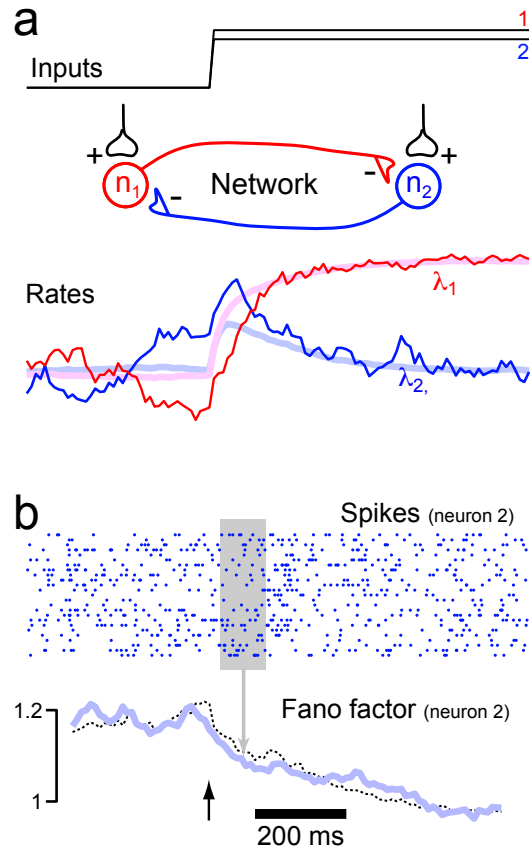# *Supplementary Materials*

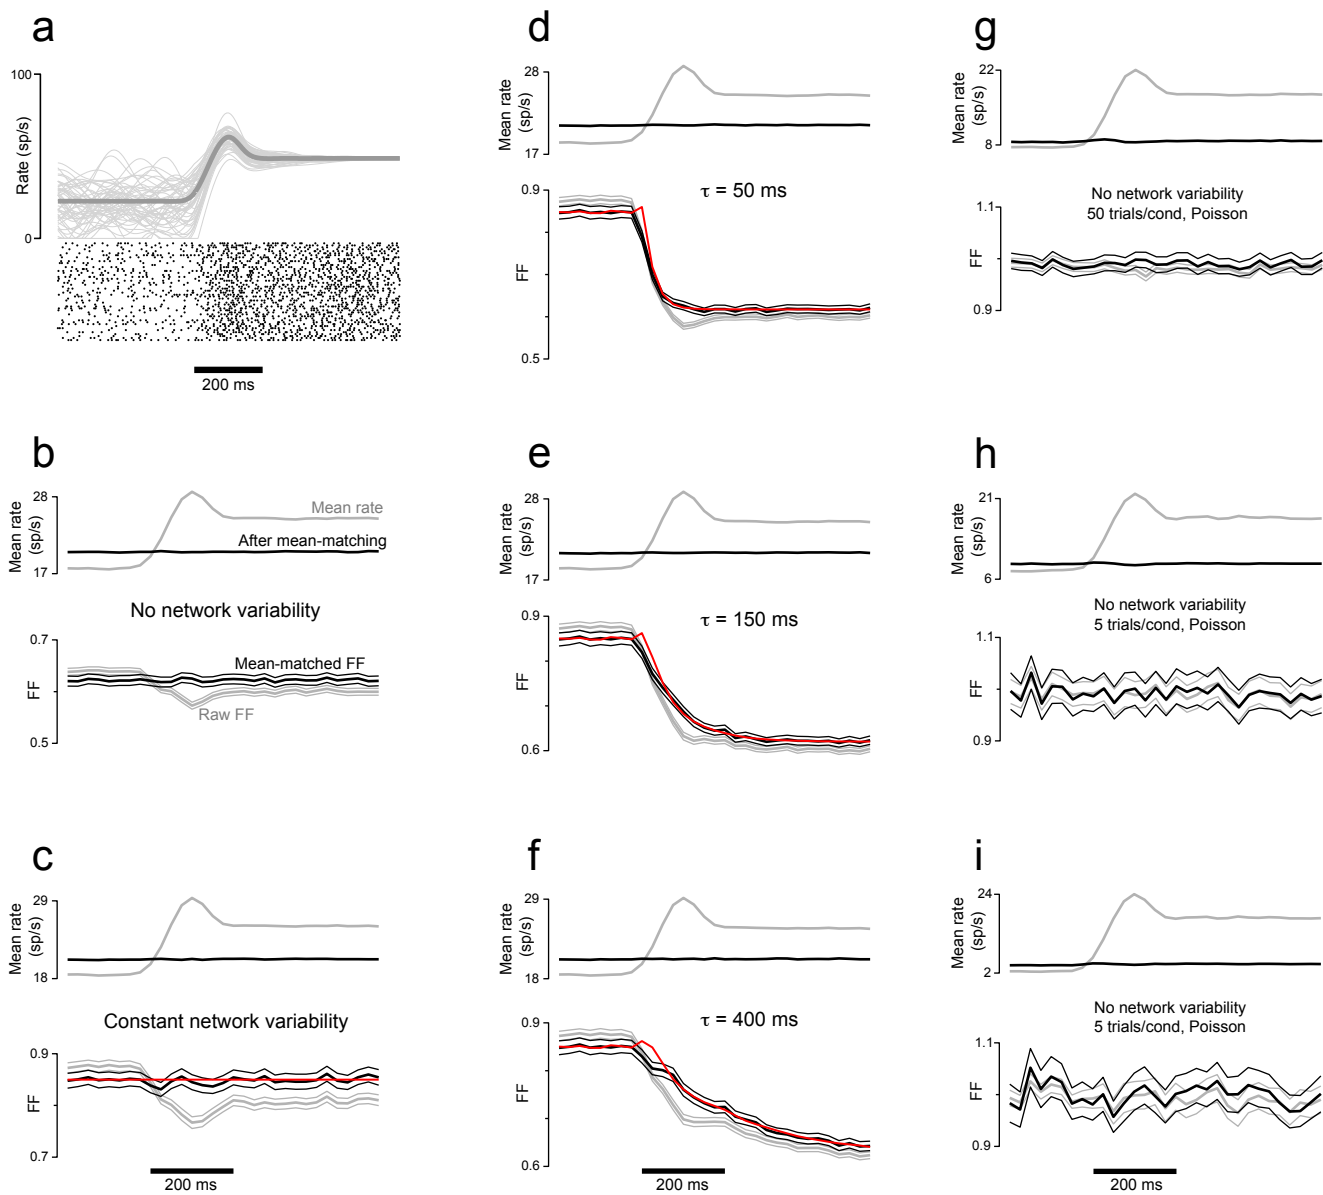## Stimulus onset quenches neural variability: a widespread cortical phenomenon

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

a

Inputs

1
2

+ n₁ — Network — n₂ +
-                          -

Rates

λ_1
λ_2,

b

Spikes (neuron 2)

1.2

Fano factor (neuron 2)

1

200 ms

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 singe-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.

a

Rate (sp/s)

100

0

200 ms

b

Mean rate (sp/s)

28

17

Mean rate

After mean-matching

No network variability

FF

0.7

0.5

Mean-matched FF

Raw FF

c

Mean rate (sp/s)

29

18

Constant network variability

FF

0.9

0.7

200 ms

d

Mean rate (sp/s)

28

17

τ = 50 ms

FF

0.9

0.5

e

Mean rate (sp/s)

28

17

τ = 150 ms

FF

0.9

0.6

f

Mean rate (sp/s)

29

18

τ = 400 ms

FF

0.9

0.6

200 ms

g

Mean rate (sp/s)

22

8

No network variability
50 trials/cond, Poisson

FF

1.1

0.9

h

Mean rate (sp/s)

21

6

No network variability
5 trials/cond, Poisson

FF

1.1

0.9

i

Mean rate (sp/s)

24

2

No network variability
5 trials/cond, Poisson

FF

1.1

0.9

200 ms

supplementary figure 2
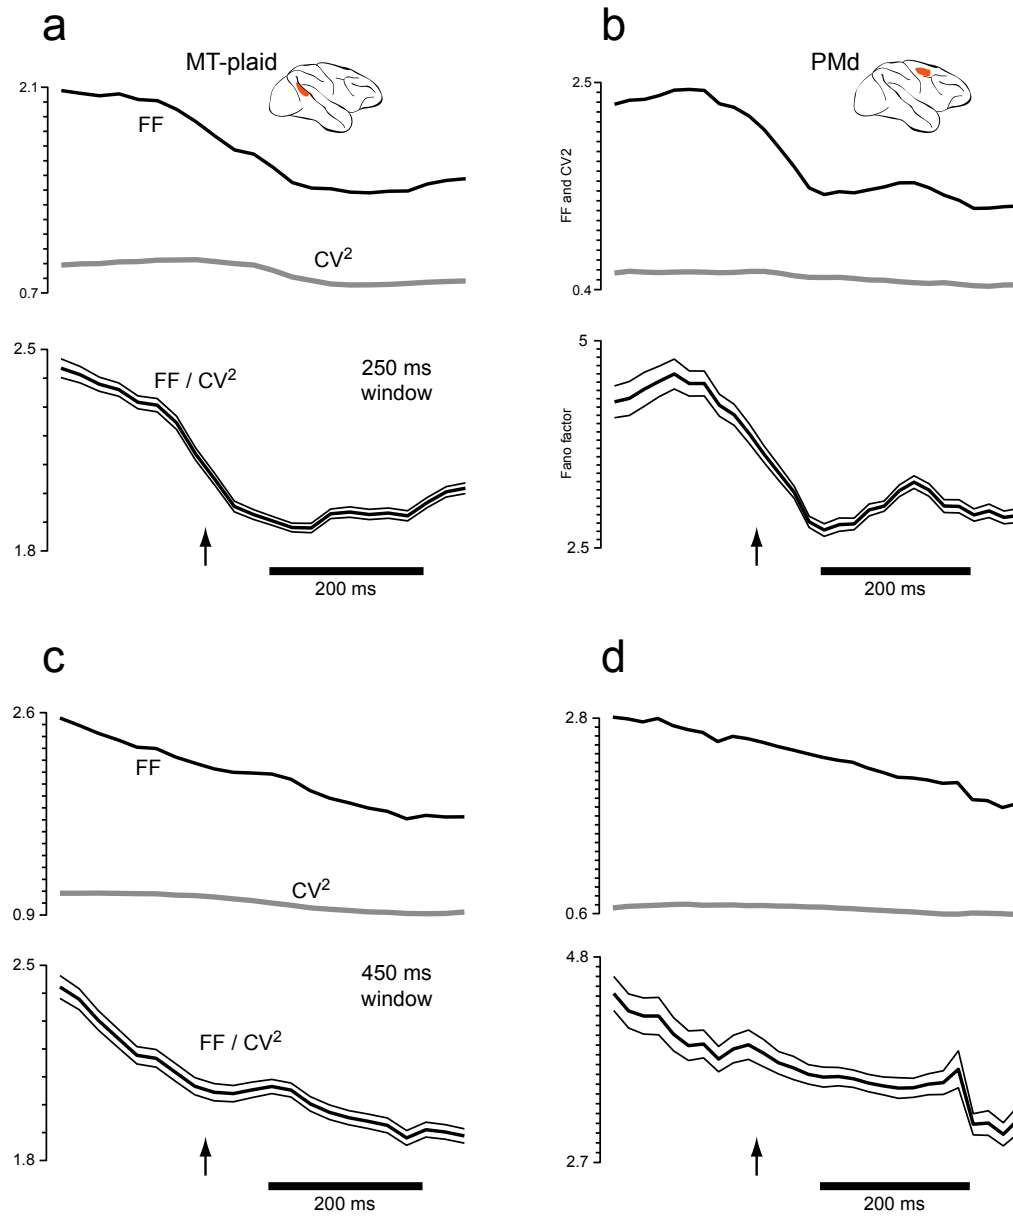(the Fano factor for simulated data)

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).

a MT-plaid

2.1
FF
CV²
0.7

2.5
FF / CV²   250 ms
window
1.8
200 ms

b PMd

2.5
FF and CV2
0.4

5
Fano factor
2.5
200 ms

c

2.6
FF
CV²
0.9

2.5
450 ms
window
FF / CV²
1.8
200 ms

d

2.8
0.6

4.8
2.7
200 ms

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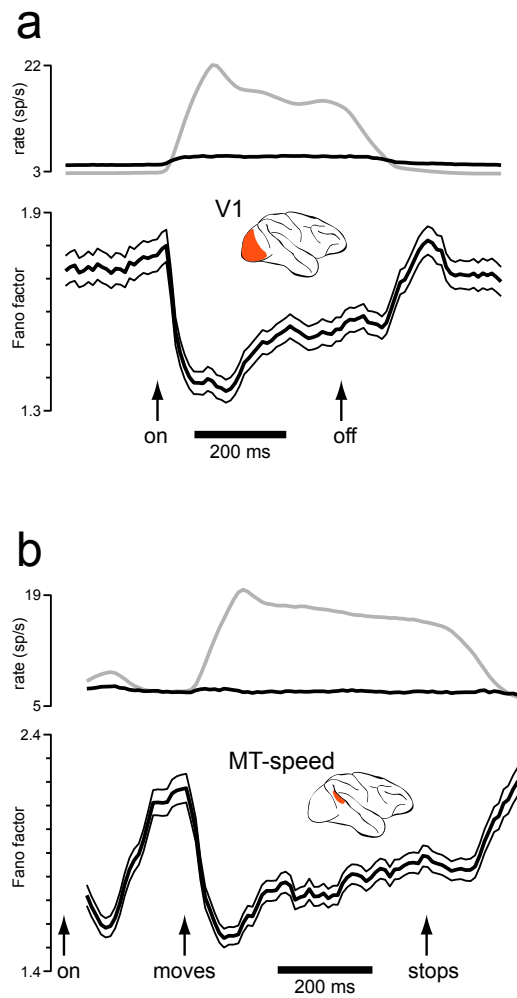(FF and the CV² of the ISIs)

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).
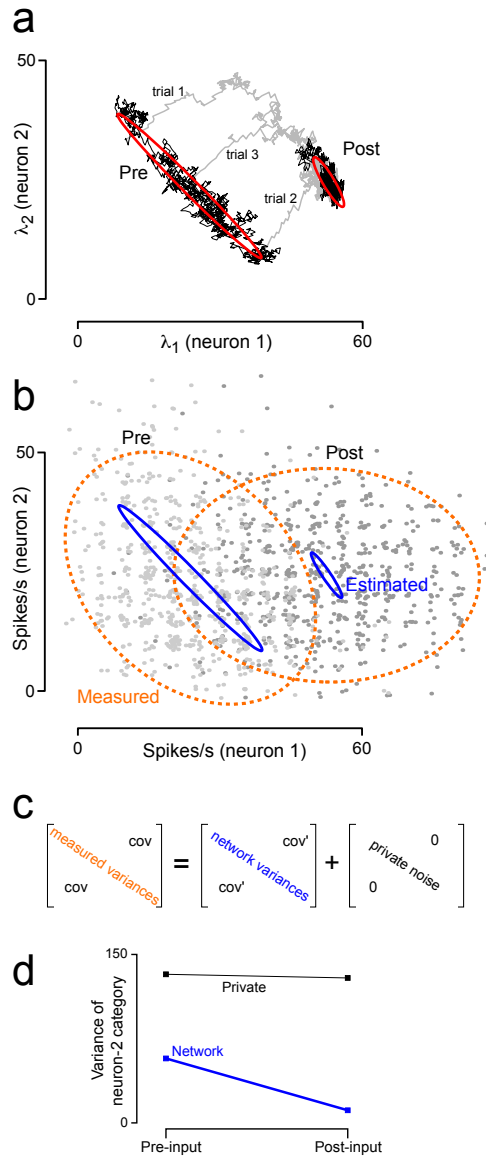
For all the above analyses, the FF/CV^2 fell following stimulus onset. This decline was the result of a strongly declining FF, normalized by a CV^2 that was both low and fairly constant (consistent with Figure 11d of Nawrot et al. 2008[4]). This confirms that the change in the FF is due to a change in 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.
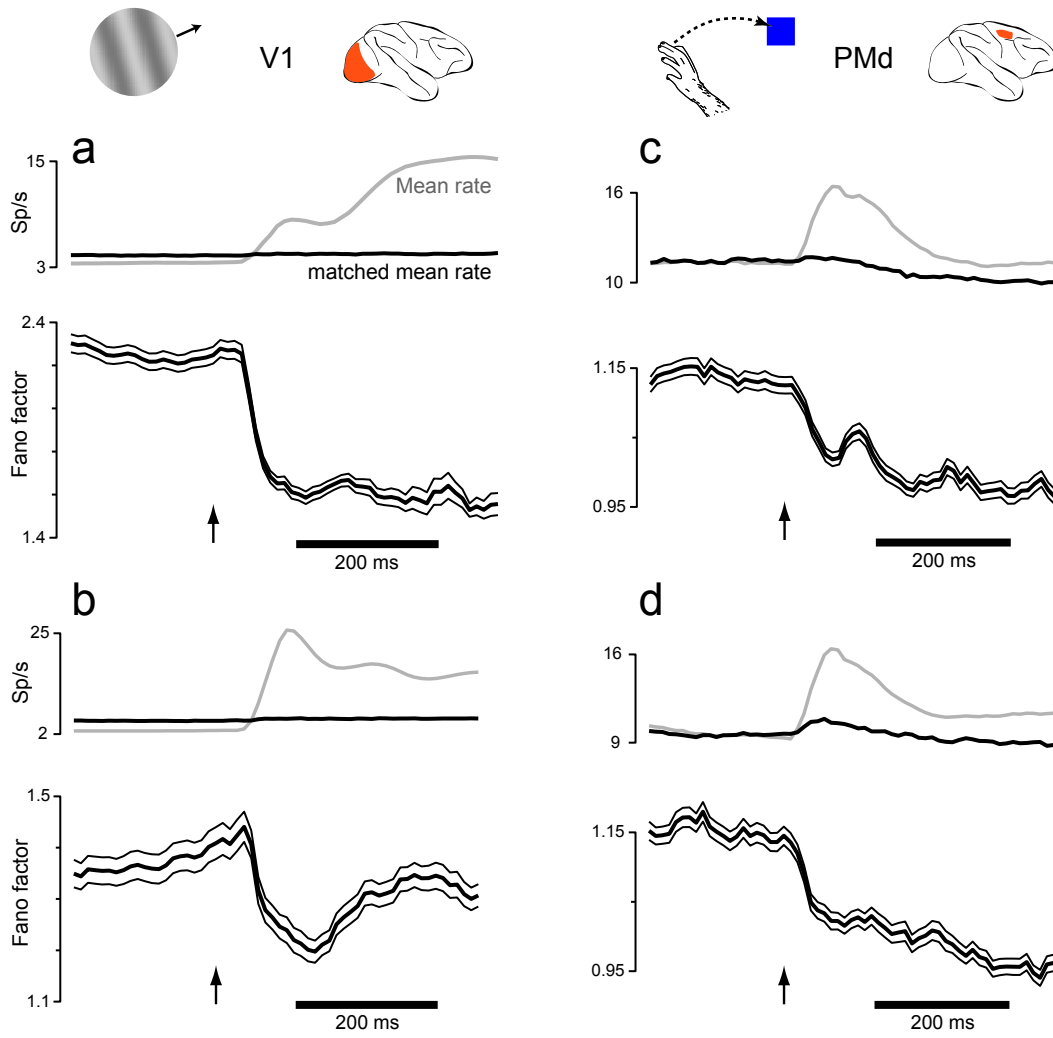
a

rate (sp/s)

22

3

Fano factor

1.9

1.3

V1

on

200 ms

off

b

rate (sp/s)

19

5

Fano factor

2.4

1.4

MT-speed

on

moves

200 ms

stops

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 6*c*, 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.

a

λ_2 (neuron 2)

50

0

trial 1

trial 3

Pre

Post

trial 2

0

λ_1 (neuron 1)

60

b

Spikes/s (neuron 2)

50

0

Pre

Post

Estimated

Measured

0

Spikes/s (neuron 1)

60

c

$$\begin{bmatrix} \textit{measured variances} & \text{cov} \\ \text{cov} & \end{bmatrix} = \begin{bmatrix} \textit{network variances} & \text{cov'} \\ \text{cov'} & \end{bmatrix} + \begin{bmatrix} \textit{private noise} & 0 \\ 0 & \end{bmatrix}$$

d

Variance of neuron-2 category

150

0

Private

Network

Pre-input

Post-input

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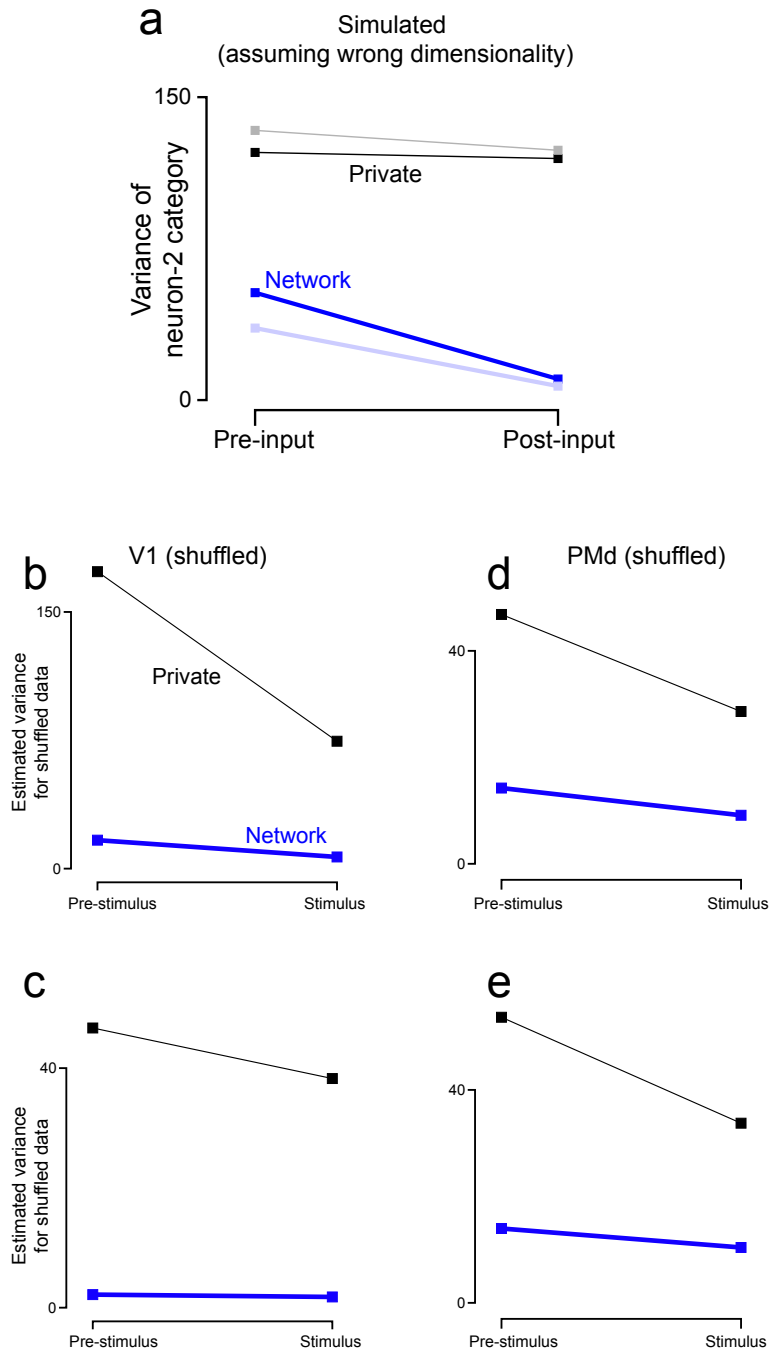(application of factor analysis to simulated data)

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 1*b*. 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 1*b*). 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).

V1

a

Sp/s
15
3

Mean rate

matched mean rate

Fano factor
2.4
1.4

200 ms

b

Sp/s
25
2

Fano factor
1.5
1.1

200 ms

PMd

c

16
10

Fano factor
1.15
0.95

200 ms

d

16
9

1.15
0.95

200 ms

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.

a Simulated
(assuming wrong dimensionality)

Variance of neuron-2 category

150

Private

Network

0

Pre-input   Post-input

b V1 (shuffled)

Estimated variance for shuffled data

150

Private

Network

0

Pre-stimulus   Stimulus

c

Estimated variance for shuffled data

40

0

Pre-stimulus   Stimulus

d PMd (shuffled)

40

0

Pre-stimulus   Stimulus

e

40

0

Pre-stimulus   Stimulus

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 is 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.

a

b

c

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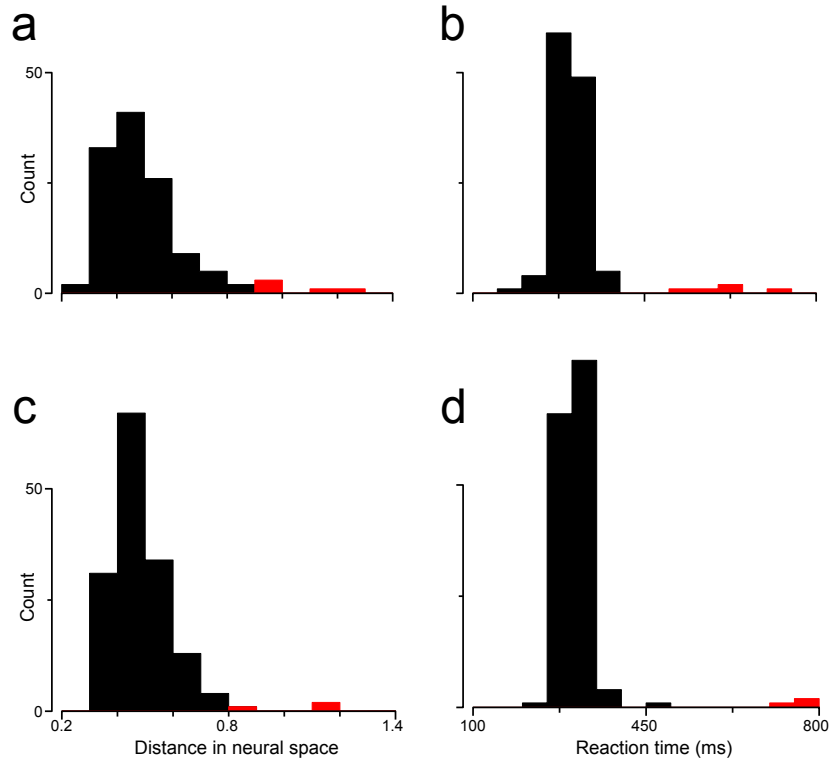
supplementary figure 8
(comparison of behavioral and neural outlier trials)

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 7*b* 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 7*c*). 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 control[5], 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 V1[6].

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 times[7]. 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 once 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 7*a*. 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 7*c*).

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

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, 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).