. Spontaneous conditions referred to the activity during

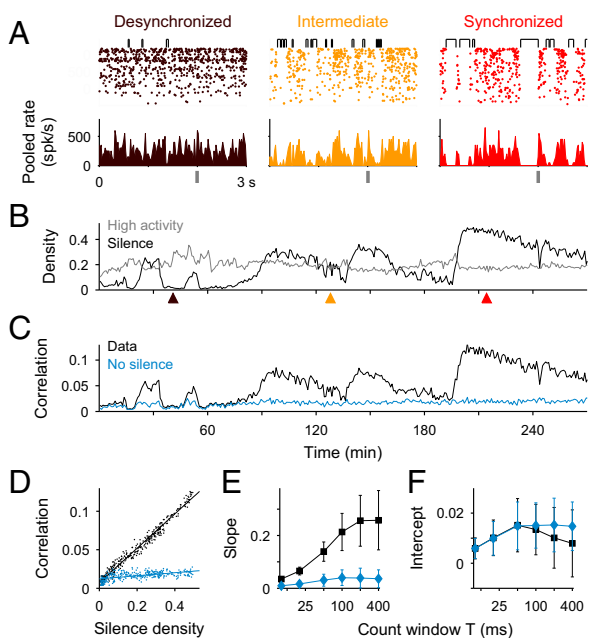


Fig. 1. The relation between correlations and silence density during spontaneous activity in A1. (A) Spike rastergrams (Top) and pooled population rate (Bottom; bin, 20 ms) from single and multiunit spike trains ($n = 84$) during epochs of desynchronized (brown), intermediate (orange), and synchronized (red) brain state. Top brackets indicate silent periods, i.e., consecutive 20-ms bins with no spikes. Vertical gray bar marks stimulus onset. (B and C) Silence density S (B, black) and averaged spontaneous spike count correlations ρ (C) obtained across 50-s epochs in one recording session. High activity density, i.e., fraction of bins with more than six spikes, is shown for comparison (B, gray; the threshold six was chosen to match the averages of the two densities). Correlations were averaged over pairs of recorded single units (black; $n = 3,240$ pairs; count window $T = 100$ ms) and surrogate data (blue) obtained by removing all silent periods. Arrowheads indicate fragments shown in A. (D) ρ vs. S for the two sets in C. Each dot represents a 50-s epoch. Linear fits from experimental and surrogate data (lines) have slopes 0.22 and 0.019, and intercepts 0.007 and 0.012, respectively. (E and F) Slope (E) and correlation intercept (F) from the linear fit of ρ vs. S as a function of window size T (mean \pm SD over $n = 6$ animals). Colors as in C and D.

1.5-s intervals preceding each stimulus. We divided each recording session into adjacent 50-s epochs, sufficiently short to capture fast brain state transitions and long enough to obtain good estimates of the spiking statistics. In each epoch, we computed S as the fraction of 20-ms bins during which the spontaneous population activity had zero spikes (Fig. 1A, brackets above rasters). The epochs with relatively high values of S were classified as the synchronized brain state (Fig. 1A, Right). During those epochs, silent periods, obtained by merging consecutive empty bins, captured mostly down periods. The epochs with S close to zero showing shorter and less frequent silent periods, were classified as the desynchronized brain state (Fig. 1A, Left). We chose 20-ms bins as the discretization yielding the maximal discriminability between brain states (Fig. S1): finer binning saturated the S to 1, whereas longer bins missed brief silences yielding near-zero S for all epochs. Moreover, we controlled that the variations of S across epochs were not explained by changes in the pooled population rate (Fig. S2 A and B; SI Methods) and that they mimicked the state changes derived independently from the spectral density of local field potential (LFP) (8) (Fig. S2C).

Relationship Between Spontaneous Correlations and Silence Density in A1.

We first asked whether variations in silence density could explain the state dependence of correlations between pairs of single units. For each epoch, we obtained the spontaneous spike count correlation coefficient ρ computed across time during spontaneous activity using a count window $T = 100$ ms and averaged across single-unit pairs (Fig. 1C). As expected, ρ was larger during synchronized state epochs where the population activity fluctuated between silent and active periods (9, 34). Interestingly, the relationship between correlation and S was linear and had a very small intercept ρ_0 (Fig. 1D). This behavior was robustly observed across experiments (Fig. S3) and over a broad range of count windows T (Fig. 1 E and F). Moreover, the density of high-activity periods, defined as bins with spike counts above a certain threshold, did not show the same covariation with ρ (Fig. 1B), showing that the density S was particularly predictive of correlation magnitude. We next asked whether correlations were solely due to the occurrence of silent periods among otherwise uncorrelated activity. To test this, we generated a surrogate dataset in which, for each epoch, silent periods were removed and the remaining periods with spikes were concatenated to form a continuous recording with $S = 0$ (SI Methods). Correlations in this dataset were weak in all epochs (Fig. 1 C–F), suggesting that the increase in ρ accompanying brain state synchronization in the original data was mediated by an increase in the silent-period probability and that the additional mechanism producing the offset correlation ρ_0 remained relatively constant.

The Dynamics of Correlation During Stimulus-Evoked Responses.

We next studied the dynamics of the population in response to short acoustic clicks (duration, 5 ms; interclick interval, 2.5 or 3.5 s). We used a sliding spike count window ($T = 50$ ms) and computed the averaged instantaneous rate, spike count correlation $\rho(t)$ (2, 14, 15), and spike count Fano factor (11–13) by performing the statistics across repeated stimulus presentations and averaged over single units or single-unit pairs (Methods). Similarly, we computed the instantaneous silence density $S(t)$ using 20-ms bins. To reveal the impact of brain state, we grouped trials depending on whether they occurred in epochs of low ($S < 0.05$), intermediate ($0.05 \leq S \leq 0.2$), or high ($0.2 < S$) brain state synchronization (Figs. 1A and 2). Despite a difference in spontaneous baseline rate across states, the averaged peak response was roughly independent of state (Fig. 2A). The dynamics of $S(t)$ showed a fast drop to a near-zero value in all states followed by a fast rebound whose magnitude depended on the cortical state. Except in the most synchronized epochs, the rebound showed an overshoot above the baseline $S(t)$, revealing that, after the initial

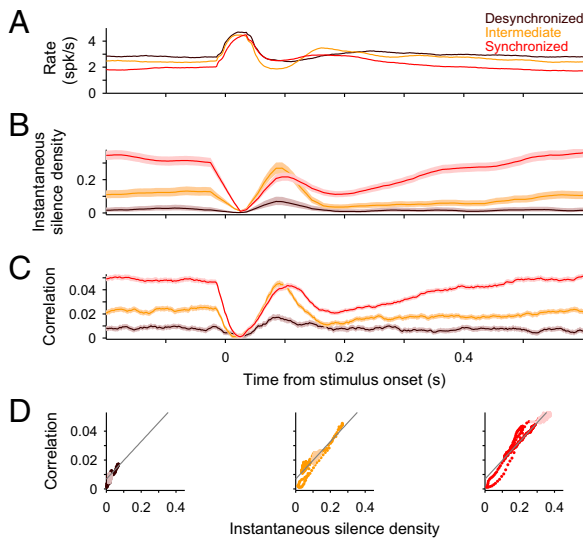


Fig. 2. Evoked dynamics of correlations in A1 across brain states. (A–C) Mean population-averaged rate $r(t)$ (A; $n = 81$ single units), instantaneous silence density $S(t)$ (B), average instantaneous correlation $\rho(t)$ (C; $n = 3,240$ pairs) in response to a click stimulus during the desynchronized (brown), intermediate (orange), and synchronized (red) epochs observed in one experiment (same as in Fig. 1 A–D). Statistics were obtained across stimulus repetitions within the given brain state. Rate and correlation used $T = 50$ -ms sliding windows (time step, 2 ms). (D) Instantaneous correlation $\rho(t)$ vs. $S(t)$ for each brain state in B and C. Darker dots correspond to the interval (–25, 515) ms, with zero being the stimulus onset. Gray lines show the linear fit to the spontaneous ρ vs. S relation (Fig. 1D). Shaded areas in C illustrate 95% confidence bands (bootstrap).

increase in population rate, the stimulus tended to increase the probability of the circuit to go into silence (Fig. 2B and Fig. S4B). After the rebound, the more synchronized epochs showed a second drop followed by a slow recovery to baseline (Fig. 2B and Fig. S4B). Interestingly, instantaneous average correlation exhibited almost identical time course as $S(t)$ in each of the different brain states (Fig. 2C and Fig. S4C). Thus, the linear relation between correlation and S found during spontaneous activity, was approximately maintained instantaneously during evoked responses (Fig. 2D). The instantaneous Fano factor followed qualitatively similar dynamics and brain state dependence to correlation except that the range of spanned values was narrower (Fig. S5).

Computational Rate Model Reproduces ρ – S Relation Across Brain State Changes. To understand the mechanisms underlying correlations and their relation with silence density, we analyzed a model with two sources of neuronal variability: the first reflecting variations in the firing rate r , and the second reflecting the spiking stochasticity existent at constant rate (1, 11, 35–37). Under this assumption, spike count correlations could be explained, at least in part, if the rate variability was correlated across neurons (35, 37). For example, it has been shown that individual spike count statistics can be described by a Poisson process in which the rate varies from trial to trial following a broad unimodal distribution (37). We tested this model in spontaneous conditions across different brain states and compared it with a model in which the rate followed a bimodal distribution reflecting network transitions between a silent ($r = 0$) and an active state ($r > 0$). We found that the statistics of population spike counts and particularly pairwise covariance, were better described by the bimodal model, especially during synchronized epochs when ρ and S were largest (Fig. S6). This suggests that silent periods reflected separable events in the circuit dynamics rather than periods of very low rate drawn from a unimodal distribution.

Assuming that rate variations are mainly produced by the network transitions into silence, we initially considered that neurons fired conditionally independent Poisson spike trains with the same varying rate $r(t)$. Under these assumptions, the spike count correlation coefficient $\rho(t)$ reads as follows (SI Methods):

$$\rho(t) = \frac{\text{Var}[R(t; T)]}{\text{Var}[R(t; T)] + \langle R(t; T) \rangle}, \quad [1]$$

where $R(t; T) = \int r(t') dt'$, named integrated rate, is the expected number of spikes elicited by each neuron in the interval $(t - T/2, t + T/2)$,




Fig. 3. Bistable rate model with adaptation captures the correlation vs. silence density relationship. (A) Network model composed of a self-coupled excitatory population with rate $r(t)$, exhibiting rate adaptation $a(t)$ of strength β and receiving a fluctuating external input $l(t)$. (B) Phase plane showing rate (cubic curves) and adaptation (straight line) nullclines and stable fixed points (filled circles). Increasing l shifts the rate nullcline (arrow), whereas increasing β decreases the slope of adaptation nullcline. (C) Stability analysis in the l – β plane shows four regimes: limit cycle (light gray area), single active attractor (gray area), single silent attractor (black area), and bistable regime with silent and active attractors (dark gray area). Example nullclines are superimposed on each regime (white curves). (D) Baseline mean rate r (Left), silence density S (Middle), and correlation coefficient ρ (Right; count window $T = 100$ ms; Eq. 1) as a function of l and β (black lines mark borders of bistable regime). Increasing l at fixed β causes an increase in r and a decrease in S and ρ , mimicking the transition toward the desynchronized state. (E) ρ vs. S relationship obtained along the desynchronization axis $\beta = 1$ and l varying from 0.4 to 4 for conditionally independent spiking ($c_0 = 0$; gray line) and nonzero spiking covariability ($c_0 = 0.01$; black line). Shaded area shows the ρ – S pairs obtained in the model when sampling the region of the (l, β) plane shown in C. (F) Network spontaneous activity for three (β, l) pairs (squares in D and E) chosen to reproduce the different brain state categories (Fig. 2): desynchronized (Top), intermediate (Middle), and synchronized (Bottom). For each state, phase plane (Left; as in B), example rate trace (Middle), and rate histogram (Right) are shown. Gray area in the phase planes shows the basin of attraction of active fixed point. Dotted lines indicate silence detection threshold. Values for each brain state were $\beta = 1$ s and $l = 1.1, 1.6,$ and 2 a.u.

$t + T/2$) given $r(t)$. The mean (angle brackets) and variance of $R(t;T)$ in Eq. 1 are taken with respect to the rate variability. In the absence of transitions, $r(t)$ is approximately constant and both $\text{Var}[R(t;T)]$ and $\rho(t)$ are zero. This is consistent with the spontaneous activity data if the small offset correlation ρ_0 at $S = 0$ is neglected (Fig. 1D and Fig. S3). To describe the dynamics of network transitions, we used a 2D dynamical model (22, 25) that allowed a mechanistic interpretation of the rate variability during spontaneous and evoked conditions (Fig. 3A). The model can be considered a mean-field description of a population of recurrently coupled excitatory neurons with mean rate $r(t)$ that receive an external input $I(t)$ and exhibit rate adaptation $a(t)$ (Methods). Because of the concavity of the transfer function and the recurrent coupling, the network can exhibit bistability with a silent attractor and an active attractor with low rate (22) (Fig. 3B). To determine whether the model could reproduce the state dependence of spontaneous correlations, we simulated its dynamics using a fluctuating external input and computed the mean rate, the silence density [defined as the fraction of time with $r(t)$ below a silence threshold] and the correlation (Eq. 1; $T = 100$ ms) for a range of values of the baseline input I and the adaptation strength β (Fig. 3D). The statistics were performed across time. We chose an axis in the bistable region of the (I, β) plane (squares in Fig. 3D) where the model reproduced the approximately linear ρ - S relation found in the data (compare Figs. 3E and 1D). Beyond this region, the model yielded much higher S values revealing that, as S approaches 1 (complete silence), correlation tends to zero (Fig. 3E, gray area). Thus, the ρ - S relation is generally nonmonotonic, but it can be approximately linear for the range of S values found in the data (i.e., $S < 0.5$). Finally, we relaxed the condition that spiking was conditionally independent across neurons and explored how the correlation changed when a small constant spiking covariability term was introduced (SI Methods). This caused a shift in the ρ - S relationship (Fig. 3E) that mimicked the offset ρ_0 observed in the experimental data (Fig. 1D and Fig. S3).

Although in the model both active and silent attractors were stable across brain states, the transition rate varied due to changes in the effective size of their basins of attraction. In the desynchronized state, the active basin of attraction was effectively large because the active fixed point was far from the basin's border, whereas the silent basin of attraction was small (Fig. 3F, Top). Therefore, the system remained active most of the time and the external fluctuations triggered sporadic and short transitions into silence (Fig. 3F, Top), yielding low S and small ρ (Fig. 3D). Decreasing I took the system toward the synchronized state (Fig. 3D, squares) where the active attractor was shifted closer to the basin's border, whereas the silent attractor was moved further (Fig. 3F, Bottom). Thus, excursions to the silent branch were more frequent and lasted longer, yielding high S and larger ρ (Fig. 3D and F, Bottom). Because transitions were triggered by fluctuations, the duration of silent and active periods was very irregular and the temporal structure of the rate did not show an oscillatory behavior. This nonrhythmic pattern during synchronized epochs was also observed in the data.

As an alternative to the bistable model, we considered a monostable dynamical model that produced Gaussian-like fluctuations in $r(t)$. Silent periods were not caused by network transitions but simply reflected large downward deflections of $r(t)$ that were mirrored by upward deflections of a similar magnitude (Fig. S7E). In this unimodal model, states of high correlation are associated with an increased population rate variance, yielding larger density of both silent and high-activity events. In contrast, high activity density in spontaneous data did not systematically increase with ρ , a feature better captured by the bistable model (Fig. 1B and Fig. S7).

Rate Model Reproduces the Evoked Dynamics of Correlation. The network model also reproduced the state-dependent dynamics of correlation during click-evoked responses. A brief current step

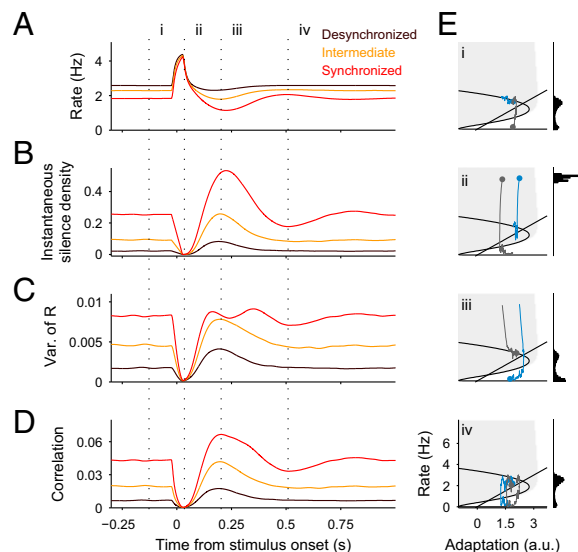


Fig. 4. Rate model stimulus-evoked dynamics across brain states. (A–D) Stimulus-evoked mean instantaneous rate $r(t)$ (A), instantaneous silence density $S(t)$ (B), variance of integrated rate $R(t)$ (C), and instantaneous correlation $\rho(t)$ (D) for the three brain states defined in Fig. 3 (same color code). Statistics were obtained across repeated presentations of the stimulus (square pulse of 10 ms). (E) Phase plane trajectories obtained from single-trial evoked responses in the synchronized state. Each phase plane (convention as in Fig. 3F, Left) shows a snapshot of two example trajectories (blue and gray traces) taken at successive times i , ii , iii , and iv (dotted lines in A–D). Histograms obtained over multiple trials (Right) correspond to $r(t)$ at the current time (dots). Independently of the network state at stimulus onset, the stimulus reliably elicits a stereotyped increase of rate that quenches the rate variability and correlations (times i and ii). Due to increased adaptation following stimulation, the system moves closer to the border of the basin of attraction of the active equilibrium point (shaded area). This increases the probability of falling into the silent branch with respect to prestimulus baseline (blue trace, time iii). In the synchronized state, the network shows a second transient decrease in $S(t)$ before returning to equilibrium (time iv).

was presented for each of the three (I, β) combinations representing different brain states (Fig. 3F), and the mean instantaneous rate $r(t)$, integrated rate variance, $S(t)$, and $\rho(t)$ were computed across repeated trials (Fig. 4). Because the stimulus produced a stereotyped response independently of whether the network was in the silent or active branch (Fig. 4E, ii), $S(t)$ dropped to zero and the mean peak rate was the same for all brain states. This caused that the rate variance, and in turn $\rho(t)$, also dropped to near-zero values (Fig. 4C and D, ii , and Movies S1 and S2). After the peak response, the rate was suppressed below baseline due to a rebound of the silence density (Fig. 4A and B, iii). This was a consequence of the response-evoked increase in adaptation: adapted trajectories run closer to the border of the active basin of attraction, which increased their probability to fall into the silent branch (Fig. 4E, iii , blue traces). The rebound in $S(t)$ led to an increase in rate variance, which, combined with a reduced $r(t)$, produced a prominent rebound in correlation (Fig. 4A–D, iii , and Movies S1 and S2). In the synchronized state, the rebound in $\rho(t)$ was followed by a second smaller drop (Fig. 4D, iii and iv , red trace), reflecting that the system exhibited a weak oscillatory behavior in relaxing back to equilibrium. The model predicted that $\rho(t)$ in response to stimuli recruiting more adaptation should show a larger rebound and a more pronounced second drop (Fig. S8).

Last, we searched for the minimal model reproducing the key aspects of the correlation evoked dynamics. A simplified bistable model with no adaptation was sufficient to reproduce the strong drop in $\rho(t)$ (Fig. S9). An additional mechanism that transiently

increased the poststimulus silence probability (e.g., short-term depression of the feedforward synaptic afferents) was necessary to generate the rebound (Fig. S9). Negative-feedback mechanisms, such as rate adaptation or synaptic short-term depression in the recurrent connections (20), could generate the rebound and, in addition, cause the second small drop in $\rho(t)$ observed in the synchronized state. The monostable Gaussian model behaved qualitatively different from the bistable model (Fig. S9), suggesting that the nonlinearity underlying the transitions into silence was necessary to reproduce the dynamics of $\rho(t)$.

Discussion

We have shown that average spike count correlations among neighboring neurons can be largely explained by transient excursions of the cortical circuit into silence. A rate network model with adaptation showing stochastic transitions between two attractors reproduced the ρ - S relation observed during spontaneous and evoked conditions. Two recent studies have proposed that fluctuations in neuronal excitability cause noise correlations in monkey visual cortex (34, 37). Whereas in one study correlated fluctuations of excitability only accounted for a fraction of the total correlation (37), in the other fluctuations “resembling up and down states” explained almost all measured correlations (34). We extended these findings by showing that, during cortical synchronization, a bimodal distribution of the excitability captures the statistics of population spike counts in the rat auditory cortex more accurately than the previously proposed unimodal model (37) (Fig. S6). We showed that stochastic transitions into silence seem to be the mechanism generating bimodal fluctuations in excitability. Modeling these transitions allowed us to understand the dynamics of correlation caused by sensory stimulation. We showed that correlations measured for spontaneous activity across brain states exhibited a tight relationship with silence density (Fig. 1D) and that, unexpectedly, this relationship was preserved instantaneously during evoked responses (Fig. 2D). The correlation ρ_0 found during spontaneous activity in the absence of silent periods (Fig. 1C and D) could be due to a number of factors (e.g., very slow rate fluctuations) but does not seem to reflect the irreducible effect of the connectivity as correlation dropped to near-zero values during evoked responses (Fig. 2C and Fig. S4C).

We hypothesized that cortical networks can be described as a bistable system with an active and a silent attractor. During active periods, the network produces Poisson-like, low-rate, asynchronous activity due to balance between excitation and inhibition (9, 38). Fluctuations arising from external inputs or from the internal spiking activity during active periods would generate transitions between the two attractors. It is not clear, however, how to build such bistable network. Balanced networks at low rates show a linear input-output relation (38), whereas bistability requires nonlinearities (22, 24, 27, 39). Extensive theoretical work has aimed to reconcile multistability and irregular firing, mainly in the context of persistent activity circuits (40, 41). We simplified the problem and, assuming a balanced state during active periods, built a rate network model to investigate the transition dynamics between the two attractors. In contrast to previous network models of up/down transitions that, due to a strong fatigue mechanism (e.g., spike frequency adaptation), operate in the oscillatory regime (Fig. 3C, *Top Right*) (24, 28, 39), our model displays weak adaptation and relies on fluctuations to escape from the otherwise-stable attractors (26, 27). Moreover, by increasing the external input I , plus optionally decreasing the adaptation strength β , the network becomes desynchronized (Fig. 3D-F), as shown in previous models (23, 25, 39, 42). Desynchronization is also thought to decrease the synaptic efficacy of intracortical connections. In our model, this would decrease the curvature of the rate nullcline and remove the silent branch, effectively linearizing the system (25). Our findings suggest that desynchronization implies moving away from the bifurcation, which increases the region around the fixed point where

the system behaves linearly (Fig. 4). However, they also show that certain features, such as correlation rebound, cannot be captured by a linear system and require the existence of a silent branch (Fig. S9).

Several studies have recently modeled (13, 16–18) the dynamics of stimulus-evoked neuronal variability (11). The proposed mechanisms range from suppression of an ongoing chaotic state (16), to the quenching of variability produced by the transitions of the network among multiple discrete states (17, 18) or along a line attractor (13). Our model shares the idea that the stimulus puts the system away from the bifurcation where network state transitions occur. There are, however, two main differences. First, except in ref. 13, the average correlation in these networks was negligible. This is because either neuronal variability across the network was independent (16) or because the transitions between network discrete states involved the coordination of small subpopulations of cells (17, 18). Second, the nature of multistability in these models followed from a specific connectivity structure such as clustered connections (17, 18) or local connectivity following a ring topology (13). Our bistable network does not rely on a particular structured connectivity, which could explain why the effect is ubiquitously observed across cortical areas with different connectivity schemes (11).

Although our experiments were performed under urethane anesthesia, silent periods also occur during waking. Global dynamics resembling up and down switching have been observed in rodents during quiescent wakefulness (43, 44) or during a perceptual task (45) as well as in awake primates (46). Previous studies have hypothesized that correlations could impact the encoding of information in large networks. Our results contribute to build a mechanistic framework for recent findings showing that, depending on variables such as sleep pressure, task engagement, locomotion, or sensory stimulation, circuits exhibit different dynamics that shape the structure of correlations. Whether these correlations ultimately impact information encoding will depend on how efficiently animals process sensory information under these different brain states (45, 47), a question that needs to be further investigated.

Methods

Experimental Techniques. All experiments were carried out in accordance with protocols approved by the Animal Ethics Committee of the University of Barcelona and by the Rutgers University Animal Care and Use Committee. Six rats (Sprague-Dawley; 250–400 g) were anesthetized with urethane (1.5 g/kg body weight) and silicon microelectrodes (Neuronexus) with 32 or 64 channels were inserted in deep layers (depth, 600–1,200 μm) of the primary auditory cortex (A1). We simultaneously recorded the spiking activity from many single units and multiunits (means, 86 and 45, respectively) and the LFP in response to acoustic “clicks” (5-ms square pulses; interstimulus interval, 2.5 or 3.5 s). Details on the techniques and spikes sorting procedures are described in *SI Methods*.

Data Analysis. Long continuous recordings (mean, ~ 2 h) were divided into 50-s epochs, and brain state was estimated in each epoch based on spontaneous pooled population activity, i.e., the merge of single and multiunit spike trains during the 1.5-s intervals preceding each stimulus presentation. Brain state was quantified using silence density (S) defined as the fraction of 20-ms time bins with no population activity (i.e., zero spikes; Fig. 1B and D-F, black). Silent and active periods were obtained from the merge of consecutive empty and nonempty bins, respectively. High activity density (Fig. 1B, gray) was computed, similarly to S , as the fraction of time bins with a spike count above a given fixed threshold. We computed averaged spontaneous correlation $\rho(T)$ as the Pearson correlation coefficient between the spike counts of neuronal pairs computed across time (count window $T = 100$ ms) and averaged over all single-unit pairs (Fig. 1C-F). A surrogate dataset was created to assess the amount of correlation during active periods (Fig. 1C-F, blue): silent periods were removed from spontaneous activity and the remaining active periods were concatenated to form, for each epoch, a continuous recording with $S = 0$.

To analyze evoked activity, we used S to classify epochs into three brain state categories: desynchronized ($S < 0.05$), intermediate ($0.05 \leq S \leq 0.2$), and synchronized ($S > 0.2$). We computed the mean population-averaged instantaneous rate, correlation $\rho(t)$ (Pearson correlation coefficient) and

spike count Fano factor across stimulus repetitions within each state category, using sliding windows (2, 11–15) ($T = 50$ ms; Fig. 2 and Figs. S4 and S5). The instantaneous silence density $S(t)$ in each category was defined as the fraction of trials with no spikes in the bin $(t, t + \Delta t)$ with $\Delta t = 20$ ms (Fig. 2 C and D and Fig. S4C). Details are described in *SI Methods*.

Computational Model. We assumed that neuronal variability had two sources (1, 11, 35–37): the first resulted from the variations in the population rate $r(t)$ mainly caused by transitions between silent and active network attractors, and the second arising from spiking stochasticity existent at constant rate. Neurons fired statistically identical Poisson spike trains with rate $r(t)$. We considered two cases: when the Poisson spike trains are conditionally independent, i.e., the only source of correlation is the fluctuations in $r(t)$, the spike count Pearson correlation coefficient $\rho(t)$ is given by Eq. 1. We also considered covariability introduced by the spiking stochasticity. This implies that the Poisson spike trains, conditioned on the rate, had an instantaneous covariance of amplitude c_0 , a coefficient that is added in the numerator of Eq. 1 to yield the expression of $\rho(t)$. The term $R(t; T) = \langle n(t; T) | r(t) \rangle$ in Eq. 1 is the expected number of spikes in the interval $(t - T/2, t + T/2)$ given $r(t)$, and can be numerically obtained from the integral of $r(t)$ in that interval (we refer to this as the integrated rate).

To describe the fluctuations of $r(t)$, we used a rate model with adaptation $a(t)$ where the dynamics were given by the following (22, 25, 28):

$$\begin{aligned} \tau_r \frac{dr}{dt} &= -r(t) + \phi(ar(t) - a(t) + I(t) - \theta), \\ \tau_a \frac{da}{dt} &= -a(t) + \beta r(t), \end{aligned} \quad [2]$$

where $\theta = 2$ was the activation threshold and the external input $I(t) = I + \text{stim}(t) + \sigma \xi(t)$ was composed of constant term I (range, 0–4 a.u.), the stimulus step function (amplitude, 60 a.u.; duration, 10 ms), and a noise term modeled as an Ornstein–Uhlenbeck process $\xi(t)$ (mean, 0; SD, 4.5; and time constant, 0.5 ms). The transfer function was $\phi(x) = gx^2$, if $0 < x \leq 1$, $\phi(x) = g\sqrt{4x-3}$, if $x > 1$, and zero otherwise. The time constants were set to $\tau_r = 5$ ms and $\tau_a = 250$ ms. Other parameters were $\alpha = 4.6$ s and $\beta = 0.3$ –3 s. Silence density was defined as the fraction of time $r(t) < 0.9$ Hz. Additional details are provided in *SI Methods*.

ACKNOWLEDGMENTS. We thank D. Jercog, E. Montbrío, and A. Renart for discussions; A. Amir for help during experiments; L. Hollender for sharing her data; and K. Wimmer, E. Kublik, and A. Compte for comments on the manuscript. This work was supported by the Polish Ministry of Science and Higher Education “Mobility Plus” Program Grant 641/MOB/2011/0 (to G.M.); the Spanish Ministry of Economy and Competitiveness together with the European Regional Development Fund Grants BES-2011-049131 (to A.H.-M.), and SAF2010-15730, SAF2013-46717-R, and RYC-2009-04829 (to J.R.); European Union Marie Curie Grant IRG PIRG07-GA-2010-268382 (to J.R.); and Wellcome Trust Grant 095668 (to K.D.H.). Part of the work was carried out at the Esther Koplowitz Centre, Barcelona.

- Shadlen MN, Newsome WT (1998) The variable discharge of cortical neurons: Implications for connectivity, computation, and information coding. *J Neurosci* 18(10):3870–3896.
- Smith MA, Kohn A (2008) Spatial and temporal scales of neuronal correlation in primary visual cortex. *J Neurosci* 28(48):12591–12603.
- Kohn A, Smith MA (2005) Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J Neurosci* 25(14):3661–3673.
- Cohen MR, Maunsell JHR (2009) Attention improves performance primarily by reducing interneuronal correlations. *Nat Neurosci* 12(12):1594–1600.
- Mitchell JF, Sundberg KA, Reynolds JH (2009) Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron* 63(6):879–888.
- Cohen MR, Newsome WT (2008) Context-dependent changes in functional circuitry in visual area MT. *Neuron* 60(1):162–173.
- Ecker AS, et al. (2010) Decorrelated neuronal firing in cortical microcircuits. *Science* 327(5965):584–587.
- Goard M, Dan Y (2009) Basal forebrain activation enhances cortical coding of natural scenes. *Nat Neurosci* 12(11):1444–1449.
- Renart A, et al. (2010) The asynchronous state in cortical circuits. *Science* 327(5965):587–590.
- Poulet JFA, Petersen CCH (2008) Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. *Nature* 454(7206):881–885.
- Churchland MM, et al. (2010) Stimulus onset quenches neural variability: A widespread cortical phenomenon. *Nat Neurosci* 13(3):369–378.
- Mochol G, Wójcik DK, Wypych M, Wróbel A, Waleszczyk WJ (2010) Variability of visual responses of superior colliculus neurons depends on stimulus velocity. *J Neurosci* 30(9):3199–3209.
- Ponce-Alvarez A, Thiele A, Albright TD, Stoner GR, Deco G (2013) Stimulus-dependent variability and noise correlations in cortical MT neurons. *Proc Natl Acad Sci USA* 110(32):13162–13167.
- Oram MW (2011) Visual stimulation decorrelates neuronal activity. *J Neurophysiol* 105(2):942–957.
- Middleton JW, Omar C, Doiron B, Simons DJ (2012) Neural correlation is stimulus modulated by feedforward inhibitory circuitry. *J Neurosci* 32(2):506–518.
- Rajan K, Abbott LF, Sompolinsky H (2010) Stimulus-dependent suppression of chaos in recurrent neural networks. *Phys Rev E Stat Nonlin Soft Matter Phys* 82(1 Pt 1):011903.
- Deco G, Hugues E (2012) Neural network mechanisms underlying stimulus driven variability reduction. *PLoS Comput Biol* 8(3):e1002395.
- Litwin-Kumar A, Doiron B (2012) Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nat Neurosci* 15(11):1498–1505.
- Amit DJ, Brunel N (1997) Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cereb Cortex* 7(3):237–252.
- Tsodyks M, Uziel A, Markram H (2000) Synchrony generation in recurrent networks with frequency-dependent synapses. *J Neurosci* 20(1):RC50.
- Brunel N, Hakim V (1999) Fast global oscillations in networks of integrate-and-fire neurons with low firing rates. *Neural Comput* 11(7):1621–1671.
- Latham PE, Richmond BJ, Nelson PG, Nirenberg S (2000) Intrinsic dynamics in neuronal networks. I. Theory. *J Neurophysiol* 83(2):808–827.
- Bazhenov M, Timofeev I, Steriade M, Sejnowski TJ (2002) Model of thalamocortical slow-wave sleep oscillations and transitions to activated States. *J Neurosci* 22(19):8691–8704.
- Compte A, Sanchez-Vives MV, McCormick DA, Wang X-J (2003) Cellular and network mechanisms of slow oscillatory activity (<1 Hz) and wave propagations in a cortical network model. *J Neurophysiol* 89(5):2707–2725.
- Curto C, Sakata S, Marguet S, Itskov V, Harris KD (2009) A simple model of cortical dynamics explains variability and state dependence of sensory responses in urethane-anesthetized auditory cortex. *J Neurosci* 29(34):10600–10612.
- Deco G, Marti D, Ledberg A, Reig R, Sanchez-Vives MV (2009) Effective reduced diffusion-models: A data driven approach to the analysis of neuronal dynamics. *PLoS Comput Biol* 5(12):e1000587.
- Mejias JF, Kappen HJ, Torres JJ (2010) Irregular dynamics in up and down cortical states. *PLoS One* 5(11):e13651.
- Mattia M, Sanchez-Vives MV (2012) Exploring the spectrum of dynamical regimes and timescales in spontaneous cortical activity. *Cogn Neurodyn* 6(3):239–250.
- Okun M, et al. (2012) Population rate dynamics and multineuron firing patterns in sensory cortex. *J Neurosci* 32(48):17108–17119.
- Sanchez-Vives MV, McCormick DA (2000) Cellular and network mechanisms of rhythmic recurrent activity in neocortex. *Nat Neurosci* 3(10):1027–1034.
- Steriade M, Nuñez A, Amzica F (1993) A novel slow (<1 Hz) oscillation of neocortical neurons in vivo: Depolarizing and hyperpolarizing components. *J Neurosci* 13(8):3252–3265.
- Cowan RL, Wilson CJ (1994) Spontaneous firing patterns and axonal projections of single corticostriatal neurons in the rat medial agranular cortex. *J Neurophysiol* 71(1):17–32.
- Steriade M, Timofeev I, Grenier F (2001) Natural waking and sleep states: A view from inside neocortical neurons. *J Neurophysiol* 85(5):1969–1985.
- Ecker AS, et al. (2014) State dependence of noise correlations in macaque primary visual cortex. *Neuron* 82(1):235–248.
- Stauder B, Rotter S, Grün S (2008) Can spike coordination be differentiated from rate covariation? *Neural Comput* 20(8):1973–1999.
- Churchland AK, et al. (2011) Variance as a signature of neural computations during decision making. *Neuron* 69(4):818–831.
- Goris RLT, Movshon JA, Simoncelli EP (2014) Partitioning neuronal variability. *Nat Neurosci* 17(6):858–865.
- van Vreeswijk C, Sompolinsky H (1996) Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* 274(5293):1724–1726.
- Hill S, Tononi G (2005) Modeling sleep and wakefulness in the thalamocortical system. *J Neurophysiol* 93(3):1671–1698.
- Renart A, Moreno-Bote R, Wang X-J, Parga N (2007) Mean-driven and fluctuation-driven persistent activity in recurrent networks. *Neural Comput* 19(1):1–46.
- Mongillo G, Hansel D, van Vreeswijk C (2012) Bistability and spatiotemporal irregularity in neuronal networks with nonlinear synaptic transmission. *Phys Rev Lett* 108(15):158101.
- Destexhe A (2009) Self-sustained asynchronous irregular states and Up-Down states in thalamic, cortical and thalamocortical networks of nonlinear integrate-and-fire neurons. *J Comput Neurosci* 27(3):493–506.
- Petersen CCH, Hahn TTG, Mehta M, Grinvald A, Sakmann B (2003) Interaction of sensory responses with spontaneous depolarization in layer 2/3 barrel cortex. *Proc Natl Acad Sci USA* 100(23):13638–13643.
- Luczak A, Barthó P, Marguet SL, Buzsáki G, Harris KD (2007) Sequential structure of neocortical spontaneous activity in vivo. *Proc Natl Acad Sci USA* 104(1):347–352.
- Sachidhanandam S, Sreenivasan V, Kyriakatos A, Kremer Y, Petersen CCH (2013) Membrane potential correlates of sensory perception in mouse barrel cortex. *Nat Neurosci* 16(11):1671–1677.
- Tan AYY, Chen Y, Scholl B, Seidemann E, Priebe NJ (2014) Sensory stimulation shifts visual cortex from synchronous to asynchronous states. *Nature* 509(7499):226–229.
- Vyazovskiy VV, et al. (2011) Local sleep in awake rats. *Nature* 472(7344):443–447.