

High-Dimensional Dynamics in Low-Dimensional Networks

Yue Wan¹ and Robert Rosenbaum^{1*}

Department of Applied and Computational Mathematics and Statistics, University of Notre Dame, Notre Dame, Indiana 46556, USA



(Received 12 November 2025; accepted 28 May 2026; published 3 August 2026)

Many networks in nature and applications have an approximate low-rank structure in the sense that their connectivity structure is dominated by a few dimensions. It is natural to expect that dynamics on such networks would also be low-dimensional. Indeed, theoretical results show that low-rank networks produce low-dimensional dynamics whenever the network is isolated from external perturbations or input. However, networks in nature are rarely isolated. Here, we study the dimensionality of dynamics in recurrent networks with low-dimensional structure driven by high-dimensional inputs or perturbations. We find that dynamics in such networks can be high- or low-dimensional, and we derive mathematical conditions on the network structure under which linearized dynamics are high-dimensional. In many low-rank networks, dynamics are suppressed in directions aligned with the network's low-rank structure, a phenomenon we term "low-rank suppression." We show that several low-rank network structures arising in nature satisfy the conditions for generating high-dimensional dynamics and low-rank suppression. Our results clarify important but counterintuitive relationships between a recurrent network's connectivity structure and the structure of its response to external input.

DOI: [10.1103/79h2-36r3](https://doi.org/10.1103/79h2-36r3)

Subject Areas: Biological Physics, Nonlinear Dynamics

I. INTRODUCTION

Recent work shows that many networks arising in nature and applications have low-rank structure [1]. This observation raises an important question: What is the relationship between the dimensionality of a network's *structure* and the dimensionality of the *dynamics* on the network?

In neuroscience, for example, the implications of low-rank recurrent connectivity on neural dynamics is a topic of intense research [2–12]. Theoretical studies demonstrating low-dimensional dynamics in low-rank neuronal network models are consistent with some neural recordings showing low-dimensional neural activity [13–15]. However, many of these recordings are made in the context of low-dimensional stimuli or behavior. A growing number of more recent studies have shown that neural activity can be high-dimensional, particularly in response to high-dimensional stimuli or behavior [15–20]. These observations raise the important point that the dimensionality of dynamics on a network depends on the structure of the network's external input in addition to its internal structure.

There is a large body of theoretical work on the dimensionality of dynamics on low-rank networks [1–12, 21–25]. Some theoretical and computational studies demonstrating low-dimensional dynamics in low-rank networks focus on networks that are self-contained in the sense that they operate in the absence of noise or external inputs or perturbations [1, 2, 6]. Other work considers external input but assumes that the input is perfectly aligned with the network's low-dimensional structure [4]. In reality, networks in nature are rarely isolated from internal or external perturbations or inputs, and these perturbations are not necessarily aligned with the network's low-dimensional structure. Some [3, 5] consider the case of general external input to networks that are weakly low-rank in the sense that the largest singular values are $\mathcal{O}(1)$ in magnitude, consistent with connection weights that scale like $1/N$ where network size N is assumed to be large.

In this study, we consider networks with strongly low-rank structure in the sense that they have a small number of asymptotically large singular values, while the rest are small or moderate in magnitude. This is consistent with connection weights that are asymptotically larger than $1/N$ in magnitude, often taken to be $\mathcal{O}(1/\sqrt{N})$ in theoretical work. Importantly, we study the dimensionality of the dynamics on these networks when they are driven by external input. The dimensionality of activity driven by intrinsic dynamics is a fundamentally different question treated in previous work [1, 2, 6]. We also focus on dynamics near linearly stable

*Contact author: Robert.Rosenbaum@nd.edu

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

equilibrium. We first show by example that networks under these conditions can produce high- or low-dimensional dynamics in response to high-dimensional inputs, depending on the network's structure. Perhaps counterintuitively, we also find that networks under these conditions *suppress* inputs aligned to their dominant low-dimensional structure, an effect we call “low-rank suppression.”

We use mathematical analysis to derive conditions under which low-rank suppression and high-dimensional dynamics can occur in the linearized dynamics of strongly low-rank networks close to a stable equilibrium. Our analysis shows that low-rank suppression arises generally in these networks. High-dimensional responses to high-dimensional external inputs, on the other hand, arise when the dominant low-rank part of the network's connectivity matrix has a nonvanishing component with the EP property. EP matrices are matrices for which the column space and row space are equal (the abbreviation stands for “equal principal” or “equal projectors”) [26,27]. Low-dimensional dynamics can arise in the linearized dynamics of strongly low-rank networks through a form of non-normal amplification [28–30], but our mathematical analysis shows that this form of non-normal amplification can occur only when the network is non-EP. Non-normality is not sufficient. To our knowledge, this connection between the EP property of matrices, non-normal amplification, and the dimensionality of recurrent network dynamics has not been described in previous work.

We show that many common structural features such as biased weights, modularity, and spatial connectivity structure naturally produce strongly low-rank connectivity matrices with low-rank parts satisfying the EP condition, implying that these networks produce low-rank suppression and high-dimensional responses to high-dimensional inputs. We also draw connections between low-rank suppression and the mathematical theory of balanced networks [31], extending previous work in this direction [4]. Finally, we demonstrate our conclusions in dynamics on a real epidemiological network.

Our conclusions have important implications for the interpretation of low-dimensional network structure. In neuroscience, our results can explain why neural populations generate high-dimensional responses to high-dimensional stimuli and tasks [16]. Our results also generalize and extend the theory of excitatory-inhibitory balance [4,31] and amplification arising from breaks in this balance [32–34]. Beyond neuroscience, our results imply that perturbations misaligned to low-rank network structure are most effective at driving responses. This counterintuitive observation could be used to design more effective interventions to epidemiological, biological, social, and other networks.

II. HIGH-DIMENSIONAL DYNAMICS AND LOW-RANK SUPPRESSION IN A NETWORK WITH RANK-ONE STRUCTURE

For illustrative purposes, we begin with a simple, linear recurrent network model [Fig. 1(a)]:

$$\tau \frac{d\mathbf{z}}{dt} = -\mathbf{z} + W\mathbf{z} + \mathbf{x}, \quad (1)$$

where $\mathbf{x}(t)$ is an external perturbation or input, $\mathbf{z}(t)$ is the network response, and $\tau = 1$. The $N \times N$ recurrent connectivity matrix W takes the form

$$W = W_0 + W_1,$$

where

$$W_0 = c\mathbf{u}\mathbf{u}^T$$

is a normal, rank-one matrix and

$$W_1 = \frac{\rho}{\sqrt{N}}Z$$

is a full-rank random matrix. Specifically, $\mathbf{u}$ is a random vector with $\|\mathbf{u}\| = 1$, Z is an $N \times N$ matrix with entries drawn independent identically distributed (i.i.d.) from a standard normal distribution, and $\rho > 0$. If

$$|c| \gg \rho,$$

then W is “effectively low-rank” [1] in the sense that it has one large singular value near $|c|$ and the remaining singular values are much smaller [approximately bounded by 2ρ ; Fig. 1(b)]. Similarly, W has one large eigenvalue near c and the remaining eigenvalues lie approximately within a circle of radius ρ in the complex plane [Fig. 1(c)] [35,36].

Importantly, stability of the network dynamics requires that $\rho < 1$ and $c < 1$ [37]. If the dynamics were unstable, then $\|\mathbf{z}(t)\|$ would grow exponentially toward ∞ , so we consider only parameter regimes with stable dynamics. Since strong low-rank structure also requires $|c| \gg \rho$, stability requires $c < 0$ when $\rho = \mathcal{O}(1)$. In simulations here, we take $\rho = 0.5$ and $c = -10$. We consider more general models and network structures later.

We first provided a high-dimensional input $\mathbf{x}(t)$. Specifically, for our first simulation, each $\mathbf{x}_j(t)$ was an i.i.d., smooth, stationary Gaussian process, which models internal noise, external perturbations, or external input to the network. Conventional wisdom might lead us to expect a low-dimensional network response dominated by variability in the direction of $\mathbf{u}$. Indeed, a *feedforward* network with the same connectivity matrix and same input produces approximately one-dimensional dynamics, because it




FIG. 1. Response properties of a recurrent network with rank-one structure. (a) Schematic of model: The connectivity matrix W quantifies connections between nodes $\mathbf{z}$, which receive external perturbations or input $\mathbf{x}$. (b) The singular values of W have one dominant term, indicating approximate rank-one structure. (c) The eigenvalues of W have a dominant, negative term. (d) The distribution of variance across principal components of a Gaussian stochastic input $[\mathbf{x}(t); \text{green}]$ and the response $[\mathbf{z}(t); \text{blue}]$. (e),(f) The network response (e) and input (f) projected onto the plane determined by $\mathbf{u}$ and a random vector $\mathbf{u}_{\text{rand}}$ demonstrates low-rank suppression along $\mathbf{u}$. (g) The network response $[\mathbf{z}(t), \text{top}]$ and its norm $[\|\mathbf{z}(t)\|, \text{bottom}]$ given an input aligned to the low-rank structure of the network ($\mathbf{x}_{\text{aligned}}$) and a random input ($\mathbf{x}_{\text{rand}}$). The norm of $\mathbf{z}$ is more than 11 times smaller in response to the aligned versus random perturbation. (h) Local network input (purple) cancels with external input (green) to produce suppressed network responses (blue) in the direction of $\mathbf{u}$.

amplifies inputs aligned to $\mathbf{u}$ (Fig. 7 in Appendix F). However, this conclusion does not necessarily carry over to recurrent networks.

In the recurrent network, the variance explained by the principal components of $\mathbf{z}(t)$ decayed similarly to those of $\mathbf{x}(t)$ [Fig. 1(d)], indicating that $\mathbf{z}(t)$ was high-dimensional like $\mathbf{x}(t)$. The only exception was the principal component that explained the *least* variance in $\mathbf{z}(t)$, which was much weaker than the other principal components [last blue dot in Fig. 1(d)]. Perhaps surprisingly, this weakest principal component direction was closely aligned to $\mathbf{u}$ (the angle was less than 8°). Consistent with this finding, the variance of $\mathbf{z}(t)$ in the direction of $\mathbf{u}$ was more than 132 times *smaller* than the variance of $\mathbf{z}(t)$ along a random direction [Fig. 1(e)] even though the variability of $\mathbf{x}(t)$ was similar in each direction [Fig. 1(f)].

To better understand these results, we next simulated the same network with two different external input patterns. One pattern was aligned with the low-dimensional structure of W :

$$\mathbf{x}_{\text{aligned}} = \mathbf{u},$$

while the other had a random direction:

$$\mathbf{x}_{\text{random}} = \mathbf{u}_{\text{rand}},$$

which was generated identically to but independently from $\mathbf{u}$. Intuitively, we might expect the network to respond more strongly to the aligned input than to the random input, as in a feedforward network (Fig. 7 in Appendix F). In reality, we observed exactly the opposite [Fig. 1(g)]: The response to $\mathbf{x}_{\text{random}}$ was more than 11 times larger than the response to $\mathbf{x}_{\text{aligned}}$.

We use the term “low-rank suppression” to refer to this phenomenon in which inputs aligned to the low-rank structure of a network are suppressed by the network’s dynamics. In the absence of other directions that are amplified by the network, low-rank suppression leads to high-dimensional responses to high-dimensional inputs [as in Fig. 1(d)].

The results in Fig. 1 contrast with some previous modeling work demonstrating low-dimensional dynamics in low-rank recurrent networks [1,3–6]. In Sec. VIII and in

Appendix A, we provide a more detailed review of this previous work and its relation to ours.

On closer inspection, low-rank suppression in this example is not surprising, because the Jacobian at the fixed point has a large negative eigenvalue (near $c - 1$ with $c < 0$ and $|c| \gg 0$) with associated eigenvector near $\mathbf{u}$ so the dynamics in Eq. (1) are highly compressive along $\mathbf{u}$. However, as we show in the next section, low-rank suppression occurs even in the absence of a large, negative eigenvalue.

A second explanation for low-rank suppression generalizes to examples without a large, negative eigenvalue: Consider the steady-state solution to Eq. (1):

$$\mathbf{z} = [I - W]^{-1} \mathbf{x}. \quad (2)$$

Convergence to the steady state requires that $\mathbf{x}(t) = \mathbf{x}$ is static, but, as long as $\mathbf{x}(t)$ varies more slowly than $\tau = 1$, solutions approximately track the quasisteady state given by Eq. (2). Because of the matrix inverse in Eq. (2), the *large* singular value of W near $|c|$ produces a *small* singular value of $[I - W]^{-1}$ near $1/(1 - |c|) \approx 1/|c|$ with left and right singular vectors near $\mathbf{u}$. Hence, the recurrent network suppresses inputs in the direction of $\mathbf{u}$.

Both of the explanations above rely on the assumption that W_0 is a symmetric matrix, or at least a normal matrix, but in the next section, we next consider a case in which W_0 is non-normal and nonsymmetric.

An interesting consequence of low-rank suppression is that external perturbations cancel nearly perfectly with recurrent inputs in the direction of $\mathbf{u}$. More precisely, note that $W\mathbf{z}(t)$ in Eq. (1) can be interpreted as a vector of internal input to each node, whereas $\mathbf{x}(t)$ is external input and $\mathbf{z}(t)$ is the network response. In the quasisteady state, we have

$$\mathbf{z} \approx W\mathbf{z} + \mathbf{x}. \quad (3)$$

In other words, $\mathbf{z}$ tracks the sum of internal and external inputs in the quasisteady state.

Since W has a large singular value with left and right singular values aligned to $\mathbf{u}$, multiplication by W amplifies the direction $\mathbf{u}$. Therefore, the internal input $W\mathbf{z}$ is large in the direction of $\mathbf{u}$ whenever $\mathbf{z}$ is moderate in the direction of $\mathbf{u}$. In other words, $|\mathbf{u} \cdot W\mathbf{z}| \gg |\mathbf{u} \cdot \mathbf{z}|$. This fact might appear to present a paradox, because the direction $\mathbf{u}$ is *amplified* in the product $W\mathbf{z}$, but $W\mathbf{z}$ is one component of $\mathbf{z}$ [Eq. (3)] and the direction $\mathbf{u}$ is *suppressed* in $\mathbf{z}$, as we see in Figs. 1(e) and 1(g).

This apparent paradox is resolved by a cancellation between $W\mathbf{z}$ and $\mathbf{x}$ in the direction of $\mathbf{z}$. Specifically, under low-rank suppression, internal input cancels nearly perfectly with external input in the direction of $\mathbf{u}$ (i.e., $\mathbf{u} \cdot W\mathbf{z} \approx -\mathbf{u} \cdot \mathbf{x}$) so that the response $\mathbf{z}$ is weak in the direction of $\mathbf{u}$ [Fig. 1(h)]. We refer to this effect as

“low-rank cancellation,” which is closely related to the theory of excitatory-inhibitory balance in neural circuits [31,38] as we show later.

III. LOW-DIMENSIONAL DYNAMICS AND LOW-RANK SUPPRESSION IN A NON-NORMAL NETWORK WITH RANK-ONE STRUCTURE

Our observation of low-rank suppression in the example above might seem unsurprising due to the presence of a large, negative eigenvalue. It is widely known that an eigenvalue with a large, negative real part produces a suppressive dynamic. Likewise, our observation of high-dimensional dynamics might seem like a simple consequence of high-dimensional input. We next address these points with an example in which W_0 is highly non-normal (however, we later show that the EP property, not normality, is key to determining the dimensionality of dynamics). In this example, low-rank suppression arises in the absence of a large negative eigenvalue, and the network’s response is low-dimensional even though the input is high-dimensional.

We specifically consider an example in which the model and all parameters are the same as in Fig. 1 except that

$$W_0 = c\mathbf{u}\mathbf{v}^T,$$

where $\mathbf{u}$ and $\mathbf{v}$ are random, orthogonal vectors with $\|\mathbf{u}\| = \|\mathbf{v}\| = 1$. This matrix is highly non-normal in the sense that its left and right singular vectors are orthogonal, and its nonzero eigenvalue does not have a simple relationship to its singular values.

The distribution of singular values of $W = W_0 + W_1$ are identical to those from our previous example [Fig. 2(b); compare to Fig. 1(b)], but the eigenvalues are different [Fig. 2(c); compare to Fig. 1(c)]. Specifically, the eigenvalues no longer have a dominant negative outlier. They are all $\mathcal{O}(1)$ in magnitude and clustered in a single bulk.

Repeating the simulations from above for this network, we found that the response to high-dimensional input was low-dimensional in the sense that variability was much greater in one direction than all others: Over 40% of the variance in $\mathbf{z}(t)$ was captured by a single principle component [first blue dot in Fig. 2(d)], compared to less than 2% for any orthogonal direction (the remaining blue dots). This dominant direction was partially aligned to $\mathbf{u}$ (the angle between the dominant direction and $\mathbf{u}$ was 31°). Consistent with this observation, variability was 100 times larger in the direction of $\mathbf{u}$ than in a random direction [Fig. 2(e)]. Hence, low-dimensional dynamics were produced by high-dimensional input.

In addition, there was a direction in which variability was especially weak [last blue dot in Fig. 2(d)], similar to our previous example. This last principal component direction $\mathbf{u}_N$ was partially aligned to $\mathbf{v}$ (the angle between $\mathbf{u}_N$ and $\mathbf{v}$




FIG. 2. Response properties of a non-normal recurrent network with rank-one structure. The same as Fig. 1 except the low-rank part of W is non-normal. (a) Schematic of the model. (b) The singular values of W have one dominant term. (c) The eigenvalues of W lack a dominant, negative term. (d) The distribution of variance across principal components of a Gaussian stochastic input $[x(t); \text{green}]$ and the response $[z(t); \text{blue}]$. (e),(f) The network response projected onto the plane determined by (a) u and (b) the last principal component direction u_N , compared to projections onto a random vector u_{rand} . (g) The network response $[z(t), \text{top}]$ and its norm $[\|z(t)\|, \text{bottom}]$ given two different input directions (aligned and random). The norm of z is more than 11 times smaller in response to the aligned versus random perturbation. (h) Local network input (purple) projected onto u_N cancels with external input (green) to produce suppressed network responses (blue) in the direction of u_N .

was 13°). While the variability along v itself was not much smaller than a random direction (due to imperfect alignment between u_N and v), the variability along u_N was 190 times smaller than the variability in a random direction [Fig. 2(f)]. Also, low-rank cancellation was observed in the direction of u_N [Fig. 2(h)].

Low-rank suppression was observed in the sense that an input in the direction of $x_{aligned} = u + (W_1 v - v)/c$ produced a much weaker response than input in a random direction [Fig. 2(g)]. The choice of $x_{aligned}$ is discussed in Sec. IV E and Appendix D, but note that $x_{aligned}$ is closely aligned to u whenever c is large. Hence, low-rank suppression and low-rank cancellation arise even in the absence of a dominant negative eigenvalue.

To summarize our observations from the two examples above (Figs. 1 and 2), both networks produced low-rank suppression and low-rank cancellation, but only the network with a non-normal connectivity matrix (Fig. 2) produced low-dimensional dynamics. In the next section, we derive more precise conditions for low-rank suppression and low-dimensional dynamics.

IV. GENERAL CONDITIONS FOR LOW-RANK SUPPRESSION AND HIGH-DIMENSIONAL DYNAMICS

Above, we considered a simple network with rank-one structure. Here, we consider more general classes of networks and specify a list of assumptions that we make about the networks and precise conditions for low-rank suppression and low-dimensional dynamics under these assumptions. We first state two assumptions that we make throughout this section. Then, we show that low-rank suppression occurs under these assumptions alone. Finally, we analyze the conditions for high-dimensional dynamics under various assumptions on $x(t)$.

A. Assumption 1: Dynamics are linear or near a stable equilibrium

We consider a general class of dynamics of the form

$$\tau \frac{dz}{dt} = -z + F(z, x) \quad (4)$$

for some smooth vector field, $\mathbf{F}: \mathbb{R}^N \times \mathbb{R}^N \rightarrow \mathbb{R}^N$, where $\tau > 0$ sets the timescale of the dynamics. The dependence of $\mathbf{F}(\mathbf{z}, \mathbf{x})$ on $\mathbf{z}$ defines interactions between nodes, but it is difficult to define a fixed network of interactions W from the general dynamics in Eq. (4). In nonlinear systems, the effective connectivity between nodes can change with the state of the system. Moreover, nonlinear recurrent networks can produce any dynamics (high- or low-dimensional) in response to any input [39,40], so it is difficult to make any general statements about the dimensionality and structure of dynamics in recurrent networks far from a stable equilibrium. To resolve this difficulty, we assume that dynamics represent a small deviation away from a stable equilibrium so that we can linearize them. A change of coordinates under this linearization gives Eq. (1), where W is defined as the Jacobian of $\mathbf{F}$ at the fixed point [$W = \partial_{\mathbf{z}} \mathbf{F}(\mathbf{z}_0, \mathbf{x}_0)$; see Appendix B for details]. Our assumption of stability guarantees that all eigenvalues of W have a real part less than 1. To avoid systems close to instability, we additionally assume that the eigenvalues are bounded away from 1; in other words, W does not have any eigenvalues close to 1. For our examples, this implies that the spectral radius ρ of the random part W_1 is not close to 1. Eigenvalues near 1 or eigenvalues larger than 1 can produce low-dimensional dynamics through a different mechanism than the mechanisms that we are studying [1,2,21–23] (see Figs. 10–12 in Appendix F for examples).

B. Assumption 2: The connectivity matrix is strongly low rank in an asymptotic sense

We additionally assume that the network W has a strongly low-rank structure. To make this assumption precise, we consider networks in the limit of large N and assume that a small number of singular values diverge with N while the rest remain $\mathcal{O}(1)$ or smaller (see Appendix C for a more precise statement).

This assumption has two components. First is that the connectivity matrix has a small number of dominant singular values. In later sections, we show that this property arises naturally in several natural network structures such as networks with biased weights, modular structure, or smooth spatial structure.

Second, this assumption implies that the dominant singular values grow large with N . This contrasts to other work on “weakly” low-rank networks in which the largest singular values are $\mathcal{O}(1)$ [3,5,7–10]. Weakly low-rank structure arises naturally when mean weights scale like $\mathcal{O}(1/N)$, whereas strongly low-rank structure arises when mean weights are asymptotically larger than $1/N$, for example, when they scale like $\mathcal{O}(1/\sqrt{N})$ as in balanced networks [4,31,41–43]. See Appendix A for a more in-depth comparison between weakly and strongly low-rank networks.

Strongly low-rank networks can be decomposed as [1]

$$W = W_0 + W_1,$$

where W_0 has rank $r \sim \mathcal{O}(1)$ and large singular values, while the singular values of W_1 are $\mathcal{O}(1)$ at most and W_1 can be full rank. We specifically assume that W_1 is a random matrix and independent from W_0 . Since W_0 has rank r , we can write its singular value decomposition as

$$W_0 = U \Sigma V^T,$$

where Σ is a diagonal, $r \times r$ matrix of singular values

$$\sigma_k = \Sigma_{kk} \gg 1$$

while $U = [\mathbf{u}_1 \cdots \mathbf{u}_r]$ and $V = [\mathbf{v}_1 \cdots \mathbf{v}_r]$ are $N \times r$ orthonormal matrices with columns $\mathbf{u}_k$ and $\mathbf{v}_k$ that define the left and right singular vectors of W_0 .

C. Low-rank suppression occurs under Assumptions 1 and 2

Our first result is that Assumptions 1 and 2 are sufficient for low-rank suppression [as observed in Figs. 1(d) and 2(d)]. Specifically, we claim that low-rank suppression is realized under a static input of the form

$$\mathbf{x}_a = \mathbf{u}_k + \frac{W_1 \mathbf{v}_k - \mathbf{v}_k}{\sigma_k}, \quad (5)$$

where $k \leq r$ so that $\sigma_k \gg 1$ is one of the dominant singular values. We claim that inputs of this form are strongly suppressed by the network. To see why this is true, first note that $\|\mathbf{x}_a\| = 1 + o(1)$, since $\sigma_k \gg 1$, $\|\mathbf{u}_k\| = \|\mathbf{v}_k\| = 1$, and the maximum singular value of W_1 is $\mathcal{O}(1)$. Now note that the steady-state response to the input $\mathbf{x}_a$ is given by

$$\mathbf{z}_a = -\frac{\mathbf{v}_k}{\sigma_k}.$$

This can be checked by substituting the equations for $\mathbf{x}_a$ and $\mathbf{z}_a$ into the steady-state equation $\mathbf{z} = W\mathbf{z} + \mathbf{x}$. Since $\|\mathbf{v}_k\| = 1$ and $\sigma_k \gg 1$, we have that $\|\mathbf{z}_a\| \ll 1$. Therefore, the network strongly suppresses the input $\mathbf{x}_a$ from Eq. (5) in the sense that $\|\mathbf{z}_a\| \ll \|\mathbf{x}_a\|$.

Note from Eq. (5) that $\mathbf{x}_a$ is closely aligned to $\mathbf{u}_k$ since $\sigma_k \gg 1$. Hence, the suppressed input $\mathbf{x}_a$ is nearly (but not exactly) parallel to the left singular vector $\mathbf{u}_k$. Also note that the response $\mathbf{z}_a$ is perfectly aligned to $\mathbf{v}_k$. Hence, strongly low-rank recurrent networks suppress inputs aligned near the dominant left singular vectors ($\mathbf{u}_k$), and the suppressed response is aligned with the dominant right singular vectors ($\mathbf{v}_k$).

Interestingly, these input-response directions along which suppression occurs are reversed from the input-response directions along which amplification occurs in a feedforward network with the same connectivity (for example, the network in Fig. 7 in Appendix F for which $\mathbf{z} = W\mathbf{x}$ in the steady state). More specifically, in a feedforward network with connectivity W , inputs aligned to the dominant *right* singular vectors ($\mathbf{v}_k$) produce *amplified* responses in the direction of the corresponding *left* singular vectors ($\mathbf{u}_k$). This is a left-right reversal and a suppression-amplification reversal of the behavior for recurrent networks, for which inputs aligned to $\mathbf{u}_k$ produce suppressed responses in the direction of $\mathbf{v}_k$.

In summary, recurrent networks with strongly low-rank structure satisfying Assumptions 1 and 2 admit a small number of input directions along which inputs are strongly suppressed. We call this form of suppression low-rank suppression. Low-rank suppression occurs regardless of whether connectivity is normal, explaining why low-rank suppression was observed in both Figs. 1 and 2. Later, we explore consequences of low-rank suppression in common network structures. First, we discuss conditions for the separate but related phenomenon of high-dimensional dynamics.

D. Assumption 3: External inputs are stationary, high-dimensional, and slowly varying

To begin discussing the dimensionality of $\mathbf{z}(t)$, we need to make assumptions that allow us to define precisely what we mean by high- and low-dimensional dynamics. To this end, we assume that $\mathbf{x}(t)$ is a stationary ergodic stochastic process. In this case, $\mathbf{z}(t)$ is also a stationary ergodic process whenever dynamics are stable [44] as in Assumption 1. This allows us to quantify the dimensionality of $\mathbf{z}(t)$ in terms of its principal components or, equivalently, by the singular values of its stationary covariance matrix. Specifically, the variance explained by the first k principal components represents the amount of variability constrained to a k -dimensional linear manifold. For this reason, we define $\mathbf{z}(t)$ to have low-dimensional dynamics when there exists an $r \sim \mathcal{O}(1)$ such that $\sigma_j/\sigma_k \rightarrow 0$ as $N \rightarrow \infty$ whenever $k \leq r$ and $j > r$. This implies that the vast majority of variability lies on a low-dimensional linear manifold. Otherwise, we say that $\mathbf{z}(t)$ has high-dimensional dynamics.

We assume that $\mathbf{x}(t)$ is high-dimensional. If $\mathbf{x}(t)$ were low-dimensional, then low-dimensional dynamics in $\mathbf{z}(t)$ could be inherited from $\mathbf{x}(t)$ instead of being generated intrinsically (see Fig. 8 in Appendix F for an example). In assuming that $\mathbf{x}(t)$ is high-dimensional, we are focusing on the question of whether low-dimensional dynamics arise from network interactions, not from low-dimensional input alone.

For simplicity, within this section, we specifically assume that the components $\mathbf{x}_j(t)$ are i.i.d. This implies

that $\mathbf{x}(t)$ is high-dimensional, is a stronger assumption than necessary, but simplifies the exposition. In Appendix D, we consider more general inputs that are not assumed to be i.i.d.

For now, we also assume that $\mathbf{x}(t)$ varies in time more slowly than the intrinsic timescale τ of the network interactions. To make this assumption more precise, note that, whenever $\mathbf{x}(t) = \mathbf{x}$ is constant, solutions $\mathbf{z}(t)$ converge to the steady state given by $\mathbf{z} = [I - W]^{-1}\mathbf{x}$. We assume that $\mathbf{x}(t)$ varies sufficiently slowly that $\mathbf{z}(t)$ is approximated by the quasisteady-state approximation,

$$\mathbf{z}(t) \approx [I - W]^{-1}\mathbf{x}(t). \quad (6)$$

In Appendix D, we give more details about this assumption. In Appendix E, we extend our analysis and results to the case that $\mathbf{x}(t)$ is not slowly varying.

E. Conditions for low-dimensional dynamics under Assumptions 1–3

Under the quasisteady-state approximation in Eq. (6) and the assumption that elements of $\mathbf{x}(t)$ are i.i.d., it is easy to show that the covariance matrix of $\mathbf{z}(t)$ can be written as [44]

$$\bar{C}^{\mathbf{z}} \propto [I - W]^{-1}[I - W]^{-T}.$$

As a result, the question of whether $\mathbf{z}(t)$ is low-dimensional becomes a question of whether $[I - W]^{-1}$ has a small number of dominant singular values whenever W has a small number of dominant singular values. More details are given in Appendix D, where we prove that the answer to this question depends on the recurrent alignment matrix [4] of W_0 , defined by

$$P = V^T U.$$

This $r \times r$ matrix measures the alignment between the left and right singular vectors of W_0 . Specifically, $P_{jk} = \mathbf{v}_j \cdot \mathbf{u}_k$ measures the alignment between the j th right singular vector and the k th left singular vector of W_0 . This matrix is also sometimes referred to as the “overlap matrix” [12]. Our main result from Appendix D can be stated as follows.

Claim 1—Consider the dynamics defined by Eq. (4) under Assumptions 1–3. If the dynamics of $\mathbf{z}(t)$ are low-dimensional, then $P = V^T U$ has at least one asymptotically small singular value.

Note that this result is equivalent to the following.

Claim—Consider the dynamics defined by Eq. (4) under Assumptions 1–3. If all singular values of $P = V^T U$ are $\mathcal{O}(1)$, then the dynamics of $\mathbf{z}(t)$ are high-dimensional.

Recall that all singular values of P are bounded by 1, so the condition that all singular values of P are $\mathcal{O}(1)$ simply says that no singular values of P are asymptotically small; hence, the two versions of Claim 1 are logical

contrapositives. We were unable to prove the converse of these claims (that dynamics are low-dimensional whenever P has a small singular value), even though it is consistent with all of our simulations. Proving the converse or finding a counterexample could be a fruitful direction for future research.

Our proof of Claim 1 relied on an assumption that $\mathbf{x}(t)$ varies more slowly than the timescale τ of network interactions. In Appendix E, we consider the case where $\mathbf{x}(t)$ is not necessarily slow. In that case, the dimensionality can be broken down into the contribution from different frequency modes. When $\mathbf{x}(t)$ varies at both high and low frequencies, the results in Claim 1 are still valid (see Fig. 14 for an example).

To better understand the conditions on P in Claim 1, we consider simulations of a widely used model from computational neuroscience [3,37,45,46]:

$$\tau \frac{d\mathbf{z}}{dt} = -\mathbf{z} + W \tanh(\mathbf{z}) + \mathbf{x} \quad (7)$$

with $\tau = 1$. In these simulations, we take each $\mathbf{x}_j(t)$ to be a smooth, i.i.d. Gaussian process with correlation time $\tau_x = 10$. We take W_0 to have rank $r = 2$ so W has two dominant singular values [Fig. 3(a)], and we fix

$$U = [\mathbf{u}_1 \quad \mathbf{u}_2]$$

while exploring different choices of V .

We first consider a case in which W_0 is normal and symmetric by setting

$$V = [-\mathbf{u}_1 \quad -\mathbf{u}_2].$$

This is an extension of the example from Fig. 1 to rank $r = 2$. In this case, P is diagonal with $P_{kk} = -1$ [Fig. 3(b)] so all singular values of P are $\sigma_P = 1$ and the network produces high-dimensional dynamics [Fig. 3(c)]. Like the example from Fig. 1, W_0 has eigenvalues with large, negative real part [Fig. 3(d)]. In this example, W_0 is both normal and symmetric. Results are similar when W_0 is normal and asymmetric (obtained, for example, by applying a rotation to V in $r = 2$ dimensions) except the conclusion $|P_{kk}| = 1$ holds only when the nonzero eigenvalues of W_0 are real, which is necessarily the case when the σ_k are distinct.

High-dimensional dynamics do not require that W_0 is normal. Instead, it is sufficient that U and V share a column space, $\text{col}(U) = \text{col}(V)$. Low-rank matrices $W_0 = U\Sigma V^T$, for which U and V share a column space are called EP matrices [26]. The term “EP” is sometimes said to refer to “equal projectors.”

All normal matrices are EP, but an EP matrix is not necessarily normal. If W_0 is an EP matrix, then $\sigma_P = 1$ for

all singular values of P , so the condition in Claim 1 is satisfied.

An example of a non-normal, EP matrix W_0 is given by taking

$$V = [\mathbf{u}_2 \quad -\mathbf{u}_1].$$

The resulting matrix $W_0 = U\Sigma V^T$ is highly non-normal, because $\mathbf{u}_1 \cdot \mathbf{v}_1 = \mathbf{u}_2 \cdot \mathbf{v}_2 = 0$. For this example, P is zero along the diagonal [Fig. 3(e)], but all singular values of P are $\sigma_P = 1$, so the network still produces high-dimensional dynamics [Fig. 3(f)]. Note that W does not have any eigenvalues with a large, negative real part [Fig. 3(g)]. Hence, this example combined with the previous example [Figs. 3(b)–3(d)] show that high-dimensional dynamics can arise in the presence or absence of eigenvalues with a large, negative real part.

To obtain an example of a network that produces *low-dimensional* dynamics (see also the example in Fig. 2), we set

$$V = [-\mathbf{u}_1 \quad \mathbf{v}_\perp],$$

where $\mathbf{v}_\perp$ is orthogonal to $\mathbf{u}_1$ and $\mathbf{u}_2$, so P has one singular value at $\sigma_P = 1$ and another at $\sigma_P = 0$ [Fig. 3(h)]. Because of the singular value at zero, the condition in Claim 1 is not met, and the network produces low-dimensional dynamics in which one principal component captures an outsized proportion of the variability [Fig. 3(i)]. For this example, W has an eigenvalue with a large, negative real part [Fig. 3(j)]. However, low-dimensional dynamics also arise in the case [Figs. 3(k) and 3(l)]

$$V = [-\mathbf{u}_2 \quad \mathbf{v}_\perp],$$

in which case W lacks eigenvalues with a large, negative real part [Fig. 3(m)]. Hence, these two examples demonstrate that low-dimensional dynamics can arise in the presence or absence of eigenvalues with a large, negative real part.

So far, in all of the examples with high-dimensional dynamics, W_0 is an EP matrix. However, W_0 need not be fully EP for dynamics to be high-dimensional. We now consider the case in which left and right singular vectors are not perfectly aligned but have some nonvanishing overlap:

$$V = \sqrt{1-c}[-\mathbf{u}_1 \quad \mathbf{v}_\perp] + \sqrt{c}[-\mathbf{u}_1 \quad -\mathbf{u}_2],$$

where $0 < c < 1$. This is a linear combination of the EP and non-EP matrices from the examples above. In this case, W_0 is not EP, but P has a singular value at $\sigma_P = \sqrt{c}$ and the other at $\sigma_P = 1$ [Fig. 3(n)]. As long as c is not close to zero, the condition in Claim 1 is satisfied, and the network produces high-dimensional dynamics [Fig. 3(o)]. Also, W has eigenvalues with a large, negative real part [Fig. 3(p)].




FIG. 3. Conditions for high-dimensional dynamics in a network with rank-two structure. (a) Singular values of W demonstrate an effective low-rank structure. (b) The alignment matrix P when W_0 is normal and symmetric. (c) The variance explained by each principal component of the network dynamics, $\mathbf{z}(t)$. (d) The eigenvalues of W . (e)–(g) The same as (b)–(d) except W_0 is EP, but non-normal. (h)–(j) The same as (b)–(d) except W_0 is not EP. (k)–(m) The same as (h)–(j) but using a different non-EP structure. (n)–(p) The same as (b)–(d) except W_0 is non-EP with a nonvanishing EP component. See the text for the precise definition of W_0 in each case.

From this example, we may conclude that the low-rank part of the connectivity matrix needs to have only a nonvanishing EP component to produce high-dimensional dynamics.

The numerical results in Fig. 3 provide an interpretation of Claim 1: If W_0 has a nonvanishing component with the EP property, then dynamics are high-dimensional. To our knowledge, this connection between the EP property of

matrices and the dimensionality of network dynamics has not been made before.

The examples in Fig. 3 also demonstrate that the eigenvalue spectrum is generally not useful for determining whether dynamics are high- or low-dimensional in strongly low-rank networks near a stable equilibrium with high-dimensional input. Instead, as Claim 1 shows, the singular values of the recurrent alignment matrix, P , are key to

determining the dimensionality of dynamics in these networks.

When low-dimensional dynamics are present, they arise from an effectively feedforward mechanism that takes variability from $\mathbf{x}(t)$ and $W_1 \mathbf{z}(t)$ aligned with the column space of V and maps it to strong variability along the column space of U (see Appendix D for more details). This mechanism is a form of non-normal amplification [28–30]. Our results add some new results to the literature on non-normal amplification under the conditions in Assumptions 1–3. First, our results show that this form of non-normal amplification requires that W is highly non-EP, not just highly non-normal. See, for example, the example in Figs. 3(e)–3(g), in which W_0 is highly non-normal, but EP so that amplification does not occur. Second, our results on low-rank suppression show that the amplified responses in the direction of U associated with non-normal amplification are accompanied by suppressed responses approximately aligned to V [see Figs. 2(e) and 2(f)]. Finally, Claim 1 shows that non-normal amplification is the *only* way to generate low-dimensional dynamics under Assumptions 1–3.

It is common in theoretical work to consider models in which the entries of U and V are random and independent with zero mean. In this case, $\mathbf{u}_j$ and $\mathbf{v}_k$ are nearly orthogonal, so $\sigma_p \ll 1$ and the network produces low-dimensional dynamics (Fig. 9 in Appendix F; compare to Fig. 2). However, many networks arising in nature do not have purely random structure. Next, we show that many naturally arising network structures satisfy our conditions for high-dimensional dynamics.

V. BIASED WEIGHTS AND MODULAR NETWORKS

So far, we considered networks with unbiased weights $E[W_{jk}] = 0$, which is common in modeling studies, but many networks in nature have weights with nonzero mean. Biased weights can produce low-rank structure. As a simple example, consider a random network with independent weights satisfying

$$E[W_{jk}] = m \quad \text{and} \quad \text{std}(W_{jk}) = s.$$

If m and s scale similarly, then the largest singular value of W is near $\sigma_1 = |m|N$, while the next-largest singular value scales like $\sigma_2 = \mathcal{O}(s\sqrt{N})$, implying an effective rank-one structure when N is large. The dominant rank-one part has constant entries, so it is normal and the network exhibits low-rank suppression and high-dimensional dynamics (Fig. 15 in Appendix F; see Appendix B in Ref. [47] for related results).

More generally, modular structure arises in many natural settings [1,48]. Specifically, many networks in nature represent the interaction between n populations, and mean

connection weights between populations are often nonzero. The adjacency matrices of modular networks can be arranged to have a block structure

$$W = \begin{bmatrix} W^{1,1} & \dots & W^{1,n} \\ W^{2,1} & \dots & W^{2,n} \\ \vdots & \dots & \vdots \\ W^{n,1} & \dots & W^{n,n} \end{bmatrix},$$

where $W^{a,b}$ is a $N_a \times N_b$ submatrix quantifying connections from population b to population a . In general, each submatrix can have a different, nonzero mean and standard deviation:

$$E[W_{jk}^{a,b}] = m_{ab} \quad \text{and} \quad \text{std}(W_{jk}^{a,b}) = s_{ab}.$$

If we assume that each population has $N_a \sim \mathcal{O}(N)$ members and that the m_{ab} scale similarly to the s_{ab} , then W has up to n dominant singular values that scale like $\mathcal{O}(m_{ab}N)$ and the remaining singular values scale like $\mathcal{O}(s_{ab}\sqrt{N})$. Hence, when $s_{ab} \sim m_{ab}$, large modular networks with biased weights naturally produce a low-rank structure in the sense that a small number of singular values are asymptotically larger than the rest [1]. If, additionally, $m_{ab} \gg \mathcal{O}(1/N)$, then the dominant singular values are asymptotically large, so the network is strongly low-rank in the sense defined in Assumption 2.

Networks of this form can be decomposed as $W = W_0 + W_1$, where $W_0 = E[W]$ is constant within each block, so W_0 has rank at most n . In general, W_0 is not a normal matrix, but it is an EP matrix because $\text{col}(U)$ and $\text{col}(V)$ are each spanned by the indicator vectors of the n populations:

$$\text{col}(U) = \text{col}(V) = \text{span}\{\mathbf{1}_1, \mathbf{1}_2, \dots, \mathbf{1}_n\},$$

where each $\mathbf{1}_k$ is an N -dimensional indicator vector with entries defined by

$$[\mathbf{1}_k]_j = \begin{cases} 1 & \text{if index } j \text{ is in population } k, \\ 0 & \text{otherwise.} \end{cases}$$

Therefore, modular networks with biased weights exhibit low-rank suppression and high-dimensional dynamics.

As a specific example of a modular network, we consider a model of a local neuronal network in the cerebral cortex. Cortical neurons obey Dale's law: All outgoing connections from a particular neuron have the same polarity, positive for excitatory neurons and negative for inhibitory neurons, and mean connection weights also depend on the postsynaptic neuron type [49,50]. These properties produce a modular structure in which the columns of the adjacency matrix corresponding to excitatory neurons are




FIG. 4. Low-rank suppression and excitatory-inhibitory balance in a modular network. (a),(b) A modular network with biased blocks modeling excitatory and inhibitory neurons has low-rank structure. (c) The alignment matrix shows that the network is EP but not normal. (d) An input that is constant within each block (purple) is aligned to the low-rank part, but a random input (green) is not. (e) The norm of $\mathbf{z}$ is more than 13 times smaller in response to the aligned versus random perturbation. (f) The excitatory (positive; red) component of mean input (local and external combined) balances with the inhibitory (negative; blue) component to produce a much smaller total input (gray), a widely observed phenomenon in neural circuits. Network dynamics obey Eq. (7).

non-negative, while the columns corresponding to inhibitory neurons are nonpositive. Without loss of generality, we can order the neurons so that the first N_e neurons are excitatory and the remaining N_i are inhibitory (with $N = N_e + N_i$). In this case, W has a 2×2 block structure [Fig. 4(a)]. We also assume that $N_e, N_i \sim \mathcal{O}(N)$, consistent with the fact that around 80% of neurons in cortex are excitatory. We consider the case in which m_{ab} and s_{ab} scale like $1/\sqrt{N}$, consistent with experiments [51] and theoretical work [4,24,31,41–43].

Together, these biologically justified assumptions imply that the network has a strongly low-rank structure, as defined in Assumption 2, with rank $r = 2$ [Fig. 4(b)]. Specifically, two singular values scale like $\sqrt{N}$ while the others are $\mathcal{O}(1)$. Moreover, the low-rank part of the connectivity matrix is EP [Fig. 4(c)]. Hence, the network produces high-dimensional dynamics and low-rank suppression despite the fact that W is low-dimensional and non-normal. The column space of U and V consists of all vectors that are uniform within each population:

$$\begin{aligned} \text{col}(U) = \text{col}(V) &= \text{span}\{\mathbf{1}_e, \mathbf{1}_i\} \\ &= \{[a \cdots ab \cdots b]^T | a, b \in \mathbb{R}\}, \end{aligned}$$

where $\mathbf{1}_e$ and $\mathbf{1}_i$ are indicator vectors for the excitatory and inhibitory populations, respectively. Therefore, external inputs $\mathbf{x}$ that are uniform within each population are aligned to the low-rank part and suppressed relative to random perturbations [Figs. 4(d) and 4(e)]. In other words, external inputs that stimulate all excitatory neurons equally and all inhibitory neurons equally are suppressed relative to inputs that are inhomogeneous within one or both populations, consistent with previous work on balanced network

models [24,32–34]. As predicted, simulations also show high-dimensional responses to high-dimensional stimuli (Fig. 16 in Appendix F), consistent with observations that neural responses in monkey visual cortex are high-dimensional when visual stimuli are high-dimensional [16].

Averaging over the excitatory and inhibitory populations in this network represents a projection onto $\text{col}(U) = \text{col}(V)$. As a result, the cancellation mechanism illustrated in Fig. 1(h) manifests as a tight balance between mean excitatory (positive) and inhibitory (negative) input to neurons [Fig. 4(f)], a phenomenon that is widely observed in neural recordings [52–58] and widely studied in computational models [4,31,34,41]. Hence, the widely studied theory of balanced networks can be interpreted as a special case of the low-rank suppression and low-rank cancellation studied here, specific to the particular modular network structure arising from Dale’s law.

VI. AMPLIFICATION OF DISORDERED PERTURBATIONS IN NETWORKS WITH SPATIAL STRUCTURE

Many networks in nature exhibit connection strength that depends smoothly on the distance between nodes in physical or other spaces, resulting in an effective low-rank structure [1,42]. As a simple example, we consider a model in which each node is assigned a spatial location, $\theta \in [0, 1]$, and connection weights in W_0 decay like a Gaussian function of the distance, $d\theta$, between nodes computed with periodic boundaries. We assume strong coupling in the sense that individual weights scale like $1/\sqrt{N}$. In this case, only nearby nodes are strongly connected [Fig. 5(a)]. Note that stability requires that connections in W_0 are negative since coupling is strong. Connectivity is also perturbed by a




FIG. 5. Low-rank suppression in a network with spatial structure. (a) Connectivity matrix W . Connection strength is a Gaussian function of distance. (b) Singular values of W demonstrate effective low-rank structure. (c) A spatially smooth perturbation (purple) is aligned to the low-rank part of W , while a spatially disordered perturbation (green) is not. (d),(e) Therefore, the network response to a smooth perturbation is much weaker than the response to a disordered perturbation (compare vertical tick marks). The norm of $\mathbf{z}$ is more than 7 times smaller in response to the smooth versus disordered perturbation. Network dynamics obeyed Eq. (7).

random component W_1 as above. This connectivity structure is effectively low-rank [Fig. 5(b)] and the low-rank part is normal, with left and right singular vectors spanned by the low-frequency Fourier basis vectors (Fig. 17 in Appendix F). Therefore, networks of this form satisfy our conditions for low-rank suppression and high-dimensional dynamics (Fig. 18 in Appendix F).

Perturbations that are smooth in space are aligned to the low-rank part of the connectivity matrix because they are formed by sums of low-frequency spatial Fourier modes. Hence, perhaps surprisingly, strongly connected networks with spatial structure are more sensitive to spatially disordered perturbations than to spatially smooth perturbations [Figs. 5(c)–5(e); compare tick labels on vertical axes].

VII. LOW-RANK SUPPRESSION AND HIGH-DIMENSIONAL DYNAMICS IN AN EPIDEMIOLOGICAL NETWORK

We next consider a real epidemiological network, specifically a network of high school social contacts [59], which was used in recent theoretical work on low-rank network dynamics [1]. In that work, the authors considered a quenched mean-field reduction of the susceptible-infected-susceptible model:

$$\tau \frac{d\mathbf{z}}{dt} = -\mathbf{z} + \gamma(1 - \mathbf{z}) \circ Q\mathbf{z}, \quad (8)$$

where $\gamma > 0$ and $\circ$ denotes elementwise multiplication. Each $\mathbf{z}_j(t)$ models the probability that an individual is infected. Here, Q is the proportional to the proximity matrix of 637 high school students, indicating whether each pair of students were in proximity of each other during a

specific school week [59]. This matrix is effectively low-rank in the sense that it has a small number of dominant singular values [Fig. 19(a) in Appendix F]. In Ref. [1], it was shown that dynamics generated by Eq. (8) on this network are effectively low-dimensional.

However, Eq. (8) is completely self-contained without any perturbations. Perturbations arise in epidemiological dynamics through interactions with individuals from outside of the modeled network and through the natural stochasticity of infection. Specifically, $\mathbf{z}(t)$ represents a probability of infection. Actual infection events will be random, which can be approximated by adding a stochastic forcing term to Eq. (8). Moreover, students interact with siblings and parents at home (who are not included in the network represented by Q), which produces additional external forcing. As a simple model of these perturbations, we modified the model by adding an external forcing term to obtain

$$\tau \frac{d\mathbf{z}}{dt} = -\mathbf{z} + \gamma(1 - \mathbf{z}) \circ [Q\mathbf{z} + \mathbf{x}(t) \circ \mathbf{z}]. \quad (9)$$

Equation (9) admits a fixed point when $\mathbf{x}(t) = 0$. Following the linearization discussed in Sec. IV A and Appendix B, the effective connectivity W is defined by the Jacobian matrix at the fixed point. We find that, like Q , the effective connectivity matrix W has a small number of dominant singular values, so our theory predicts that this network exhibits low-rank suppression. Moreover, the recurrent alignment matrix P defined from W satisfies the conditions for high-dimensional dynamics from Claim 1 [Fig. 19(a) in Appendix F]. Simulations of Eq. (9) confirm that the model exhibits low-rank suppression [Figs. 6(a)




FIG. 6. Low-rank suppression and high-dimensional dynamics in a real epidemiological network. (a) The response to a perturbation aligned with the dominant low-rank part is weaker than the response to a perturbation of the same magnitude in a random direction. The distance of z from its fixed point is more than 10 times smaller in response to the aligned versus random perturbation. (b) In the presence of high-dimensional random perturbations, the variance explained by each principal component of the dynamics is similar to the input, indicating high-dimensional dynamics. Network dynamics obeyed Eq. (7).

and 19(b)–19(d) in Appendix F] and high-dimensional dynamics [Fig. 6(b)] in contrast to the model without external input studied in Ref. [1].

Importantly, the “aligned” input direction that we used to demonstrate low-rank suppression in Fig. 6(a) is defined in terms of the dominant left singular vector of the *effective* connectivity matrix W , i.e., the Jacobian at the fixed point. Using the left singular vector of the explicit connectivity, Q , in place of the effective connectivity W does not produce suppression [see Figs. 19(c) and 19(d) in Appendix F]. This observation validates our linearization argument and highlights the necessity of using effective connectivity in place of explicit connectivity in nonlinear networks.

Taken together, our results highlight the importance of accounting for external perturbations and internal noise when studying the dimensionality of epidemiological dynamics. Moreover, the results imply that epidemiological networks can be more sensitive to random perturbations than to perturbations aligned with the network’s low-rank structure. For example, if the forcing term is taken to represent interventions like masking or inoculation, then our results predict that it is more effective to apply these interventions randomly to individuals than to apply interventions to individual subpopulations or other structures reflected in the dominant structure of the proximity matrix.

VIII. DISCUSSION

We presented theory and examples showing that strongly low-rank recurrent networks close to a stable equilibrium suppress inputs or perturbations aligned with the dominant directions of their connectivity matrices. We also derived conditions under which these networks generate high-dimensional dynamics. We showed that many low-rank structures that arise in nature are consistent with low-rank suppression and high-dimensional dynamics.

Recent parallel work [11] also demonstrated that low-rank networks can produce suppression and high-dimensional dynamics under some assumptions. In that work, connectivity matrices had a small number of nonzero singular values which were $\mathcal{O}(1)$ in magnitude, in contrast to the large singular values assumed in our models (see Appendix A for a more complete comparison). Other related work showed that trained neural networks can exploit suppression to cancel noise [25].

Our results rely on a linearization of the dynamics around a stable fixed point. The dimensionality of nonlinear network dynamics away from a stable equilibrium is also an important question. However, nonlinear networks can produce any dynamics in response to any inputs [39,40] and effective connectivity can change with network state, so it is difficult to draw broad, general conclusions about the dimensionality of dynamics in nonlinear networks. Previous work [1] circumvented this difficulty by parametrizing the model in terms of the vector of local inputs from the network, $y = Wz$, and then used bounds on the Jacobian matrix of the dynamics with respect to y , but they did not consider external inputs or perturbations. Combining our approaches by incorporating external inputs into their analytical approach is a promising direction for future work.

We focused on the response of low-rank networks to external input, whereas much previous work focuses on the intrinsic dynamics of low-rank networks in the absence of external drive. Landau and Sompolinsky [4] consider a similar model to ours but look at the interaction between externally driven responses and intrinsic, chaotic dynamics. However, they restrict to low-dimensional external inputs in the column space of the low-rank part of the connectivity matrix. Future work could combine these approaches to better understand the interaction between chaotic, intrinsic dynamics and high-dimensional external input.

It is tempting to ascribe low-rank suppression to the presence of an eigenvalue with a large, negative real part [as in Figs. 1 and 3(b)–3(d)]. A large, negative eigenvalue is well known to produce a suppressive dynamic. However, the results in Fig. 2 and Sec. IV E demonstrate that the eigenvalue spectrum is not generally relevant to the presence or absence of low-rank suppression. In other words, in recurrent networks driven by external input near a stable equilibrium, low-rank suppression depends only on the presence of large singular values of the connectivity

matrix W . High-dimensional dynamics, on the other hand, depend additionally on the singular values of the recurrent alignment matrix $P = V^T U$.

Our results have implications for networks in nature and applications. For example, in neuroscience, the implications of low-rank recurrent connectivity on neural dynamics is a topic of intense research [3–12]. Our results are consistent with an observation that networks of neurons produce high-dimensional activity in the context of high-dimensional stimuli or tasks but low-dimensional dynamics in response to low-dimensional stimuli or tasks [16].

Beyond neuroscience, our results imply that perturbations to a low-rank network are more effective when they are delivered nonuniformly across subpopulations or space. More generally, perturbations to a low-rank network are more effective when they are not aligned to the network's low-rank structure. This result can be leveraged to design and test more effective interventions, for example, to epidemiological, ecological, or social networks.

IX. METHODS

All simulations and analyses were performed in PYTHON using a combination of custom written PYTORCH and NUMPY code. All differential equations, except for Figs. 6 and 19 in Appendix F, were solved using a simple forward Euler scheme with a step size of $dt = 0.01$. The simulations in Figs. 6 and 19 in Appendix F were solved using a Runge-Kutta scheme adapted from the approach used in previous work [1]. Networks in Figs. 1 and 2 used the linear model from Eq. (1). Networks in Figs. 3–5 used the nonlinear model from Eq. (7).

ACKNOWLEDGMENTS

We thank Ashok Litwin-Kumar and David G. Clark for helpful conversations and comments on drafts of the manuscript. This work was supported by the Air Force Office of Scientific Research (AFOSR) Awards No. FA9550-21-1-0223 and No. FA9550-26-1-0004 and the National Science Foundation under Neuronex Award No. DBI-1707400.

DATA AVAILABILITY

Code to produce all figures is available [60].

APPENDIX A: COMPARISON TO PREVIOUS WORK ON DYNAMICS OF LOW-RANK NETWORKS

Here, we focus on the response of networks to high-dimensional external input. Some previous work on low-rank networks and low-dimensional dynamics focuses on the dimensionality of *intrinsic* dynamics generated by the network in the absence of external input or in the presence of low-dimensional external input. For example, chaotic

and other nonlinear dynamics in untrained and trained networks can be low-dimensional (see Refs. [1,2,21–23] and Figs. 10 and 11 in Appendix F). In contrast, we are here primarily interested in the network's response to high-dimensional external input. In our model, the network's activity is driven primarily by the external input (which is filtered by the network's internal structure) and not by intrinsic dynamics. When the network is strongly low-rank, we showed that low-dimensional dynamics can arise only when the recurrent alignment matrix has small singular values (see Claim 1).

In this section, we review the relationship between our models and the models from previous theoretical work on low-rank networks, especially work by Ostojic *et al.* [3,5,7–10], work by Thibault, Allard, and Desrosiers [1], and work by Landau and Sompolinsky [4,6].

All of these studies consider recurrent networks with connectivity of the form $W = W_0 + W_1$, where W_0 is low-rank and W_1 is a full-rank random matrix, just like our model. We next describe the salient properties that distinguish these models from ours.

Thibault, Allard, and Desrosiers [1] assume that their networks are self-contained and do not receive any time-varying external input. Specifically, they explicitly restrict their dynamics to models of the form

$$\frac{d\mathbf{z}}{dt} = F(\mathbf{z}, W\mathbf{z}),$$

which excludes the possibility of modeling a time-varying external input, like $\mathbf{x}(t)$ in our model. Models of this form cannot describe networks that are part of a larger network or networks that are modulated by time-varying, external factors. Similar assumptions were made in one study by Landau and Sompolinsky [6]. In other work by Landau and Sompolinsky [4], external input was included, but this input was assumed to be perfectly aligned to the low-rank structure of the connectivity matrix and, therefore, low-dimensional.

The salient differences between our model and the models considered by Ostojic and colleagues [3,5,7–10] are more subtle. Like us, they consider external inputs that are not aligned to the low-rank part of W . Also, like us, they take $W = W_0 + W_1$, where W_1 is full rank with random entries and the variance of the entries in W_1 scale like $\mathcal{O}(1/N)$ so that the maximum singular value of W_1 scales like $\mathcal{O}(1)$.

However, in contrast to our models, the low-rank components of the networks considered by Ostojic and colleagues take the form

$$W_0 = \sum_{\mu=1}^r \frac{\mathbf{m}_\mu \mathbf{n}_\mu^T}{N}, \quad (\text{A1})$$

where r is the rank and each $\mathbf{m}_\mu$ and $\mathbf{n}_\mu$ are $N \times 1$ vectors with entries that scale like $\mathcal{O}(1)$. Specifically, they are taken

to be random vectors, and the variance of their entries scales like $\mathcal{O}(1)$. Often, they are taken to be biased random vectors with a nonzero mean that also scales like $\mathcal{O}(1)$. Because of the $1/N$ factor in Eq. (A1), the variance of the entries in W_0 scales like $\mathcal{O}(1/N^2)$, and, when the entries are biased, the mean entry in W_0 scales like $\mathcal{O}(1/N)$. Regardless of whether entries are biased, the singular values of W_0 scale like $\mathcal{O}(1)$ in the models considered by Mastrogiuseppe and Ostojic [3]. Hence, the singular values of the low-rank part (W_0) and the random part (W_1) have the same scale whenever the spectral radius of W_1 is $\mathcal{O}(1)$, in contrast to our models in which the singular values of W_0 are considered to be asymptotically larger than the singular values of W_1 . This difference in scaling is acknowledged by Ostojic *et al.* [3,5,7–10], who refer to their networks as “weakly low-rank.”

To summarize, when W_1 is a random matrix with $\mathcal{O}(1)$ spectral radius and when the entries of W_0 are biased, the primary difference between many strongly and weakly low-rank networks in practice is that the entries of W_0 scale like $\mathcal{O}(1/N)$ in weakly low-rank networks while they are larger [typically $\mathcal{O}(1/\sqrt{N})$] in many strongly low-rank networks, for example, in work on balanced networks [4,6]. The transition between these two regimes can be visualized by gradually scaling the magnitude of the low-rank part from weak to strong. Figure 22 in Appendix F shows that low-rank suppression arises gradually in this case.

The distinction between strongly and weakly low-rank networks is more nuanced when the full-rank, random part W_1 of the connectivity matrix is assumed to be weak or absent. For example, recent work by Mastrogiuseppe, Carmona, and Machens [11] considered the case in which the full-rank part is absent ($W_1 = 0$) and the low-rank connectivity matrix $W = W_0$ has $\mathcal{O}(1)$ singular values. While their use of $\mathcal{O}(1)$ singular values is similar to weakly low-rank networks from other work [3], the absence of a random part means that their connectivity matrices are truly low-rank (not just approximately low-rank). In that work [11], Mastrogiuseppe, Carmona, and Machens showed that the emergence of suppression and high-dimensional dynamics depends on the degree of overlap between the left and right singular vectors in ways that are not captured by our asymptotic theory.

For the sake of comparison to the models in Ref. [11], we consider a rank-one network structure like the one in Fig. 1. Specifically, we take W_1 to be random with spectral radius $\rho > 0$ and take

$$W_0 = c\mathbf{u}\mathbf{u}^T.$$

In Fig. 1 and its analysis, we assume that $\rho \sim \mathcal{O}(1)$. In this case, a strong (i.e., asymptotically dominating) low-rank structure requires that $|c| \gg 1$. Therefore, stability in this case requires that $c < 0$, which produces strong, negative eigenvalues [as in Fig. 1; however, consider also the

examples in Figs. 2, 3(k)–3(m), and 9, which are strongly low-rank without large eigenvalues]. We could alternatively take the random, full-rank part to be weak: $\rho \ll 1$ (or it could be absent, $\rho = 0$, as in Ref. [11]). In this case, we can take the low-rank part to be moderate in magnitude [$|c| = \mathcal{O}(1)$] while maintaining an asymptotically dominant low-rank structure because $|c| \gg \rho$ when $|c| = \mathcal{O}(1)$ and $\rho \ll 1$. Repeating the simulation from Fig. 1 demonstrates a weakened form of low-rank suppression with high-dimensional dynamics when $c < 0$ and an amplification of aligned inputs with weakly low-dimensional dynamics when $c > 0$ (Figs. 20 and 21 in Appendix F). We also repeated all of the simulations from Fig. 3 in networks with $\rho \ll 1$ and $\sigma_1, \sigma_2 \sim \mathcal{O}(1)$, so that the network is still low-rank in the sense that the first two singular values dominate but not strongly low-rank because they are $\mathcal{O}(1)$ [see Fig. 23(a)]. Under these conditions, the dynamics were high-dimensional regardless of the structure of W_0 [Figs. 23(b)–23(f)]. To understand this result, note that low-dimensional dynamics in our models [Figs. 2, 3(i), and 3(l)] arise from strong amplification of a small number of directions by W . When the low-rank part of W is weaker, these low-dimensional dynamics will be weaker or absent. In other words, when the dominant singular values of W_0 are $\mathcal{O}(1)$ and smaller, the identity in $[I - W]^{-1}$ from the quasisteady-state analysis more strongly dominates the W (or, at least, the W does not dominate the I). More directly, the $\mathbf{x}(t)$ in Eq. (1) dominates the $W\mathbf{z}$. Therefore, high-dimensional inputs are simply transferred to high-dimensional responses when W is weaker.

When $c < 0$, the network exhibits a weaker form of low-rank suppression along with high-dimensional responses (Fig. 20 in Appendix F). Similar examples and results are considered in Ref. [11]. In conclusion, some of our theoretical results on low-rank suppression and high-dimensional dynamics require that the low-rank part is large in *absolute* terms ($|c| \gg 1$), not just *relative* terms ($|c| \gg \rho$).

When $|c| = \mathcal{O}(1)$, we can take $c > 0$ while keeping the network stable with a dominant low-rank part (in contrast to strongly low-rank networks like Fig. 1, where stability requires $c < 0$). When $c > 0$, inputs aligned to $\mathbf{u}$ are *amplified* compared to random perturbations (Fig. 21 in Appendix F), reversing the trend of low-rank suppression. However, the amplification is weak. In response to high-dimensional input, the network dynamics have a dominant principal component direction, but the dominance is weaker than in other examples we have considered [Fig. 21(d); compare to Figs. 2 and 9] because stability imposes a bound on c (specifically, $c < 1$) and, therefore, on $\|W_0\|$.

Landau and Sompolinsky [4,6] consider strongly low-rank networks. However, as noted above, their external input $\mathbf{x}(t)$ is aligned to the low-rank part of W and is, therefore, low-dimensional itself. Thibault, Allard, and

Desrosiers [1] consider several network models, but their “rank-perturbed Gaussian” model is strongly low-rank and equivalent to the network structure we study in Fig. 1. Complicating matters, Thibeault, Allard, and Desrosiers directly compare their rank-perturbed Gaussian model to the networks in Ostojic *et al.*, despite the fact that the scaling of their low-rank parts differ by a magnitude of $\sqrt{N}$. Specifically, in Thibeault, Allard, and Desrosiers [1], the low-rank part is defined by Eq. (A1), where $\mathbf{m}$ and $\mathbf{n}$ are unbiased Gaussian random vectors. The entries of $\mathbf{m}$ have $\mathcal{O}(1/N)$ variance while the entries of $\mathbf{n}$ have $\mathcal{O}(1)$ variance, so the entries of W_0 have $\mathcal{O}(1/N)$ variance, in contrast to the $\mathcal{O}(1/N^2)$ variance used by Mastrogiuseppe and Ostojic [3]. Importantly, this means that the singular values of W_0 are $\mathcal{O}(\sqrt{N})$ in the rank perturbed Gaussian model analyzed by Thibeault, Allard, and Desrosiers [1] but $\mathcal{O}(1)$ in the models by Ostojic *et al.* Hence, the rank perturbed Gaussian models considered by Thibeault, Allard, and Desrosiers are strongly low-rank, in contrast to the weakly low-rank networks considered by Ostojic *et al.*

APPENDIX B: DETAILS FOR THE LINEARIZATION OF DYNAMICS AROUND A STABLE EQUILIBRIUM

Here, we make precise our assumption that dynamics represent a small perturbation around a stable equilibrium and give a more precise derivation of the linearized dynamics.

We begin by considering nonlinear dynamics of the form

$$\tau \frac{d\mathbf{z}}{dt} = \mathbf{F}(\mathbf{z}, \boldsymbol{\eta}),$$

where $\boldsymbol{\eta}(t)$ is the external input before any change of coordinates is applied. We assume that there exists a constant in time, $\boldsymbol{\eta}(t) = \boldsymbol{\eta}_0$, that produces a stable steady-state solution $\mathbf{z}_0$. In other words, there exist $\boldsymbol{\eta}_0$ and $\mathbf{z}_0$ such that

$$\mathbf{F}(\mathbf{z}_0, \boldsymbol{\eta}_0) = \mathbf{z}_0.$$

We additionally assume that this fixed point is hyperbolically stable. In other words, the Jacobian matrix

$$J = -I + \partial_{\mathbf{z}} \mathbf{F}(\mathbf{z}_0, \boldsymbol{\eta}_0)$$

has eigenvalues with a strictly negative real part. Here, I is the identity matrix, and $\partial_{\mathbf{z}} \mathbf{F}(\mathbf{z}_0, \boldsymbol{\eta}_0)$ is the Jacobian matrix of $\mathbf{F}$ with respect to $\mathbf{z}$ evaluated at the fixed point. We then consider a small perturbation around this fixed point driven by a perturbation to the forcing term:

$$\boldsymbol{\eta}_p(t) = \boldsymbol{\eta}_0 + \epsilon \boldsymbol{\eta}(t).$$

The response $\mathbf{z}_p(t)$ of the network to the perturbed forcing term $\boldsymbol{\eta}_p(t)$ can be written to linear order in ϵ as

$$\mathbf{z}_p(t) = \mathbf{z}_0 + \epsilon \mathbf{z}(t) + \mathcal{O}(\epsilon^2).$$

The perturbation $\mathbf{z}(t)$ obeys the linearized equation

$$\tau \frac{d\mathbf{z}}{dt} = -\mathbf{z} + W\mathbf{z} + \mathbf{x}(t), \quad (\text{B1})$$

where

$$W = \partial_{\mathbf{z}} \mathbf{F}(\mathbf{z}_0, \boldsymbol{\eta}_0)$$

is interpreted as the effective connectivity matrix and

$$\mathbf{x}(t) = \partial_{\boldsymbol{\eta}} \mathbf{F}(\mathbf{z}_0, \boldsymbol{\eta}_0) \boldsymbol{\eta}(t)$$

is interpreted as the effective input. Our assumption that the equilibrium is stable implies that all eigenvalues of W have a real part less than 1. We additionally assume that the real part of all eigenvalues of the Jacobian are bounded away from 0, meaning that the real part of all eigenvalues of W are bounded below 1. In other words, there is a $d > 0$ with $d \sim \mathcal{O}(1)$ such that all eigenvalues of W satisfy $\text{Re}(\lambda_W) < 1 - d$. In particular, this implies that $\rho < 1 - d$, where ρ is the spectral radius of the random part W_1 of W .

Note that Eq. (B1) is equivalent to the linear dynamics in Eq. (1). Hence, the linear network dynamics in Eq. (1) are a sort of normal form for the linearized dynamics described here. Therefore, we are justified in focusing on the linear dynamics of Eq. (1) in the analysis in the results.

In many models, the perturbation is purely additive, $\mathbf{F}(\mathbf{z}, \boldsymbol{\eta}) = \mathbf{H}(\mathbf{z}) + \boldsymbol{\eta}$, so that the effective input is equal to the raw input, $\mathbf{x} = \boldsymbol{\eta}$. In other models, the raw input can have a different structure and dimensionality to the effective input. For illustrative purposes, we consider a contrived example of such a model in Fig. 8 in Appendix F.

If a model is fully linear, i.e., if a model is already written in the form of Eq. (B1), then we do not need to assume that dynamics remain close to the equilibrium, so inputs and perturbations need not be weak when dynamics are linear.

Some models are parametrized in terms of a connectivity matrix that is not equal to the Jacobian matrix. For example, the system might be written as $d\mathbf{z}/dt = \mathbf{G}(W\mathbf{z}, \mathbf{x})$, where W is meant to quantify a connectivity matrix, but $\partial_{\mathbf{z}} \mathbf{G}(W\mathbf{z}_0, \mathbf{x}_0) \neq W$. In this case, our interpretation of the effective connectivity matrix is not consistent with the original parametrization, but they will often share structural properties such as their approximate dimensionality. We consider one such example from epidemiology at the end of the results. In that example, we found that the use of effective connectivity in place of the explicit connectivity matrix is important for our results to hold (see Fig. 19 in Appendix F).

APPENDIX C: PRECISE DEFINITIONS OF STRONGLY LOW-RANK NETWORKS AND LOW-DIMENSIONAL DYNAMICS

As mentioned above, we focus our attention on strongly low-rank networks. Intuitively, this means that the connectivity matrix has a small number of large singular values and the rest are not large. To make this definition mathematically precise, we need to consider an asymptotic limit of large N , where N is the network size. Let σ_k be the k th largest singular value of W . In strongly low-rank networks, there exists an $r \sim \mathcal{O}(1)$ such that $\lim_{N \rightarrow \infty} \sigma_k = \infty$ for $k = 1, \dots, r$. In addition, $\sigma_k \leq \mathcal{O}(1)$ for $k > r$.

We additionally need to specify what we mean by “low-dimensional” and “high-dimensional” dynamics. Let λ_k be the k th largest eigenvalue of the covariance matrix of $\mathbf{z}(t)$, equivalently the k th largest singular value since covariance matrices are symmetric and positive definite. Note that λ_k is a non-negative real number representing the amount of variance explained by the k th principal component of $\mathbf{z}(t)$. Intuitively, $\mathbf{z}(t)$ is low-dimensional whenever a large portion of the variance is explained by a small number of singular values. To make this precise, we say that the dynamics of $\mathbf{z}(t)$ are low-dimensional whenever there is an $r \sim \mathcal{O}(1)$ such that $\lim_{N \rightarrow \infty} \lambda_j / \lambda_k = 0$ whenever $j > r$ and $k \leq r$.

APPENDIX D: CONDITIONS FOR HIGH-DIMENSIONAL DYNAMICS IN RESPONSE TO STATIONARY, STOCHASTIC PERTURBATIONS

We now consider conditions for the emergence of high-dimensional dynamics. We assume that $\mathbf{x}(t)$ is a stationary, ergodic stochastic process. Define the cross-spectral matrix $\tilde{C}^{\mathbf{x}}(f)$ of $\mathbf{x}(t)$ at frequency f as the Fourier transform of the matrix of cross-covariance matrix:

$$\tilde{C}_{jk}^{\mathbf{x}}(f) = \int_{-\infty}^{\infty} C_{jk}^{\mathbf{x}}(\tau) e^{-2\pi i f \tau} d\tau,$$

where $C_{jk}^{\mathbf{x}}(\tau) = \text{cov}[\mathbf{x}_j(t), \mathbf{x}_k(t + \tau)]$ is the stationary cross-covariance. When each $\mathbf{x}_j(t)$ is i.i.d., then $\tilde{C}^{\mathbf{x}}(f) = I c_x(f)$ is a multiple of the identity matrix, where $c_x(f)$ is the power spectral density of each $\mathbf{x}_j(t)$. The assumption of i.i.d. elements is true for all of the examples we consider, but we do not apply this simplification until later in our calculation.

The cross-spectral density of the network response $\mathbf{z}(t)$ is defined analogously, and it can be derived under the dynamics in Eq. (B1) to get [42–44,61,62]

$$\tilde{C}^{\mathbf{z}}(f) = [r(f)I - W]^{-1} \tilde{C}^{\mathbf{x}}(f) [r(f)I - W]^{-*}, \quad (\text{D1})$$

where

$$r(f) = 1 - 2\pi\tau f i$$

is scalar.

Equation (D1) quantifies the covariance structure of $\mathbf{z}(t)$ at any given frequency mode f , but we are often specifically interested in the zero-lag temporal covariance:

$$\tilde{C}_{jk}^{\mathbf{z}} = \text{cov}[\mathbf{z}_j(t), \mathbf{z}_k(t)] = \int_{-\infty}^{\infty} \tilde{C}_{jk}^{\mathbf{z}}(f) df. \quad (\text{D2})$$

From Eq. (D1), we therefore have

$$\tilde{C}^{\mathbf{z}} = \int_{-\infty}^{\infty} [r(f)I - W]^{-1} \tilde{C}^{\mathbf{x}}(f) [r(f)I - W]^{-*} df. \quad (\text{D3})$$

If the timescale of fluctuations in $\mathbf{x}(t)$ is much slower than the timescale τ of network dynamics, then $\tilde{C}^{\mathbf{x}}(f) \approx 0$ at frequencies much higher than $1/(2\pi\tau)$ (see Fig. 13 in Appendix F). Indeed, we can take this to be the definition of the statement that the fluctuations in $\mathbf{x}(t)$ are much slower than τ . Since $r(f) \approx 1$ whenever $f \ll 1/(2\pi\tau)$, we therefore have that $r(f) \approx 1$ whenever $\tilde{C}^{\mathbf{x}}(f)$ is not close to zero (see Fig. 13 in Appendix F). As a result, the only parts of the integrand that contribute to the integral in Eq. (D3) are the low-frequency components, $f \approx 0$. In this case, we can replace $r(f)$ with 1 to obtain

$$\tilde{C}^{\mathbf{z}} \approx [I - W]^{-1} \tilde{C}^{\mathbf{x}} [I - W]^{-T}. \quad (\text{D4})$$

Note that Eq. (D4) can also be derived directly from the quasisteady-state approximation $\mathbf{z}(t) \approx [I - W]^{-1} \mathbf{x}(t)$. Our assumption above that $\tilde{C}^{\mathbf{x}}(f) \approx 0$ for $f > 1/\tau$ is essentially equivalent to the assumption that the quasisteady-state approximation is accurate. Either approach can be used to derive Eq. (D4). In Appendix E below, we provide more details about the frequency dependence of variability and how it relates to the speed at which $\mathbf{x}(t)$ varies.

When principal component analysis is applied to $\mathbf{z}(t)$, the variance explained by each principal component is given by the ordered list of eigenvalues of the covariance matrix, $\tilde{C}^{\mathbf{z}}$. From Eq. (D4), we see that these eigenvalues are proportional to the squares of the singular values of the matrix

$$R = [I - W]^{-1} \sqrt{\tilde{C}^{\mathbf{x}}},$$

where $\sqrt{\tilde{C}^{\mathbf{x}}}$ is the matrix square root of the symmetric positive definite covariance matrix $\tilde{C}^{\mathbf{x}}$ (distinct from the entrywise square root, in general). Therefore, the decay of the variance explained by each principal component of $\mathbf{z}$ [as in the blue dots in Fig. 1(d)] are described by the squared singular values of R .

In the models we consider, each $\mathbf{x}_j(t)$ is an i.i.d. stochastic process, so $\tilde{C}^{\mathbf{x}} = vI$ is a multiple of the identity where $v = \text{var}[\mathbf{x}_j(t)]$ is the stationary covariance of $\mathbf{x}_j(t)$.

Therefore, for our models,

$$R = [I - W]^{-1} \sqrt{v}$$

and

$$\bar{\mathbf{C}}^{\mathbf{z}} \approx [I - W]^{-1} [I - W]^{-T} v,$$

where recall that $v > 0$ is a scalar.

In more general classes of models in which $\bar{\mathbf{C}}^{\mathbf{x}}$ is not a multiple of the identity matrix, the dimensionality of $\mathbf{z}(t)$ might be reduced when $\bar{\mathbf{C}}^{\mathbf{x}}$ is effectively low-rank (see Fig. 8 in Appendix F). However, in this situation, low-dimensional dynamics are inherited from the input, not generated by interactions within the network. By restricting to high-dimensional $\mathbf{x}(t)$, we are focused on the question of whether low-dimensional dynamics arise from network interactions.

In conclusion, when $\mathbf{x}(t)$ is high-dimensional, the dimensionality of the dynamics of $\mathbf{z}(t)$ is determined by the effective rank of the matrix

$$A = [I - W]^{-1}.$$

More specifically, the variance explained by each principal component of $\mathbf{z}(t)$ is proportional to the square of the singular values of A ,

$$\text{var explained by } k\text{th PC of } \mathbf{z}(t) \approx \sigma_{A,k}^2 v,$$

where $\sigma_{A,k}$ is the k th singular value of A (assuming singular values and principal components are sorted in decreasing order) and $v = \text{var}[x_j(t)]$ is the stationary variance of each $\mathbf{x}_j(t)$.

Therefore, $\mathbf{z}(t)$ is low-dimensional whenever A has a small number of large singular values. As above, we can use the fact that singular values commute with matrix inverses to write our conclusions in terms of

$$Q = I - W.$$

Specifically,

$$\text{var explained by } k\text{th PC of } \mathbf{z}(t) \approx \frac{v}{\sigma_{Q,N-k}^2}, \quad (\text{D5})$$

where $\sigma_{Q,N-k}$ is the $(N - k)$ th singular value of Q , i.e., the k th from the last singular value.

To demonstrate these analytical results, we repeated the simulation from Fig. 1(d) and added the predictions from Eq. (D5) as red dots [Fig. 24(a) in Appendix F]. Surprisingly, the theory did not closely match the simulations. We suspected that this was due to finite sampling: The simulation was performed over the time interval $t \in [0, T]$, where $T = 5 \times 10^3$ [for comparison, $\tau = 1$ and

the correlation timescale of $\mathbf{x}(t)$ was $\tau_x = 10$]. We suspected that Eq. (D4) would be accurate when $\bar{\mathbf{C}}^{\mathbf{x}}$ is replaced by the empirical covariance matrix of $\mathbf{x}(t)$. Under this substitution, Eq. (D5) would be replaced by

$$\text{var explained by } k\text{th PC of } \mathbf{z}(t) \approx \frac{1}{\sigma_{\hat{B},N-k}^2},$$

where

$$\hat{B} = [I - W] \sqrt{\hat{\mathbf{C}}^{\mathbf{x}}}$$

and $\hat{\mathbf{C}}^{\mathbf{x}}$ is the empirical covariance matrix. Or, equivalently and more simply,

$$\text{var explained by } k\text{th PC of } \mathbf{z}(t) \approx \sigma_{\hat{R},k}^2, \quad (\text{D6})$$

where

$$\hat{R} = [I - W]^{-1} \sqrt{\hat{\mathbf{C}}^{\mathbf{x}}}$$

is the sampled value of R . Using Eq. (D6) gives a more accurate prediction [Fig. 24(a) in Appendix F, green dots]. This confirms that the error in the red dots from Fig. 24(a) is due largely to undersampling of $\mathbf{x}(t)$. We next increased the simulation time tenfold to $T = 5 \times 10^4$. In this case, the original Eq. (D5) was more accurate [Fig. 24(b) in Appendix F, red dots], further confirming that the errors in Fig. 24(a) are due largely to sampling error.

As concluded previously, $\mathbf{z}(t)$ exhibits low-rank suppression [as demonstrated by the last blue dot in Fig. 1(d)] whenever Q has a small number of asymptotically large singular values (equivalently, whenever A has a small number of asymptotically small singular values). Now, we may conclude that $\mathbf{z}(t)$ has low-dimensional dynamics whenever Q has some asymptotically *small* singular values (equivalently, whenever A has a small number of asymptotically large singular values). Conversely, *high-dimensional* dynamics can occur only whenever Q does not have any asymptotically small singular values (equivalently, whenever A does not have any asymptotically large singular values). We can summarize as follows.

- (i) Low-rank suppression [as in Fig. 1(g)] occurs whenever $Q = I - W$ has at least one asymptotically large singular value.
- (ii) High-dimensional dynamics [as in Fig. 1(d)] can occur only whenever $Q = I - W$ lacks any asymptotically small singular values.

Previously, in Sec. IV E, we effectively showed that $Q = I - W$ has large singular values (and, therefore, low-rank suppression occurs) whenever W is strongly low-rank. Conditions under which $Q = I - W$ lacks *small singular values* (and, therefore, high-dimensional dynamics occur) are not so simple. Specifically, the lack or presence of small

singular values depends on the recurrent alignment matrix

$$P = V^T U,$$

which measures the alignment between the left and right singular vectors $\mathbf{u}_k$ and $\mathbf{v}_k$, respectively. Specifically, $P_{jk} = \mathbf{v}_j \cdot \mathbf{u}_k$ so that $|P_{jk}| = 1$ whenever $\mathbf{v}_k = \pm \mathbf{u}_k$ and $P_{jk} = 0$ whenever $\mathbf{v}_k$ is orthogonal to $\mathbf{u}_k$. Note that singular values of P are bounded by unity, $\sigma_P \leq 1$. We next show that if all singular values of P are $\mathcal{O}(1)$, then the network exhibits high-dimensional responses to high-dimensional inputs.

We now restate and prove Claim 1 from Sec. IV E.

Claim—Under the model and Assumptions 1–3 in the main text, if $P = V^T U$ does not have any asymptotically small singular values, then the dynamics of $\mathbf{z}(t)$ are high-dimensional.

Proof. We prove this claim by proving its contrapositive: If the dynamics of $\mathbf{z}(t)$ are low-dimensional, then $P = V^T U$ has at least one asymptotically small singular value.

The dynamics of $\mathbf{z}(t)$ are low-dimensional whenever $\mathbf{z}(t)$ has a small number of dominant principal components as in Fig. 2(d) [conversely, $\mathbf{z}(t)$ is high-dimensional whenever there is no such dominant principal component, as in Fig. 1(d)]. From the discussion above, we know that $\mathbf{z}(t)$ is low-dimensional whenever $A = [I - W]^{-1}$ has at least one asymptotically large singular value or, equivalently, whenever $Q = I - W$ has an asymptotically small singular value. Therefore, our original claim is equivalent to the following: If $Q = I - W$ has at least one asymptotically small singular value, then $P = V^T U$ also has at least one asymptotically small singular value.

We prove this version of the claim directly. Assume that Q has an asymptotically small singular value. Then, there is a $\mathbf{z}$ satisfying $\|\mathbf{z}\| = 1$ and

$$(I - W)\mathbf{z} = o(1),$$

where the notation $o(1)$ means that $\|(I - W)\mathbf{z}\| \rightarrow 0$ as $N \rightarrow \infty$. It is sufficient to show that there is a $\mathbf{y}$ with $\|\mathbf{y}\| = 1 + o(1)$ satisfying

$$\|P\mathbf{y}\| = o(1).$$

We have that

$$W_0 \mathbf{z} + W_1 \mathbf{z} = \mathbf{z} + o(1).$$

Multiplying both sides on the left by U^T gives

$$\Sigma V^T \mathbf{z} + U^T W_1 \mathbf{z} = U^T \mathbf{z} + o(1).$$

Since W_1 is random and independent from W_0 and U and since U^T projects from N to $r \ll N$ dimensions, the term $U^T W_1 \mathbf{z}$ represents the projection of a random vector onto the low-dimensional column space of U , so $\|U^T W_1 \mathbf{z}\| = o(1)$. We, therefore, have

$$\Sigma V^T \mathbf{z} = U^T \mathbf{z} + o(1).$$

We next claim that $\mathbf{z} = U U^T \mathbf{z} + o(1)$. To see why this is true, note that $U U^T \mathbf{z}$ is the orthogonal projection of $\mathbf{z}$ onto the column space of W_0 . Since W_0 dominates $W = W_0 + W_1$ (and, therefore, the column space of W is dominated by that of W_0) and $W\mathbf{z} \approx \mathbf{z}$, we may conclude that $\mathbf{z}$ lies predominantly in the column space of W and, therefore, of W_0 . In other words, $\mathbf{z} = U U^T \mathbf{z} + o(1)$. Hence, we can rewrite the equality above as

$$\Sigma V^T U U^T \mathbf{z} = U^T \mathbf{z} + o(1),$$

which reduces to

$$\Sigma P \mathbf{y} = \mathbf{y} + o(1),$$

where $P = V^T U$ and $\mathbf{y} = U^T \mathbf{z}$. Note again that $\|\mathbf{y}\| = \|U^T \mathbf{z}\| = 1 + o(1)$ since $\|\mathbf{z}\| = 1$ and $\mathbf{z} = U U^T \mathbf{z} + o(1)$. We then have that

$$\|P \mathbf{y}\| = \|\Sigma^{-1} \mathbf{y}\| + o(1) = o(1),$$

because Σ^{-1} is a diagonal matrix with $o(1)$ terms on the diagonal. ■

The proof of Claim 1 also tells us the dominant directions of variability whenever dynamics are *low-dimensional* as in Figs. 2 and 3(g). Since $\mathbf{z}$ is approximately aligned to the U in the proof, we may conclude that low-dimensional dynamics are caused by excess variability along the column space of U , i.e., the column space of W_0 . This explains the observations in Fig. 2(e). More specifically, low-dimensional dynamics arise when external input and the random part of internal input (i.e., $W_1 \mathbf{z}$) have variability that overlaps with the column space of V . Because these right singular vectors are associated with large singular values of W , this variability gets mapped to strong variability in the column space of U . Since P has small singular values, this variability in the column space of U does not overlap with the column space of V , so it cannot be canceled by recurrent interactions. This creates an effectively feedforward mapping that amplifies variability along the directions $\mathbf{v}_k$ and maps it to variability in the directions of $\mathbf{u}_k$. This mechanism is a form of non-normal amplification [28–30] (see the discussion in Sec. IV E for a more in-depth comparison).

These results can also help us understand some subtleties of low-rank suppression. Our theory from Sec. IV C shows that the input $\mathbf{x}_a = \mathbf{u}_k + (W_1 \mathbf{v}_k + \mathbf{v}_k)/\sigma_k$ is suppressed. Since σ_k is large, this input is close to $\mathbf{u}_k$ but has a correction term $\mathbf{x}_c = (W_1 \mathbf{v}_k + \mathbf{v}_k)/\sigma_k$. In Fig. 1(g), the “aligned” input that produced a suppressed response was chosen simply as $\mathbf{x} = \mathbf{u}$, without the correction term (using $\mathbf{x}_a$ with the correction term also produced a suppressed response; not pictured). For the example in Fig. 2, we used

$\mathbf{x} = \mathbf{x}_a$ with the correction term because using $\mathbf{x} = \mathbf{u}$ without a correction term did not produce a suppressed response. To understand why this happens, note that networks like the one in Fig. 2 (for which $P = V^T U$ is singular or has a small singular value) take inputs from the directions $\mathbf{v}_k$ and amplify them in proportion to σ_k (according to our proof of Claim 1). Hence, the $\mathbf{v}_k/\sigma_k$ term in $\mathbf{x}_a$ is small by a factor of σ_k , but the response is also amplified by a factor of σ_k , so the correction term produces a $\mathcal{O}(1)$ response and cannot be ignored. In summary, the correction term $\mathbf{x}_c$ is necessary to include when P is singular or has small singular values (as in Fig. 2).

APPENDIX E: DIMENSIONALITY OF DYNAMICS WHEN INPUTS ARE NOT SLOWLY VARYING

In Sec. IV E and Appendix D, we assumed that $\mathbf{x}(t)$ varied on a slower timescale than the timescale τ of the network dynamics, justifying the quasisteady-state approximation in Eq. (6). While this assumption simplifies our analysis and conclusions, we can understand the dimensionality of dynamics without it. Here, we analyze the dimensionality of $\mathbf{z}(t)$ under all of the conditions from Assumptions 1–3 except the assumption that $\mathbf{x}(t)$ is slowly varying.

Beginning with Eq. (D3), we see that the dimensionality of $\mathbf{z}(t)$, in general, depends on the interaction between $\tilde{\mathbf{C}}^{\mathbf{x}}(f)$, W , and $r(f)$ across different frequencies f . Here, we again focus on the special case that each $\mathbf{x}_j(t)$ is i.i.d. In this case, $\tilde{\mathbf{C}}^{\mathbf{x}}(f) = c_x(f)I$, where $c_x(f)$ is the scalar power spectral density of each $\mathbf{x}_j(t)$. Therefore, the covariance matrix of $\mathbf{z}(t)$ can be written as [44]

$$\tilde{\mathbf{C}}^{\mathbf{z}} = \int_{-\infty}^{\infty} [r(f)I - W]^{-1} [r(f)I - W]^{-*} c_x(f) df, \quad (\text{E1})$$

where $\cdot^{-*}$ is the inverse conjugate transpose and $r(f) = 1 - 2\pi\tau f i$.

First, consider $\tilde{\mathbf{C}}^{\mathbf{z}}(f)$ at frequencies lower than the timescale of the dynamics, $f \ll 1/(2\pi\tau)$. At these frequencies, we have that $r(f) \approx 1$ so that

$$\tilde{\mathbf{C}}^{\mathbf{z}}(f) \approx [I - W]^{-1} [I - W]^{-T} c_x(f)$$

at sufficiently low frequencies. Since $c_x(f)$ is a scalar, we have that

$$\tilde{\mathbf{C}}^{\mathbf{z}} \propto [I - W]^{-1} [I - W]^{-T}$$

for $f \ll 1/(2\pi\tau)$. Therefore, the arguments and conclusions from Sec. IV E, Appendix D, and Claim 1 apply, in general, to the dimensionality of dynamics at low frequencies. Specifically, the dimensionality of the dynamics at low frequencies (i.e., the slow dynamics) can be low-dimensional only when $P = V^T U$ has a small singular value. This is true even when $\mathbf{x}(t)$ has power at high frequencies.

However, the stationary covariance matrix of $\mathbf{z}(t)$ accounts for variability at all frequencies [see Eqs. (E1)]. When $\mathbf{x}(t)$ is slowly varying, $c_x(f) \approx 0$ at higher frequencies (see Fig. 13), so the low-frequency analysis above is sufficient (consistent with the results obtained in Appendix D and Sec. IV E). When $\mathbf{x}(t)$ has power at high frequencies, these higher frequencies need to be accounted for.

At high frequencies satisfying $f \gg \|W\|/(2\pi\tau)$, the $r(f)I$ terms in Eq. (E1) dominate the W terms so that

$$\tilde{\mathbf{C}}^{\mathbf{z}}(f) \approx \frac{c_x(f)}{1 + 4\pi^2 f^2 \tau^2} I,$$

which is a multiple of the identity and, therefore, full rank. Hence, when $\mathbf{x}(t)$ has variability at these high frequencies [i.e., when $c_x(f)$ is not close to zero for $f \gg \|W\|/(2\pi\tau)$; see Fig. 14(a)], these high-frequency components can add only high-dimensional variability to $\mathbf{z}(t)$ because $\mathbf{z}(t)$ tracks $\mathbf{x}(t)$ at high frequencies.

If $\mathbf{x}(t)$ has variability *only* at these high frequencies [$f \gg \|W\|/(2\pi\tau)$], then only these frequencies would contribute to the integral in Eq. (E1), and, therefore, $\mathbf{z}(t)$ would be high-dimensional regardless of the structure of W . However, it is somewhat rare that natural stochastic processes have power only at higher frequencies.

More commonly, when $\mathbf{x}(t)$ varies quickly, it has power at both low and high frequencies. In other words, $c_x(f)$ is far from zero for both small and large f . For example, white noise has equal power at all frequencies and Ornstein-Uhlenbeck processes have more power at low than high frequencies even when they have fast timescales. When $\mathbf{x}(t)$ has power at low and high frequencies, the integral in Eq. (E1) has contributions from both low and high frequencies. Low-dimensional variability can be introduced only at low frequencies and, by Claim 1 and our arguments above, can be introduced only when $P = V^T U$ has a small singular value. Since these low-dimensional dynamics represent an asymptotically large singular value, this combination of low-dimensional variability at low frequencies and high-dimensional variability at high frequencies results in low-dimensional variability overall, since the low-dimensional variability is strong while the high-dimensional variability is $\mathcal{O}(1)$. As a result, when $\mathbf{x}(t)$ has power at both low and high frequencies, the conclusions of Claim 1 apply. Specifically, low-dimensional variability arises only when P has small singular values. To test this conclusion, we repeated the simulations from Fig. 3 with timescales swapped. Specifically, we set $\tau = 10$ and $\tau_x = 1$ so that $c_x(f)$ is far from zero at both low frequencies [$f \ll 1/(2\pi\tau)$] and high frequencies [$f \gg \|W\|/(2\pi\tau)$; see Fig. 14(a) in Appendix F]. As predicted, low-dimensional variability arises only in the case that $P = V^T U$ has a singular value at zero [Figs. 14(b)–14(f)], identical to the case when $\mathbf{x}(t)$ varies slowly (compare Fig. 14 to Fig. 3).

APPENDIX F: ADDITIONAL SIMULATIONS AND DATA

This appendix contains results from additional simulations and data. The simulations themselves and data are described in the respective captions of Figs. 7–24 and contextualized within the main text and appendices.




FIG. 7. Response properties of a feedforward network with rank-one structure. The same as Fig. 1 except the recurrent network was replaced by a feedforward network with dynamics satisfying $\tau d\mathbf{z}/dt = -\mathbf{z}' + \mathbf{W}\mathbf{x}$. Unlike the recurrent network in Fig. 1, the feedforward network is most sensitive to inputs aligned with its low-rank structure, and its dynamics are dominated by one-dimensional variability. (a) Schematic of the feedforward model. (b) The singular values of $\mathbf{W}$ have one dominant term. (c) The eigenvalues of $\mathbf{W}$ have a dominant, negative term. (d) The distribution of variance across principal components of a Gaussian stochastic input $[\mathbf{x}(t); \text{green}]$ and the response $[\mathbf{z}(t); \text{blue}]$. (e),(f) The network response (e) and input (f) projected onto the plane determined by $\mathbf{u}$ and a random vector, $\mathbf{u}_{\text{rand}}$, demonstrates amplification along $\mathbf{u}$. (g) The network response $[\mathbf{z}(t), \text{top}]$ and its norm $[\|\mathbf{z}(t)\|, \text{bottom}]$ given an input aligned to the low-rank structure of the network ($\mathbf{x}_{\text{aligned}}$) and a random input ($\mathbf{x}_{\text{rand}}$). (h) The input $\mathbf{x}$ and the response $\mathbf{z}$, projected onto $\mathbf{u}$.




FIG. 8. Dynamics of a network with rank-one structure driven by an input perturbation with one dimensional structure. The same as Fig. 1 except the dynamics obeyed $\tau d\mathbf{z}/dt = -\mathbf{z} + \mathbf{W}\mathbf{z} + \mathbf{W}_x\boldsymbol{\eta}(t)$, where $\boldsymbol{\eta}(t)$ is a realization of the same Gaussian stochastic process used for $\mathbf{x}(t)$ in Fig. 1 and $\mathbf{W}_x$ is an effectively low-rank matrix generated identically to but independently from $\mathbf{W}$. (a) The distribution of variance across principal components of the effective input [$\mathbf{x}(t) = \mathbf{W}_x\boldsymbol{\eta}(t)$; green] and the response [$\mathbf{z}(t)$; blue]. The effective input is low-dimensional and the network dynamics $\mathbf{z}(t)$ inherit these low-dimensional dynamics from $\mathbf{x}(t)$. (b),(c) The network response (b) and input (c) projected onto the plane determined by $\mathbf{u}$ and a random vector $\mathbf{u}_{rand}$. (d) Local network input (purple) cancels with external input (green) to produce suppressed network responses (blue) in the direction of $\mathbf{u}$. Low-rank suppression and cancellation are still exhibited along the vector $\mathbf{u}$, defining the low-rank structure of $\mathbf{W}$ [(b)–(d)], which is distinct from the dominant direction of $\mathbf{x}(t)$.




FIG. 9. Response properties of a network with independent left and right singular vectors. The same as Fig. 2 except $W_0 = c\mathbf{u}\mathbf{v}^T$, where $\mathbf{u}$ and $\mathbf{v}$ are independent random unit vectors. The network behaves similarly to the network in Fig. 2 because $\mathbf{u}$ and $\mathbf{v}$ are nearly orthogonal. (a) Schematic of model. (b) The singular values of W have one dominant term. (c) The eigenvalues of W lack a dominant, negative term. (d) The distribution of variance across principal components of a Gaussian stochastic input [$\mathbf{x}(t)$; green] and the response [$\mathbf{z}(t)$; blue]. (e),(f) The network response projected onto the plane determined by (a) $\mathbf{u}$ and (b) the last principal component direction $\mathbf{u}_N$, compared to projections onto a random vector $\mathbf{u}_{\text{rand}}$. (g) The network response [$\mathbf{z}(t)$, top] and its norm [$\|\mathbf{z}(t)\|$, bottom] given two different input directions (aligned and random). The norm of $\mathbf{z}$ is more than 11 times smaller in response to the aligned versus random perturbation. (h) Local network input (purple) projected onto $\mathbf{u}_N$ cancels with external input (green) to produce suppressed network responses (blue) in the direction of $\mathbf{u}_N$.




FIG. 10. Dynamics of an unstable, chaotic network with rank-one structure. The same as Fig. 1 except $\rho = 2$ and the magnitude of c was increased by the same factor ($c = -40$) to maintain the same degree of low-rank structure (the same ratio between the spectral radius of the random part ρ and the magnitude of the low-rank part $|c|$), and the dynamics obey Eq. (7). The instability produced by taking $\rho > 1$ generates chaotic dynamics. (a) Schematic of model. (b) The singular values of W have one dominant term. (c) The eigenvalues of W have a dominant, negative term. (d) The distribution of variance across principal components of a Gaussian stochastic input [$\mathbf{x}(t)$; green] and the response [$\mathbf{z}(t)$; blue]. (e),(f) The network response (e) and input (f) projected onto the plane determined by $\mathbf{u}$ and a random vector $\mathbf{u}_{\text{rand}}$. (g) The network response [$\mathbf{z}(t)$, top] and its norm [$\|\mathbf{z}(t)\|$, bottom] given an input aligned to the low-rank structure of the network ($\mathbf{x}_{\text{aligned}}$) and a random input ($\mathbf{x}_{\text{rand}}$). (h) Local network input (purple), external input (green), and network response (blue) in the direction of $\mathbf{u}$.




FIG. 11. Dynamics of an unstable, nonchaotic network with rank-one structure. The same as Fig. 1 except $c = 10$ and the dynamics obey Eq. (7). The fixed point at $\mathbf{z} = 0$ is unstable because $c > 1$. (a) Schematic of model. (b) The singular values of W have one dominant term. (c) The eigenvalues of W have a dominant, negative term. (d) The distribution of variance across principal components of a Gaussian stochastic input $[\mathbf{x}(t); \text{green}]$ and the response $[\mathbf{z}(t); \text{blue}]$. (e),(f) The network response (e) and input (f) projected onto the plane determined by $\mathbf{u}$ and a random vector $\mathbf{u}_{\text{rand}}$. (g) The network response $[\mathbf{z}(t), \text{top}]$ and its norm $[\|\mathbf{z}(t)\|, \text{bottom}]$ given an input aligned to the low-rank structure of the network ($\mathbf{x}_{\text{aligned}}$) and a random input ($\mathbf{x}_{\text{rand}}$). (h) Local network input (purple), external input (green), and network response (blue) in the direction of $\mathbf{u}$.




FIG. 12. Dynamics of a stable network with rank-one structure close to criticality. The same as Fig. 1 except $\rho = 0.95$. The fixed point at $\mathbf{z} = 0$ is stable but close to instability because the maximum real part of the eigenvalues of the Jacobian is close to 0. (a) Schematic of model. (b) The singular values of W have one dominant term. (c) The eigenvalues of W have a dominant, negative term. (d) The distribution of variance across principal components of a Gaussian stochastic input $[\mathbf{x}(t); \text{green}]$ and the response $[\mathbf{z}(t); \text{blue}]$. (e),(f) The network response (e) and input (f) projected onto the plane determined by $\mathbf{u}$ and a random vector $\mathbf{u}_{\text{rand}}$. (g) The network response $[\mathbf{z}(t), \text{top}]$ and its norm $[\|\mathbf{z}(t)\|, \text{bottom}]$ given an input aligned to the low-rank structure of the network ($\mathbf{x}_{\text{aligned}}$) and a random input ($\mathbf{x}_{\text{rand}}$). (h) Local network input (purple), external input (green), and network response (blue) in the direction of $\mathbf{u}$.




FIG. 13. Power spectral density of $\mathbf{x}(t)$ used in simulations. Average power spectral density of $\mathbf{x}_j(t)$ (green) compared to $|r(f)|^2 = 1/(1 + 4\pi^2 f^2 \tau^2)$ (gray) when $\tau_x = 10$ and $\tau = 1$. These parameters are used in all figures except for Fig. 14. The vertical red line shows $f = 1/(2\pi\tau)$.




FIG. 14. High- and low-dimensional dynamics when inputs are not slow. Simulations of a network with rank-two structure (the same networks as in Fig. 3) when inputs are slower than the intrinsic timescale of the network, specifically, $\tau_x = 1$ and $\tau = 10$. (a) Average power spectral density of $\mathbf{x}_j(t)$ (green) compared to $|r(f)|^2 = 1/(1 + 4\pi^2 f^2 \tau^2)$ (gray). The vertical red line shows $f = 1/(2\pi\tau)$. Since $\mathbf{x}(t)$ is faster than intrinsic network dynamics ($\tau_x < \tau$), $c_x(f)$ has power at higher frequencies where $r(f) \approx 0$. (b) The variance explained by each principal component of the network dynamics, $\mathbf{z}(t)$. The same as Fig. 3(c) except that $\tau_x = 1$ and $\tau = 10$. (b)–(f) The variance explained by each principal component of the network dynamics, $\mathbf{z}(t)$, for five different network structures. The same as Figs. 3(c), 3(f), 3(i), 3(l), and 3(o), respectively, except that we set $\tau_x = 1$ and $\tau = 10$. Variability in $\mathbf{z}(t)$ is high- and low-dimensional under the same conditions.




FIG. 15. Response properties of a network with biased weights. The same as Fig. 1 except we used the dynamics in Eq. (7), and W_{jk} are drawn i.i.d. from a normal distribution with mean $-0.5/\sqrt{N}$ and standard deviation $0.5/\sqrt{N}$, where $N = 200$. (a) Schematic of model. (b) The singular values of W have one dominant term. (c) The eigenvalues of W have a dominant, negative term. (d) The distribution of variance across principal components of a Gaussian stochastic input $[\mathbf{x}(t)$; green] and the response $[\mathbf{z}(t)$; blue]. (e),(f) The network response (e) and input (f) projected onto the plane determined by $\mathbf{u}$ and a random vector $\mathbf{u}_{\text{rand}}$. (g) The network response $[\mathbf{z}(t)$, top] and its norm $[\|\mathbf{z}(t)\|]$, bottom] given an input aligned to the low-rank structure of the network ($\mathbf{x}_{\text{aligned}}$) and a random input ($\mathbf{x}_{\text{rand}}$). (h) Local network input (purple), external input (green), and network response (blue) in the direction of $\mathbf{u}$.




FIG. 16. Low-dimensional dynamics in a modular network. The same as Fig. 1 except we used the dynamics in Eq. (7) and the modular network from Fig. 4. (a) The distribution of variance across principal components of a Gaussian stochastic input $[\mathbf{x}(t)$; green] and the response $[\mathbf{z}(t)$; blue]. (b),(c) The network response (b) and input (c) projected onto the plane determined by $\mathbf{u}$ and a random vector $\mathbf{u}_{\text{rand}}$.




FIG. 17. Dominant singular vectors of a spatial network are Fourier modes. The singular vectors corresponding to the ten largest singular values of a network with spatial structure (the same network in Fig. 5).




FIG. 18. Low-dimensional dynamics in a spatial network. The same as Fig. 1 except we used the dynamics in Eq. (7) and a network with spatial structure (the same network from Fig. 5). (a) The distribution of variance across principal components of a Gaussian stochastic input [$\mathbf{x}(t)$; green] and the response [$\mathbf{z}(t)$; blue]. (b),(c) The network response (b) and input (c) projected onto the plane determined by $\mathbf{u}$ and a random vector $\mathbf{u}_{\text{rand}}$.




FIG. 19. Singular values and dynamics in an epidemiological network. (a) Singular values of Q (red) and W (blue). (b),(c) Projections of z onto the dominant left singular vectors of W and Q , respectively, for the epidemiological model (the same model as in Fig. 6). (d) The response to a perturbation aligned with the dominant left singular vector of Q compared to the response to a perturbation of the same magnitude in a random direction. The same as Fig. 6(a) except the dominant singular vector of Q was used instead of the dominant singular vector of W . The lack of suppression demonstrates the necessity of using the effective connectivity matrix W in place of Q , which is the connectivity matrix used to define the dynamics.




FIG. 20. Response properties of a weakly low-rank network with $\rho \ll 1$ and $c < 0$. The same as Fig. 1 except $c = -0.5$ and $\rho = 0.01$. (a) Schematic of model. (b) The singular values of W have one dominant term, but it is not large (compare to Fig. 1). (c) The eigenvalues of W have a dominant, negative term, but it is not large in magnitude (compare to Fig. 1). (d) The distribution of variance across principal components of a Gaussian stochastic input $[x(t); \text{green}]$ and the response $[z(t); \text{blue}]$. (e),(f) The network response (e) and input (f) projected onto the plane determined by u and a random vector u_{rand} . (g) The network response $[z(t), \text{top}]$ and its norm $[\|z(t)\|, \text{bottom}]$ given an input aligned to the low-rank structure of the network ($x_{aligned}$) and a random input (x_{rand}). (h) Local network input (purple), external input (green), and network response (blue) in the direction of u .




FIG. 21. Response properties of a weakly low-rank network with $\rho \ll 1$ and $c > 0$. The same as Fig. 1 except $c = 0.5$ and $\rho = 0.01$. Note that stability requires that $c < 1$ when $c > 0$. (a) Schematic of model. (b) The singular values of W have one dominant term, but it is not large (compare to Fig. 1). (c) The eigenvalues of W have a dominant, negative term, but it is not large in magnitude (compare to Fig. 1). (d) The distribution of variance across principal components of a Gaussian stochastic input $[\mathbf{x}(t); \text{green}]$ and the response $[\mathbf{z}(t); \text{blue}]$. (e),(f) The network response (e) and input (f) projected onto the plane determined by $\mathbf{u}$ and a random vector $\mathbf{u}_{\text{rand}}$. (g) The network response $[\mathbf{z}(t)]$, top] and its norm $[\|\mathbf{z}(t)\|]$, bottom] given an input aligned to the low-rank structure of the network ($\mathbf{x}_{\text{aligned}}$) and a random input ($\mathbf{x}_{\text{rand}}$). (h) Local network input (purple), external input (green), and network response (blue) in the direction of $\mathbf{u}$.




FIG. 22. Quantifying low-rank suppression as the maximum singular value is scaled. We simulated a network with symmetric low-rank structure (the same model as in Fig. 1) while changing the value of $c < 0$ to change the value of the maximum singular value of W , $\sigma_{\text{max}} = |c|$. This demonstrates a smooth transition between weakly and strongly low-rank structure. For each value of c , we computed the ratio of $\|\mathbf{z}\|$ in response to an aligned input versus a random input, i.e., the height of the two bumps in Fig. 1(g) as a measure of low-rank suppression. Low-rank suppression occurs when this ratio is small.




FIG. 23. High-dimensional dynamics in weakly low-rank networks. Simulations of networks with effective rank-two structure (similar to networks in Fig. 3) except the low-rank structure is weak. Specifically, $\rho = 0.01$, $\sigma_1 = 0.5$, and $\sigma_2 = 0.25$. (a) The singular values of W have two dominant terms, but they are not large [compare to Fig. 3(a)]. (b)–(f) The variance explained by each principal component of the network dynamics, $\mathbf{z}(t)$, for five different network structures. The same as Figs. 3(c), 3(f), 3(i), 3(l), and 3(o), respectively, except that $\rho = 0.01$, $\sigma_1 = 0.5$, and $\sigma_2 = 0.25$ (weak low-rank structure).




FIG. 24. Comparison between theory and simulations for shorter and longer simulations. (a) The distribution of variance across principal components of the network response $[\mathbf{z}(t); \text{blue}]$ for simulations of a network with effective rank-one structure (the same model as in Fig. 1). Red dots show theoretical values from a direct analytical expression [i.e., using Eq. (D5)]. Green dots show values obtained by using the empirical principal components of $\mathbf{x}(t)$ [i.e., using Eq. (D6)]. The network was simulated for $T = 5 \times 10^3$ time units. (b) The same as (a) except we increased the simulation time to $T = 5 \times 10^4$.

- [1] V. Thibault, A. Allard, and P. Desrosiers, *The low-rank hypothesis of complex systems*, *Nat. Phys.* **20**, 294 (2024).
- [2] J. Aljadeff, D. Renfrew, M. Vugué, and T. O. Sharpee, *Low-dimensional dynamics of structured random networks*, *Phys. Rev. E* **93**, 022302 (2016).
- [3] F. Mastrogiuseppe and S. Ostojic, *Linking connectivity, dynamics, and computations in low-rank recurrent neural networks*, *Neuron* **99**, 609 (2018).

- [4] I. D. Landau and H. Sompolinsky, *Macroscopic fluctuations emerge in balanced networks with incomplete recurrent alignment*, *Phys. Rev. Res.* **3**, 023171 (2021).
- [5] A. Dubreuil, A. Valente, M. Beiran, F. Mastrogiuseppe, and S. Ostojic, *The role of population structure in computations through neural dynamics*, *Nat. Neurosci.* **25**, 783 (2022).
- [6] I. D. Landau and H. Sompolinsky, *Coherent chaos in a recurrent neural network with structured connectivity*, *PLoS Comput. Biol.* **14**, e1006309 (2018).
- [7] F. Schuessler, A. Dubreuil, F. Mastrogiuseppe, S. Ostojic, and O. Barak, *Dynamics of random recurrent networks with*

- correlated low-rank structure*, *Phys. Rev. Res.* **2**, 013111 (2020).
- [8] A. Valente, J. W. Pillow, and S. Ostoic, *Extracting computational mechanisms from neural data using low-rank rnns*, *Adv. Neural Inf. Process. Syst.* **35**, 24072 (2022).
- [9] L. Cimeša, L. Ciric, and S. Ostoic, *Geometry of population activity in spiking networks with low-rank structure*, *PLoS Comput. Biol.* **19**, e1011315 (2023).
- [10] M. Beiran, N. Meirhaeghe, H. Sohn, M. Jazayeri, and S. Ostoic, *Parametric control of flexible timing through low-dimensional neural manifolds*, *Neuron* **111**, 739 (2023).
- [11] F. Mastrogiuseppe, J. Carmona, and C. K. Machens, *Stochastic activity in low-rank recurrent neural networks*, *PLoS Comput. Biol.* **21**, e1013371 (2025).
- [12] D. G. Clark, O. Marschall, A. van Meegen, and A. Litwin-Kumar, *Connectivity structure and dynamics of nonlinear recurrent neural networks*, *Phys. Rev. X* **15**, 041019 (2025).
- [13] M. T. Kaufman, M. M. Churchland, S. I. Ryu, and K. V. Shenoy, *Cortical activity in the null space: Permitting preparation without movement*, *Nat. Neurosci.* **17**, 440 (2014).
- [14] J. A. Gallego, M. G. Perich, L. E. Miller, and S. A. Solla, *Neural manifolds for the control of movement*, *Neuron* **94**, 978 (2017).
- [15] C. Stringer, M. Pachitariu, N. Steinmetz, C. B. Reddy, M. Carandini, and K. D. Harris, *Spontaneous behaviors drive multidimensional, brainwide activity*, *Science* **364**, eaav7893 (2019).
- [16] C. Stringer, M. Pachitariu, N. Steinmetz, M. Carandini, and K. D. Harris, *High-dimensional geometry of population responses in visual cortex*, *Nature (London)* **571**, 361 (2019).
- [17] F. Lanore, N. A. Cayco-Gajic, H. Gurnani, D. Coyle, and R. A. Silver, *Cerebellar granule cell axons support high-dimensional representations*, *Nat. Neurosci.* **24**, 1142 (2021).
- [18] L. Avitan and C. Stringer, *Not so spontaneous: Multidimensional representations of behaviors and context in sensory areas*, *Neuron* **110**, 3064 (2022).
- [19] A. Markanday, S. Hong, J. Inoue, E. De Schutter, and P. Thier, *Multidimensional cerebellar computations for flexible kinematic control of movements*, *Nat. Commun.* **14**, 2548 (2023).
- [20] J. Manley, S. Lu, K. Barber, J. Demas, H. Kim, D. Meyer, F. M. Traub, and A. Vaziri, *Simultaneous, cortex-wide dynamics of up to 1 million neurons reveal unbounded scaling of dimensionality with neuron number*, *Neuron* **112**, 1694 (2024).
- [21] V. Mante, D. Sussillo, K. V. Shenoy, and W. T. Newsome, *Context-dependent computation by recurrent dynamics in prefrontal cortex*, *Nature (London)* **503**, 78 (2013).
- [22] D. Sussillo and O. Barak, *Opening the black box: Low-dimensional dynamics in high-dimensional recurrent neural networks*, *Neural Comput.* **25**, 626 (2013).
- [23] C. Huang, D. A. Ruff, R. Pyle, R. Rosenbaum, M. R. Cohen, and B. Doiron, *Circuit models of low-dimensional shared variability in cortical networks*, *Neuron* **101**, 337 (2019).
- [24] D. J. O'Shea, L. Duncker, W. Goo, X. Sun, S. Vyas, E. M. Trautmann, I. Diester, C. Ramakrishnan, K. Deisseroth, M. Sahani *et al.*, *Direct neural perturbations reveal a dynamical mechanism for robust computation*, *bioRxiv* 2022 (2022).
- [25] F. Schuessler, F. Mastrogiuseppe, S. Ostoic, and O. Barak, *Aligned and oblique dynamics in recurrent neural networks*, *eLife* **13**, RP93060 (2024).
- [26] H. Schwerdtfeger, *Introduction to Linear Algebra and the Theory of Matrices* (P. Noordhoff, Groningen, 1950).
- [27] D. S. Bernstein, *Scalar, Vector, and Matrix Mathematics: Theory, Facts, and Formulas-Revised and Expanded Edition* (Princeton University Press, Princeton, NJ, 2018).
- [28] L. N. Trefethen, *Pseudospectra of matrices*, *Numer. Funct. Anal. Optim.* **91**, 234 (1991).
- [29] B. K. Murphy and K. D. Miller, *Balanced amplification: A new mechanism of selective amplification of neural activity patterns*, *Neuron* **61**, 635 (2009).
- [30] G. Hennequin, T. P. Vogels, and W. Gerstner, *Non-normal amplification in random balanced neuronal networks*, *Phys. Rev. E* **86**, 011909 (2012).
- [31] C. Van Vreeswijk and H. Sompolinsky, *Chaos in neuronal networks with balanced excitatory and inhibitory activity*, *Science* **274**, 1724 (1996).
- [32] I. D. Landau, R. Egger, V. J. Dercksen, M. Oberlaender, and H. Sompolinsky, *The impact of structural heterogeneity on excitation-inhibition balance in cortical networks*, *Neuron* **92**, 1106 (2016).
- [33] R. Pyle and R. Rosenbaum, *Highly connected neurons spike less frequently in balanced networks*, *Phys. Rev. E* **93**, 040302(R) (2016).
- [34] C. Ebsch and R. Rosenbaum, *Imbalanced amplification: A mechanism of amplification and suppression from local imbalance of excitation and inhibition in cortical circuits*, *PLoS Comput. Biol.* **14**, e1006048 (2018).
- [35] V. L. Girko, *Circular law*, *Theory Probab. Appl.* **29**, 694 (1985).
- [36] T. Tao, V. Vu, and M. Krishnapur, *Random matrices: Universality of ESDs and the circular law*, *Ann. Probab.* **38**, 2023 (2010).
- [37] H. Sompolinsky, A. Crisanti, and H.-J. Sommers, *Chaos in random neural networks*, *Phys. Rev. Lett.* **61**, 259 (1988).
- [38] C. van Vreeswijk and H. Sompolinsky, *Chaotic balanced state in a model of cortical circuits*, *Neural Comput.* **10**, 1321 (1998).
- [39] W. Maass, T. Natschläger, and H. Markram, *Real-time computing without stable states: A new framework for neural computation based on perturbations*, *Neural Comput.* **14**, 2531 (2002).
- [40] H. Jaeger and H. Haas, *Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication*, *Science* **304**, 78 (2004).
- [41] A. Renart, J. De La Rocha, P. Bartho, L. Hollender, N. Parga, A. Reyes, and K. D. Harris, *The asynchronous state in cortical circuits*, *Science* **327**, 587 (2010).
- [42] R. Rosenbaum, M. A. Smith, A. Kohn, J. E. Rubin, and B. Doiron, *The spatial structure of correlated neuronal variability*, *Nat. Neurosci.* **20**, 107 (2017).
- [43] C. Baker, C. Ebsch, I. Lampl, and R. Rosenbaum, *Correlated states in balanced neuronal networks*, *Phys. Rev. E* **99**, 052414 (2019).

- [44] C. W. Gardiner, *Handbook of Stochastic Methods for Physics, Chemistry and the Natural Sciences*, 2nd ed. (Springer-Verlag, Berlin, 1985).
- [45] D. Sussillo and L. F. Abbott, *Generating coherent patterns of activity from chaotic neural networks*, *Neuron* **63**, 544 (2009).
- [46] R. Rosenbaum, *Modeling Neural Circuits Made Simple with PYTHON* (MIT Press, Cambridge, MA, 2024).
- [47] D. J. MacKay and K. D. Miller, *Analysis of Linsker's application of Hebbian rules to linear networks*, *Network* **1**, 257 (1990).
- [48] S. Gu, M. G. Mattar, H. Tang, and G. Pan, *Emergence and reconfiguration of modular structure for artificial neural networks during continual familiarity detection*, *Sci. Adv.* **10**, eadm8430 (2024).
- [49] R. B. Levy and A. D. Reyes, *Spatial profile of excitatory and inhibitory synaptic connectivity in mouse primary auditory cortex*, *J. Neurosci.* **32**, 5609 (2012).
- [50] C. K. Pfeffer, M. Xue, M. He, Z. J. Huang, and M. Scanziani, *Inhibition of inhibition in visual cortex: The logic of connections between molecularly distinct interneurons*, *Nat. Neurosci.* **16**, 1068 (2013).
- [51] J. Barral and A. D. Reyes, *Synaptic scaling rule preserves excitatory–inhibitory balance and salient neuronal network dynamics*, *Nat. Neurosci.* **19**, 1690 (2016).
- [52] M. Okun and I. Lampl, *Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities*, *Nat. Neurosci.* **11**, 535 (2008).
- [53] B. V. Atallah and M. Scanziani, *Instantaneous modulation of gamma oscillation frequency by balancing excitation with inhibition*, *Neuron* **62**, 566 (2009).
- [54] H. Adesnik and M. Scanziani, *Lateral competition for cortical space by layer-specific horizontal circuits*, *Nature (London)* **464**, 1155 (2010).
- [55] Y. Shu, A. Hasenstaub, and D. A. McCormick, *Turning on and off recurrent balanced cortical activity*, *Nature (London)* **423**, 288 (2003).
- [56] M. Wehr and A. M. Zador, *Balanced inhibition underlies tuning and sharpens spike timing in auditory cortex*, *Nature (London)* **426**, 442 (2003).
- [57] B. Haider, A. Duque, A. R. Hasenstaub, and D. A. McCormick, *Neocortical network activity in vivo is generated through a dynamic balance of excitation and inhibition.*, *J. Neurosci.* **26**, 4535 (2006).
- [58] A. L. Dorn, K. Yuan, A. J. Barker, C. E. Schreiner, and R. C. Froemke, *Developmental sensory experience balances cortical excitation and inhibition.*, *Nature (London)* **465**, 932 (2010).
- [59] R. Mastrandrea, J. Fournet, and A. Barrat, *Contact patterns in a high school: A comparison between data collected using wearable sensors, contact diaries and friendship surveys*, *PLoS One* **10**, e0136497 (2015).
- [60] <https://github.com/RobertRosenbaum/HighDimLowDimCode>.
- [61] J. Trousdale, Y. Hu, E. Shea-Brown, and K. Josić, *Impact of network structure and cellular response on spike time correlations*, *PLoS Comput. Biol.* **8**, e1002408 (2012).
- [62] C. Baker, E. Froudarakis, D. Yatsenko, A. S. Tolia, and R. Rosenbaum, *Inference of synaptic connectivity and external variability in neural microcircuits*, *J. Comput. Neurosci.* **48**, 123 (2020).