The trade-off between neural computation and oscillatory coordination.

Erik J Peterson[a,1] **and Bradley Voytek**[a,b,c]

[a]Department of Cognitive Science; [b]Neurosciences Graduate Program; [c]Halıcıoğlu Data Science Institute, University of California, San Diego, 9500 Gilman Drive La Jolla, CA, 92093

This manuscript was compiled on May 3, 2018

Oscillations can improve neural coding by grouping action potentials into synchronous windows of activity, but this same effect can harm coding when action potentials become over-synchronized. Diseases ranging from Parkinson's to epilepsy suggest that over-synchronization leads to pathology, but the precise boundary separating healthy from pathological synchrony remains an open theoretical problem. Here we study a simple model that shows how error in individual cells' computations is traded for population-level synchronization. To put the in biological terms accessible to the cell we conceive of a "voltage budget" where instantaneous moments of membrane voltage can be partitioned into oscillatory and computational terms. By comparing these budget terms we derive a new set of biologically measurable inequalities that separate healthy from pathological synchrony. Finally, we derive an optimal non-biological algorithm for exchanging computational error with population synchrony.

Rhythmic entrainment is a common feature of biological systems, but complete synchronization is often undesirable. This can be conceptually illustrated in the case of neural oscillations, where a totally unsynchronized neuronal population might lack communication capacity whereas a perfectly synchronized population might lack computation capacity (1). The biological reality lies in between, where moderate oscillations coordinate the firing of many individual neurons, creating synchronous windows of population communication (2). Temporally grouping action potentials in this manner improves signal to noise (3) and increases the number of coincident firing events (4, 5), driving learning at individual synapses (6–8). Complete independence between neurons dramatically reduces these temporal coincidences, whereas complete synchronization eliminates each neuron's individual firing characteristics, negating any possible computational contribution (9).

To understand this problem more intuitively we imagine a population where each neuron receives the same spiking input but otherwise is independent from its neighbors. In this model each cell's response depends only on its immediate synaptic weight and its long-term membrane dynamics. Even in this simple situation the membrane response of each neuron can be complex, spanning chaotic irregular activity, bursting, accelerating, and a decay in rate driven by adaptation (10, 11). Even in the simplest case of regular-firing, with uniform sampling of synaptic weights, a population can exhibit substantial response variability (as shown in Figure 1**a**).

From a theoretical perspective, the high-dimensional nature of an independent neural response is a powerful potential computational resource (12, 13). However if there are many populations, all trying to communicate at once, we can again imagine how individually useful high-dimensional responses begin to feedback onto each other—with variability amplifying variability. Allowing a population to fall under the sway of a oscillator is one way to stabilize communications (14) while simultaneously, as already noted, improving signal-to-noise and increasing the synaptic learning rate.

Brittian *et al* has considered the problem of pathological synchrony in a model of Parkinson's disease. In line with our illustration, they suggested synchrony can tune the complexity of a population's response. They use this to explain why increased beta (13-30 Hz) synchrony in Parkinson's patients can lead to a symptom-causing loss of computational capacity (9). While inspirational to our approach, theirs does not offer quantitative predictions and its generality is limited by a strong assumption they make about the nature of neural computation.

In our formalization, we only assume that neurons that act independently have the highest possible computational expression. We study this ideal in the simple model described above: a population of independent neurons entrained by a homogeneous global oscillator. We treat any correlation induced between independent neurons as a perturbation from its computational ideal—as an error. As the strength of the oscillation grows, then, error increases as neurons are forced to synchronize with one other.

To study the trade-off between synchrony and error in biological terms accessible to the cell, we examine instantaneous moments of membrane voltage. For a small window of time it is reasonable to separate out the computational drive from the oscillatory influence, creating what we call a "voltage budget". We use these budgets as a neuron might: for estimating the budget at one moment in time to predict how a later cycle of an oscillation will affect spiking.

Our main contribution is the introduction of a voltage budget analysis, which allows us to define a mathematical, biologically observable criterion to distinguish between healthy and pathological oscillatory synchronization. We note that even in the healthy range there is a continuous trade-off between the induction of synchrony and the introduction of error. Finally, we describe a new non-biological algorithm that achieves an optimal trade-off between error and synchrony.

Results

We model a simple neural network: a population of N neurons entrained by a single global oscillator, governed by an amplitude (A) and frequency (f) of the form $A/2(1+\sin(2\pi f t)$

EJP and BV designed the study and wrote manuscript. EJP implemented the model and conducted the analysis.

The authors have no conflicts of interest to declare.

[1]To whom correspondence should be addressed. E-mail: Erik.Exists@gmail.com

DRAFT

with $A \geq 0$. When oscillatory amplitude is zero, each neuron is completely independent. As oscillatory amplitude grows, each neuron's firing is increasingly perturbed by the strength of the oscillation. We illustrate this in Figure 1 where, in the rightmost panel, oscillatory amplitude is relatively high and the population's firing pattern bears little resemblance to the original population activity as driven by the input (shown in the leftmost panel).

We define computation here in the most basic mathematical sense: an input mapped to an output. We formalize this first in the abstract general form of, $f : x \rightarrow y$.. Here f is a function, synonymous with computation, and x and y are any set of inputs and outputs. In practice we implement f as an adaptive exponential (AdEx) integrate-and-fire model (15) (see *Methods*), and limit input and output to binary time-series, i.e. action potentials.

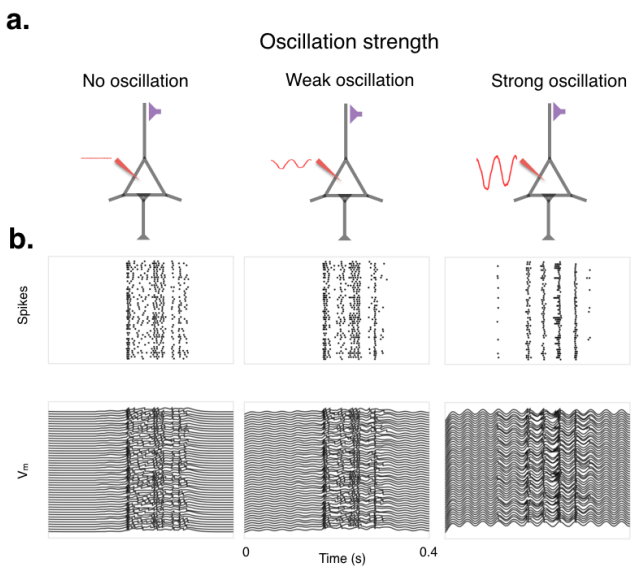

Fig. 1. The trade-off between computation and coordination. **a**. Illustration of a single model neuron receiving synaptic input (purple synapse), subjected to increasing levels oscillatory power (red waveform). **b.** A population of N independent neurons subjected to three levels of oscillatory power ($A = 0.0$, $A = 0.1$, $A = 0.6$ nS; $f = 30$ Hz). The top panel is a plot of action potentials per neuron and the bottom is membrane voltages of those same neurons.

The voltage budget. To understand the boundary between healthy and pathological synchrony, we analyze computation from a neuronal perspective by examining changes in the membrane potential of individual neurons. To simplify this analysis we examine extremely small, nearly instantaneous, windows of time (w). For a window of time much less than the membrane time constant ($w << \tau_m$) it is reasonable to treat the resting potential (V_r) and the firing threshold (V_t) as constant values. With these terms fixed, the total amount of voltage becomes fixed and the membrane potential available in each neuron becomes a physically conserved quantity. This means that no energy is allowed in or out, closing the system.

To aid our thinking about this conserved system we use economic metaphors. The total amount of voltage available to a neuron in w is termed a "budget" that can be "spent" by the neuron. It can be spent to either minimize computational errors—a perturbation of spiking—caused by the global oscillator, or to better align itself with the population (Fig. 2). With this approach the voltage "cost" of increases in oscillatory power can be explicitly analyzed in terms of cellular physiology.

We decompose each budget V_b into three terms: 1) the computational voltage V_c, 2) the oscillatory voltage V_o, and, 3) the open voltage V_n (Fig 1). The open voltage serves the same role as the potential energy term common in basic analyses of the physics of baseball thrown into the air. For example: as the ball rises and is slowed by gravity, kinetic energy is traded for potential energy. Put another way, V_n represents the unused capacity of the system—the difference between $V_c + V_o$ and the threshold potential.

The basic budget relationships are shown in equations 1 and 2. In Eq 2, computation and communication explicitly compete for influence on the eventual spike that happens when the threshold is reached (when $V_n \rightarrow 0$).

$$V_b = V_t - V_r \quad [1]$$

$$V_b = V_o + V_c + V_n \quad [2]$$

In practice though we study budget terms as ratios, because working in unitless quantities is simpler mathematically and eases empirical comparisons. Specifically, we study the oscillatory power normalized by the total size of each neuron's budget (V_o/V_b), and most importantly we study the ratio of oscillation to computation (V_o/V_c).

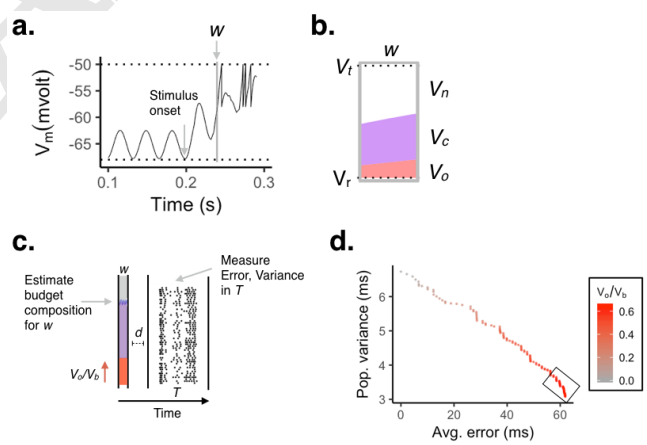

Fig. 2. The voltage budget. **a**. Membrane voltage of a single neuron. Horizontal dotted lines indicate the voltage threshold (V_t; top line) and the resting potential (V_r; bottom line). The vertical gray box depicts the 2 ms budget window (w; A = 0.1 nA; f = 20 Hz). Stimulus onset occurs at 0.2 seconds. **b.** Example of a voltage budget decomposition in time window w in panel *a.*. Red represents V_o; purple depicts V_c; white space is V_n. **c**. Diagram of budget prediction. A budget estimate is formed at t, then d seconds later a single cycle of oscillation begins. Over the period of that cycle's length T we estimate the error and variance of the population spiking. **d.** Plot of average error versus population variance over period T. Increases in V_o/V_b are denoted in red. Boxed area highlights the error plateau phase. In all plots: w = 2 ms; d = 2 ms; oscillation frequency f = 20 Hz.

Budget analysis. We imagine each neuron seeks to spend its budget as prudently as possible *for each cycle of oscillation*. We model this by using a voltage budget at time t to predict the effect of a later ($t + d$) cycle of oscillation on spiking,

where $t + d$ is the time at which we induce a burst of oscillation. We then measure spiking activity for T seconds after the oscillation begins, where $T = 1/f$ (illustrated in Figure 2**c**). Tuning oscillations requires the cell make a temporal predictions, because once a cycle of oscillation begins the dynamics of the network are committed to completing that cycle. Mathematically, oscillations begin as an initial value problem; how the oscillation starts largely determines where it ends.

To demonstrate the broader trend we first examine the effect of a 20 Hz rhythm. Later we show that the key trends for this rhythm is robust to changes to frequency (4-30 Hz), noise (0 - 2.5 mV), and synaptic weights (0.3 -45 nS) (Fig 5), as well as to neuron type (Fig 6).

Each experiment begins with a reference run, where the neurons are initialized with their unique synaptic weights, capacitances, and recovery variables. After 200 milliseconds of settling in time, each neuron to subjected to the stimulus—a pulse of 8 Hz Poisson activity lasting 50 milliseconds. The entire spiking response to this brief stimulus is recorded for later reference. The time of the first spike serves as a locking point E. The reference run is followed by the budget run, where we measure the level of computational influence V_c. In each neuron, we then calculate V_o for many possible levels of peak oscillatory power A, and store these for later reference. We do not imagine that a real biological system would need to explicitly conduct either the reference or the budget phases. This would instead be learned by the system, becoming implicit in the synaptic weights and delay times of the circuit. Finally, in the experimental phase, we sweep over a large range of A values corresponding to V_o changes ranging from < 0.1 mV to greater than 5 mV. In each of these, a single cycle of oscillation introduced at $E+d$ seconds, and spiking behavior is monitored for T seconds 2**c**). In all our simulations $d = 2$ ms.

A model of bursts. We study single cycle "bursts" of oscillation for two reasons. The first reason is empirical: when individual experimental trials of real oscillatory data are examined, oscillations often appear as bursts (16, 17) even though averaging many trials gives what appears to be sustained rhythm (18). As a result of this averaging, most theoretical work on oscillations assume a sustained rhythm though there are exceptions (18–20). Still, it remains unclear what, in general, a single cycle of oscillation can accomplish, which is why we focus on the this case here.

The second reason is practical. Bursts, by definition, are accompanied by aperiodic windows of neuronal activity that serve as a natural reference for doing error calculations. By comparing firing during an aperiodic trial to a later periodic trial, we can measure how synchronization perturbs single neuron computations in a simple and direct way.

The trade-off between error and variance. As V_o rises in proportion of the total budget V_b, population variance declines consistent with previous models (21, 22). We also observe that while variance decreases computational errors increase. This trade-off is roughly linear for low power oscillations (Fig2**d**). This early trend implies a nearly one-to-one trade-off between coordination at the population level and errors introduced at the neuronal level. However as relative oscillatory power grows eventually error plateaus at the inflection point of the error-variance curve (boxed region, Fig. 2**d**). Separating the pseudo-linear from the plateau phase is the basis for how we come to separate healthy from pathological synchrony.

Defining pathological oscillation. The fundamental idea we wish to test: can a pathological oscillation be identified using only computational activity as a reference? To do this we consider the ratio of oscillation to computation, V_o/V_c. Before getting to our numerical results though it is worth considering V_o/V_c on its own. When $V_o/V_c < 1$, computation dominates the system in that the majority of the voltage driving each spike reflects the computational dynamics of each neuron. But beyond $V_o/V_c = 1$, the majority of spikes reflect the *modulatory* oscillation. Intuitively then, once a modulator dominates, it is in effect no longer a modulator: the neurons downstream gain information mostly about the oscillatory entrainment "signal" and not the population's normal computation. This idea leads to our first new definition: oscillations that exceed $V_o/V_c = 1$ are *intrinsically pathological*.

Numerical analysis of V_o/V_c confirms this prediction: at $V_o/V_c = 1.0$, the average error and variance of the population plateau and are incrementally—asymptotically—approaching their maximum values (for error see Figure 4**a** and **c**). At the individual level, error for all neurons plateaus by $V_o/V_c = 1.0$, however the steepness and curvature of this rise is neuron dependent (Figure 4**b**). These cell-level differences suggest that global oscillation is not optimal. That is, better trade-offs may be possible if the oscillation is tuned to each neuron's particular computational curve; an idea we return to below.

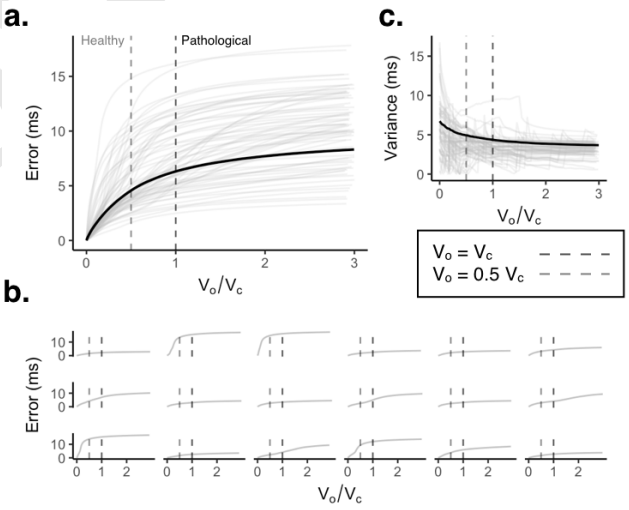

Fig. 3. Voltage budget spent on oscillation increases error and reduces variance. **a.** Plot of error predicted by changes to V_o/V_c. Grey lines indicate individual neurons. Black is the population average. Dashed lines suggest the end points for the pathological (dark) and healthy (light) inequalities. **b.** Randomly chosen individual examples taken from *a*. **c.** Variance predicted by the ratio V_o/V_c. Grey lines indicate individual neurons. Black represents the population average.

Defining healthy oscillation. In a later section we prove that the optimal trade-off between error and variance is linear. This proof leads to our second criterion: on normative grounds we suggest that healthy oscillations have $V_o/V_c < 0.5$. This point demarcates the average transition from pseudo-linear to plateau (Figure 3**a** and **b**). At $V_o/V_c = 0.5$, this point also separates the transition in variance from linear to its

DRAFT

(even sharper) plateau phase, after which increasing oscillatory power has an increasingly marginal effect on error or variance, as those are close to maximal and minimal respectively (3**c**).

Variance and V_o. Variance and oscillatory budget (V_o) have a more complex relationship than error and oscillation. In individual neurons we see that as oscillatory budget increases variance exists as two phenomena, which act in opposition. For sufficiently small increases in oscillatory budget, synchrony declines along a logistic path (see light orange traces in Figure 4**a** and **e**). However once the oscillatory budget reaches a critical value, specific to each neuron, variance discontinuously increases, denoted by the dark orange in Figure 4**a** and **e**. These sudden jumps happen every time a new, extra action potential is generated. New action potentials tend to be at the extremes of an oscillatory cycle, as illustrated in Figure 4**f**. In contrast to variance, however, individual neuron's errors change smoothly (Figure 4**e**).

On average, increases in oscillatory budget increase synchrony. Counter-intuitively, however, very weak oscillations can generate as much synchrony as strong oscillations. In fact, the quantiles analysis in Figure 4**c** suggests that weak oscillations, when targeted at a select sub-set of neurons, can have lower overall variance, while showing much greater capacity for error/variance trade-off.

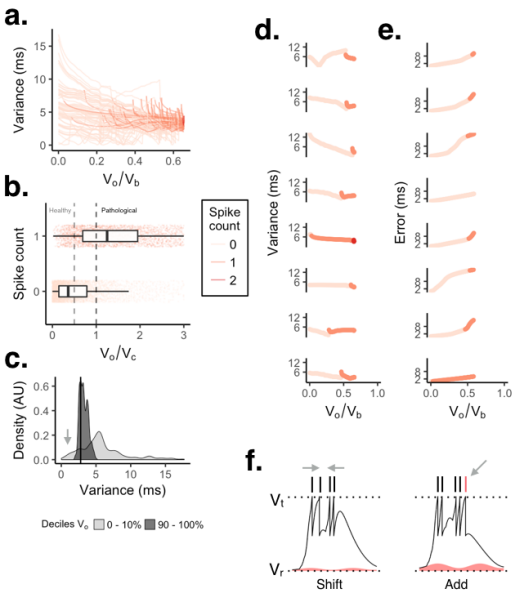

Fig. 4. Oscillations both decrease and increase variance. **b.** A plot of variance and V_o, where traces from individual neurons are colored by the number of new spikes introduced by V_o. Light orange indicates that the oscillation introduced zero new spikes; dark orange is when the oscillation introduced one spike; red is two spikes. **b.** Spike count as a function of V_o/V_c. The left side of the light dashed line denotes the healthy ratio of V_o to V_c while the dark line denotes the start of the pathological region. **c.** Histogram density plots for the top and bottom variance deciles. The black line indicates the minimum population variance. The arrow indicates neurons that show a lower total variance for weaker, rather than stronger, oscillations. **d.** Randomly chosen examples of variance for individual neurons. **e.** Examples of error for individual neurons. **f.** Illustration of the opposing possible effects an increase in V_o can have on on variance. Left panel depicts an oscillation acting to shift spike times closer together. Right panel depicts a stronger oscillation that adds an additional spike, shown in red.

Testing robustness to frequency. Frequency has little effect on error (Fig. 5**a**) but profoundly affects the synchrony (Fig. 5**b**). Oscillations slower than 20 Hz show similar trends outlined in Figures 2-5 (yellow arrow, Fig. 5**b**), while oscillations faster than 20 Hz generate less total synchrony and plateau faster (green arrow, Fig. 5**a**).

Influence of membrane noise and synaptic weight. Biologically-relevant levels of noise have little effect on normalized variance or error (Fig. 5**c**-**d**). The range of synaptic weights, in combination with the input firing rate, modestly alter the normalized error profile (Fig. 5**e**) but can strongly affect variance (Fig. 5**f**). As shown in Fig. 5**f**, when operating in the healthy range, where $V_o/V_c \leq 0.5$, all examined synaptic weights show a consistent monotonic decrease in synchrony as V_o rises. However in the pathological range $V_o/V_c \geq 1$, weaker synapses and high V_o generate a strong increase in variance—a change in variance 5-fold larger than any other observed in our simulations. The strength of this effect might tempt us to speculate that weak synapses are particularly prone to generating pathological symptoms. When placed under what might otherwise be considered a low level of power and weak level of entrainment, neurons with weaker synapses may generate unusually asynchronous responses.

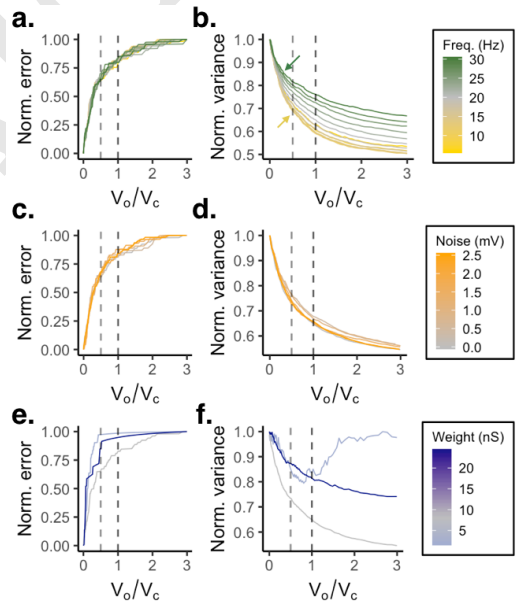

Fig. 5. Testing oscillatory budget robustness to frequency, noise, and synaptic weight. The gray colored curve on each plot represents the "standard" value otherwise used throughout. Light and dark dashed lines represent healthy and pathological V_o/V_c ratios, respectively. **a.** Simulated effect of oscillation frequency (4-30 Hz) on average error, normalized against each neuron's maximum error. **b.** Population variance decreases more quickly for stronger oscillations when those oscillations are lower frequency. **c.** Injection of noise into membrane voltage between 0 and 2.5 mV shows that noise has little effect on either error or, **d.**, variance as a function of oscillatory power. Oscillation frequency was 20 Hz (gray curve). **e.** Simulating changes in the synaptic weight range and its impact on average error. Plots are colored by the average weight. Individual ranges for the standard weight value in gray were 1.5-15 nS. Light blue illustrates weak synapses in the 0.3-3 nS range, while dark blue represents strong synapses sampled from 4.5-45 nS. Individual synaptic weights were sampled independently and uniformly for each condition. Oscillation frequency was 20 Hz (gray curve). Synaptic weights have relatively weaker effects on population error, but, **f.**, show stronger effects on population variance.

Influence of cell-type. So far we have focused on regular firing neurons. Now we consider a heterogeneous population of neurons displaying a wide range of firing modes. Our neuron dynamics are governed by the AdEx model, which is known to fit a wide range of real firing properties including fast firing, adaption, accelerating activity, bursting, and chaotic irregular firing (11). However, rather than simulating preset classes of neurons, we opted instead to smoothly sample a large range of the neural space, generating a diverse pool of firing behaviors. The space is illustrated by a few randomly selected examples in Figure 6**a**.

On average, the same budget ratios that predict healthy synchrony in regular firing neurons hold true in simulations of heterogeneous neuronal populations Fig. 6**b-c**). All models show the same one-to-one trade-off between variance and error below $V_o/V_c = 0.5$ and a error plateau by $V_o/V_c = 1.0$.

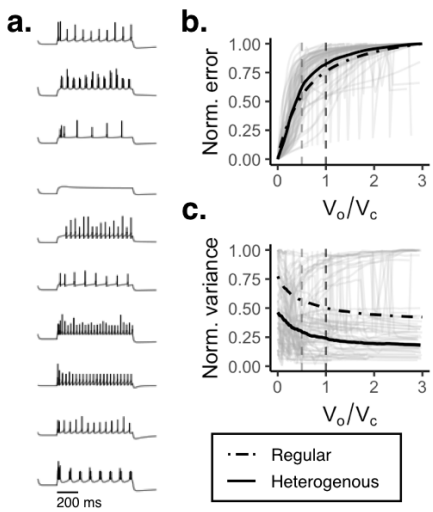

Fig. 6. Influence of cell type on voltage budget analysis. **a.** Examples of heterogeneous firing modes, driven by a 0.8 second, 0.25 nA square-wave pulse. **b.** Average (black) and individual errors (gray) for $N = 250$ heterogeneous neurons. Average error from a regular firing network is redrawn from Figure 3*a* (dot-dash). The left side of the light vertical line denotes the healthy oscillation ratio ((V_o/V_c)=0.5) while the dark light denotes the start of the pathological region ((V_o/V_c)=1.0). **c.** Variance as a function of the V_o/V_c ratio. Black is the population average. Grey traces represent individual neurons, and the dot-dashed line represents variance from the regular population.

Optimal synchronization. Our model shows that oscillations increase computational error and reduce population variance in a roughly linear fashion such that there is a trade-off between the two: oscillations reduce the population variance at the cost of computational accuracy.

To measure population synchrony and average individual neuron computational error, we use two related metrics. The first is the mean absolute error (E), which measures the average computational error. The second is the mean absolute deviation (D), which measures the variance of each neuron's spiking relative to the population average. Exploring the mathematical connections between these metrics in the abstract offers insight into the structure of the computation/communication problem, and allows us to prove an optimal algorithm for oscillatory synchronization.

$$E = \frac{1}{K}\sum_{i=1}^{K} |\hat{y} - y_i| \qquad [3]$$

$$D = \frac{1}{K}\sum_{i=1}^{K} |\bar{y} - y_i| \qquad [4]$$

Here y are spike times from individual neurons, reflecting an example when oscillatory amplitude A is greater than zero. $\hat{y}$ is the set of reference times acquired without oscillations, *i.e.*, when $A = 0$. We denote examples from y as y_i, and $\bar{y}$ is the average of y.

What's the best algorithm to shift spike times? To formalize an answer, assume we wish to change spiking variance D by some amount $\epsilon \in \mathbb{R}^1$. When ϵ is 0 there is no change in synchrony and so by definition $y = \hat{y}$ and $E = 0$. As we increase ϵ the question becomes how should we distribute that perturbation, or error, among the K spikes in the spiking population y? That is, how do we set each neuron's spiking error, ϵ_i, for the series $y = (\hat{y}_1 + \epsilon_1, \hat{y}_2 + \epsilon_2) \ldots (\hat{y}_K + \epsilon_K)$?

A naive approach, similar in character to a global oscillator, is to spread the error uniformly among the entire spiking population, y. Formally, if we decompose total error into i equal error partitions, then we have a uniform error distribution case where $\epsilon = |\epsilon_1| + |\epsilon_2| + \ldots + |\epsilon_K|$. If we wish to use error to minimize variance we must set the sign of each perturbation ϵ_i to oppose the sign of y_i; if y_i is negative, ϵ_i is positive, and vice versa (for example see Figure 7**a**, *bottom panel*). When implemented over a range of variances, this uniform approach gives rise the blue error-variance curve in Figure 7**b**. The question then becomes is this uniform strategy the optimal algorithm to balance the trade-off between computational error and population variance? That is, is there a smaller value of D for a given level of error, E?

To explore optimality, we introduce a single degree of freedom. We hold all errors equal as before, except for two neurons m and n. This lets us ask the question: by introducing a single degree of freedom can we generate more synchrony than the simple uniform error distribution strategy? If we can do so, we know that the uniform strategy is *not* optimal. To simplify the analysis, first we center all y and $\hat{y}$, by subtracting $\bar{y}$ from all, and remove the normalization term $1/K$, leading to equations 5 and 6.

$$E' = \sum_{i=1}^{K} |\hat{y}' - y_i'| \qquad [5]$$

$$D' = \sum_{i=1}^{K} |y_i'| \qquad [6]$$

An instructive, but extreme, use of a single degree of freedom is to assign all the values from one free perturbation to the other. That is we set ϵ_n to 0, and ϵ_m to $\frac{-2\epsilon}{K}$. From here it becomes clear that if we apply ϵ_m to $\max(|y|)$ and ϵ_n to $\min(|y|)$ this will produce the largest possible decrease in adjusted variance, D', and represents a $\max(|y|)$ - $\min(|y|)$ decrease in D' compared to the uniform approach.

DRAFT

If we free ourselves from the contrived example of uniformly distributing ϵ, we can see now that the optimal approach is to apply ϵ to $\max(|y|)$; $\max(|y|) \rightarrow \max(|y|) \pm \epsilon$. This means that for any given arrangement of spike times, the optimal trade-off between computation and population synchronization introduced by an oscillation is to shift the spike farthest from the mean. We implement this as an incremental algorithm, which can change the variance by ϵ by intuitively ranking y, shifting the spike $\max(y)$ by some very small amount in time δ until a running sum of all these tiny perturbations is equal to ϵ. In theory, the smaller the δ the better the algorithm approximates the true optimal solution. In practice, a δ below 0.00001 seconds is sufficient. See Figure 7**a** (*top panel*) for an example of a synchronized spike train and **b** for a depiction of the optimal error-variance curve.

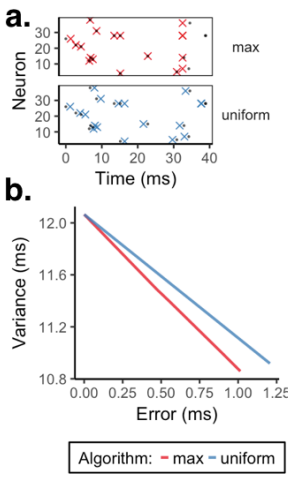

Fig. 7. A comparison of two algorithmic strategies for inducing synchrony. **a.** In these examples, two populations of synchronized neurons have their original spikes (•) shifted to new positions (X) based on two different algorithms: either by iteratively shifting the spikes furthest from the mean (top, red) or by shifting each spiking uniformly (bottom, blue). **b.** Error-variance trade-off curves for the strategies illustrated in *a*. The smaller the error for a given level of variance, the more optimal the algorithm is. Note that the $max(|y|)$ algorithm (red) optimally minimizes the trade-off between computational error and population variance.

Discussion

We build a simple model of external oscillatory entrainment. We mimic the real biological case where one "pacemaker" population coordinates another, aperiodic, population, such as in the case of top-down oscillatory influence (23–26). Our model is the simplest case we could devise that allows for the precise, biologically testable, predictions of oscillatory over-entrainment.

Limits of the model. We studied a toy model designed as a best-case scenario for understanding the trade-off between computation and communication. Our purpose was to create an initial (the first, as far as we are aware) quantitative model of healthy versus pathological oscillations, and so begin by using the simplest model that embodies the problem: uncoupled neurons subjected to a global oscillator. Real biological systems, and more complete simulations, feature extensive yet sparse connections between neurons. These connections naturally create dependencies between the activity of one neuron and the others in its population—a basic phenomenon we do not capture here. Theoretical analysis of neural coding, and decades of empirical research, however, suggest that, despite this extensive connectivity, real neurons act with a high degree of independence, which is optimal for computational efficiency. The biological implementation of this independence relies on a combination of independent dendritic computation and the precise arrangement of excitatory and inhibitory circuits. So, despite the simplicity of our model, it may act as a reasonable approximation of real complex networks which maintain a high degree of independence between neurons.

Oscillations in a network can also arise by self-organization, where rhythmicity is driven by interactions *within* the entrained population. Our model is not well-suited to this case; previous theoretical work suggests that, even in idealized cases, there is a minimum level of voltage budget needed to initiate and sustain an intrinsic oscillation (27). That is, in self-organized systems the order parameter *can't* be expected to smoothly vary in the biologically relevant 1-5 mV range, which is a requirement for our analysis to hold. Understanding the interaction between computation and coordination within a self-organized population remains an open question.

Real synchrony is often weak. It is not clear how oscillatory entrainment can be weak, yet also be an important general feature of nervous system function, as is frequently supposed (1, 2). That is, to observe oscillations in real local field potentials often requires little more than placing an electrode in the appropriate region, as field potential oscillations are relatively ubiquitous (24, 28). Observing the same oscillation in the the spiking behavior of neurons, however, often requires recording from many—even hundreds—of neurons, especially in cortical areas (29–31). In these recordings about half the neurons show no preference for an oscillation's phase (32). Those that are entrained are often weakly entrained, synchronized by at most few percent.

Our model suggests that the most effective oscillations are precisely those that are both weak and sparse. When the oscillatory power remains below the that of all other "computational" inputs, the system can exchange errors in a single neuron's computation for group-level synchrony, measured in the voltage budget analysis as the quantitative ratio $V_o/V_c \leq 0.5$. This ratio's predictions are relatively invariant to oscillation frequency, noise in the membrane potential, and variations synaptic weight; this ratio also predicts the firing properties of a large range of heterogeneous cell-types. Further, strong oscillations offer only marginal improvements in synchrony: once an oscillation grows too strong it induces new action potentials in the population. These extra action potentials tend to be at the trailing end of the neuron's response to input, increasing variance rather than gathering spikes together. Finally, our new strategy for provably-optimal coordination targets a only small fraction of the population. Targeting all neurons for coordination has a larger error cost than targeting only a few of the more extreme action potentials in a given cycle.

Oscillations as epiphenomenon. Oscillations could be a side effect, or epiphenomenon, of neural physiology. Mathematical and experimental analysis of both simple (30, 31) and complex biological (10) structures suggests a relatively large portion of the neural parameter-space generates oscillations. As a

DRAFT

Table 1. Model parameters.

Symbol	Range (unit)	Description
A	0 - 0.15 (namp)	Peak current (order parameter)
ϕ	0	Oscillation phase
f	0 - 30 (Hz)	Oscillation frequency
r_s	6 (Hz)	Stimulus firing rate
C	100-200 (pfarad)	Membrane capacitance
w_{in}	1.5 - 15 (nsiemens)	Synaptic weight
τ_e	5 (msecond)	Excitatory synaptic time constant
V_e	0 (mvolt)	Excitatory synaptic reversal potential
I_{bias}	10 (namps)	Bias current
σ	0 - 5 (namps)	Injected membrane noise
g_l	10 - 18 (nsiemens)	Leak conductance
V_l	-50 - -70 (mvolt)	Leak reversal potential
a	-11 - 4 (nsiemens)	Slow recovery term
τ_w	30 - 300 (msecond)	Slow time constant
b	0 - -12 (psiemens)	Instantaneous recovery term
V_{rheo}	-46 - -58 (mvolt)	Instantaneous recovery voltage
Δt	2.0 (msecond)	Recovery voltage rate term

result, oscillations may be only a nuisance. An artifact of other biological factors. An epiphenomenon.

On the other hand, oscillations may have arisen early on during nervous system evolution, initially as an artifact. However, over time these oscillations were co-opted and put to, perhaps several distinct, uses. After more than 80 years of study, separating these two possibilities remains an open problem.

We offer a new approach to the oscillations-as-epiphenomenon debate. By deriving *a priori* quantitative bounds between the healthy and pathological ranges of oscillation, and in defining an optimal algorithmic approach to synchrony, we suggest that these normative constraints can help in finally separating functional oscillations from physiological epiphenomena. To see how, recall that our analysis suggests that oscillatory input—and all the other neuronal inputs (collected into the "computational" term in our model)—into the neuron exist in equilibrium. At one end of this equilibrium are neurons whose action potentials are independent. At the other end are neurons who are completely synchronized, and therefore redundant. Oscillations that are *just* an artifact would be expected to explore both extremes. On the other hand, oscillations that track strict normative bounds must be functional, rather than epiphenomenal.

Materials and Methods

The network. Our model was a network of $N = 250$ Adaptive Exponential (AdEx) neurons (15) sharing a common oscillatory drive, $I_{\text{osc(t)}} = \frac{A}{2}(1 + \sin(t2\pi f + \phi))$; where f is the frequency of the oscillation, and ϕ denotes the phase offset. Each neuron was driven by an identical Poisson process. Variability in computational output depended solely on the variations in synaptic weight, and 'slow" membrane recovery parameters a and τ_w, all of which were independently sampled from a uniform distribution (15) All synapses were excitatory, and governed by a single exponential decay term with time constant τ_e. Membrane noise currents were driven by a Ornstein–Uhlenbeck process, with a 5 ms time constant.

Neuron-type. Initially we modeled regular firing neurons, whose parameters were piblihsed in shown (11). To generate the heterogeneous population in Figure 6, we sampled uniformly sampled from within a large range of AdEx membrane parameters, whose values are found in Table 1. These were in turn drawn from several neuron types previously described by (11).

ACKNOWLEDGMENTS. B.V is supported by a Sloan Research Fellowship, the Whitehall Foundation (2017-12-73), and the National Science Foundation (1736028). EJP wishes to thank Rachel L. Storer for invaluable discussions during development the budget analysis approach.

DRAFT

1. Voytek B, Knight RT (2015) Dynamic network communication as a unifying neural basis for cognition, development, aging, and disease. *Biological Psychiatry*.
2. Fries P (2005) A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends in Cognitive Sciences* 9(10):474–480.
3. Chen Y, Zhang H, Wang H, Yu L, Chen Y (2013) The Role of Coincidence-Detector Neurons in the Reliability and Precision of Subthreshold Signal Detection in Noise. *PLoS ONE* 8(2).
4. Zhou P, Burton SD, Snyder AC, Smith MA, Nathan N (2015) Establishing a Statistical Link Between Network Oscillations and Neural Synchrony. pp. 1–22.
5. Schaefer AT, Angelo K, Spors H, Margrie TW (2006) Neuronal oscillations enhance stimulus discrimination by ensuring action potential precision. *PLoS Biology* 4(6):1010–1024.
6. Brette R (2012) Computing with neural synchrony. *PLoS computational biology* 8(6):e1002561.
7. Song S, Miller KD, Abbott LF (2000) Competitive Hebbian Learning through Spike-Time Dependent Synaptic Plasticity. *Nature Neuroscience* 3:919–926.
8. Markram H, Lübke J, Frotscher M, Sakmann B (1997) Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* 275(5297):213–215.
9. Brittain JS, Sharott A, Brown P (2014) The highs and lows of beta activity in cortico-basal ganglia loops. *European Journal of Neuroscience* 39(11):1951–1959.
10. Markram H (2015) Reconstruction and Simulation of Neocortical Microcircuitry. *Cell* 163(2):456–492.
11. Naud R, Marcille N, Clopath C, Gerstner W (2008) Firing patterns in the adaptive exponential integrate-and-fire model. *Biological Cybernetics* 99(4-5):335–347.
12. Advani M, Ganguli S (2016) Statistical Mechanics of High-Dimensional Inference. pp. 1–12.
13. Fusi S, Miller EK, Rigotti M (2016) Why neurons mix: High dimensionality for higher cognition. *Current Opinion in Neurobiology* 37:66–74.
14. Izhikevich EM (2010) *Dynamical Systems in Neuroscience.* (The MIT Press), pp. 1–464.
15. Brette R, Gerstner W (2005) Adaptive Exponential Integrate-and-Fire Model as an Effective Description of Neuronal Activity. *J Neurophysiol* 94:3637–3642.
16. Feingold J, Gibson DJ, DePasquale B, Graybiel AM (2015) Bursts of beta oscillation differentiate postperformance activity in the striatum and motor cortex of monkeys performing movement tasks. *Proceedings of the National Academy of Sciences* 112(44):201517629.
17. Alekseichuk I, et al. (2016) Spatial Working Memory in Humans Depends on Theta and High Gamma Synchronization in the Prefrontal Cortex. *Current Biology* pp. 1–9.
18. Jones SR (2016) When brain rhythms aren't 'rhythmic': implication for their mechanisms and meaning. *Current Opinion in Neurobiology* 40:72–80.
19. Lundqvist M, et al. (2016) Gamma and Beta Bursts Underlie Working Memory. *Neuron* 90(1):152–164.
20. Shin H, Law R, Tsutsui S, Moore CI, Jones SR (2017) The rate of transient beta frequency events predicts behavior across tasks and species. *eLife* 6:1–31.
21. Needleman DJ, Tiesinga PHE, Sejnowski TJ (2001) Collective enhancement of precision in networks of coupled oscillators. *Physica D: Nonlinear Phenomena* 155:324–336.
22. Fries P (2009) Neuronal gamma-band synchronization as a fundamental process in cortical computation. *Annual review of neuroscience* 32:209–24.
23. Engel aK, Fries P, Singer W (2001) Dynamic predictions: oscillations and synchrony in top-down processing. *Nature reviews. Neuroscience* 2(10):704–716.
24. Buzsáki G, Anastassiou Ca, Koch C (2012) The origin of extracellular fields and currents–EEG, ECoG, LFP and spikes. *Nature reviews. Neuroscience* 13(6):407–20.
25. Fries P (2015) Rhythms for Cognition: Communication through Coherence. *Neuron* 88(1):220–235.
26. Heeger DJ (2017) Theory of cortical function. *Proceedings of the National Academy of Sciences* 114(8):1773–1782.
27. Börgers C, Franzesi G, LeBeau FEN, Boyden ES, Kopell NJ (2012) Minimal size of cell assemblies coordinated by gamma oscillations. *PLoS Computational Biology* 8(2).
28. Siegel M, Donner TH, Engel AK (2012) Spectral fingerprints of large-scale neuronal interactions. *Nature Reviews Neuroscience* 13(February):20–25.
29. Wang XJ (2010) Neurophysiological and computational principles of cortical rhythms in cognition. *Physiological reviews* 90(3):1195–268.
30. Börgers C, Kopell N (2003) Synchronization in networks of excitatory and inhibitory neurons with sparse, random connectivity. *Neural computation* 15:509–538.
31. Brunel N (1999) Fast Global Oscillations in Networks of Integrate-and-Fire Neurons with Low Firing Rates. 1671:1621–1671.
32. Jia X, Tanabe S, Kohn A (2013) Gamma and the Coordination of Spiking Activity in Early Visual Cortex. *Neuron* 77(4):762–774.