

Stimulus onset quenches neural variability: a widespread cortical phenomenon

Churchland MM**^{1,2}, Yu BM**^{1,2,13}, Cunningham JP¹, Sugrue LP^{2,3}, Cohen MR^{2,3}, Corrado GS^{2,3}, Newsome WT^{2,3,4}, Clark AM⁶, Hosseini P⁶, Scott BB⁶, Bradley DC⁶, Smith MA⁷, Kohn A^{8,9}, Movshon JA⁹, Armstrong KM^{2,4}, Moore T^{2,4}, Chang SW¹⁰, Snyder LH¹⁰, Priebe NJ¹¹, Finn IM¹², Ferster D¹², Ryu SI^{1,5}, Santhanam G¹, Sahani M¹³, and Shenoy KV^{1,2}

¹Dept. of Electrical Engineering, ²Neurosciences Program, ³Howard Hughes Medical Institute, ⁴Department of Neurobiology, Stanford University School of Medicine, ⁵Department of Neurosurgery, Stanford University School of Medicine, Stanford University, Stanford CA, 94305, USA. ⁶Department of Psychology and Committee on Computational Neuroscience, University of Chicago, Chicago, IL 60637, USA. ⁷Department of Neuroscience and Center for the Neural Basis of Cognition, University of Pittsburgh, Pittsburgh, PA, USA. ⁸Department of Neuroscience, Albert Einstein College of Medicine, Bronx, NY, USA. ⁹Center for Neural Science, New York University, New York, NY 10003, USA. ¹⁰Department of Anatomy and Neurobiology, Washington University School of Medicine, St. Louis, MO, USA. ¹¹Section of Neurobiology, School of Biological Sciences, University of Texas at Austin, Austin, Texas 78705, USA. ¹²Department of Neurobiology and Physiology, Northwestern University, Evanston, IL 60208, USA. ¹³Gatsby Computational Neuroscience Unit, University College London, London, United Kingdom.

Author Contributions. MMC and BMY contributed equally to this work. MMC wrote the manuscript, performed the FF and FA analyses, and made the figures. GPFA was developed by BMY, JPC, MS and KVS. This application of FA was devised by MMC and BMY. The mean-matched Fano factor was developed by MMC and KVS. The conception for the study arose from conversations between MMC, KVS, BMY, DCB, MRC, WTN and JAM. V1 data (extracellular) were collected in the laboratory of JAM by MAS and AK and in the laboratory of AK. V4 data were collected in the laboratory of TM by KMA. MT (plaid) data were collected in the laboratory of DCB by AMC, PH and BBS. MT (dots) data were collected in the laboratory of WTN by MRC. LIP and OFC data were collected in the laboratory of WTN by LPS, using an experimental design developed by LPS and GSC. PRR data were collected in the laboratory of LHS by SWC. PMd data were collected in the laboratory of KVS by BMY, SIR, GS and MMC. MT (direction/area and speed) data were collected in the laboratory of Stephen G. Lisberger by NJP and MMC. Intracellularly-recorded V1 data were collected by NJP and IMF in the laboratory of DF. All authors contributed to manuscript revisions and editing, particularly JAM, WTN, LPS, DF, JPC, BMY and KVS.

Neural responses are typically characterized by computing the mean firing rate. Yet response variability can exist across trials. Many studies have examined the impact of a stimulus on the mean response, yet few have examined the impact on response variability. We measured neural variability in 13 extracellularly-recorded datasets and one intracellularly-recorded dataset from 7 areas spanning the four cortical lobes. In every case, stimulus onset caused a decline in neural variability. This occurred even when the stimulus produced little change in mean firing rate. The variability decline was observable in membrane potential recordings, in the spiking of individual neurons, and in correlated spiking variability measured with implanted 96-electrode arrays. The variability decline was observed for all stimuli tested, regardless of whether the animal was awake, behaving, or anaesthetized. This widespread variability decline suggests a rather general property of cortex: that its state is stabilized by an input.

A fundamental approach of systems neuroscience is to probe the brain with repeated stimulus ‘trials’ and infer neural mechanism from the recorded responses. Extracellularly-recorded responses are typically analyzed by computing the average spike rate across trials. By averaging, the experimenter hopes to overcome the apparent noisiness of spiking and estimate the true change in the neuron’s underlying firing rate. It is likely true that much of the recorded spiking variability is effectively noise, and doesn’t reflect fundamentally different responses on different trials. Yet it is nevertheless clear that the neural response can vary meaningfully across trials. For example, the neural state may be initially similar across trials, but become variable in response to a stimulus, as in¹. Alternately, sensory cortex can be restless and active^{2,3} prior to stimulus onset. A central question is whether the stimulus-driven response suppresses such ongoing variability^{4,5,6,7}, superimposes with it^{2,8,9}, or yields even greater variability due to non-linear interactions¹⁰? In general, does stimulus onset drive variability up (due to the variable responses themselves) or down (due to suppression of variable ongoing activity)?

In general, the mean rate provides an incomplete characterization of the neural response. A fuller characterization requires – at the least – knowing whether rate variability is present and how it changes with time. For example, the responses in **Figure 1a** and **b** have similar means, yet one would infer different things about the neural circuits that gave rise to them. The mean in **Figure 1c** erroneously suggests little stimulus-driven response. **Supplementary Figure 1** illustrates a similar scenario using a simulated network. Because such situations may be common, it is important to characterize not only the stimulus-driven change in mean rate, but also the stimulus-driven change in rate variance.

The impact of a stimulus on variability could, of course, depend on the brain area, stimulus, and task. However, stimulus onset reduces both membrane potential variability in anaesthetized cat V1^{4,5} and firing-rate variability in premotor cortex of reaching monkeys^{11,12}. The presence of similar effects in two very different contexts suggests that a decline in variability could be a common feature of the cortical response. This would

agree with recent theoretical work^{13,14} indicating that such an effect may be a general property of large recurrent networks.

To address this issue, we analyzed recordings from many cortical areas, driven via a variety of stimuli. A measure of firing-rate variability (the Fano factor) revealed a stimulus-driven decline in variability that was similar in timecourse to the decline in V1 membrane-potential variability. This decline was present not only for anaesthetized V1, but for all cortical areas tested regardless of the stimulus or behavioral state. The decline was also present in the correlated firing-rate variability of neurons recorded using implanted multi-electrode arrays. Finally, we demonstrate how recently developed methods, applied to simultaneous multi-electrode recordings, can reconstruct the variable evolution of firing rates on individual trials.

Results.

Across-trial variability in the membrane potential

Stimuli and task events can alter the structure and correlation¹⁵ of membrane-potential variability. In particular, visual stimuli drive a reduction in membrane potential (V_m) variability in cat primary visual cortex (V1) that is largely independent of stimulus orientation^{4,5}. We re-analyzed data from⁵ to illustrate the timecourse of this effect (**Fig. 2**). Stimulus onset drives an immediate decline in V_m variability. This decline occurs even for non-preferred stimuli that elicit little change in mean V_m (see also⁴ and Fig. 7c,d of⁵). Average variability (across all neurons/conditions) declined rapidly following stimulus onset and then remained at a rough plateau (**Fig. 2c**). The variability in question is across-trial variability, with a fairly long autocorrelation. When V_m is low (or high), it tends to stay low (or high) for 10's to 100's of ms.

The relationship between intracellularly-recorded V_m variability and extracellularly-recorded firing-rate variability is likely to be complex, given the nonlinear and dynamic relationship between V_m and firing rate (e.g., considerable V_m variability occurs below threshold). One nevertheless expects across-trial V_m variability to produce across-trial firing rate variability. Consistent with this expectation, firing-rate variability declines in parallel with V_m variability, as we will show below. A larger question is whether the observed decline in variability is specific to V1 or whether it reflects a broader phenomenon. The latter is suggested both by the presence of a similar effect in a premotor cortex¹¹ and by recent theoretical work^{13,14}.

Addressing this issue requires quantifying firing-rate variability in extracellular recordings. While quantifying V_m variability is straightforward, quantifying firing-rate variability is more complicated. Extracellularly-recorded spike trains are usually described in terms of an underlying firing rate (often termed λ) observed via a noisy point process (e.g., Poisson) that produces spikes. It should be stressed that this conception captures the statistics of neurons embedded in a network^{16,17}: spike generation at the axon hillock is not responsible for the noisy spiking-process statistics¹⁸, nor is firing rate synonymous with membrane-potential. Rather, the underlying firing rate can be thought

of as the average response of many similarly tuned neurons, or as the average response of one neuron across truly identical trials. Of course, repeated trials are not guaranteed to be truly identical: the underlying firing rate may differ somewhat. It is precisely this variability that we wish to capture, while ignoring variability arising from the roughly-Poisson spiking. Spiking variability may have interesting structure of its own, but for present purposes acts as noise.

Poisson spiking-process ‘noise’ can severely mask underlying firing-rate variability (**Supplementary Fig. 1**). It is thus rarely possible to discern changes in firing-rate variability by eye. We use two approaches to isolate underlying firing-rate variability from the variability contributed by spiking noise. The first approach employs a modified method for computing the Fano factor. This method is applicable to conventionally recorded single-neuron data, allowing analysis of a large number of existing datasets. The second approach uses factor analysis to assess covariance in large-scale simultaneous recordings. These methods are technically very different, but both are intended to assess the same thing: the degree of across-trial firing rate variability, independent of the contribution of noisy spiking statistics.

A variability decline across multiple cortical areas

We first employed the Fano factor (FF): the spike-count variance divided by the spike-count mean. Counts were made in a sliding window over the duration of the trial. The FF has been used extensively to characterize neural variability^{19,20,21,22,23}. The FF is influenced both by variability arising from spiking noise, and by across-trial variability in the underlying firing rate²⁴. Most prior work assumes that the underlying firing rate is similar across trials, and uses the FF to assess the statistics of spiking noise, which are roughly Poisson ($FF \approx 1$) for most of cortex. In the present work, we begin with the assumption that spiking noise is roughly Poisson, and we use the FF to assess across-trial variability in the underlying rate. An FF greater than one will be interpreted as an indication of across-trial firing-rate variability. Changes in the FF will be interpreted as reflecting changes in across-trial firing-rate variability^{11,24,25}. Although this approach begins by assuming Poisson spiking noise, it is reasonably robust to violations of that assumption (it is sufficient that spiking-noise variance scale linearly with the mean; the

slope needn't be unity). We will also present critical controls for potential artifacts related to non-Poisson spiking noise. But before doing so we examine how the FF behaves across a variety of cortical areas.

We computed the mean firing rate and the FF for ten datasets from seven cortical areas of the macaque monkey (**Fig. 3**): V1, V4, MT, the lateral intra-parietal area (LIP), the parietal reach region (PRR), dorsal premotor cortex (PMd) and orbitofrontal cortex (OFC). Responses were to various visual stimuli or, for OFC, to juice reward. For each area the FF was averaged across neurons/conditions (see below). This is similar to what was done for the membrane potential analysis, and reflects both a desire for statistical power and the expectation that variability may change for both preferred and non-preferred stimuli (as in **Fig. 2a,b**).

In every case, stimulus onset drove a decline in firing-rate variability as assessed by the FF (all p-values <0.02). This is remarkable, given the diversity of areas, stimuli, and behavioral states. Variability declined during responses to simple visual stimuli, during operantly-conditioned responses (PRR and PMd), and during reward-driven responses (OFC). The variability decline was present regardless of whether the animal was anaesthetized (V1 and the MT datasets at bottom), passively viewing (V4) or performing a task (the other six datasets). For the MT datasets at bottom, stimulus onset occurred in two stages: pattern onset followed by motion onset. Both events drove a decline in variability, although only the more effective moving stimulus drove a sustained decline.

We have previously proposed that declining variability in premotor cortex is related to the progress motor preparation¹¹. The changes in V1 membrane potential variability have been linked to shunting inhibition^{4,26} and contrast-invariant orientation tuning⁵. These explanations are no less likely given the current data. Yet it is clear that the decline in variability is not limited to those cases, but constitutes a cortex-wide effect. This suggests a very general property of cortex: that its state is stabilized by an input.

The mean-matched Fano factor controls for rising rates

The FF declines as mean firing rates are rising. This is expected: in the intracellular data it was similarly true that V_m variability dropped as the mean V_m rose. It

is nevertheless critical to rule out any possibility that the declining FF is trivially related to rising rates. At higher rates, refractory periods will tend to regularize spiking (a violation of the Poisson assumption) which could reduce the FF simply because the variability of spiking-process noise is reduced. A less obvious concern is that across-trial firing-rate variability could be constant, but becomes normalized by a higher mean after stimulus onset (the FF is the spike-count variance divided by the mean).

To address these issues, the FF traces above (**Fig. 3**) were computed using a method we term ‘mean-matching.’ The motivation is as follows: interpretation would be aided if we could record from a population of neurons that did not change its overall distribution of mean firing rates. If some neurons responded to the stimulus with an increased mean rate, and others with a decreased mean rate, then the overall distribution of mean rates (across neurons and conditions) could stay conveniently constant. For many areas (e.g., PRR) this is not far from what was observed. The overall mean rate always increased post-stimulus, yet this was often driven by a small percentage of neurons that strongly preferred the stimulus. Most neurons exhibited more modest changes in mean rate, with declines being almost as common as increases. For other areas (especially anaesthetized V1) it was the case that the vast majority of neurons increased their mean rate following stimulus onset. Nevertheless, for all datasets there was considerable overlap, before and after stimulus onset, in the distribution of mean rates. Mean-matching exploits this overlap.

To apply mean-matching, for each neuron/condition we compute the mean and variance of the spike count in a sliding window (**Fig. 4a**). At each time, we plot the variances versus their respective means, with one data-point for every neuron/condition (**Fig. 4b**, gray and black points). The ‘raw FF’ is simply the slope of the regression relating the variance to the mean. The concern being addressed is that the raw FF might be lower following stimulus onset simply because rates are higher. Inspection reveals this to be unlikely: even for the range of mean counts that is shared across times, there are noticeably fewer high-variance points for the later times (compare -100 and $+100$ ms). Mean-matching formalizes this observation. We compute distributions of mean counts (*gray* marginal distributions) and take the greatest common distribution (*black*). At each

time, individual points are randomly excluded until the actual distribution matches the common distribution. The mean-matched FF is based on the remaining (*black*) points.

All FFs reported in **Figure 3** employed mean-matching. Very similar results were obtained for the raw FF (as in **Fig. 4c**). Thus, the decline is not somehow produced by mean-matching. Mean-matching simply controls for a potential artifact. Also note that, even following mean-matching, most of the analyzed neurons still showed stimulus-driven changes in mean rate (both increases and decreases). Only the *overall* distribution of firing rates is held constant. Before application to real data, mean-matching was applied to many simulated datasets, a few of which are shown in **Supplementary Figure 2**. The mean-matched FF successfully avoids artifacts related to changing mean rates over a wide range of circumstances, including low trial-counts, non-Poisson spiking, low mean rates, and large changes in mean rate.

Further controls for the FF

We employed a complementary method to further control for potential changes in spiking-noise statistics. We computed both the FF (for all neurons/conditions) and the square of the coefficient of variation of the inter-spike intervals (CV^2) using the method of²⁴. While the FF declined following stimulus onset, the CV^2 changed little (**Supplementary Fig. 3**). This is not consistent with a large change in spiking-process statistics (which should produce a large decline in both the FF and CV^2) but is consistent with a decline in across-trial rate variability, a similar conclusion to that drawn in²⁴. One can also control for rising mean rates (and resulting refractoriness) by focusing on conditions/neurons that *individually* exhibit little change in mean rate (e.g., null-directions and other non-responsive conditions). This necessitates that there exist a sufficient number of such cases to allow further analysis. Of the datasets in **Figure 3**, four (MT-plaids, PRR, PMd, and MT-speed) could be analyzed in this manner. In each case (**Fig. 5**), the FF decline is as clear for the non-responsive conditions ($p < 10^{-6}$, $p < 10^{-6}$, $p < 0.002$, $p < 10^{-4}$) as it had been for the original analyses. This finding also addresses a second concern: low-rate neurons (which are preferentially preserved by mean-matching following stimulus onset) might tend to have low intrinsic FFs, resulting in an artifactual FF decline. However, such an artifact cannot affect the analysis in **Figure 5**: exactly the

same neurons/conditions contribute at all times. Furthermore, the raw FF always showed the same decline as the mean-matched FF. In summary, the decline in the FF cannot be a simple consequence of refractoriness, changing spiking-noise statistics, or normalization. Rather it appears to reflect a widespread decline in the variance of the underlying firing rate.

Variability declines even for ‘non-responsive’ conditions

Analyses of non-responsive conditions (**Figs. 2,5**) act a control, but also make a deeper point. The firing rate – averaged across trials – may change little, yet neurons are not truly non-responsive. It appears that underlying rates *do* change: from a variable range of rates before stimulus onset to a more consistent rate (with a similar mean) following stimulus onset. As in **Figure 1c**, it may be only the mean rate that changes little. The critical factor is not whether a neuron appears responsive from its mean, but whether the population as a whole is responsive. This is illustrated by the data from anaesthetized MT (**Fig. 3, bottom**). The onset of a static dot pattern evokes a weak and transient increase in mean rate at the population level, and a transient (though large) decline in the FF. In contrast, dot motion onset evokes a sustained increase in the mean population rate and a sustained FF decline. This FF decline is present and sustained even for those neurons where the mean rate is changing little (**Fig. 5d**). Thus, the FF decline depends on the stimulus effectively driving the population as a whole, but not on individual-neuron mean responsiveness. Consistent with this, the FF rises again fairly rapidly when the population ceases to respond following stimulus offset (**Supplementary Fig. 4**).

A decline in correlated variability

Does each neuron undergo its own decline in variability, or is the above effect related to variability that is shared? To assess this, our second approach exploits the covariance of simultaneous recordings, obtained via 96-electrode array recordings from two of the above areas: V1 and PMd. The measured covariance matrix (assessed across trials) reflects both shared firing-rate variability and ‘private’ spiking-process noise. It is thus necessary to decompose the measured covariance matrix into shared (network-level)

and private components (**Fig. 6a**). A number of applicable methods exist, all of which exploit the assumption that shared variability should be relatively low dimensional (i.e., that neurons are tuned to a modest number of shared factors). Intuitively, if one can average across similarly tuned neurons to remove spiking noise, the underlying rate (and its variance) becomes more accessible. More precisely, if neurons are tuned for 8 shared factors, the covariance matrix describing the underlying rate will be of rank 8. The measured covariance matrix should then be decomposed into a rank-8 ‘network’ covariance matrix, and a diagonal matrix capturing private spiking-process noise. Spiking-process variability is not necessarily physiologically cell-intrinsic^{16,18,20}, but is weakly correlated between most neurons and thus ‘private’ for present purposes, especially given the long analysis windows being used. Of the applicable methods²⁷, factor analysis²⁸ has the most appropriate noise model: different neurons may have different levels of private noise. This is essential, as Poisson spiking ‘noise’ scales with the mean rate. **Supplementary Figure 5** further illustrates how FA separates network variability from private spiking noise.

We applied FA to four datasets: two V1 datasets from different animals and two PMd datasets collected from one animal on separate days. After isolating the network variances (blue diagonal in **Fig. 6a**) we computed the average network variance (across neurons/conditions) before and after stimulus onset. The average variance was computed following a mean-matching procedure, just as was performed for the FF. Mean-matching aids interpretation by allowing direct comparison of the variance present for a given mean rate. It also insures that any contamination of network variability by private noise (inevitable for finite trial-counts) does not increase with rising rates. Non-mean-matched data are presented subsequently.

A decline in network variability was observed in all datasets (**Fig. 6b,c,e,f**). For V1, network variability was 66% and 40% lower following stimulus onset. For PMd, network variability was 53% and 45% lower. This decline was significantly greater for network variability than for private noise, in both absolute and relative terms (**Fig. 6d,g** and legend). Thus, the FF and FA analyses agree: stimulus onset drives a decline in firing-rate variability that is not simply due to a change in the regularity of spiking itself.

The decline in rate covariance is consistent with²⁹ (from which one of the V1 datasets is drawn), in which correlations were lower overall after stimulus onset. However, the present result does not simply follow from that result. The mean pairwise correlation can decline even while overall covariance grows (e.g., if negative correlations become more common or if private spiking variability climbs). Interestingly, in three of the four cases, FA also indicates a small reduction in private noise. Limited data make it difficult to determine whether this indicates an actual change in spiking-process statistics³⁰ or occurs because some network variability is incorrectly assigned to private noise. The latter can occur if the true dimensionality of the data is higher than the dimensionality assumed by the analysis (**Supplementary Fig. 7a**). This possibility is of minor concern: it would imply we are underestimating the network-variance decline. The converse (mis-assignment of a private change to a network change) is potentially concerning yet unlikely: for trial-shuffled data FA correctly assigns the variance change to private noise with only a small ‘leak’ into network variance (**Supplementary Fig. 7b–e**). Furthermore, the decline in network variability was the larger effect, both in absolute (**Fig. 6b,c,e,f**) and percentage (**Fig. 6d,g**) terms. In summary, there may or may not be modest changes in private spiking-process variability following stimulus onset. In either case, the measured decline in variability is principally a decline in network-level variability, shared among neurons.

As an aside, the shared variability being assessed is indeed variability in the firing rate across trials, not variability in precise spike patterns. Rerunning FA and FF analyses after ‘jittering’ each spike within a 20 ms window produced essentially identical results. This is unsurprising given the large (400 ms) analysis window.

A qualification regarding the decline in variability

For both the FF and FA, we compared identical distributions of mean rates before and after stimulus onset. Our central result must thus be stated carefully: for a given mean rate there is less variability around that mean following stimulus onset (e.g., the variance around a 10 spike/s mean might shrink from 15 to 8 spikes²/s²). This will often imply that overall variability declines, but that can fail to occur when rates are initially very low. In the limit, when initial firing rates are zero, variability is zero and will

necessarily rise when rates rise. For area PMd, overall firing rates change modestly following stimulus onset and overall variability does indeed decline. However, in anaesthetized V1 most neurons have baseline rates near zero, and thus near-zero variability. Overall variability is thus essentially guaranteed to rise following stimulus onset. This can be seen by plotting network-level rate variability versus mean rate (**Fig. 6h,i**). In this analysis, spiking variability has been approximately factored out, leaving primarily underlying-rate variability. The effect of stimulus onset is essentially the same for both PMd and V1: the black curve lies well below the gray curve. For any given mean rate there is less variability around that rate after stimulus onset. However, for V1 mean rates are low before stimulus onset, necessitating low rate variance. Overall variability thus rises following stimulus onset (average variance rises from 77 ± 5 SE spikes²/s² to 190 ± 8 spikes²/s²) even though it is lower for any given mean rate. For PMd, the overall change in mean rates is modest, and variance is lower after stimulus onset both for a given rate and in overall terms (average variance drops from 28 ± 2 spikes²/s² to 16 ± 1 spikes²/s²). Thus, the central effect is a decline in the variability present around a given mean rate. This may or may not translate into an overall variance decline, depending on whether overall mean rates rise modestly or sharply. This qualification is restricted to the analysis of firing rate, whose variability must be zero when the mean is zero. By contrast, the membrane voltage shows considerable variability even when firing rates are near zero (**Fig. 2**).

Single trial analysis of variability

The ‘neural response’ is incompletely characterized by its mean. A fuller characterization is provided by also assessing firing-rate variance. Yet ideally one would directly observe the evolution of firing rates on individual trials, much as the membrane potential was directly observed on individual trials in Figure 2. Such visualization is likely important not only for the initially-high variance, but also for the variance that survives stimulus onset. The presence of remaining firing-rate variability after stimulus onset is underscored by its correlation with pre-stimulus variability^{2,8,10,31} (**Supplementary note 1**). The FA analysis (**Fig. 6h,i**) also indicates that considerable firing-rate variability remains following stimulus onset. For PMd, such variability is

behaviorally meaningful: it correlates with subsequent reaction time¹¹ and reach speed³². Estimating the evolution of single-trial firing rates is difficult or impossible based on single-neuron recordings, but becomes plausible for large-scale simultaneous recordings. To do so we employ an extension of FA termed ‘Gaussian-process factor analysis’ (GPFA)³³. GPFA projects the responses of many neurons into a low-dimensional space where each trial’s ‘neural trajectory’ can be traced through time. Essentially, each axis captures the smoothed and averaged response of similarly-tuned neurons. More precisely, each neuron’s underlying rate is a linear combination of the responses represented on each axis (much as for principal component analysis). Similar visualization has recently been employed in invertebrate systems^{1,34,35}. A full characterization of single-trial variability across all conditions is beyond the scope of the present study. Still, the analyses below serve to illustrate the power, and in some cases the limitations, of viewing neural responses on a single-trial basis.

The collapse in variance in PMd is readily visualized via GPFA (**Fig. 7a**). The neural state is initially variable across trials, even though the monkey’s external state – including fixation – is tightly controlled. Following reach-target onset, the neural state becomes less variable and changes its center of mass (i.e., mean rates change). These features are particularly clear in the supplementary movies. Across both datasets and all 14 target locations, variability was always lower by the go cue than it had been pre-stimulus (for the two days, variance was an average of 27% and 37% smaller per dimension, relative to 100 ms pre-stimulus, $p < 10^{-6}$ and $p < 10^{-8}$; t-tests across 14 conditions). Variability was lower still by movement onset (36% and 46% smaller than pre-stimulus variability; $p < 10^{-10}$, $p < 10^{-11}$).

Despite the above results, the principal goal of GPFA is not to further quantify variability and its decline (that is better accomplished by the FF and FA analyses). Rather, we would like to know *how* individual trials deviate from their mean. In particular, while most neural trajectories hewed close to their mean after target onset there were intriguing exceptions. For example, the red trace in **Figure 7b** follows the usual trajectory up until the go cue (presumably indicating successful target detection and initial planning). Yet following the go cue the trajectory wanders before falling back on track. On this trial the monkey did indeed have an unusually long reaction time (629 ms;

all others ranged from 240–375 ms). Other trials, for other target locations/days, contained similar outliers associated with long reaction times (**Fig. 7c, Supplementary Fig. 8**). Thus, across-trial averaging obscures not only the variance decline but also other interesting features, including events that may only have happened once.

We also applied GPFA to data from V1. Neural trajectories for one example condition (a 45° drifting grating) are shown in **Figure 8**. The initial overshoot corresponds to the onset transient; the looping trajectory is driven by the periodic stimulus. Across-trial variability is, if anything, more pronounced than in PMd. Individual-trials showed considerable variability in the stimulus-driven trajectory, and also in the trajectory before stimulus onset and after its offset. Of course, given a limited number of neurons, GPFA cannot provide noise-free estimates of single-trial firing rates. Yet much of the observed variability is real: shuffling the data to break neuron-neuron correlations greatly reduced variability relative to the original data (**Fig. 8b,c**). This was similarly true for every stimulus condition for all datasets tested. For the two V1 datasets, the original data had on average 81% and 89% more variance ($p < 10^{-10}$ for both, t-test across conditions) than the shuffled data. The same was true for PMd (42% and 43% more variance; $p < 10^{-15}$ for both). Thus, the variable individual-trial trajectories do indeed reflect trial-to-trial departures from the mean behavior. In particular, for V1 there is a surprising degree of variability in the initial state (**Fig. 8c, blue**), given the low average initial rates of most neurons. Because mean-matching cannot be applied to GPFA, total variability certainly did not decline following stimulus onset. Indeed it was higher, in agreement with **Figure 6h**. In summary, GPFA is in some ways less flexible than the FF or FA analyses (mean-matching cannot be applied, and the noise model is fixed across time). Still, the ability to reconstruct single-trial responses, and observe events that may only have occurred once, lends it a unique utility.

Discussion

Stimulus onset drives a widespread decline in the variability of the neural state. Consistent with prior reports, this effect is directly observable at the level of the membrane potential^{4,5}. At the level of firing rates, the effect is not initially obvious. Much of the measured variability at the level of a single neuron is due to quasi-Poisson spiking statistics, whose variance rises with mean rate²⁰. Nevertheless, appropriately-controlled metrics reveal that firing-rate variability declines sharply following stimulus onset. Specifically, for a given mean rate there is much less variability around that rate after stimulus onset. The ability to assess this effect in extracellular records was critical, as it allowed examination of datasets from many cortical areas, employing many stimuli, in both anesthetized and behaving animals. Stimulus onset might have been expected to impact neural variability differently depending on the area, stimulus, task, and behavioral state. Yet we consistently observed a decline in neural variability across all 14 datasets, from occipital to frontal cortex. This finding argues that a stimulus-driven decline in variability is a widespread feature of the cortical response.

The FF and FA methods both indicate that the central effect is a decline in the across-trial variability of the underlying firing rate, rather than a change in the spiking-statistics applied to that rate. This decline in firing-rate variability is likely a network property, for two reasons. First, the decline occurs even when the mean firing rate of a given neuron changes little (e.g., for non-preferred stimuli) provided that the network as a whole responds to the stimulus. Second, the variability in question is shared among neurons within an area, as revealed by FA. This finding, and the similar decline across areas, suggests that the relevant network could span much of cortex. Alternately, it may be that different areas exhibit the same effect simply because it is a common property of cortex. Distinguishing these possibilities will require large-scale simultaneous recordings from multiple areas.

Many prior studies have examined correlations between firing rate variability and behavior (e.g.,^{31,32,36,37,38}). A few have related across-trial firing-rate variability with task difficulty,³⁹ attention²⁵ or motor learning^{40,41}. The present study is the first to fully document a simpler effect: the change in variability driven by the stimulus itself. That said, related effects can certainly be found in prior studies. This is true not only for the

membrane potential^{4,5}, but also for firing-rate variability. We previously reported (using methodology related to the FF) a variability decline in premotor cortex¹¹. Variability also undergoes changes in motor cortex^{24,42}. In retina, LGN and V1, the FF was lowest when rates were highest⁴³. In a recent study of V4²⁵, the FF dropped after the stimulus entered the receptive field. There is also a small FF drop in previously-published MT data⁴⁴. Variability declined throughout the trial for multiple areas during a probabilistic choice task⁴⁵. Zebrafish mitral cell responses become more reliable following odor presentation⁴⁶, and both cerebellar and FEF variability reach minima before movement^{47,48}. A stimulus-driven decline in correlated activity^{7,29} was found for V1, consistent with a decline in (correlated) rate variability (also see⁴⁹ for a similar effect r.e., LFP). Finally, there are reports⁵⁰ of lower variability in driven versus spontaneous thalamic activity. In some of these studies variability changes were of passing interest, and controls for changing spiking-process statistics were not always applied (although see^{11,24}). Yet our results argue that most of these effects likely reflect a decline in network-level firing-rate variance.

Mechanistically, the variance decline implies that cortical circuits become more stable when driven. Many network types can be stabilized by an input. Natural candidates are recurrent networks with attractor dynamics⁵¹. More broadly, a stimulus-driven decline in variability may be a general feature of large recurrent networks^{13,14}. The widespread effects we observed may thus simply reflect the widespread presence of recurrent circuitry. These ‘network level’ explanations are consistent with the shared nature of the firing-rate variability, and with our previous finding that a drop in PMd variability relates to reaction time¹¹. Yet other data indicate a bottom-up source; in particular, shunting inhibition appears to be the source of the decline in V1 membrane-potential variability^{4,26}. One hopes these two levels of explanation will be reconciled, although it is not yet clear how to do so.

The functional interpretation of converging rates likely varies by area. For PMd, the decline may relate to motor preparation^{11,12}. In area V4, it may relate to attention²⁵. Often it is tempting to relate the variability decline to cognitive-level phenomena: stimulus onset is a natural candidate to focus not only attention and planning, but other forms of mental processing as well. On the other hand, the decline was observed even

under anesthesia (all V1 datasets and two of the MT datasets) consistent with lower-level explanations⁴. See **Supplementary note 2** for further discussion of the possible origins – local versus cognitive – of the initial variability. The anesthetized data additionally rule out effects related to microsaccades or other unobserved changes in behavior. Finally, effects in the anaesthetized data appear unrelated to long-timescale drift of anesthesia: the decline is undiminished when computed more locally in time (**Supplementary note 3**).

Despite being widespread, the changes in variability are unlikely to constitute a ‘coding channel’. It is unlikely that the brain could measure or interpret variability that is present only across trials. It is the experimenter who, when interpreting the recorded responses, needs to know when variability is substantial and how it is changing. Most basically, we need to know when the mean is representative. More intriguingly, different classes of dynamics – single-attractor, multi-attractor, integrator – can have different variability ‘signatures’⁵². Variability measurements may thus allow hypotheses regarding circuit-level dynamics to be tested at the level of the single neuron. More generally, while the mean and the variance are indispensable metrics, one also wishes to observe the individual-trial responses themselves. GPFA and related methods can approximately reconstruct individual-trial neural trajectories from simultaneous extracellular recordings⁵³, and thus have the potential to uncover telling features of the neural response that are normally lost to averaging.

Supplementary material is provided.

Acknowledgements

This work was supported by a Helen Hay Whitney postdoctoral fellowship (MMC), Burroughs Welcome Fund Career Awards in the Biomedical Sciences (MMC, KVS), NIH Director's Pioneer Award 1DP1OD006409 (KVS), Gatsby Charitable Foundation (MS, BMY), National Institute of Health NINDS Collaborative Research in Computational Neuroscience Grant 5-R01-NS054283-02 (KVS, MS), the Michael Flynn Stanford Graduate Fellowship and NIH-NINDS-CRCNS-R01 (JPC), the Howard Hughes Medical Institute and NIH grant EY05603 (WTN, LPS, MRC, GSC), an HHMI predoctoral fellowship (MRC), NIH grant NIH EY014924 (TM, KMA), Sloan Foundation (TM, KVS), Pew Charitable Trust (TM), NIH EY015958 and EY018894 (MAS), NIH EY02017 and EY04440 (JAM), NIH EY016774 (AK), NIH 1 EY13138-01 (DCB, AMC, PH, BBS), NIH EY019288 (NJP), the Pew Charitable Trust (NJP), EY04726 (DF), NDSEG Fellowships (BMY, GS), NSF Graduate Research Fellowships (BMY, GS), the Christopher and Dana Reeve Foundation (KVS, SIR), and the following awards to KVS: Stanford Center for Integrated Systems, National Science Foundation Center for Neuromorphic Systems Engineering at Caltech, Office of Naval Research, and Whitaker Foundation. We thank Stephen G. Lisberger, in whose laboratory two of the datasets (MT-direction/area and MT-speed) were collected.

Figure 1. Schematic illustration of possible types of across-trial firing rate variability. In each panel, we suppose that the same stimulus is delivered four times (four trials) yielding four different responses. Panels *a* and *b* were constructed to have the same mean response across the four trials. *a.* Stimulus-driven decline in variability. *b.* Stimulus-driven rise in variability. *c.* Stimulus-driven decline in variability with little change in mean rate.

Figure 2. Analysis of intracellularly-recorded membrane potential from cat V1. Stimuli were drifting sine-wave gratings presented at different orientations / frequencies. Analysis employed a 50 ms sliding window (box filter) to match the 50 ms window used for the FF analysis. Similar results were obtained with a shorter (5 ms) or longer (100 ms) window. *a.* Data from one example neuron. V_m for individual trials (*black*) is plotted on top of the mean (*gray*). Data are shown when no stimulus was delivered, for a non-preferred stimulus, and for a preferred stimulus. The arrow marks stimulus onset. *b.* Similar plot for a second example neuron. *c.* The mean and variance of V_m across all 52 neurons and all stimuli. Flanking traces give SEMs.

Figure 3. Changes in firing-rate variance for eight datasets (one per panel). Insets indicate stimulus type. Data are aligned on stimulus onset (arrow). For the two bottom panels (MT-area/direction and MT-speed) the dot pattern appeared at time zero (first arrow) and began moving at the second arrow. The mean rate (gray) and the FF (black with flanking SEs) were computed using a 50 ms sliding window. For OFC, where response amplitudes were small, a 100 ms window was used to gain statistical power. Analysis included all conditions, including non-preferred. The FF was computed after mean-matching (**Fig. 4**). The resulting stabilized means are shown in black. The mean number of trials/condition was 100 (V1), 24 (V4), 15 (MT plaids), 88 (MT dots), 35 (LIP), 10 (PRR), 31 (PMd), 106 (OFC), 125 (MT-direction/area), and 14 (MT-speed).

Figure 4. Illustration of how the mean-matched FF was computed. Data are from the MT-plaids dataset. **a.** Spike rasters for the 46 trials (one per line) recorded from one MT neuron (127) for one stimulus condition (upwards-moving plaid). Shaded areas show 4 locations of the sliding window, which moved in 10 ms increments. For each window location, the spike count was computed for each trial. The mean and variance (across trials) of that count then contributed one data-point to the subsequent analysis (*next panel*). **b.** The FF was computed based on scatterplots of the spike-count variance versus mean. Each scatterplot corresponds to a window in *a*. Each point plots data for one neuron/condition (red indicates the neuron/condition from *a*). The orange line has unity slope: the expected variance-to-mean relationship for Poisson spiking. Data above the orange line is consistent with the presence of underlying-rate variability. Grey dots show all data. Gray lines are regression fits to all data (constrained to pass through zero, weighted according to the estimated SE of each variance measurement). Gray distributions are of mean counts. These appear to have different areas because of the vertical log scale. Black points are those preserved by mean-matching (*methods*). Black distributions are thus identical within the bin resolution. Black lines are regression slopes for the mean-matched data. **c.** The FF versus time. Arrows indicate time-points from the panels above. The gray trace (with flanking SEs) plots the ‘raw’ FF: the slope of the grey lines from panel *b*. The black trace plots the ‘mean-matched’ FF: the slope of the black lines.

Figure 5. Changes in the FF after restricting the analysis to neurons / conditions with little change in mean rate (e.g., non-preferred conditions). Of the original ‘neuron-conditions’ (the response of one neuron to one condition), this analysis preserved 28% (MT, panel *a*), 49% (PRR, panel *b*), 27% (PMd, panel *c*), and 41% (MT-speed, panel *d*). A 100 ms window (rather than 50 ms) was employed to regain lost statistical power. The trace at the top of each panel shows the mean rate, averaged across all included neurons/conditions. The trace with flanking standard errors shows the FF, computed with no further mean matching. Arrows indicate stimulus onset. For the MT-speed dataset (panel *d*) the stimulus appeared at the very start of the record (*first arrow*) and began moving 256 ms later (*second arrow*).

Figure 6. Application of FA to data from V1 and PMd. See *methods* and **Supplementary Figure 6** for a description of datasets. **a.** FA was applied to covariance matrices ($\#neurons \times \#neurons$) of spike counts, taken in an analysis window that either ended at stimulus onset (pre-stimulus) or began just after stimulus onset (stimulus). The measured covariance matrix was approximated as the sum of a ‘network’ covariance matrix and a diagonal matrix of private noise. To produce the plots in *b–g*, network variances were averaged across the subset of neurons/conditions (48% and 30% for V1, 74% and 79% for PMd) whose distribution of mean rates was matched before and after stimulus onset (similar to **Fig. 4**). **b.** Estimated variances for one V1 dataset. Network variability declined more than private variability in both absolute ($p < 10^{-7}$) and relative (percent of initial value, $p < 10^{-7}$) terms (paired t-tests across conditions). **c.** similar plot for a V1 dataset from a second monkey: $p < 0.002$ (absolute) and $p < 0.002$ (relative). **d.** Summary comparison for V1. Changes in variability (stimulus minus pre-stimulus) were expressed in percentage terms. Data to the left of zero indicate that network variability underwent the larger decline. The distribution includes all conditions and both datasets. The mean and standard error is given by the black symbol at top ($p < 10^{-7}$ compared to zero, paired t-test). Gray symbols give individual means for each dataset. **e.** Same as *b,c* but for one PMd dataset (G20040123). Network variability declined more in absolute ($p < 0.005$) and relative ($p < 0.001$) terms. **f.** Similar plot for a second PMd dataset (G20040122): $p < 0.05$ (absolute) and $p < 0.02$ (relative). **g.** Summary comparison for PMd (distribution mean < 0 ; $p < 10^{-4}$). **h.** Relationship between mean firing rate and network-level (shared firing-rate) variance. Data (same dataset as panel *b*) were binned by mean rate, and the average network variance ($\pm SE$) was computed for each bin. This was done both pre-stimulus (gray) and after stimulus onset (black). The average was taken across neurons and conditions (each datum being averaged was, for one condition, one element of the blue diagonal in panel *a*). Distributions of mean rates are shown at bottom. The analysis in *b* was based on the overlapping (mean-matched) portion of these distributions. **i.** Similar plot for PMd (same dataset as panel *e*).

Figure 7. Individual-trial neural trajectories, computed using GPFA. **a.** Projections of PMd activity into a 2-dimensional state space. Each black point represents the location of neural activity on one trial. Gray traces show trajectories from 200 ms before target onset until the indicated time. The stimulus was a reach target (135°, 60mm distant), with no reach allowed until a subsequent go cue (see *b,c*). Fifteen (of 47) randomly-selected trials are shown. Same dataset as in **Figure 6e**. **b.** Trajectories are plotted until movement onset. Blue dots indicate 100 ms before stimulus (reach target) onset. No reach was allowed until after the go cue (green dots), 400–900 ms later. Activity between the blue and green dots thus relates to movement planning. Movement onset (black dots) was ~300 ms after the go cue. For display, eighteen randomly-selected trials are plotted, plus one hand-selected trial (red, trialID211). Covariance ellipses were computed across all 47 trials. This is a 2-dimensional projection of a 10 dimensional latent space. In the full space, the black ellipse is far from the edge of the blue ellipse. This projection was chosen to accurately preserve the relative sizes (on a per dimension basis) of the true 10-dimensional volumes of the ellipsoids. G20040123 dataset. **c.** Same presentation and target location as in *b*, but for data from another day’s dataset (G20040122, red trial: trialID793).

Figure 8. Projections of V1 activity into a 2-dimensional space. Blue, black and red traces show activity before, during and after stimulus presentation (a drifting 45° grating). Same dataset as **Figure 6c**. **a.** The mean trajectory and three trials picked by hand. The grey spot shows the average location of pre-stimulus activity. In a few cases (e.g., upper left portion of the rightmost panel), traces were moved *very* slightly apart to make it clear that they traveled in parallel rather than crossed. **b.** Trajectories after data were shuffled to remove correlated variability. 25 randomly-selected trials are plotted (lighter traces) along with the mean (saturated traces). **c.** Same as in **b**, but for the original unshuffled data.

Methods

Datasets

All methods were approved in advance by the relevant institutional animal care and use committees. Data were collected from seven cortical areas of the macaque monkey and one cortical area of the cat by members of nine laboratories. For most datasets, methods have been described previously^{5,11,29,54,55,56,57,58,59,60}. Briefly, for the intracellularly-recorded V1 datasets, stimuli were 100% contrast sine-wave gratings presented at different orientations. For the extracellularly-recorded V1 dataset, stimuli were 100% contrast sine-wave gratings (6.25 Hz, 1.3 cycles/°) drifting in one of 12 directions. For V4, stimuli were one or two oriented bars situated in the neuron's receptive field. For some conditions, similar bars appeared in the opposite hemifield. For the first MT dataset, stimuli were square-wave gratings superimposed to produce a plaid. These patterns contained a small amount of variable-contrast texture, which was identical for all trials of a given type within a recording session (making it sensible to compute the response variance across those trials). For the second MT dataset, stimuli were 0%-coherence random dots. These had exactly the same dot pattern for repeated trials within a session. For LIP, stimuli were two colored saccade targets, between which the monkey chose. Analysis was further segregated by the eventual choice and the estimated value of that choice to the monkey. E.g., one of the conditions corresponded to rightwards saccades to a red target of high reward probability. For PRR, stimuli were reach/saccade targets (modality instructed by color) both within and outside the neuron's receptive field. For PMd, stimuli were reach targets at one of 2 distances and 7 directions. For OFC, the task was similar to that used for LIP, but analysis was locked to reward delivery (after the saccade). Analysis was segregated based on the features that OFC neurons seemed most sensitive to: choice, whether that choice repeated the last choice, and whether reward was delivered. E.g., one of the conditions corresponded to rewarded saccades to a red target when that target was chosen last time. For the third and fourth MT datasets, stimuli were random dot patterns. For 'MT-direction/area', stimulus direction and dot-patch area varied across conditions. For the 'MT-speed' dataset, stimuli were presented at different speeds (1–128°/s) in the preferred and null directions. For these two datasets, the exact location of the individual dots within the aperture was

different on each trial. This presumably did inflate the FF by injecting additional across-trial variability. However, this appears to be a modest concern: robust FF declines were still observed. These two datasets were also unusual in that the stimulus (the dot patch) appeared at the moment the trial began (and data collection commenced). Motion onset occurred 256 ms later.

Extracellular V1, MT-direction/area, and MT-speed datasets were recorded from opioid-anaesthetized monkeys (*macaca fascicularis* or *macaca nemestrina*). Intracellular V1 recordings were made from anaesthetized cats. All others were made from awake monkeys (*macaca mulatta*) performing a task with controlled fixation. For the V4 dataset, the task was simply to fixate while stimuli were presented. For PMd, the task required a reach. For PRR the task required a reach or a saccade. For the remaining datasets the task required a saccade indicating a choice or perceptual discrimination. Tasks involved a delay between stimulus onset and when a response was allowed, with one exception. For the MT-dots dataset, the monkey was allowed to respond as soon as he had made the discrimination. Analysis was thus limited to reaction times longer than 250 ms to allow a sufficient analysis window. All tasks involved a period of enforced fixation prior to stimulus onset. For LIP this was not true for all trials, and analysis was thus restricted to those that did. Fixation enforcement makes it unlikely that changes in variability are indirectly due to saccades. Additionally, datasets recorded from anaesthetized/paralyzed animals showed the same changes in variability. Finally, it is unlikely that changes in the FF, with a latency of ~60 ms, could be indirectly due to saccades, which would incur a latency of 80–200 ms⁶¹ plus an additional visual delay of 40–80 ms before firing rates could be affected. For all analyses, each trial's data were aligned on stimulus onset.

Most datasets were recorded across multiple days (and for LIP and MT, more than one monkey) using single electrodes. Exceptions were the V1 and PMd datasets. (The following applies to the FF analysis. See below regarding FA and GPFA.) The V1 data (dataset 565L) were recorded in one day from a 96 channel electrode array (Cyberkinetics). Only single-unit isolations (44) were analyzed. The PMd dataset was recorded using a similar array. Data were pooled across 5 days to increase trial and neuron count (although many units were likely recorded repeatedly). Single-neuron

counts for the datasets used for the FF analysis were: 44 (V1), 49 (V4), 314 (MT plaids), 82 (MT dots), 43 (LIP), 175 (PRR), 73 (PMd), 116 (OFC), 47 (MT-direction/area), and 57 (MT-speed). The mean number of trials/condition was 100 (V1), 24 (V4), 15 (MT plaids), 88 (MT dots), 35 (LIP), 10 (PRR), 31 (PMd), 106 (OFC), 125 (MT-direction/area), and 14 (MT-speed). The total numbers of ‘neuron-trials’ (number of neurons times the number of trials per neuron) were 52,800 (V1), 7063 (V4), 106,007 (MT plaids), 7179 (MT dots), 8357 (LIP), 14,160 (PRR), 35,953 (PMd), 49,389 (OFC), 61,343 (MT-direction/area) and 12,797 (MT-speed).

FA requires simultaneously recorded data (e.g., array-based recordings from a single day). Simulations similar to that in **Supplementary Figure 5** convinced us that the V1 and PMd datasets used for the FF analyses had neuron/trial-counts barely sufficient to provide FA with sufficient statistical power. To increase statistical power we included multi-unit isolations in the analysis. We used the two PMd recording days with the most ($> \sim 50$) trials per condition (G20040123 & G20040122). For V1, we employed two datasets (567R and 106R). Both used the same stimuli as the dataset used for the FF (described above) but had higher unit/trial counts. The first (567R) had only 20 single-unit isolations, but many multi-unit isolations (102) and more trials/condition (200 rather than 100) relative to the dataset used for the FF analysis. The second dataset (106R) had a very large (172) number of primarily multi-unit isolations, and even more trials/condition (400). These datasets were thus well-suited to the FA analysis. Both these datasets exhibited very clear FF declines (**Supplementary Fig. 6**). However, the decision was made to keep the original FF analysis (**Fig. 3**) restricted to the single-unit-only dataset. While the FF can be reliably computed for multi-unit recordings (the sum of independent Poisson-distributed counts is Poisson-distributed), the mean-matching procedure is not guaranteed to work.

Fano factor

The FF was computed using matlab (The Mathworks) code available at (<http://www.stanford.edu/~shenoy/GroupCodePacks.htm>). Data for each neuron and condition are initially treated separately. A condition corresponds to multiple trials of the same stimulus. Spike counts were computed in a 50 ms sliding window moving in 10 ms

steps. We then computed the variance (across trials) and mean of the spike count. Thus, the data for neuron₁, condition₁, time₁ were reduced to two numbers: the variance and mean of the spike count in a window centered on that time. For each time, a scatterplot of the variance versus the mean was compiled, with one point per neuron/condition (typically a few thousand points). The ‘raw’ FF was the slope of the regression relating the variance to the mean. This regression was weighted according to the estimated sampling error for the variance (computed based on the number of trials and the mean rate, and assuming Poisson statistics). The observed effects are not due to the use of a short (50 ms) window: effect magnitudes were generally larger with larger windows, at the cost of blurred timecourses.

To compute the ‘mean-matched Fano factor’, we first computed the distribution of mean counts (one datum per neuron/condition) for each time being analyzed (grey distributions; **Fig. 4b**). We then computed the greatest common distribution present at all times (black distribution). Each bin of this common distribution had a height equal to the smallest value for that same bin across all distributions at all times. For each time, we then matched the analyzed distribution of mean rates to this common distribution. This was accomplished by discarding, from each bin, randomly chosen datapoints (each corresponding to the variance/mean pair for one neuron/condition) until the height of that bin matched that of the common distribution. The FF was then computed. This procedure was repeated 50 times with different random seeds, with a different subset of the data preserved on each repetition, and thus a slightly different resulting FF. The reported FF is the mean of these individual FFs.

This mean-matching method discards less data than might be expected, and far less than if we restrict analysis to individual neurons whose rates were not changing for particular conditions (as in **Fig. 5**). Mean matching preserved 49% (V1), 31% (V4), 71% (MT plaids), 32% (MT dots), 48% (LIP), 84% (PRR), 75% (PMd), 86% (OFC), 46% (MT-direction/area) and 70% (MT-speed) of the ‘neuron-conditions’. Prior studies using the FF have occasionally presented controls involving informal mean-matching: drawing attention to different values of the FF over a similar range of rates^{11,25,43}. The current method removes the need for such additional controls. Mean-matching also does more

than validate the basic effect; it prevents artifacts from corrupting the timecourse (**Supplementary Fig. 2**).

The statistical significance of the FF decline was assessed by comparing the FF 100 ms before stimulus onset and 200 ms after (for the MT-direction/area and MT-speed datasets, where the FF underwent two declines, we compared the time of motion onset with 200 ms after). Significance was computed based on sampling distributions estimated from the 95% confidence intervals yielded by the regression used to compute the FF. The decline was statistically significant in all cases: $p < 0.001$ (V1), $p < 10^{-5}$ (V4), $p < 10^{-10}$ (MT plaids), $p < 0.01$ (MT dots), $p < 0.005$ (LIP), $p < 0.002$ (PRR), $p < 10^{-10}$ (PMd), $p < 0.02$ (OFC), $p < 10^{-14}$ (MT-direction/area), and $p < 10^{-7}$ (MT-speed).

Factor analysis

FA was applied to spike counts in a 400 ms window, which either ended at stimulus onset ('Pre-stimulus') or began 100 ms after stimulus onset ('Stimulus'). For the 2nd V1 dataset (**Fig. 6c**), stimulus duration was only 400 ms, so we used a shorter (350 ms) window that began 50 ms after stimulus onset. A data matrix, D , was compiled, with D_{ij} being the spike count for the i^{th} trial and the j^{th} neuron. D included only neurons with rates ≥ 1 spike/s for both periods. FA was then applied to this matrix (neurons are variables, trials are observations). FA was performed separately on the D matrices compiled before and after stimulus onset. FA was performed separately for each condition. Subsequent results were averaged. We used code written by Z. Ghahramani²⁷ to fit the FA parameters using expectation-maximization. Essentially identical results were obtained using matlab's 'factoran'. FA requires an estimate of the dimensionality of the latent space capturing shared variability. To estimate this, we exploited the fact that recordings were made for multiple stimulus conditions (14 for PMd, 12 and 8 for the two V1 datasets). We computed the mean (across time and trials) response for each condition. We also included the pre-stimulus period as a condition. We then applied PCA (to the $c \times n$ matrix) and asked how many dimensions were necessary to account for 95% of the *across-condition* response variance (PCA is acceptable because it is being applied to trial-averaged data). For PMd and V1, 5 and 4 dimensions were required. These were then the latent dimensionalities assumed by FA. These numbers likely

underestimate of the true dimensionality of the network response across conditions (our stimuli spanned only part of the space of all stimuli). This underestimate is acceptable: it implies that FA will underestimate the drop in the network variance, making our current interpretation conservative (**Supplementary Fig. 7**). Finally, similar results were obtained across a range of reasonable latent dimensionalities (2–8).

We used FA to estimate the private and shared component of each neuron’s variance (that is, we computed but did not analyze the full network covariance matrix). This kept each variance associated with a mean rate, and allowed us to apply the mean-matching procedure used for the FF. Mean-matching aids interpretation because 1) even when rates are rising dramatically, we can ask whether rate variance is changing for a given mean rate 2) there is a clear expectation that the estimated private noise should change little (assuming effects are not due to private noise), and 3) there is less concern that a large increase in private noise (which will occur with rising mean rates) might contaminate the network variance estimate. The percentage of neuron/conditions preserved following mean-matching was 31% and 48% (two V1 datasets) and 74% and 79% (two PMd datasets).

The ‘estimated variances’ plotted in **Figure 6b,c,e,f** are the mean network (shared) variances, averaged across neurons/conditions, after mean-matching. Significance was tested via t-test (n = number of conditions).

Gaussian-process factor analysis

GPFA³³ was used to extract temporally-smooth single-trial neural trajectories. Spike counts were taken in non-overlapping 20 ms bins, then square-rooted. The square-root transform⁶² allows GPFA (which assumes stationary private noise for each neuron) to handle what would otherwise be a time-varying scaling of private noise with changing rates. Unlike FA, GPFA was applied simultaneously to data from all timepoints. It was critical that FA was applied separately before and after stimulus onset, to allow it to capture any changes in private noise. For GPFA we require a visualization in which the same latent space captures the neural trajectory across time. Using all timepoints also increases the quantity of analyzable data. The cost is that GPFA, unlike the FA analysis, cannot account for changes in private noise with time. However, such changes are likely

modest (**Fig. 6b,c,e,f**) and are not the focus of the GPFA analysis (which seeks to track changes in network state, not private noise).

Figures 7 and 8 plot 2-dimensional projections of the higher-dimensional trajectories found by GPFA. For PMd, subsequent analysis quantified ellipsoid volumes at different times. This analysis employed the higher-dimensional latent space. The dimensionality of that space was estimated using cross-validation. This is reasonable for GPFA, relative to FA (where we did not use cross-validation), because of the increase in datapoints gained by including different times. Across the 14 reach targets, each analyzed separately, the latent dimensionality ranged from 8 to 12. It is expected that this dimensionality is higher than that for FA, as the latter needed only to account for variability at a single time, while GPFA must account for variability across both trials and time. GPFA produced similar results when we used a dimensionality smaller (6) or higher (15) than that chosen by cross validation. Because dimensionality impacts the covariance volume, we report that volume per dimension (i.e., $\text{volume}^{1/k}$ where k is the latent dimensionality; k was always the same pre- and post-target).

For PMd, the 2D projections in **Figure 7** were chosen to preserve the relative per-dimensional volumes of the 8–12 dimensional covariance ellipses. For V1, we simply plotted the top two dimensions returned by GPFA, which maximizes the captured variability (analogous to the first two principal components).

References

- ¹ Briggman, K. L., Abarbanel, H. D. & Kristan, W. B., Jr. Optical imaging of neuronal populations during decision-making. *Science* **307**, 896-901 (2005).
- ² Arieli, A., Sterkin, A., Grinvald, A. & Aertsen, A. Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science* **273**, 1868-1871 (1996).
- ³ Kenet, T., Bibitchkov, D., Tsodyks, M., Grinvald, A. & Arieli, A. Spontaneously emerging cortical representations of visual attributes. *Nature* **425**, 954-956 (2003).
- ⁴ Monier, C., Chavane, F., Baudot, P., Graham, L. J. & Fregnac, Y. Orientation and direction selectivity of synaptic inputs in visual cortical neurons: a diversity of combinations produces spike tuning. *Neuron* **37**, 663-680 (2003).
- ⁵ Finn, I. M., Priebe, N. J. & Ferster, D. The emergence of contrast-invariant orientation tuning in simple cells of cat visual cortex. *Neuron* **54**, 137-152 (2007).
- ⁶ Monier, C., Fournier, J. & Fregnac, Y. In vitro and in vivo measures of evoked excitatory and inhibitory conductance dynamics in sensory cortices. *J Neurosci Methods* **169**, 323-365, doi:S0165-0270(07)00554-7 [pii] 10.1016/j.jneumeth.2007.11.008 (2008).
- ⁷ Kohn, A., Zandvakili, A. & Smith, M. A. Correlations and brain states: from electrophysiology to functional imaging. *Curr Opin Neurobiol* **19**, 434-438, doi:S0959-4388(09)00063-4 [pii] 10.1016/j.conb.2009.06.007 (2009).
- ⁸ Azouz, R. & Gray, C. M. Cellular mechanisms contributing to response variability of cortical neurons in vivo. *J Neurosci* **19**, 2209-2223 (1999).
- ⁹ Fiser, J., Chiu, C. & Weliky, M. Small modulation of ongoing cortical dynamics by sensory input during natural vision. *Nature* **431**, 573-578 (2004).
- ¹⁰ Kisley, M. A. & Gerstein, G. L. Trial-to-trial variability and state-dependent modulation of auditory-evoked responses in cortex. *J Neurosci* **19**, 10451-10460 (1999).
- ¹¹ Churchland, M. M., Yu, B. M., Ryu, S. I., Santhanam, G. & Shenoy, K. V. Neural variability in premotor cortex provides a signature of motor preparation. *J Neurosci* **26**, 3697-3712 (2006).
- ¹² Rickert, J., Riehle, A., Aertsen, A., Rotter, S. & Nawrot, M. P. Dynamic encoding of movement direction in motor cortical neurons. *J Neurosci* **29**, 13870-13882, doi:29/44/13870 [pii] 10.1523/JNEUROSCI.5441-08.2009 (2009).
- ¹³ Sussillo, D. & Abbott, L. F. Generating coherent patterns of activity from chaotic neural networks. *Neuron* **63**, 544-557 (2009).
- ¹⁴ Abbott, L. F., Rajan, K. & Sompolinsky, H. in *Neuronal Variability and its Functional Significance* eds M. Ding & D. Glanzman) (in press).
- ¹⁵ Poulet, J. F. & Petersen, C. C. Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. *Nature* **454**, 881-885 (2008).
- ¹⁶ Shadlen, M. N. & Newsome, W. T. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci* **18**, 3870-3896 (1998).

- 17 van Vreeswijk, C. & Sompolinsky, H. Chaos in neuronal networks with balanced
excitatory and inhibitory activity. *Science* **274**, 1724-1726 (1996).
- 18 Mainen, Z. F. & Sejnowski, T. J. Reliability of spike timing in neocortical
neurons. *Science* **268**, 1503-1506 (1995).
- 19 Gur, M., Beylin, A. & Snodderly, D. M. Response variability of neurons in
primary visual cortex (V1) of alert monkeys. *J Neurosci* **17**, 2914-2920 (1997).
- 20 Carandini, M. Amplification of trial-to-trial response variability by neurons in
visual cortex. *PLoS Biol* **2**, E264 (2004).
- 21 Bradley, A., Skottun, B. C., Ohzawa, I., Sclar, G. & Freeman, R. D. Visual
orientation and spatial frequency discrimination: a comparison of single neurons
and behavior. *J Neurophysiol* **57**, 755-772 (1987).
- 22 Tolhurst, D. J., Movshon, J. A. & Dean, A. F. The statistical reliability of signals
in single neurons in cat and monkey visual cortex. *Vision Res* **23**, 775-785 (1983).
- 23 Vogel, A., Hennig, R. M. & Ronacher, B. Increase of neuronal response
variability at higher processing levels as revealed by simultaneous recordings. *J
Neurophysiol* **93**, 3548-3559 (2005).
- 24 Nawrot, M. P. *et al.* Measurement of variability dynamics in cortical spike trains.
J Neurosci Methods **169**, 374-390 (2008).
- 25 Mitchell, J. F., Sundberg, K. A. & Reynolds, J. H. Differential attention-
dependent response modulation across cell classes in macaque visual area V4.
Neuron **55**, 131-141 (2007).
- 26 Borg-Graham, L. J., Monier, C. & Fregnac, Y. Visual input evokes transient and
strong shunting inhibition in visual cortical neurons. *Nature* **393**, 369-373,
doi:10.1038/30735 (1998).
- 27 Roweis, S. & Ghahramani, Z. A unifying review of linear gaussian models.
Neural Comput **11**, 305-345 (1999).
- 28 Everitt, B. S. *An introduction to latent variable models.* (Chapman & Hall, 1984).
- 29 Smith, M. A. & Kohn, A. Spatial and temporal scales of neuronal correlation in
primary visual cortex. *J Neurosci.* **28**, 12591-12603 (2008).
- 30 Davies, R. M., Gerstein, G. L. & Baker, S. N. Measurement of time-dependent
changes in the irregularity of neural spiking. *J Neurophysiol* **96**, 906-918 (2006).
- 31 Williams, Z. M., Elfar, J. C., Eskandar, E. N., Toth, L. J. & Assad, J. A. Parietal
activity and the perceived direction of ambiguous apparent motion. *Nat Neurosci*
6, 616-623 (2003).
- 32 Churchland, M. M., Afshar, A. & Shenoy, K. V. A central source of movement
variability. *Neuron* **52**, 1085-1096 (2006).
- 33 Yu, B. M. *et al.* Gaussian-process factor analysis for low-dimensional single-trial
analysis of neural population activity. *J Neurophysiol* **102**, 614-635 (2009).
- 34 Stopfer, M., Jayaraman, V. & Laurent, G. Intensity versus identity coding in an
olfactory system. *Neuron* **39**, 991-1004 (2003).
- 35 Levi, R., Varona, P., Arshavsky, Y. I., Rabinovich, M. I. & Selverston, A. I. The
role of sensory network dynamics in generating a motor program. *J Neurosci* **25**,
9807-9815 (2005).
- 36 Britten, K. H., Newsome, W. T., Shadlen, M. N., Celebrini, S. & Movshon, J. A.
A relationship between behavioral choice and the visual responses of neurons in
macaque MT. *Vis Neurosci* **13**, 87-100 (1996).

- 37 Medina, J. F. & Lisberger, S. G. Variation, signal, and noise in cerebellar sensory-
motor processing for smooth-pursuit eye movements. *J Neurosci* **27**, 6832-6842
(2007).
- 38 McAdams, C. J. & Maunsell, J. H. Effects of attention on the reliability of
individual neurons in monkey visual cortex. *Neuron* **23**, 765-773 (1999).
- 39 Horwitz, G. D. & Newsome, W. T. Target selection for saccadic eye movements:
prelude activity in the superior colliculus during a direction-discrimination task. *J*
Neurophysiol **86**, 2543-2558 (2001).
- 40 Mandelblat-Cerf, Y., Paz, R. & Vaadia, E. Trial-to-trial variability of single cells
in motor cortices is dynamically modified during visuomotor adaptation. *J*
Neurosci **29**, 15053-15062, doi:29/48/15053 [pii]
10.1523/JNEUROSCI.3011-09.2009 (2009).
- 41 Kao, M. H., Doupe, A. J. & Brainard, M. S. Contributions of an avian basal
ganglia-forebrain circuit to real-time modulation of song. *Nature* **433**, 638-643
(2005).
- 42 Oram, M. W., Hatsopoulos, N. G., Richmond, B. J. & Donoghue, J. P. Excess
synchrony in motor cortical neurons provides redundant direction information
with that from coarse temporal measures. *J Neurophysiol* **86**, 1700-1716 (2001).
- 43 Kara, P., Reinagel, P. & Reid, R. C. Low response variability in simultaneously
recorded retinal, thalamic, and cortical neurons. *Neuron* **27**, 635-646 (2000).
- 44 Osborne, L. C., Bialek, W. & Lisberger, S. G. Time course of information about
motion direction in visual area MT of macaque monkeys. *J Neurosci* **24**, 3210-
3222 (2004).
- 45 Lee, D. & Seo, H. in *Neuronal variability and its functional significance* eds D.
Glanzman & M. Ding) (Oxford Press, in press).
- 46 Friedrich, R. W. & Laurent, G. Dynamics of olfactory bulb input and output
activity during odor stimulation in zebrafish. *J Neurophysiol* **91**, 2658-2669
(2004).
- 47 Fortier, P. A., Smith, A. M. & Kalaska, J. F. Comparison of cerebellar and motor
cortex activity during reaching: directional tuning and response variability. *J*
Neurophysiol **69**, 1136-1149 (1993).
- 48 Cohen, J. Y. *et al.* Difficulty of visual search modulates neuronal interactions and
response variability in the frontal eye field. *J Neurophysiol* **98**, 2580-2587 (2007).
- 49 Nauhaus, I., Busse, L., Carandini, M. & Ringach, D. L. Stimulus contrast
modulates functional connectivity in visual cortex. *Nat Neurosci* **12**, 70-76
(2009).
- 50 Werner, G. & Mountcastle, V. B. The Variability Of Central Neural Activity In A
Sensory System, And Its Implications For The Central Reflection Of Sensory
Events. *J Neurophysiol* **26**, 958-977 (1963).
- 51 Wang, X. J. Decision making in recurrent neuronal circuits. *Neuron* **60**, 215-234,
doi:S0896-6273(08)00836-2 [pii]
10.1016/j.neuron.2008.09.034 (2008).
- 52 Churchland, A. C., Kiani, R. & Shadlen, M. N. Time course of response
variability in area LIP in a two- and four-choice decision task. *Society for*
Neuroscience Abstracts (2006).

- 53 Churchland, M. M., Yu, B. M., Sahani, M. & Shenoy, K. V. Techniques for
extracting single-trial activity patterns from large-scale neural recordings. *Curr*
Opin Neurobiol **17**, 609-618 (2007).
- 54 Sugrue, L. P., Corrado, G. S. & Newsome, W. T. Matching behavior and the
representation of value in the parietal cortex. *Science* **304**, 1782-1787 (2004).
- 55 Cohen, M. R. & Newsome, W. T. Context-dependent changes in functional
circuitry in visual area MT. *Neuron* (in press).
- 56 Armstrong, K. M., Fitzgerald, J. K. & Moore, T. Changes in visual receptive
fields with microstimulation of frontal cortex. *Neuron* **50**, 791-798 (2006).
- 57 Armstrong, K. M. & Moore, T. Rapid enhancement of visual cortical response
discriminability by microstimulation of the frontal eye field. *Proc Natl Acad Sci*
U S A **104**, 9499-9504 (2007).
- 58 Chang, S. W., Dickinson, A. R. & Snyder, L. H. Limb-specific representation for
reaching in the posterior parietal cortex. *J Neurosci* **28**, 6128-6140 (2008).
- 59 Bradley, D. C. *et al.* A distributed neural correlate of visual motion segmentation.
(submitted).
- 60 Priebe, N. J., Churchland, M. M. & Lisberger, S. G. Constraints on the source of
short-term motion adaptation in macaque area MT. I. the role of input and
intrinsic mechanisms. *J Neurophysiol* **88**, 354-369 (2002).
- 61 Boch, R. & Fischer, B. Saccadic reaction times and activation of the prelunate
cortex: parallel observations in trained rhesus monkeys. *Exp Brain Res* **50**, 201-
210 (1983).
- 62 Kihlberg, J. K., Herson, J. H. & Schotz, W. E. Square root transformation
revisited. *Appl. Statist.* **21**, 76-81 (1972).




figure 1

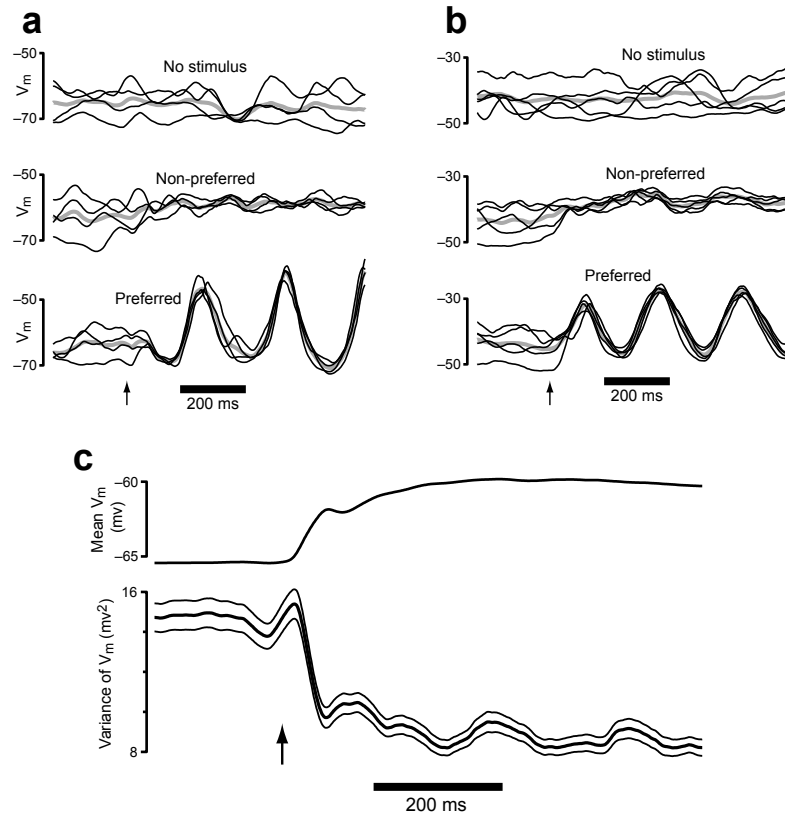


figure 2

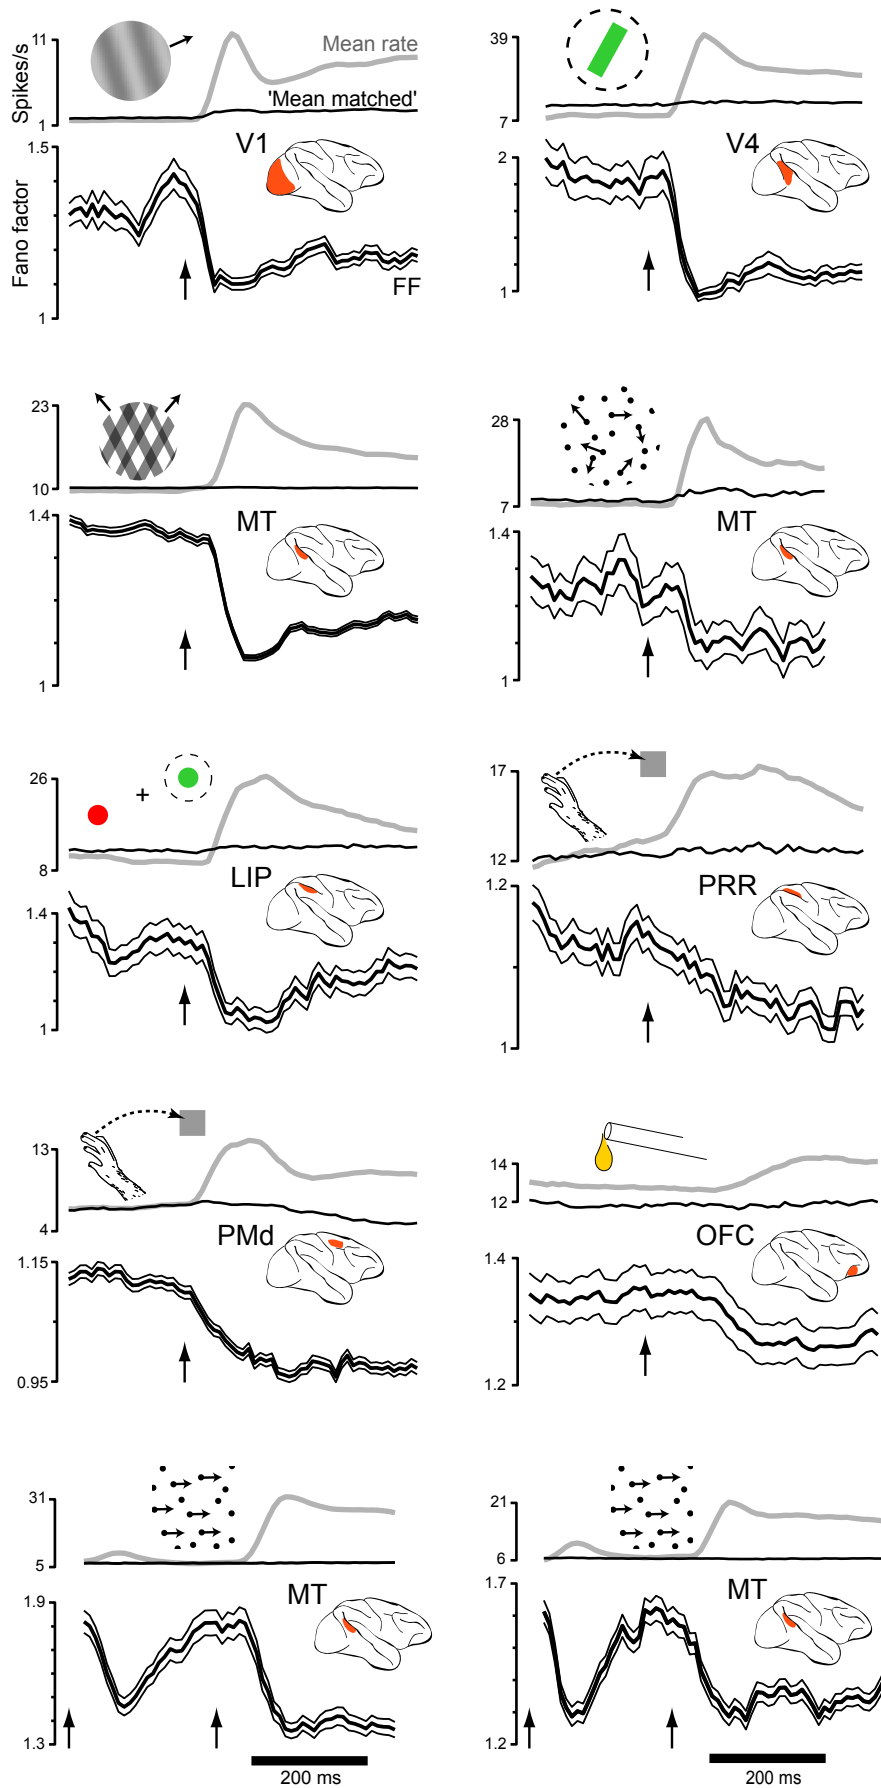


figure 3

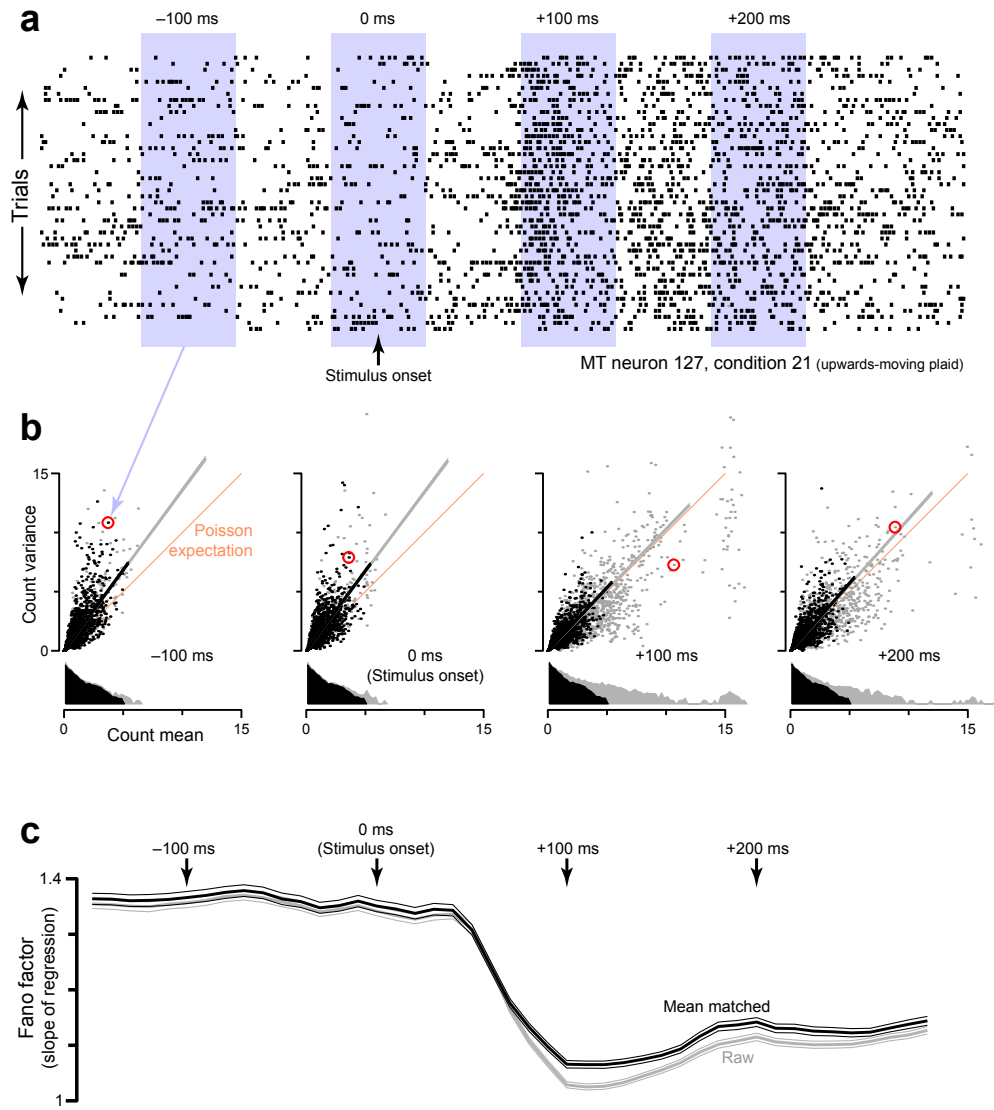


figure 4

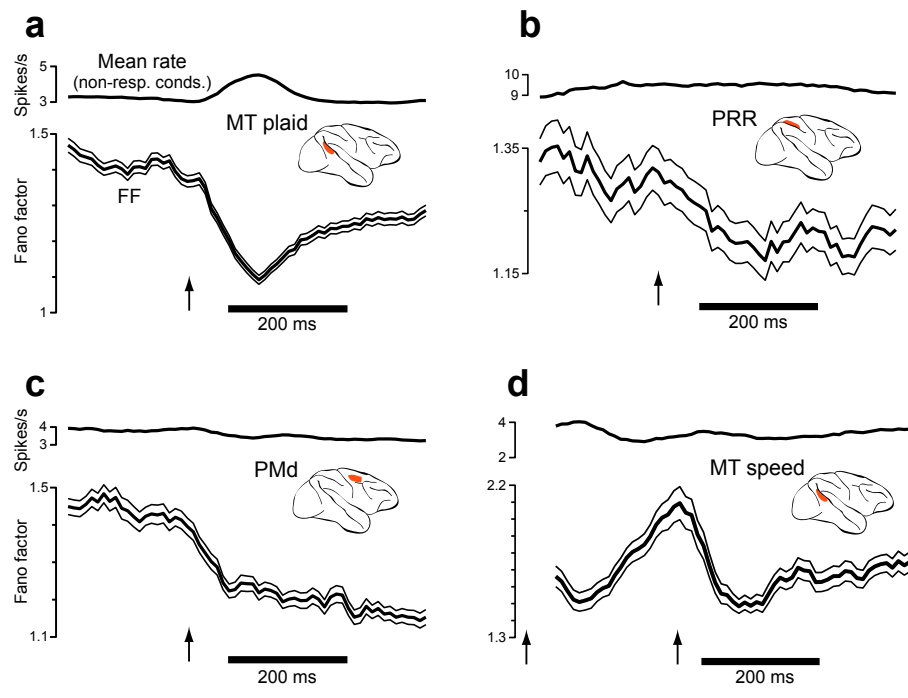


figure 5

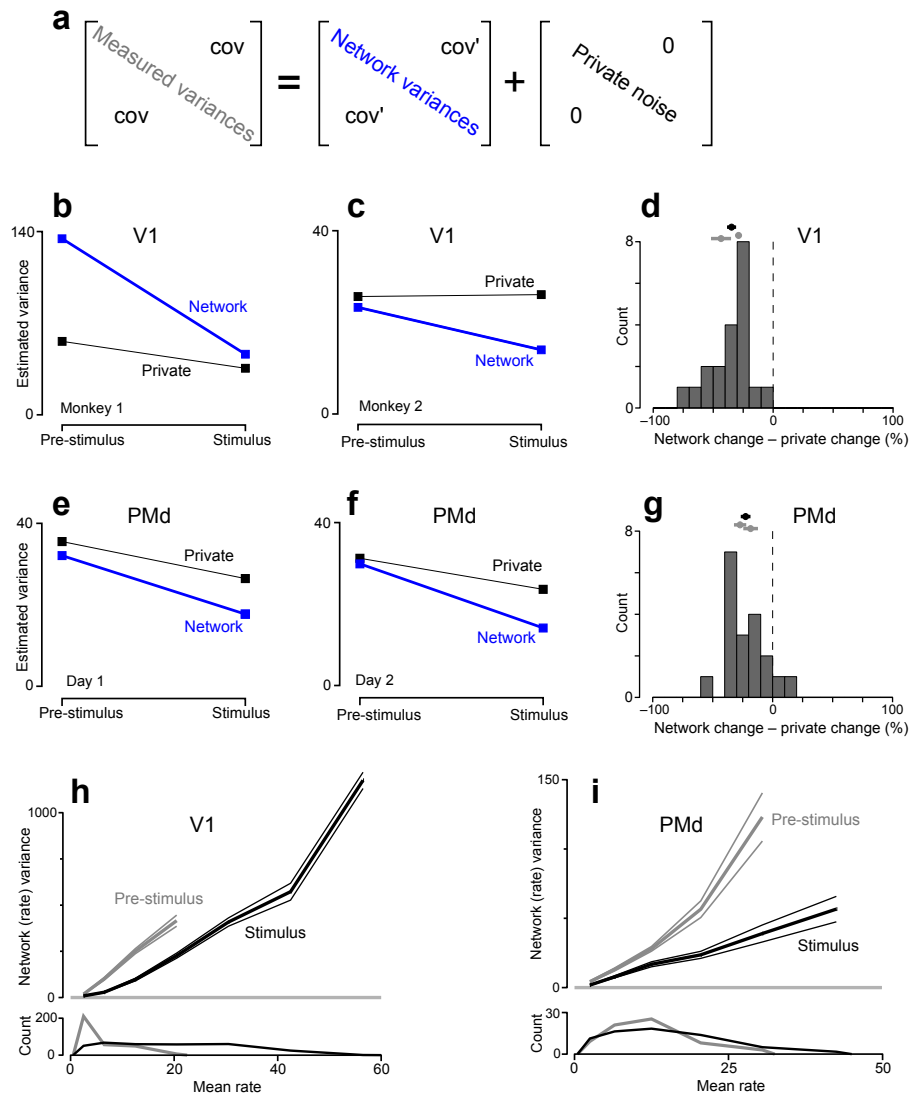


figure 6

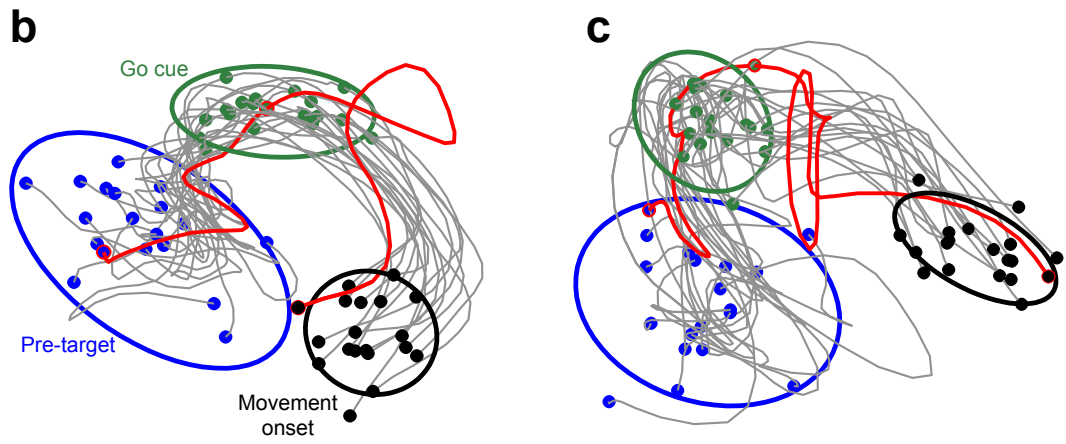



figure 7

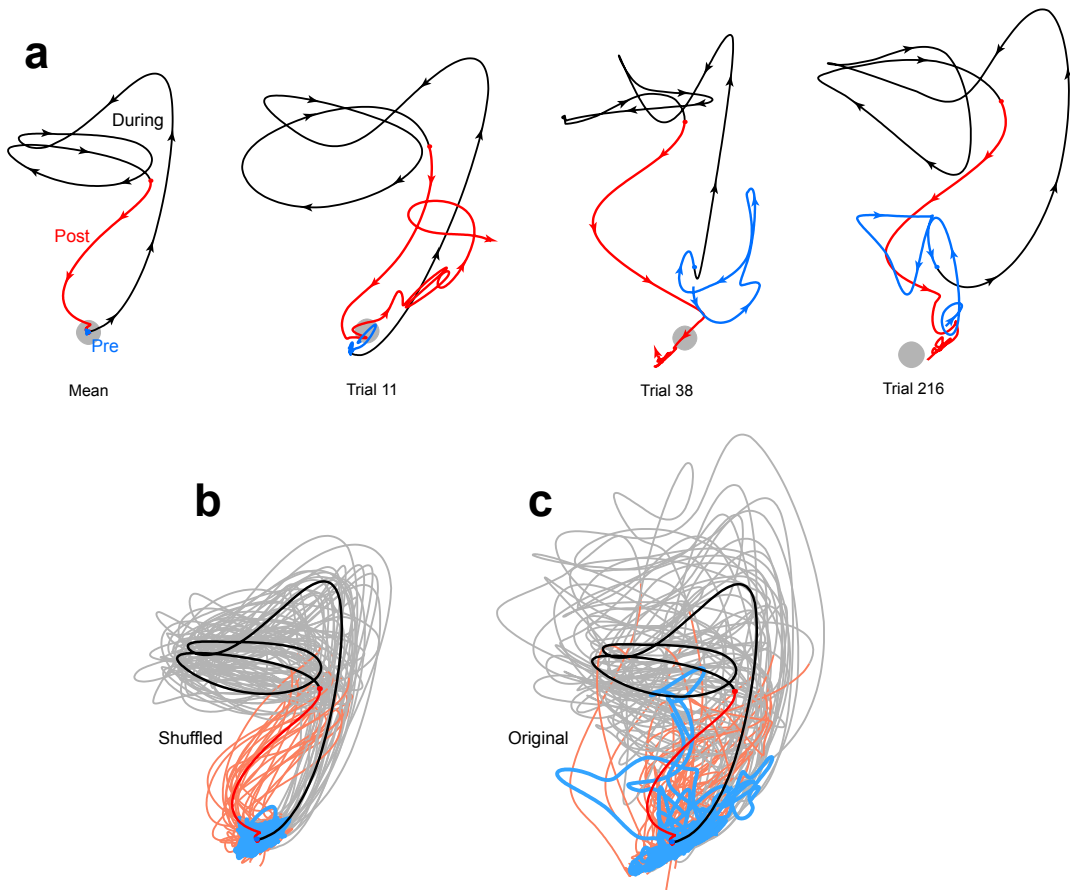
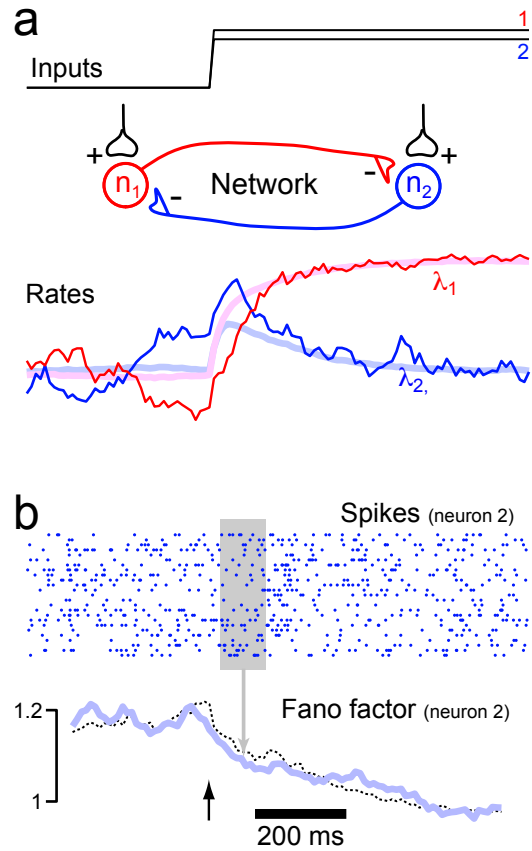


figure 8

Supplementary Materials

Stimulus onset quenches neural variability: a widespread cortical phenomenon

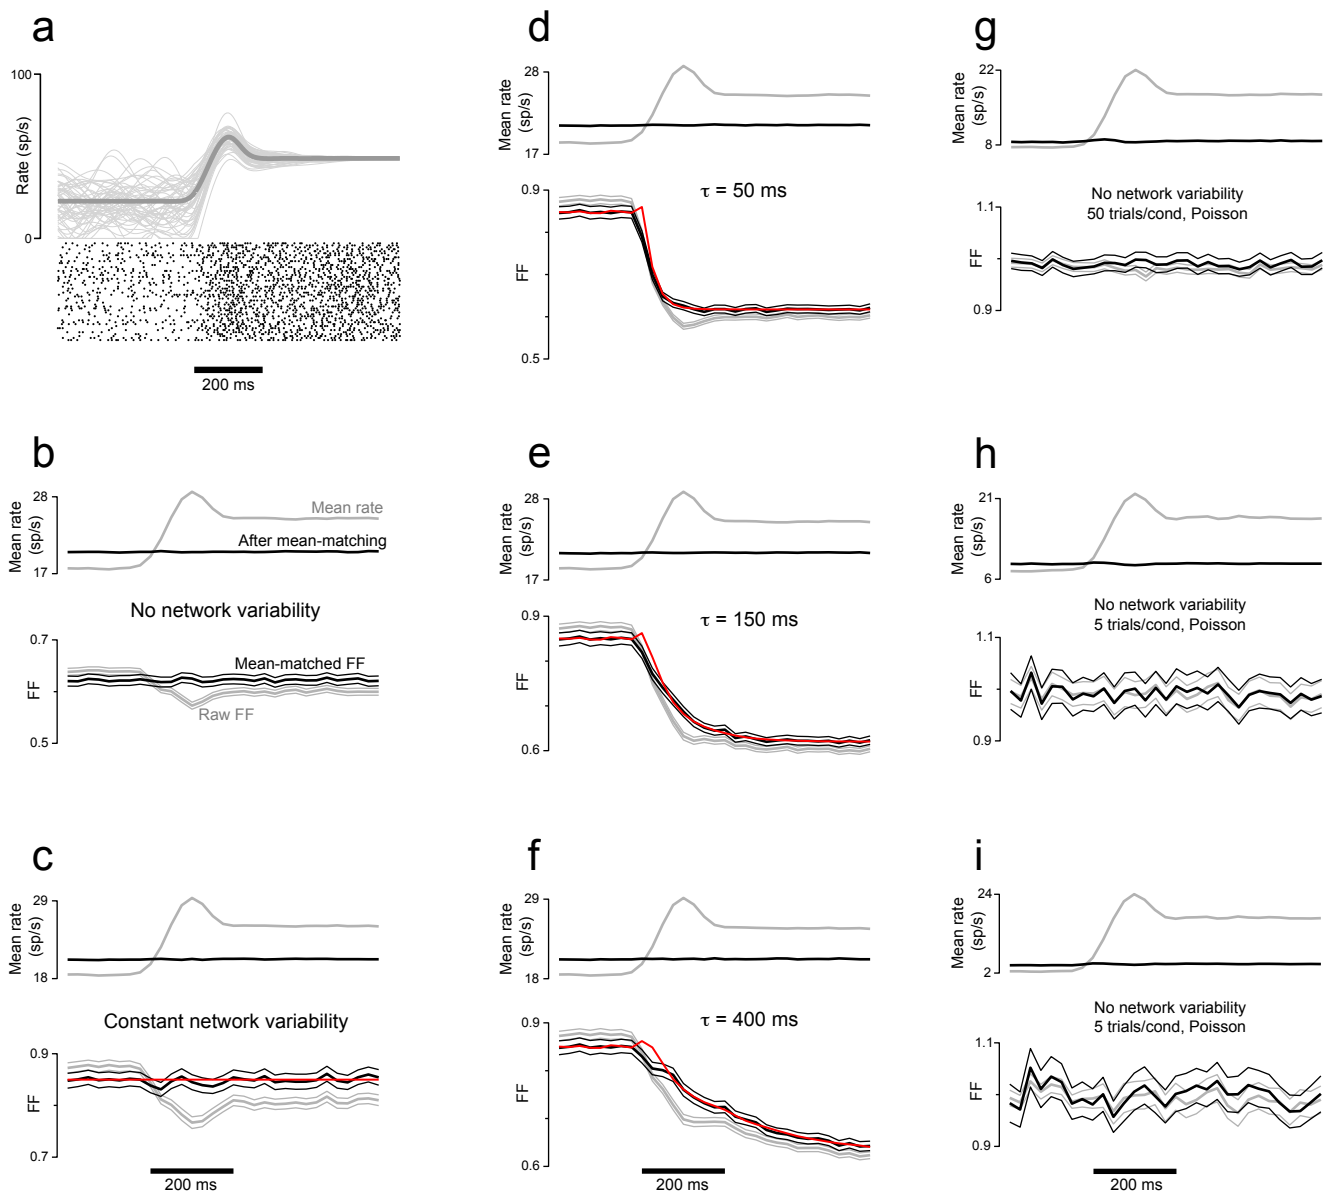
Churchland MM**, Yu BM**, Cunningham JP, Sugrue LP, Cohen MR, Corrado GS, Newsome WT, Clark AM, Hosseini P, Scott BB, Bradley DC, Smith MA, Kohn A, Movshon JA, Armstrong KM, Moore T, Chang SW, Snyder LH, Priebe NJ, Finn IM, Ferster D, Ryu SI, Santhanam G, Sahani M, and Shenoy KV



supplementary figure 1
(stimulus-driven changes in variability for a simple simulation)

Supplementary figure 1. Stimulus-driven changes in across-trial variability can be exhibited even by simple network architectures. Details of the simulations are described in the supplementary materials below. **a.** A simulated two-neuron recurrent network. The inputs were initially zero, and stepped to new values at the indicated time. The evolution of rates (λ) was different on each trial, as the network contained noise. One example trial is shown (*saturated colored traces*). Light, thicker traces show the mean λ across all trials. **b.** Simulated spike-trains were produced via an inhomogeneous Poisson process based on each trial's λ_2 . 25 trials are shown. The FF (*blue, bottom*) was computed from 10,000 such trials, using a 100 ms sliding window (*gray shading*). The dashed trace plots the true across-trial variance of λ_2 . The black arrow indicates the step input.

In the absence of an input, this network has two shallow competing attractors. It is thus unstable given small amounts of noise, but becomes more stable when driven. As a consequence, average rates (*thick light-colored traces* in *a*) are initially unrepresentative of single-trial firing rates (*colored traces* give each neuron's rate for that trial). Following the input step, single-trial rates hew more closely to their means; across-trial variability is reduced. That decline depends on the attractor dynamics of this example 'winner-take-all' network. A network with different dynamics – e.g., integrator dynamics^{1,2,3} – might produce rising variance. As in Figure 1*a* of the text, mean rate alone may be inadequately informative. For example, neuron 2 exhibits only a modest mean response (*thick blue trace*). This might (mistakenly) suggest weak participation in the network computation. In contrast, the sustained variance decline reveals the involvement of both neurons, and correctly suggests attractor dynamics.



supplementary figure 2
(the Fano factor for simulated data)

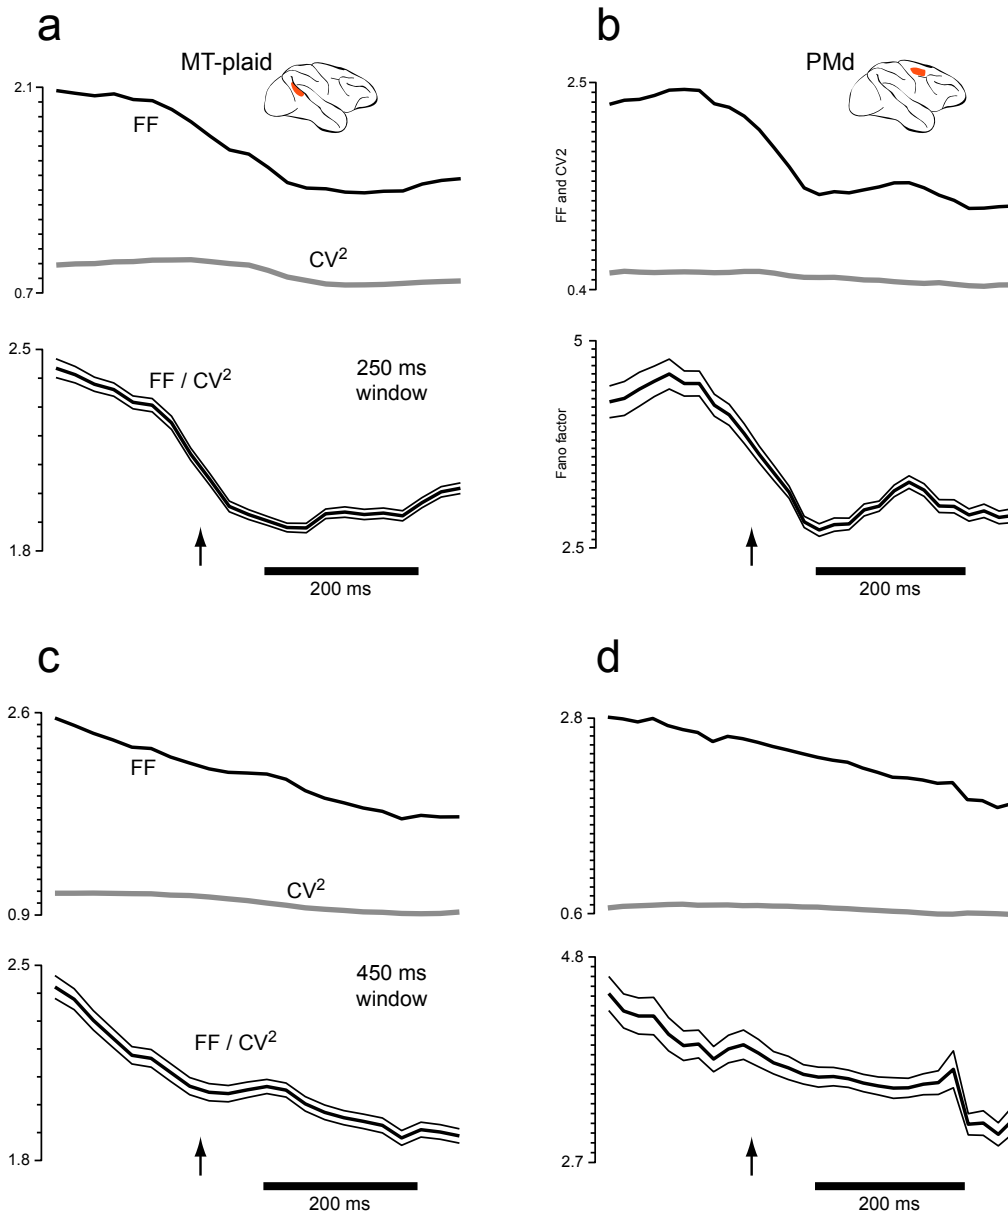
Supplementary figure 2. Verification of the mean-matching method for simulated data. **a.** Illustration of how simulated data was produced. Each neuron had a mean rate (*thick gray trace*) that underwent a change to a new value at 400 ms (left-hand edge of *calibration bar*). On each simulated trial (50 in this case), the underlying rate (thin grey traces) differed from the mean. Beginning at 400 ms, this rate variance collapsed (with a time constant of 150 ms for this simulation). On each trial, the underlying rate provided the λ for a renewal process that was used to produce spikes (rasters). That process obeyed either Poisson, Poisson-with-refractory, or gamma-interval (order 2) spiking statistics.

Subsequent panels demonstrate the behavior of the ‘raw’ FF and the mean-matched FF under a variety of circumstances. Gamma-interval spiking statistics provide the sternest test for the FF, as they most seriously violate the assumption that the spike-count variance scales linearly with the mean. Thus, to demonstrate the robustness of the method, the simulations in panels *b-f* employ gamma spiking statistics. Data similar to that in *a* were produced for 2000 simulated neuron-conditions (the equivalent of 20 conditions for 100 neurons). Each neuron’s mean rate underwent a change at 400 ms, with the initial and final mean rates being drawn from uniform distributions over the intervals [0 35] and [0 50] spikes/s. Thus, individual-neuron mean rates could either rise or fall, though the overall tendency was towards increasing rates. Additional trial-to-trial variability, of a given magnitude and with a given time-constant, was then added to the mean rate (using the procedure that produced the thin grey traces in *panel a*). This was done independently for each neuron-condition. **b.** Application of the method to a simulation with no variability in underlying rate (i.e., where the thin gray traces in *a* were all superimposed on the mean). The mean-matched FF (*black*, with flanking SEs) correctly reports little change in variance. In contrast, the raw FF (*grey*) incorrectly reports a small drop, due to the gamma-interval spiking statistics, which become more regular at higher rates. The grey trace at top plots the ‘measured’ mean rate across all simulated neurons / conditions. The black trace plots the mean rate after applying the mean-matching procedure. **c.** Application of the method to a simulation with constant variability in the underlying rate (i.e., the thin traces in *a* never converge). The mean-matched FF correctly reports little change in underlying variance. The raw FF is again plagued by a small artifact. The red trace plots the true (unchanging in this case) timecourse of underlying rate variability. **d-f.** Similar plots, but for different time-constants for the convergence of the underlying rates. The mean-matched FF approximately captures the true (*red*) timecourse. The raw FF performs reasonably well, but tends – due to the rising rates and gamma-interval spiking statistics – to overestimate the steepness of the initial drop. In summary,

even when the Poisson assumption is violated, the mean-matched FF avoids artifacts that impact the raw FF (this was also true when we used an absolute refractory period, rather than gamma spiking statistics). This allows the mean-matched FF to accurately capture underlying changes (or lack thereof) in rate variance.

Panels g-i present an exploration of how the FF is affected by trial count. For these simulations, there was no variability in the underlying rate, and spiking statistics were Poisson. If well behaved, the FF should thus remain near unity for all times. As above, firing rates changed at 400 ms, and 2000 ‘neuron-conditions’ were employed, with a range of initial and final mean firing rates. **g.** Initial and final mean-rate distributions (across neuron-conditions) were drawn from the intervals [0 15] and [0 35]. We simulated 50 trials per neuron-condition. **h.** Same as **g** but with 5 trials per neuron-condition. **i.** Same as **h** but with a lower initial range of rates, [0 5].

Inspection of panels *g-i* reveals that neither the FF (*gray*) nor the mean-matched FF (*black*) is biased by a low trial count, or by a low initial range of rates. In particular, the FF is not biased upwards for low rates / trial counts. One might have expected this, as the denominator will often be near zero, but in those cases the numerator is similarly small. That said, when both trial counts and spike rates are very low, mean-matching is not guaranteed to completely eliminate artifacts due to non-Poisson spiking. Estimates of the mean are then poor, and mean-matching becomes only approximate. In this case, the mean-matched FF will show smaller artifactual changes than will the raw FF, but it may not behave perfectly. For example, for gamma order-2 spiking, with low initial rates (0-5 spikes/s) and low trial counts (5), the raw FF will of course decline sharply with rising rates, and the mean-matched FF can actually show a slight *rise*. However, we have never seen this pattern in real data, and all the datasets we analyzed had an average of at least 10 trials per condition, and often 100 (V1) or more (the V1 datasets used for FA, the FF for which is plotted in Supplementary figure 6).



supplementary figure 3
 (FF and the CV² of the ISIs)

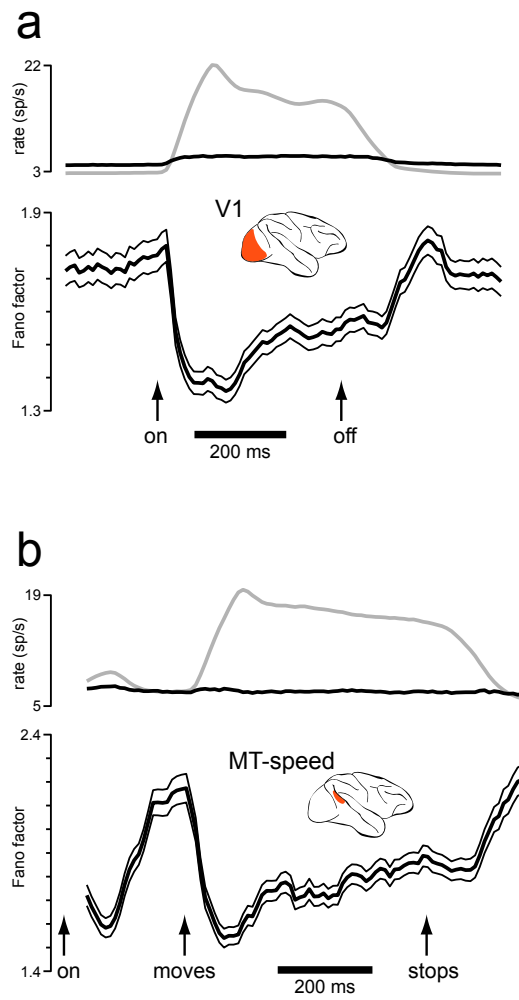
Supplementary figure 3. An additional analysis to determine whether the decline in the Fano factor (FF) is due to declining firing rate variability, or to spiking statistics that become more regular. This analysis complements the mean-matching method. Mean-matching controls for changes in spiking statistics that are due to changing rates (e.g., the impact of a refractory period) but not for spiking statistics that might change for some other reason (e.g., perhaps spiking statistics change from Poisson to fourth-order gamma-interval due to some change in the nature of the synaptic input).

To investigate this possibility, The FF was normalized by the square of the coefficient of variation (CV^2) of the inter-spike intervals. The FF/CV^2 was computed using the method of Nawrot et al.⁴ That approach is based on the equality $FF = CV^2$, which holds for any stationary renewal process in the limit of long windows. Thus, a change in underlying spiking statistics (e.g., from Poisson to gamma-interval) may cause the FF to decline, but will not impact the FF/CV^2 (both the numerator and denominator will decline similarly). However, a decline in across-trial rate variability *will* be captured by the FF/CV^2 : the FF will decline much more than the CV^2 . A hurdle to applying this method is that mean firing rates are typically very non-stationary, at which point the above equality does not hold. To overcome this, time is rescaled (separately for each neuron and condition) to produce a stationary mean rate. The FF and CV^2 are then computed. To account for underlying rates that may not always equal the mean, the CV^2 is computed on individual trials and then averaged. After computing the FF/CV^2 , time is scaled back before averaging across neurons and conditions.

Panels a,b plot the results of this analysis using a 250 ms window (250 ms of rescaled time, which can be more or less than 250 ms of actual time, depending on whether rates are low or high, respectively). Arrows indicate stimulus onset. The FF and CV^2 are plotted individually at the top of each panel. Their ratio is plotted with flanking standard errors. The 250 ms window is longer than that used in other analyses (e.g., Figure 3). A longer window is needed for the current analysis for two reasons. First because the $FF = CV^2$ equality holds only for long windows and second, to insure there are enough spikes within the window for the CV of the inter-spike intervals to be estimated. Also to this end, analysis was restricted to neurons/conditions where the mean number of spikes in the 250 ms analysis window was at least 4. We further restricted the analysis to neurons where there was an average of at least 4 spikes during the pre-target period. Following this restriction, the median number of spikes/window was 9 (MT) and 5 (PMd). **Panels c,d** plot the same analysis, based on the same datasets, but using a 450 ms window. We enforced a minimum mean number of spikes per window of 10 (MT) and 7 (PMd). The median number of spikes/window was 27 (MT) and 8.5 (PMd).

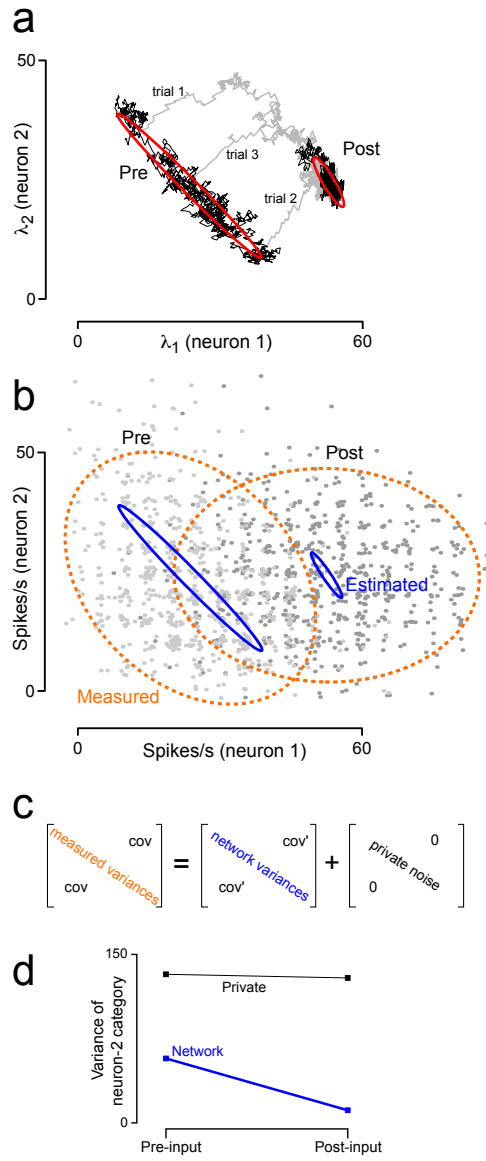
For all the above analyses, the FF/CV^2 fell following stimulus onset. This decline was the result of a strongly declining FF, normalized by a CV^2 that was both low and fairly constant (consistent with Figure 11d of Nawrot et al. 2008⁴). This confirms that the change in the FF is due to a change in across-trial underlying-rate variability, and not a change in the statistics of private spiking noise. Had spiking statistics transitioned from irregular to regular, the CV^2 of the inter-spike intervals would have undergone a large decline from an initially high value (paralleling the decline seen for the FF). In that case the FF/CV^2 ratio would have remained fairly constant. Instead the CV^2 is near or even slightly below unity (indicating spiking statistics slightly more regular than Poisson) at all times. Note that the FF decline is much larger than in the analyses elsewhere in the study, a result of using a longer window (250 or 450 ms, versus 50 ms in Figure 3). This occurs because the impact of across-trial rate variability on the FF is strongly window dependent.

Using the FF/CV^2 method, the decline in variability is somewhat smeared in time (especially in *c* and *d*). This is due to the necessity of a large analysis window. Furthermore, given the need for time rescaling, the true size of the window varies (across time, neurons, and conditions) depending on the underlying rate.



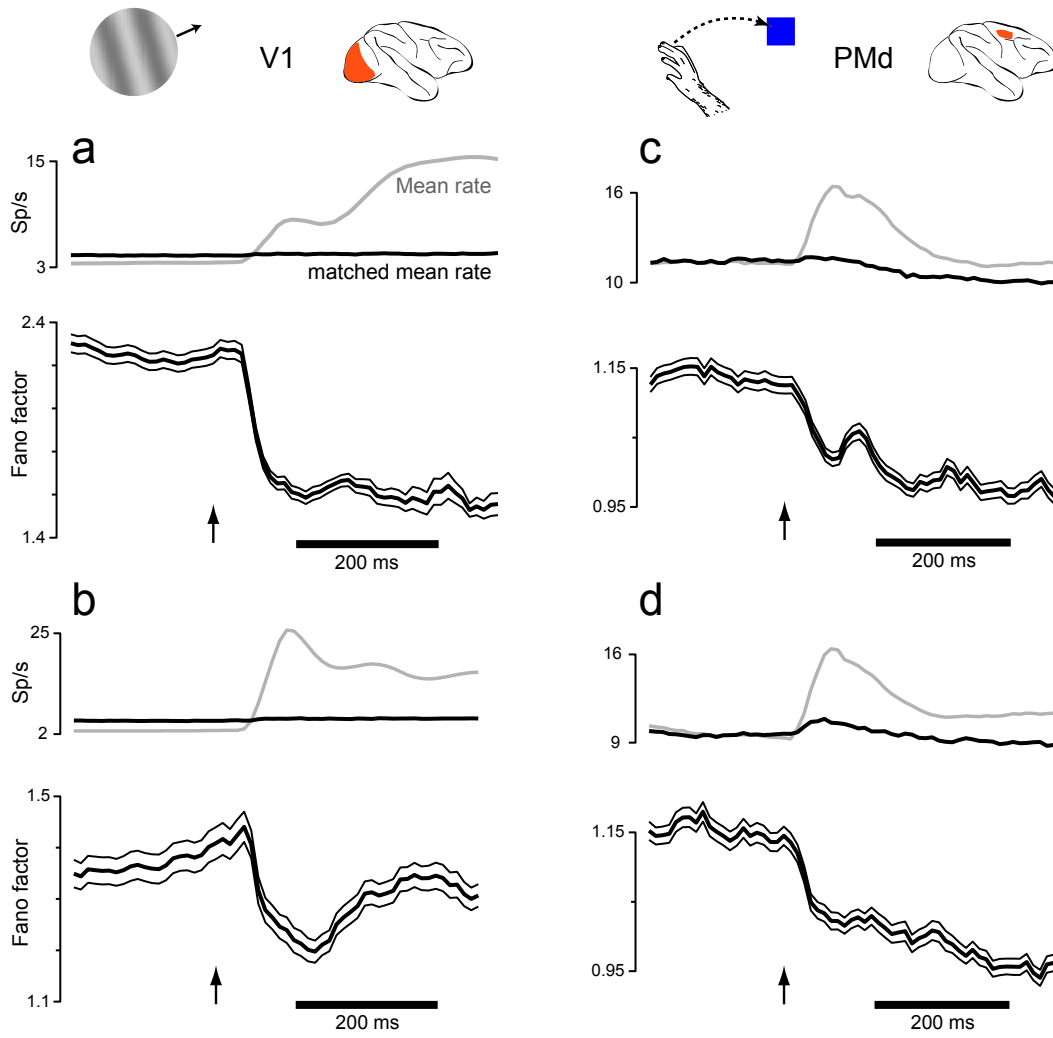
supplementary figure 4
(FF recovery after stimulus offset)

Supplementary figure 4. Timecourse of the recovery of the FF following stimulus offset. This analysis could not be performed for all datasets, as it was common for data collection to cease following stimulus offset, or for behavioral control to be poor at that time (e.g., the animal may have made its response and be consuming its reward with no fixation requirements). This analysis is thus restricted to the V1 dataset from Figure 6c, and the MT-speed dataset (Figure 3, bottom). Both these datasets involved anaesthetized preparations. For improved statistical power, a 100 ms window is used (the decline/recovery is not noticeably sharper with a 50 ms window). **a.** Recovery of the FF in V1 following stimulus offset. **b.** Recovery of the FF in MT following the offset of motion. Note that for this dataset there are three key events: the onset of a static stimulus, the onset of motion, and the cessation of motion. Offset of the static stimulus occurred after the end of the record.



supplementary figure 5
(application of factor analysis to simulated data)

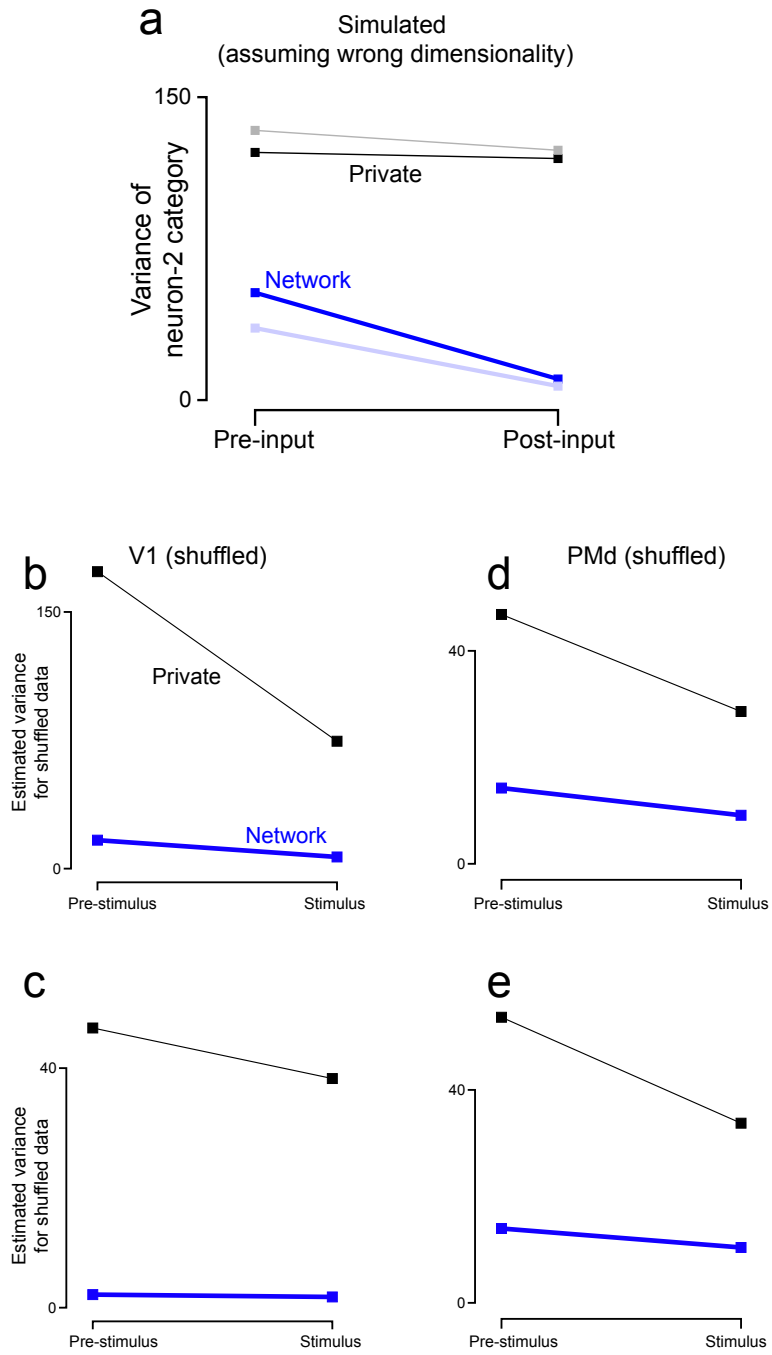
Supplementary figure 5. To illustrate the application of FA, we apply it to simulated data. **a.** Two-dimensional rate trajectories for three simulated trials from the simple artificial network shown in Supplementary figure 1b. Traces are grey for a 400 ms transition time following stimulus onset. For 10,000 simulated trials, we computed the mean state in a 200 ms window, both before (Pre) and after (Post) the onset of the input. Red ellipses plot the covariance across these measurements. The variance decline is now evident as a decline in covariance-ellipse area from pre-input to post-input. Unfortunately, in practice one measures not these underlying-rate covariances, but single-trial spike counts, simulated below by draws from Poisson distributions with rates λ_1 and λ_2 . **Panel b** plots the resulting single-trial estimates of spike rate (*grey*) and their covariance ellipses (*dashed orange*). These ‘measurements’ are dominated by the Poisson spiking-process noise. (Note that the fixed 200 ms measurement window produces discrete values. For presentation, ‘jitter’ has thus been added to make clear where density is highest.) Yet despite measurement noise, the underlying structure is recoverable (*blue ellipses*, see below), provided many trials and neurons. We simulated 10,000 trials for 60 neurons, 30 each with rates λ_1 and λ_2 . We then used FA (**panel c**) to decompose the measured covariance into a network and private components. The resulting network covariance ellipses (*blue, panel b*) accurately reflect the true underlying-rate covariance (*red ellipses in panel a*), despite non-Gaussian underlying data and spiking-process noise. Estimating the blue network covariance ellipses required abundant data, but even modestly-sized datasets allow FA to estimate the network variances of individual neurons (*blue diagonal in c*). To illustrate this, **panel d** plots the mean network variance across 200 trials for the 30 simulated units with rate λ_2 . FA accurately reports declining network (shared rate) variability, with little change in private spiking-process noise. The full network covariance matrix was computed, but we subsequently averaged variances only for units with little change in mean rate (those with rate λ_2 , whose mean changes little, see panel **a** and Supplementary figure 1b). This aids interpretation; were mean rates not stabilized, spiking-process noise would scale with rate. We similarly restricted neural analysis in Figure 6 *b–g* to matched distributions of mean rates across time (much as was done for the FF).



supplementary figure 6
 (the FF for the multi-unit array recordings)

Supplementary figure 6. The FF (computed using the mean-matching method) for the V1 and PMd datasets that were used for FA (V1: 567R and 106R; PMd: G20040123 & G20040122).

Panels a,b,c,d plot the FF for the same datasets used in *Figure 6b,c,e,f* (respectively) of the main text. Format is similar to that of Figure 3 in the main text. However, these datasets (unlike those in Figure 3) contained multi-unit isolations (see *methods*). Nevertheless, similar drops in the FF were seen. Also note that the changes in the FF (*black with flanking SEs*) do not always mirror those in the mean rate (*grey*). For example, in panel *a* the FF decline is immediate while the mean rate ramps up in two successive steps. In panel *d*, the FF initially declines as the mean rate is rising. The FF then continues declining although the mean rate is by then falling.




supplementary figure 7
(controls for effects obtained via factor analysis)

Supplementary figure 7. Exploration of when FA might mis-assign variance from one source (e.g., network-level rate variance) to the other (e.g., private spiking noise). This can happen if the specified dimensionality (rank) of the network covariance matrix is too low or too high. In the former case, some network variability will be mis-assigned as private. Thus, a decline in network variability will ‘leak’ into private noise. This possibility is of minor concern, as it would imply that our results slightly underestimate the true effect. This possibility is explored in *a* (see below). Alternately, if the rank is too high, some private noise will be mis-assigned as network variance. A decline in spiking-process variability (e.g., a change from Poisson to gamma-interval spiking statistic) could then appear as a decline in network variance. This artifact *is* of potential concern, as our interpretation is that the observed variance decline has a network source and is not principally due to more a regular spiking-process. It is perfectly acceptable if there is a real drop in spiking-process variability. However, if spiking-process variability were ‘leaking’ heavily into network variability, this would make interpretation difficult. A large leak is unlikely *a priori*: the changes in network variance seen in Figure 6 (of the main text) are larger than the changes in private noise, and remained so when we tested a range of dimensionalities (2-8, effects only become stronger for higher dimensionalities). This is inconsistent with the primary effect being a change in private noise. A further control is presented below (panels *b-e*).

Panel a. Dark traces are the same as in Supplementary figure 5*d*. Light traces plot the results of FA if the underlying network state had a higher dimensionality (6) but FA was still run assuming a dimensionality of 2. This was accomplished using a simulation in which three 2-neuron networks evolved independently. This produced data that spanned 6 dimensions. The key point is that when FA is run assuming a dimensionality that is too low, the algorithm has no choice but to assign some unexplained variance to private noise. Private noise is initially inflated, and then appears to decline. Such artifacts are possible because FA requires that we specify in advance the rank of the network covariance matrix (i.e., the dimensionality of the latent space, analogous to the number of principal components). Often there is no certain way to do so, especially as cross-validation techniques will underestimate the true network dimensionality unless the dataset is very large. We attempted to estimate the network dimensionality by assessing the dimensionality spanned by the mean population response across stimulus conditions (this yielded 4 and 5 dimensions for V1 and PMd, see *methods*). Yet given that not all possible stimuli were tested, this approach may also underestimate the true dimensionality spanned by the network state. Thus, the small decline in private spiking noise for the neural data, seen in Figure 6*b,c,e,f*, may result from an artifact similar to that shown in this panel. For finite data, this issue is difficult to resolve. Note, however, that this makes the central result conservative: the true

decline in network variance is likely *larger* than what is measured, on the assumption that we have underestimated the true dimensionality of the network space.

Panels b-e. Control for the converse concern to that addressed above. Might private noise ‘leak’ into the network variance, especially if we assumed too *high* a dimensionality for FA. Analysis was the same as for Figure 6*b,c,e,f* of the main text. However, analysis was applied to trial-shuffled data, so that there was no real covariance among neurons, only private noise. This was done by taking the data matrix, D (for which each column corresponds to one neuron, and each row to one trial), and for each column reassigning the rows randomly. This eliminates all correlation between neurons (i.e., all network variance) except that occurring by chance. FA correctly assigns the majority of the variance, and the majority of the drop in variance, to private noise (black traces undergo the largest drops), with only a small ‘leak’ into the network variance. This occurs despite the fact that FA assumed the original latent dimensionalities (see *methods*), whereas the true dimensionality of the shared network variance was actually zero due to shuffling.



supplementary figure 8
(comparison of behavioral and neural outlier trials)

Supplementary figure 8. Comparison of reaction-time outliers with ‘neural trajectory’ outliers. For the first PMd dataset (dataset G20040123) 5 trials had outlier RTs > 500 ms. Here we ask if those trials also had unusual neural trajectories, as in Figure 7b of the main text (where the red trial was one of the 5 outliers). The 5 outliers were spread across 3 targets. We applied GPFA to the recordings for each of these targets, and computed the mean neural trajectory during the 400 ms after the go cue. We then calculated the average distance, in neural space, of each trial’s trajectory from the mean trajectory. **a.** A histogram of these values. **b.** A histogram of RTs. The 5 trials that were flagged as outliers based on their RTs are colored red in both plots. It was indeed the case that the RT outliers were also outliers in terms of their neural trajectories. **c,d,** same analysis but for a different day’s dataset (G20040122, same as in Figure 7c). For this dataset there were three RT outliers (*red*), spread across three targets. These were also outliers in neural space.

Supplementary note 1. For both PMd and V1 (the datasets where FA could be applied), we found that the ‘network’ state after stimulus onset tended to correlate with the network state pre-stimulus. This was assessed, after applying FA, and finding the network state on each trial. Note that this is different from what was done in Figure 6 of the main text, where FA was used solely to estimate the network covariance matrix. Here we are using FA to find the network state on each trial (analogous to using PCA, and finding each trial’s projection onto the first few PCs). Each trial’s datum becomes a point in a 5 (V1) or 4 (PMd) dimensional latent space. On each trial, there is then one location (in one latent space) pre-stimulus, and another location (in a different latent space) during the stimulus. We regressed the network state found during the stimulus against the pre-stimulus state. For each stimulus-epoch dimension, data were regressed against all dimensions of pre-stimulus data. For example, for V1 this led to 60 total comparisons (5 stimulus-epoch dimensions, 12 conditions). Of 60 and 36 comparisons, 32% (V1) and 39% (PMd) showed a significant ($p < 0.05$) correlation.

Supplementary note 2. A fundamental question, largely unaddressed in this study, is the source and meaning of the variability that is present before target onset. On the one hand, such variability could be produced by local circuitry, which might be less stable when not driven. For example, circuitry related to gain control⁵, involving lateral inhibition, might be less stable when un-driven, much like the network in Supplementary figure 1*b*. A related hypothesis is that the decline in variability is related to the recruitment of shunting inhibition, and there is indeed good evidence for this in V1⁶.

A second possibility is that variability arises due to a lack of ‘internal’ behavioral control during the pre-stimulus period. Even if external behavior (hand and eye position) is perfectly controlled during the baseline, the monkey’s attention and ‘thoughts’ certainly are not. Variability in cognitive factors might be expected to be high during the baseline, and then be reduced following stimulus onset. In this view, variability is produced not by local circuitry, but by broader circuits involving attention or other cognitive phenomena.

In support of the local-circuitry hypothesis, the decline in variability is readily apparent in all the datasets from anaesthetized animals (all V1 datasets and the two MT datasets used for the FF). In support of the cognitive hypothesis, we have previously found that the decline in neural variability is predictive of reaction time, being faster for trials with shorter reaction times⁷. This suggests that the trials where motor preparatory becomes accurate most rapidly are also the trials where the neural state becomes consistent most rapidly. In summary, it is still quite unclear which explanation is correct, and it is quite plausible that both are at play.

Supplementary note 3. The central result of this study is that there is considerable trial-to-trial variability present before stimulus onset, and that this variability is diminished by stimulus onset. What is the timescale of this variability? In particular, for anaesthetized animals, is the variability related to slow drift in anesthesia? A slow drift in excitability would certainly create across-trial firing-rate variability. However, it isn't clear that an increase from this source would be limited to the pre-stimulus period; it would likely increase variability at all times. Still one can imagine scenarios in which this might occur. Thus, to examine whether slow drift could be a partial source of our effect, we recomputed the FF for the V1 data (from Figure 3), after splitting the 100 trials/condition into 10-trial sets. Thus, variability was being computed across a smaller set of trials, presented close to one another in time. This did indeed result in slightly lower variability overall. However, the stimulus-driven decline in variability was undiminished; the effect looked essentially identical to the original effect. Virtually identical results were also obtained if we computed variability across subsets of 5 trials each. Thus, while some slow drift is present (variability is slightly lower overall when computed locally) this appears to have little to do with the stimulus-driven decline in variability.

Supplementary video 1. A movie version of Figure 7a. Data are from PMd, and show the decline in across-trial variance after the onset of the stimulus (a reach target). The movie spans 750 ms, beginning 400 ms before stimulus onset and ending 350 ms after. The movie ends before the go cue is given. Each black dot shows the state of PMd on one trial. Fifteen randomly-chosen trials are shown. Dots turn blue for a brief moment at the time of stimulus onset. Note the subsequent drop in the variance of the dot locations (i.e., a drop in firing-rate variance). This feature of the response is at least as clear as the change in mean dot location (i.e., the change in mean firing rates). G20040123 dataset.

Supplementary video 2. Same as Supplementary video 1, but more time is shown and the trajectory of the RT-outlier trial is now included (*red*). The movie spans ~1500 ms. This time-span differs slightly across trials, as they have different go-cue and movement-onset times. At the time of the go cue, each dot turns green and further progress is halted. Progress resumes once all trials have passed the time of their respective go cues. This re-aligns the data to the go cue, much as is commonly done in PSTH's. Traces end at movement onset.

Supplementary video 3. Same as Supplementary video 1, but for the G20040122 PMd dataset (that shown in Figure 7c).

Supplementary video 4. Supplementary video 2, but for the G20040122 PMd dataset (that shown in Figure 7c).

Simulations

The simulation (Supplementary figure 1 and Supplementary figure 5*a,b*) employed two units with activities that varied continuously from zero to one. This network is merely meant to be illustrative: its details are not intended to model any real network. The network's two units form a winner-take-all network. When receiving balanced inputs (initially zero) the network is unstable and intrinsic noise has a large effect on the network state. When driven with an unbalanced input, the network settles to a state where one unit is more active. In doing so, the network becomes more stable. Network activity evolved according to,

$$\tau d \mathbf{s}(t) / dt = -\mathbf{s}(t) + f(\mathbf{W}\mathbf{s}(t) + \mathbf{b} + \mathbf{d}(t) + \mathbf{q}(t)) ,$$

where $\tau = 35$ ms, \mathbf{W} is a 2x2 anti-diagonal matrix with both entries = -0.7, \mathbf{b} is a two-dimensional bias vector with both entries = 0.7, $\mathbf{d}(t)$ is a two-dimensional input vector, and $\mathbf{q}(t)$ is two-dimensional Gaussian noise of standard deviation 0.22, independent at each time and for each neuron. At the time of stimulus onset, $\mathbf{d}(t)$ stepped from [0 0] to [0.42 0.355]. $f(x)$ is a sigmoidal transfer function: $f(x) = (\pi/2 + \text{atan}(20*(x-0.5)))/\pi$. The zero to one range of $\mathbf{s}(t)$ was mapped to a λ from zero to 60 spikes/s. These firing rates were converted to 'recorded' spike trains via an inhomogeneous Poisson process with rate λ .

To compute the red ellipses in Supplementary figure 5*a*, the mean λ on each trial was computed in a 200 ms measurement window which either ended at stimulus onset (*Pre*) or began 400 ms after stimulus onset (*Post*), by which point network 'settling' had largely finished. For Supplementary figure 5*b*, the spike count on each of 10,000 trials was drawn from a Poisson distribution, with mean equal to the λ on that trial (averaged over the measurement window) times the length of the measurement window. The spike count was transformed back to a rate for plotting purposes.

When applying Factor analysis (FA) to the simulated data, we simulated recordings from a 60 unit network in which half the units shared the underlying rate of neuron 1 and half shared the underlying rate of neuron 2. Thus, the underlying network-level variance was strongly correlated across units (perfectly within a type). Measured rates were computed using draws from a Poisson distribution that were independent for each unit. This emulates the very weak spike-spike correlations typically observed for pairs of cortical neurons. FA was applied to data for all simulated units. Supplementary figure 5*d* then restricts subsequent analysis to the neuron-2-like class.

References (supplementary)

- 1 Carpenter, R. H. & Williams, M. L. Neural computation of log likelihood in control of saccadic eye movements. *Nature* **377**, 59-62 (1995).
- 2 Hanes, D. P. & Schall, J. D. Neural control of voluntary movement initiation. *Science* **274**, 427-430 (1996).
- 3 Shadlen, M. N. & Newsome, W. T. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J Neurophysiol* **86**, 1916-1936 (2001).
- 4 Nawrot, M. P. *et al.* Measurement of variability dynamics in cortical spike trains. *J Neurosci Methods* **169**, 374-390 (2008).
- 5 Priebe, N. J., Churchland, M. M. & Lisberger, S. G. Constraints on the source of short-term motion adaptation in macaque area MT. I. the role of input and intrinsic mechanisms. *J Neurophysiol* **88**, 354-369 (2002).
- 6 Monier, C., Chavane, F., Baudot, P., Graham, L. J. & Fregnac, Y. Orientation and direction selectivity of synaptic inputs in visual cortical neurons: a diversity of combinations produces spike tuning. *Neuron* **37**, 663-680 (2003).
- 7 Churchland, M. M., Yu, B. M., Ryu, S. I., Santhanam, G. & Shenoy, K. V. Neural variability in premotor cortex provides a signature of motor preparation. *J Neurosci* **26**, 3697-3712 (2006).