The Refractory Period Matters: Unifying Mechanisms of Macroscopic Brain Waves

Corey Weistuch
corey.weistuch@stonybrook.edu
Laufer Center for Physical and Quantitative Biology and Department of Applied
Mathematics and Statistics, Stony Brook University, Stony Brook, NY 11794, U.S.A.

Lilianne R. Mujica-Parodi
lilianne.strey@stonybrook.edu
Laufer Center for Physical and Quantitative Biology, Departments of Biomedical
Engineering and of Physics and Astronomy, Program in Neuroscience, and
Department of Chemistry, Stony Brook University, Stony Brook, NY 11794, U.S.A.,
and Athinoula A. Martinos Center for Biomedical Imaging, Massachusetts General
Hospital and Harvard Medical School, Boston, MA 02129, U.S.A.

Ken Dill
dill@laufercenter.org
Laufer Center for Physical and Quantitative Biology, Department of Physics
and Astronomy, and Department of Chemistry, Stony Brook University,
Stony Brook, NY 11794, U.S.A.

The relationship between complex brain oscillations and the dynamics of individual neurons is poorly understood. Here we utilize maximum caliber, a dynamical inference principle, to build a minimal yet general model of the collective (mean field) dynamics of large populations of neurons. In agreement with previous experimental observations, we describe a simple, testable mechanism, involving only a single type of neuron, by which many of these complex oscillatory patterns may emerge. Our model predicts that the refractory period of neurons, which has often been neglected, is essential for these behaviors.

1 Introduction

A major interest in neuroscience is understanding how macroscopic brain functions, such as cognition and memory, are encoded at the microscale of neurons and their topological connectivities. One of the significant developments in this direction was the Wilson-Cowan (WC) model, describing the averaged behavior of large populations of simple excitatory and inhibitory neurons in terms of a set of coupled, mesoscale differential equations (Destexhe & Sejnowski, 2009; Wilson & Cowan, 1972, 1973). With only a few

physical parameters, WC provided one of the first mechanisms for simple (single-frequency) oscillations across the brain, such as the hypersynchronized dynamics observed during epileptic seizures (Destexhe & Sejnowski, 2009; Shusterman & Troy, 2008). More recently, generalized WC-like models have been used to describe heterogeneous populations of neurons ranging in scale from single regions to networks of activities across the whole brain (Breskin, Soriano, Moses, & Tlusty, 2006; Deco, Jirsa, Robinson, Breakspear, & Friston, 2008; Destexhe & Sejnowski, 2009; Hopfield, 1982; Schneidman, Berry, Segev, & Bialek, 2006; Tkacik, Schneidman, Berry, Michael, & Bialek, 2006; Weistuch, Mujica-Parodi, Amgalan, & Dill, 2020).

But there remain important macroscopic brain behaviors that WC-like models fail to adequately explain (Chow & Karimipanah, 2019; Muir, 1979). One example is theta oscillations in the hippocampus, which have multiple superimposed frequencies and are thought to be critical for memory formation and storage (Buzsáki, 2002; Buzsáki & Draguhn, 2004; Colgin, 2013). They are believed to be generated through recurrent feedback involving excitatory neurons only (Buzsáki, 2002). Another example is gamma oscillations, which are high-frequency chaotic firing patterns associated with a wide-range of complex brain activities (Buzsáki & Wang, 2012). They are believed to arise in networks of inhibitory neurons. While WC-like models can exhibit these complex patterns of brain activity, they require many different subtypes of neurons to do so (Goles-Chacc, Fogelman-Soulié, & Pellegrin, 1985; Keeley, Fenton, & Rinzel, 2017; Neves & Monteiro, 2016; Rajan & Abbott, 2006).

Brain oscillations can be facilitated by the refractory period of a population of neurons (Greenberg & Hastings, 1978; Sanchez-Vives & McCormick, 2000). Yet, while the WC model considers both an absolute and a relative refractory period of neurons, it can only be applied under quasi-steady-state conditions (Feldman & Cowan, 1975; Muir, 1979). As a result, the WC model only exhibits simplified, lower-dimensional behaviors. Here, we describe a model of similar simplicity to the WC model but that also accounts for the non-steady-state behaviors induced by the refractory periods of neurons (see Figure 1A compared to the effective WC model in Figure 1B for a single population of neurons). Related generalizations indeed exhibit complex dynamical features missing from Benayoun, Cowan, van Drongelen, and Wallace (2010), Pinto and Copelli (2019), and Goychuk and Goychuk (2015). However, current models make broad assumptions about neural dynamics, and their behaviors have yet to be fully determined. The novelty of this article is two-fold. First, we treat the stochastic dynamics of the model using maximum caliber (Max Cal), a principle of statistical inference that applies to systems of pathways or systems of dynamical processes, which draws more directly on data and is freer of unwarranted model assumptions (Dixit et al., 2018; Ghosh, Dixit, Agozzino, & Dill, 2020; Pressé, Ghosh, Lee, & Dill, 2013; Weistuch, Agozzino, Mujica-Parodi, & Dill, 2020). Second,
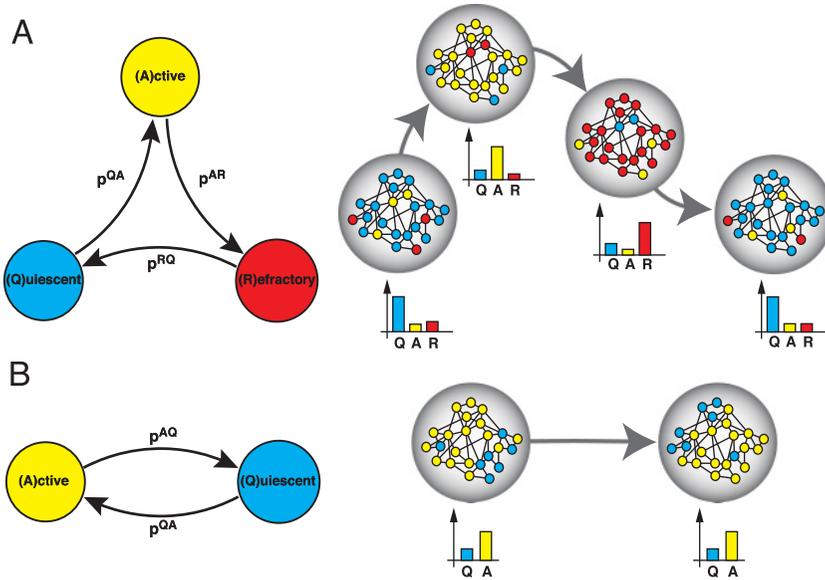
Figure 1: Neural oscillations depend on the refractory period. (A). A network representation of the model. Left: Neural activity represented as a Markov chain. Here the rates (p^{QA}, p^{AR}, and p^{RQ}) determine the occupation probabilities of the three states: Q, A, and R. Right: The fraction of neurons in each state evolves over time. Active neurons (yellow) can either excite (as shown) or inhibit neighboring quiescent neurons (blue) by modulating the average firing probability p^{QA}. Complex oscillations emerge when these neurons must wait to fire again (red). (B). Without the refractory state (left), the fraction of neurons in each state does not evolve over time (right).

the model is well characterized and provides a simple, testable mechanism for both multifrequency and chaotic oscillations.

2 The Physics of the Model

We represent a generic network of N neurons (labeled $i = 1, 2, \ldots, N$) as a graph; nodes represent each neuron, and edges are synaptic connections (see Figure 1). Each node (i) also has a time-dependent state $S_i(t)$, representing the activity of a neuron. In particular, the nodes of our network can be in any one of three states: quiescent (Q), or silent but able to fire; active (A), or firing; or refractory (R), or unable to fire. Additionally, the states of each node evolve stochastically over time: $Q \rightarrow A \rightarrow R \rightarrow Q$. The rate of each of these transitions is then chosen to reflect the biophysical dynamics of real neurons.

We use the principle of maximum caliber (Max Cal) to infer these transitions directly from the data (Dixit et al., 2018; Ghosh et al., 2020; Pressé et al., 2013). Here Max Cal, the dynamical extension of maximum entropy, provides the simplest and least-biased model consistent with a few known transition rates, such as average neuronal firing rates and correlations (Dixit et al., 2018; Ghosh et al., 2020; Pressé et al., 2013; Tkacik et al., 2006). This model takes the form of a probability distribution P over different dynamical trajectories Γ of populations of neurons.

Using Max Cal, we model how the fraction of neurons in each state (π_Q, π_A, and π_R) evolves over time. While our approach is applicable to any number of neurons, we focus on the case when this number is large. We maintain our focus here for two reasons. First, it presents an enormous simplification, as we can study the long-time behavior of our model using mean-field theory (Deco et al., 2008; Gerstner, 2000; Jirsa & Haken, 1997; Omurtag, Knight, & Sirovich, 2000). Second, it is often a reasonable approximation, as system behaviors converge to their means when their number of components N is large.

2.1 Obtaining the Stochastic Dynamics of the Model Using Maximum Caliber.

Here we ask how simple neuronal interactions might give rise to complex patterns of brain dynamics. To answer this, we use Max Cal to build a minimal model of neural dynamics. Here, the Caliber $\mathcal{C}$ is defined as the path entropy over the probability distribution of trajectories P_Γ subject to a prespecified set of constraints:

$$\mathcal{C}[P_\Gamma] = -\sum_\Gamma P_\Gamma \log P_\Gamma + \sum_{i,\Gamma} \lambda_i A_{i,\Gamma} P_\Gamma, \tag{2.1}$$

where λ_i are the Lagrange multipliers constraining generic average quantities $\langle A_i \rangle$. Here the quantities that we measure are the transitions of nodes between different states: $l_i^{QA}(t)$, $l_i^{AR}(t)$, and $l_i^{RQ}(t)$. In particular, $l_i^{QA}(t)$ is one if the ith node transitions from Q to A during the time interval $[t, t+1]$ and is otherwise zero; the other transition indicators are defined similarly. We thus want to constrain our model in such a way as to preserve the average transition rate between each pair of states:

$$r_{QA} = \frac{1}{N}\left\langle \sum_{i=1}^{N} l_i^{QA}(t) \right\rangle = \left\langle \pi_Q(t) p^{QA}(t) \right\rangle,$$

$$r_{AR} = \frac{1}{N}\left\langle \sum_{i=1}^{N} l_i^{AR}(t) \right\rangle = \left\langle \pi_A(t) p^{AR}(t) \right\rangle,$$

$$r_{RQ} = \frac{1}{N}\left\langle \sum_{i=1}^{N} l_i^{RQ}(t) \right\rangle = \left\langle \pi_R(t) p^{RQ}(t) \right\rangle. \tag{2.2}$$

Here $\langle \circ \rangle$ denotes an average over time, and the second set of equalities holds when the number of neurons N is large. The average rates r_{QA}, r_{AR}, and r_{RQ} are computed from experimental data as the time-averaged fraction of nodes transitioning from $Q \to A$, $A \to R$, and $R \to Q$, respectively. In contrast, the right-hand sides of the above equations are computed over the different trajectories that our inferred model will produce. Here, these averages are constrained using the Lagrange multipliers h^{QA}, h^{AR}, and h^{RQ}, respectively (see appendix A and B).

These Lagrange multipliers can then be incorporated into the transition probabilities p^{QA}, p^{AR}, and p^{RQ} as discussed in appendix B. Here, p^{AR} and p^{RQ} are constants and are functions of their respective Lagrange multipliers. More directly, p^{AR} (resp. p^{RQ}) can be computed as the average fraction of refracting A (resp. quiescing R) per unit time. In contrast, a key property of neurons is their ability to communicate by altering the firing activity of their neighbors. Specifically, firing neurons can either increase (excite) or decrease (inhibit) the probability that other quiescent neurons fire. Here we include this with the additional constraint

$$C = \frac{1}{N}\left\langle \sum_{i=1}^{N} l_i^{QA}(t) N_A(t) \right\rangle = N\left\langle \pi_Q(t)\pi_A(t)p^{QA}(t) \right\rangle, \tag{2.3}$$

where $N_A(t)$ is the number of active neurons at time t. A large value of C thus represents a population of excitatory neurons, as the firing probability of additional nodes increases with the number of currently active nodes. Conversely, a small (close-to-zero) value of C represents an inhibitory population, whereby the activation of a few nodes suppresses subsequent firing of additional nodes. This constraint is enforced by the Lagrange multiplier J, the coupling constant. Thus, the transition probability p^{QA} is a function of both the raw firing probability of a neuron (controlled by $h = h^{QA}$) and the feedback strength, J. This relationship is given by (see appendix B)

$$p^{QA}(\pi_A) = \frac{e^{h+J\pi_A}}{1 + e^{h+J\pi_A}}. \tag{2.4}$$

Thus, h and J can be alternatively computed by fitting the shape of p^{QA} for different values of π_A. Taken together, our model is a function of four parameters (p^{AR}, p^{RQ}, h, and J), each uniquely chosen to reproduce our four experimental constraints.

2.2 The Mean-Field Equations. Here we compute the time evolution of the average fraction of neurons in each state (π_Q, π_A, and π_R). Before proceeding, we make a few notational simplifications to enhance readability. First, we use Δ to refer to the change in a variable over time. For example, $\Delta\pi_A(t) = \pi_A(t+1) - \pi_A(t)$. And second, aside from their initial

definitions, we implicitly assume the time dependence of our variables and drop (t) when writing our equations. For example, $\pi_A(t)$ will just be written as π_A, and $\Delta\pi_A(t)$ will just be $\Delta\pi_A$. Thus, after maximizing the caliber subject to our four constraints and computing the average (mean-field) dynamics (see appendixes A and B), we find that our system can be described by two coupled equations:

$$\Delta\pi_Q = (1 - \pi_A - \pi_Q)p^{RQ} - \pi_Q p^{QA},$$
$$\Delta\pi_A = \pi_Q p^{QA} - \pi_A p^{AR}. \tag{2.5}$$

Here we have eliminated the corresponding third equation for $\Delta\pi_R$ using the constraint that the fractions of nodes of each type sum to unity (i.e., the number of neurons is fixed).

In contrast to typical modeling approaches, we have made no assumptions in deriving these equations other than the fact that our experimentally observed constraints are reasonably descriptive of neural dynamics. Thus, we expect our model to be widely applicable, even when other previous models fail.

Also, each of our parameters has a clear biological interpretation. First, p^{AR} and p^{RQ} control the average amount of time neurons spend (respectively) active and refractory. Thus, when p^{AR} is large (as might be expected of real neurons), nodes are only briefly active. On the other hand, p^{RQ} might be expected to be small, as biological neural oscillation occurs at a relatively low frequency (an action potential lasts 1 ms, but the fastest oscillations have a period of about 10 ms). Reflecting these requirements, we fix p^{AR} and p^{RQ} at 0.8 and 0.01, respectively. Additionally, h, the unit-less average firing threshold, controls the fraction of neurons that fire spontaneously. Thus, we should have $h < 0$, reflecting a low-level of baseline activity. Finally, J reflects feedback, or synaptic coupling, between neighboring neurons and can be either positive (excitatory) or negative (inhibitory).

We study two general classes of brain oscillations, corresponding to the network activities of excitatory ($J > 0$) and inhibitory ($J < 0$) neurons. Here, excitatory oscillations are characterized by high-amplitude waves of activity followed by long periods of silence during which most neurons are refractory (Buzsáki, 2002; Buzsáki & Draguhn, 2004). In contrast, networks of inhibitory neurons fire asynchronously, producing low-amplitude, high-frequency oscillations (Brunel & Hakim, 2008; Buzsáki & Wang, 2012; Korn & Faure, 2003). And unlike WC, both of these behaviors can be exhibited by our model using only a single population of neurons (Feldman & Cowan, 1975; Muir, 1979).

3 Model Properties

The formulation of WC-like models is based on quasi-steady-state dynamics and requires multiple populations of neurons to exhibit even
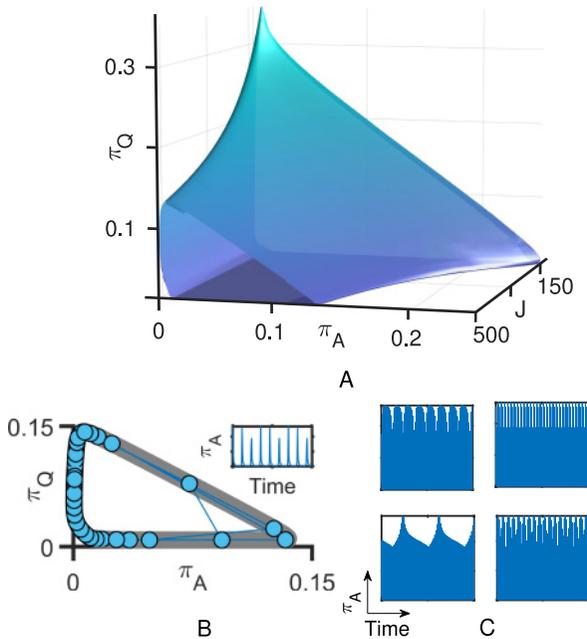
Figure 2: Excitatory couplings produce complex oscillations. (A). The phase plane (π_Q vs π_A) for different values of J (at $h = -5$), illustrating the emergence of oscillations (rings). (B). A typical trajectory (blue) in phase space (gray) and over time (inset). Because π_Q and π_A vary slightly with each cycle, the oscillatory amplitude changes over time. These changes are very sensitive to J. (C). Examples of different oscillatory patterns for different values of J.

single-frequency oscillations (see appendix C and Feldman & Cowan, 1975, and Muir, 1979). In contrast, the behaviors of real neurons are more complex, often involve few neural subtypes, and have been difficult to describe mechanistically (Buzsáki & Draguhn, 2004; Chow & Karimipanah, 2019). We next demonstrate the significant improvements of our model over these previous approaches.

Unlike WC, our model explains how a single population of excitatory neurons can generate multifrequency oscillation patterns (see Figure 2). In particular, Figure 2A, depicting the phase plane of our model, shows the emergence of oscillatory activity (rings) when the coupling $J > 0$ is nestled within a critical region. Here the amplitude of each oscillation can vary with every cycle (see Figure 2B), producing the multifrequency bands expected of real neurons (Bacak, Kim, Smith, Rubin, & Rybak, 2016; Buzsáki, 2002; Vladimirski, Tabak, O'Donovan, & Rinzel, 2008). Unlike WC, by only slightly tuning J, our model predicts the emergence of highly distinct patterns of activity (see Figure 2C). And indeed, a similar mechanism
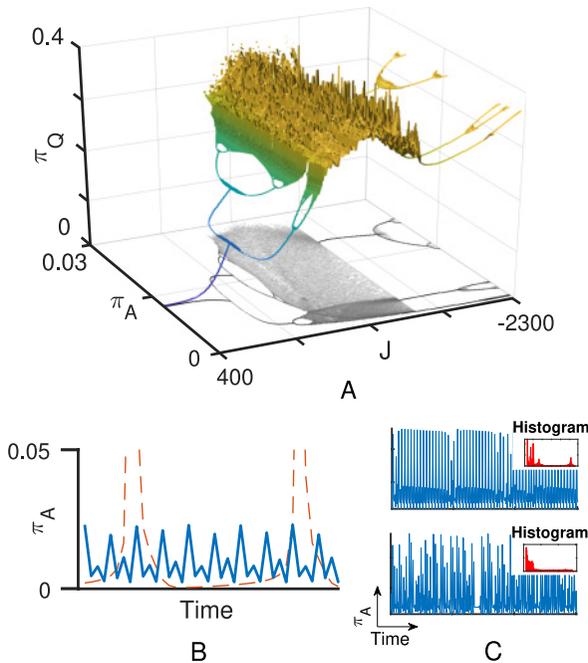
Figure 3: Inhibitory couplings produce chaotic oscillations. (A). The phase plane (π_Q vs π_A) and its projection (black) for different values of J (at $h = -1$). The number of points (for each J) corresponds to the period of the associated oscillation. As J is decreased, the oscillations become chaotic and aperiodic (orange). (B). Comparison of inhibitory (blue) to excitatory (orange) oscillations produced by our model. (C). Examples of different chaotic oscillatory patterns along with a histogram of π_A over time (inset). Here information is stored, not in the timing but in the probabilities of different amplitudes.

is thought to underlie tremendous information capacity of real networks of neurons (Averbeck, Latham, & Pouget, 2006; Panzeri, Macke, Gross, & Kayser, 2015; Vladimirski et al., 2008).

At the other extreme, recurrent inhibitory networks of neurons have been shown to produce high-frequency and sometimes chaotic firing patterns (Buzsáki & Wang, 2012; Korn & Faure, 2003). In contrast to excitatory networks, inhibitory neurons fire in small bands of only a few neurons at a time. As the strength of this inhibition is increased, these neurons fire asynchronously and chaotically (Buzsáki & Wang, 2012). And while WC-like models can produce some inhibitory oscillations, they require multiple subtypes of neurons to do so (Goles-Chacc et al., 1985; Keeley et al., 2017). Figure 3 describes how these features emerge from our model using only a single subtype of neuron. Here, Figure 3A. depicts the phase plane of
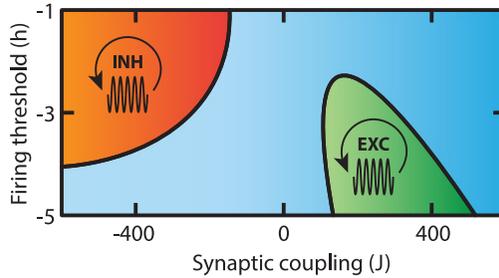
Figure 4: The phase diagram of our model, depicting the emergence of excitatory (green) and inhibitory (orange) oscillations. In the blue region, brain activity is constant over time. In contrast to WC-like models, however, oscillations can be produced by tuning h and J.

our model for different values of J ($J < 0$). The number of points for each J corresponds to the period of the inhibitory oscillations. As J is decreased, this period continually doubles until it diverges to infinity and chaos emerges. Because inhibitory neurons fire as far apart as possible, they oscillate with a much higher frequency (as well as a lower amplitude) as compared to excitatory neurons (see Figure 3B; Buzsáki & Draguhn, 2004; Buzsáki & Wang, 2012).

And despite appearing to have almost noise-like dynamics, these chaotic firing patterns robustly store information in their probability distributions of amplitudes (see Figure 3C and inset). And thus, the asynchronous oscillations in real networks of inhibitory neurons can be seen as information transmission that is fast and robust to noise (Brunel & Hakim, 2008; Buzsáki & Wang, 2012). Also, hidden within the chaotic region are occasional windows of stable oscillations that are seen when J is very negative (see Figure 3A). Whether pathological or strategic, this suggests that real networks of neurons may be able to flexibly switch between qualitatively different patterns of firing activity by only slightly changing their synaptic coupling (Korn & Faure, 2003).

Taken together, the general behavior of our model changes dramatically, in biologically expected ways, as its parameters are varied. These findings are summarized in Figure 4, illustrating how these behaviors change with h and J (with p^{RQ} and p^{AR} fixed at their previous, biologically plausible values). In particular, as long as p^{RQ} and p^{AR} are biologically appropriate, our model exhibits roughly three different behaviors (corresponding to the three colors in Figure 4): constant (equilibrium) activity and both excitatory and inhibitory oscillations (including chaos). In contrast, a single population of WC neurons exhibits only the former behavior (see appendix C). And analysis of the locations and properties of each of these regimes can be easily performed using only standard techniques (see appendix D). Thus, our model explains a huge variety of complex, natural phenomena

in a simple and practical way. In particular, h (i.e., the mean firing probability) can be manipulated experimentally by applying an external voltage to a group of neurons (Breskin et al., 2006; Eckmann et al., 2007). Also, synaptic activity (J) can be manipulated (Breskin et al., 2006). The predictions of our model (and even the phase diagram itself) can be easily tested experimentally.

4 Discussion

Here we have presented a new treatment of collective neural dynamics that starts from only the most elementary biophysical of neurons, and basic stochastic dynamics. We find a broad range of behaviors, even in the simplest case of only a single type of neuron (either excitatory or inhibitory).

Of course, many situations involve both types of neurons. Nevertheless, some situations involve only a single type. For example, theta-wave neuronal oscillations in the hippocampus are thought to play a considerable role in memory formation and spatial navigation (Buzsáki, 2002; Colgin, 2013). The currents driving these oscillations are believed to be primarily generated by recurrent excitatory-excitatory connections within the CA3 region of the hippocampus, whereby these neurons robustly synchronize using a "relaxation" mechanism akin to our model's predictions (Buzsáki, 2002; Buzsáki & Draguhn, 2004). Our model suggests how these neurons can easily toggle between and store the large number of complex oscillatory patterns required for their proper function (Buzsáki & Draguhn, 2004; Hutcheon & Yarom, 2000; Llinás, 1988).

Similarly, the emergence of chaotic neural dynamics has been seen experimentally and is believed to underlie high-frequency, gamma-band oscillations across the brain (Aihara, Takabe, & Toyoda, 1990; Brunel & Hakim, 2008; Korn & Faure, 2003). Our model generates these patterns with just inhibitory neurons (Buzsáki & Wang, 2012). And while chaotic dynamics might seem counterproductive for the brain, it has been theorized that these patterns are critical for information storage (Aihara et al., 1990; Brunel & Hakim, 2008; Jia et al., 2012; Korn & Faure, 2003). And perhaps fluctuations into the occasional window of stability within this chaos play a role in pathologies such as epilepsy (Sato et al., 2017).

It is also worth noting that our mean-field Max Cal, equation 2.5, resembles the model proposed in Goychuk and Goychuk (2015), but with two key differences: first, it makes far fewer assumptions and can be rigorously extended to small, noisy populations of neurons (see appendix B); and second, it defines time discretely and exhibits richer dynamics (e.g., chaos). At the same time, this link may provide a new avenue to examine the behaviors of more general neural networks (Pinto & Copelli, 2019; Goychuk & Goychuk, 2015).

Our model is also readily extended beyond a single type of neuron. In particular, WC provided the starting point for the Hopfield model of

associative learning, which has itself been an essential starting point for much of the recent development in artificial neural networks (Destexhe & Sejnowski, 2009; Hopfield, 1982; Hopfield & Tank, 1986). In that case, each population of neurons is assigned its own learned coupling J_{ij}, representing the storage of unique patterns of activity. In a similar vein, the present may allow generalization beyond the Hopfield model to include more complex, dynamical features (Destexhe & Sejnowski, 2009).

Appendix A: Maximizing Caliber for Markovian Processes

Here we summarize how to apply Max Cal to Markovian systems. The trajectories Γ of some variable S are defined as $\Gamma = \{S_0, S_1, \ldots, S_T\}$. Our goal is to infer P_Γ using some given information, or constraints. First, since the process is Markovian,

$$P_\Gamma = \pi(S_0) \prod_{t=1}^{T} P(S_t|S_{t-1}), \qquad (A.1)$$

where the vertical bar is used to denote the conditional probability. P denotes the transition probabilities, and π denotes a distribution over states. In particular, if the Markov chain is allowed to reach a steady-state distribution π,

$$\pi(S_t) = \sum_{S_{t-1}} P(S_t|S_{t-1})\pi(S_{t-1}), \quad \sum_{S_t} \pi(S_t) = 1. \qquad (A.2)$$

The path entropy $\mathcal{E}$ can then be written as

$$\mathcal{E}\{P\} = -\sum_{\Gamma} P_\Gamma \log P_\Gamma = -\sum_{S_0} \pi(S_0) \log \pi(S_0)$$

$$- \sum_{t=1}^{T} \sum_{S_t} \sum_{S_{t-1}} \pi(S_{t-1}) P(S_t|S_{t-1}) \log P(S_t|S_{t-1}), \qquad (A.3)$$

which for T large reduces to

$$\frac{1}{T}\mathcal{E}\{P\} = -\sum_{S_a} \sum_{S_b} \pi(S_a) P(S_b|S_a) \log P(S_b|S_a) \qquad (A.4)$$

for generic subsequent times a and b. We now write our caliber $\mathcal{C}$ as

$$\frac{1}{T}\mathcal{C}\{P\} = \sum_{S_a} \pi(S_a) \sum_{S_b} P(S_b|S_a)\left[-\log P(S_b|S_a) + \mu(S_a) + \sum_i \lambda_i A_i(S_a, S_b)\right].$$

$$(A.5)$$

Here $\mu(S_a)$ ensures that the transition probabilities $P(S_b|S_a)$ are properly normalized (sum to 1). Additionally the Lagrange multipliers λ_i enforce the constraints of $\langle A_i(S_a, S_b) \rangle$, such as the mean transition rates discussed in the main text. We find the trajectory distribution that maximizes the caliber $\mathcal{C}$:

$$\frac{\partial \mathcal{C}}{\partial P(S_b|S_a)} = 0 \implies -\log P(S_b|S_a) + \mu(S_a) + \sum_i \lambda_i A_i(S_a, S_b) - 1 = 0.$$

(A.6)

Therefore:

$$P(S_b|S_a) = e^{\mu(S_a) - 1 + \sum_i \lambda_i A_i(S_a, S_b)}.$$

(A.7)

Since the distributions need to be normalized, we have that

$$P(S_b|S_a) = \frac{e^{\sum_i \lambda_i A_i(S_a, S_b)}}{\sum_{S_b} e^{\sum_i \lambda_i A_i(S_a, S_b)}}.$$

(A.8)

Finally, using our original constraints, equation A.5, we can uniquely determine the Lagrange multipliers λ_i.

Appendix B: Deriving the Mean-Field Model from Max Cal

Here our goal is to understand how the constraints, equations 2.2 and 2.3, give rise to our mean-field model, equation 2.5. First, we use $S(t) = \{S_1(t), S_2(t), \ldots, S_N(t)\}$ to denote the states of all nodes at time t. Second, the number of nodes in each state are then given (respectively) by $N_Q(t)$, $N_A(t)$, and $N_R(t)$. And finally, transitions are indicated by the functions $l_i^{QA}(t)$, $l_i^{AR}(t)$, and $l_i^{RQ}(t)$ as indicated in the main text. We next follow the general procedure laid out in appendix A (see equation A.8) to infer the transition probabilities $P\left(l_i^{QA}|S\right)$, $P\left(l_i^{AR}|S\right)$, and $P\left(l_i^{RQ}|S\right)$. In particular, each quiescent (Q) node fires ($Q \rightarrow A$) with probability

$$P\left(l_i^{QA}|S\right) = \frac{e^{h^{QA} + J^* N_A}}{1 + e^{h^{QA} + J^* N_A}} = p^{QA}.$$

(B.1)

Similarly, each active (A) node becomes refractory ($A \rightarrow R$) with probability

$$P\left(l_i^{AR}|S\right) = \frac{e^{h^{AR}}}{1 + e^{h^{AR}}} = p^{AR},$$

(B.2)

and each refractory (R) node quiesces ($R \rightarrow Q$) with probability

$$P\left(l_i^{RQ}|S\right) = \frac{e^{h^{RQ}}}{1 + e^{h^{RQ}}} = p^{RQ}. \tag{B.3}$$

Equations B.1 to B.3 provide the rules by which our simple network of neurons evolves over time. However, here we are primarily interested in how the population dynamics of a group of neurons changes over time—in particular, N_Q, N_A, and N_R. For example, changes in N_Q can occur in two ways. First, nodes in R can quiesce ($R \rightarrow Q$), adding to the total number of Q nodes. Second, nodes in Q can fire ($Q \rightarrow A$), subtracting from the total number of Q nodes. Here we denote the number of each kind of transition as N_{RQ}, N_{QA}, and N_{AR}. The number of nodes of each type at time $t + 1$ is then given by

$$\Delta N_Q = N_{RQ} - N_{QA},$$
$$\Delta N_A = N_{QA} - N_{AR},$$
$$\Delta N_R = N_{AR} - N_{RQ}. \tag{B.4}$$

In reality, though, we only have two dynamical equations since

$$N_Q + N_A + N_R = N \tag{B.5}$$

for all t. Additionally, since each transition is independent, the number of transitions of each type is binomially distributed:

$$N_{RQ} \sim B(N_R, p^{RQ}),$$
$$N_{QA} \sim B(N_Q, p^{QA}),$$
$$N_{AR} \sim B(N_A, p^{AR}). \tag{B.6}$$

Here we use $B(N, p)$ as shorthand for the two-parameter binomial distribution; N is the number of trials, and p is the probability of each success (here, a transition of a particular node). We next ask how simple neuronal interactions might give rise to complex patterns of brain dynamics. In particular, we use mean-field theory to explore how our previous equations behave when the number of neurons is large (Deco et al., 2008). To simplify our analysis, we divide equation B.4 by the number of nodes N and instead follow how the average fraction of nodes in each state (π_Q, π_A, and π_R) changes over time. Since the mean of a binomially distributed random variable $B(N, p)$ is Np, the average dynamics of our model are given by

$$\Delta \pi_Q = \pi_R p^{RQ} - \pi_Q p^{QA},$$
$$\Delta \pi_A = \pi_Q p^{QA} - \pi_A p^{AR},$$
$$\Delta \pi_R = \pi_A p^{AR} - \pi_R p^{RQ}. \tag{B.7}$$

But we can eliminate contributions from π_R using equation B.5. In addition, to keep all variables in terms of the fractions π, we define $J = J^*N$ and thus arrive at our final relationships, equations 2.5 and 2.4.

Appendix C: Deriving the Wilson-Cowan Model from Max Cal

Here we show how the widely used Wilson-Cowan (WC) model emerges as a special case of our more general Max Cal model. For simplicity, we focus on only a single type of neuron, but the derivation (as well as our model) can almost trivially be extended to any number of neural types by adding couplings. Here we start from our mean-field model, equation 2.5. Suppose that the number of refractory neurons ($\pi_R = 1 - \pi_A - \pi_Q$) is in a quasi-steady state $\Delta \pi_R \approx 0$. Thus, adding together both parts of equation 2.5,

$$(1 - \pi_A - \pi_Q)p^{RQ} - \pi_A p^{AR} = 0 \implies \pi_Q = 1 - \pi_A \left(1 + \frac{p^{AR}}{p^{RQ}}\right). \tag{C.1}$$

Next, we define the constant $r = 1 + \frac{p^{AR}}{p^{RQ}}$. Substituting this back into our equation for $\Delta \pi_A$,

$$\Delta \pi_A = p^{QA}(1 - r\pi_A) - \pi_A p^{AR}. \tag{C.2}$$

Now, defining $p^{AR} = 1/\tau$ and rearranging, our equation turns into the exact same form as that from the WC equation (Chow & Karimipanah, 2019; Wilson & Cowan, 1972). While the original derivation of WC assumes that all of the variables are in a quasi-steady state (Chow & Karimipanah, 2019), our re-derivation suggests that it can be applied more generally. Nevertheless, WC still breaks down when the rate of change of recovering neurons is large. In other words, a single population of WC neurons cannot describe strongly coupled behaviors such as avalanches and intrinsic oscillations. In contrast, our Max Cal model provides a much more complete picture of neural dynamics while retaining the simplicity of the original WC model.

Appendix D: Bifurcation Analysis

Here we use local stability analysis to explore how our model transitions between simple equilibrium behavior and complex oscillatory dynamics as its

parameters are varied. To achieve this, we compute the equilibrium state of our model and ask how typical trajectories behave in its vicinity. In general, a system is in equilibrium if it does not change over time. Thus, the equilibrium states of our model are the coordinates where the left-hand side of equation 2.5 is 0. After standard algebraic manipulation, we find that the equilibrium point satisfies

$$\pi_A^* = \frac{p^{RQ} p^{QA}}{p^{RQ} p^{QA} + p^{QA} p^{AR} + p^{AR} p^{RQ}} = \frac{p^{RQ} p^{QA}}{p_D}, \quad \pi_Q^* = \pi_A^* \frac{p^{AR}}{p^{QA}}. \tag{D.1}$$

The behavior of trajectories near this point is then determined by the Jacobian matrix of derivatives, $\mathcal{J} = \frac{\partial(\pi_Q', \pi_A')}{\partial(\pi_Q, \pi_A)}$. Here we use ' to denote a subsequent time step $(t+1)$ and bold to denote matrices. For equation 2.5, the Jacobian is given by

$$\mathcal{J}(\pi_Q, \pi_A) = \begin{pmatrix} 1 - p^{RQ} - p^{QA} & -p^{RQ} - M \\ p^{QA} & 1 - p^{AR} + M) \end{pmatrix}. \tag{D.2}$$

Here, $M = \pi_Q J p^{QA}(1 - p^{QA})$. To describe the stability, we must compute the eigenvalues of this matrix evaluated at the equilibrium (π_Q^*, π_A^*). In particular, when the magnitudes of these eigenvalues (whether real or complex) are both less than one, all trajectories rapidly approach the equilibrium point (π_Q^*, π_A^*), that is, the dynamics are stable. But when one (or both) of these eigenvalues has magnitude greater than one, trajectories never reach equilibrium (the dynamics are unstable). Additionally, this transition can occur in several different ways, giving rise to the different types of oscillations we observe. In particular, excitatory oscillations occur when the real part of this eigenvalue is positive (leading to large oscillations between high π_A and high π_Q). In contrast, when the real part of this eigenvalue is negative, high-frequency, inhibitory oscillations occur. To determine when oscillatory behaviors occur, we thus need to determine when the eigenvalues of $\mathcal{J}$ change their stability. To simplify the expression of these eigenvalues, we define

$$F = \frac{p^{RQ} + p^{QA} + p^{AR} - M}{2}. \tag{D.3}$$

The eigenvalues, λ, are then given by

$$\lambda = 1 - F \pm \sqrt{F^2 - p_D + p^{RQ} M}. \tag{D.4}$$

We now have three cases to consider: when the unstable eigenvalue is one (excitatory), when it is -1 (inhibitory), and when it is complex with $|\lambda| = 1$ (either excitatory or inhibitory).

For the first case, we set $\lambda = 1$ and solve equation D.4 to find the critical point J_c,

$$J_c = -\frac{1}{\pi_A^*} W\left(\frac{-e^h}{p^{AR}\pi_A^*}\right), \tag{D.5}$$

where $W(x)$ is the multivalued Lambert W function. For the $\lambda = -1$ (inhibitory case), p^{QA} and p^{RQ} are both expected to be small. Solving equation D.4 after this approximation produces J_c:

$$J_c = \frac{p^{RQ} + p^{AR} - 2}{\pi_A^* p^{AR}} - \frac{1}{\pi_A^*} W\left(-\frac{p^{RQ} + p^{AR} - 1}{p^{AR}} e^{h + \frac{p^{RQ} + p^{AR} - 2}{p^{AR}}}\right). \tag{D.6}$$

And finally, solving the complex case exactly produces J_c:

$$\frac{p^{RQ} + p^{AR}}{(1 - p^{RQ})\pi_A^* p^{AR}} - \frac{1}{\pi_A^*} W\left(-\frac{p^{RQ} + p^{AR} + 1 - \frac{p^{RQ}}{\pi_A^*}}{(1 - p^{RQ})p^{AR}} e^{h + \frac{p^{RQ} + p^{AR}}{(1 - p^{RQ})p^{AR}}}\right). \tag{D.7}$$

Additionally, if $p^{RQ} \ll p^{QA}$ (which is almost always the case for biologically plausible sets of parameters),

$$\pi_A^* \approx \frac{p^{RQ}}{p^{RQ} + p^{AR}}. \tag{D.8}$$

Thus, the three scenarios describe three sets of critical transitions between different types of oscillations. The first ($\lambda = 1$) and last (λ complex) exclusively correspond to the emergence of excitatory oscillations. In contrast, the second case ($\lambda = -1$) corresponds to the emergence of inhibitory oscillations. Most important, the Lambert function $W(x)$ is only defined when $x \geq e^{-1}$. And when $-e^{-1} < x < 0$, the Lambert function has two solutions (corresponding to the beginning and end of oscillatory behavior). Thus, we have found an analytical relationship between our model parameters and the emergence of qualitatively distinct biological patterns.

Acknowledgments

the Stony Brook University Laufer Center for Physical and Quantitative Biology (K.D.).

Competing Interests

We declare no competing financial interests.

References

Aihara, K., Takabe, T., & Toyoda, M. (1990). Chaotic neural networks. *Physics Letters A*, *144*(6–7), 333–340.

Averbeck, B. B., Latham, P. E., & Pouget, A. (2006). Neural correlations, population coding and computation. *Nature Reviews Neuroscience*, *7*(5), 358–366.

Bacak, B. J., Kim, T., Smith, J. C., Rubin, J. E., & Rybak, I. A. (2016). Mixed-mode oscillations and population bursting in the pre-Bötzinger complex. *eLife*, *5*, e13403.

Benayoun, M., Cowan, J. D., van Drongelen, W., & Wallace, E. (2010). Avalanches in a stochastic model of spiking neurons. *PLOS Computational Biology*, *6*(7), e1000846.

Breskin, I., Soriano, J., Moses, E., & Tlusty, T. (2006). Percolation in living neural networks. *Physical Review Letters*, *97*(18), 188102.

Brunel, N., & Hakim, V. (2008). Sparsely synchronized neuronal oscillations. *Chaos*, *18*(1), 015113.

Buzsáki, G. (2002). Theta oscillations in the hippocampus. *Neuron*, *33*(3), 325–340.

Buzsáki, G., & Draguhn, A. (2004). Neuronal oscillations in cortical networks. *Science*, *304*(5679), 1926–1929.

Buzsáki, G., & Wang, X.-J. (2012). Mechanisms of gamma oscillations. *Annual Review of Neuroscience*, *35*, 203–225.

Chow, C. C., & Karimipanah, Y. (2019). *Before and beyond the WilsonCowan equations.* arXiv:1907.07821.

Colgin, L. L. (2013). Mechanisms and functions of theta rhythms. *Annual Review of Neuroscience*, *36*, 295–312.

Deco, G., Jirsa, V. K., Robinson, P. A., Breakspear, M., & Friston, K. (2008). The dynamic brain: From spiking neurons to neural masses and cortical fields. *PLOS Computational Biology*, *4*(8).

Destexhe, A., & Sejnowski, T. J. (2009). The Wilson–Cowan model, 36 years later. *Biological Cybernetics*, *101*(1), 1–2.

Dixit, P. D., Wagoner, J., Weistuch, C., Pressé, S., Ghosh, K., & Dill, K. A. (2018). Perspective: Maximum caliber is a general variational principle for dynamical systems. *Journal of Chemical Physics*, *148*(1), 010901.

Eckmann, J.-P., Feinerman, O., Gruendlinger, L., Moses, E., Soriano, J., & Tlusty, T. (2007). The physics of living neural networks. *Physics Reports*, *449*(1–3), 54–76.

Feldman, J., & Cowan, J. (1975). Large-scale activity in neural nets I: Theory with application to motoneuron pool responses. *Biological Cybernetics*, *17*(1), 29–38.

Gerstner, W. (2000). Population dynamics of spiking neurons: Fast transients, asynchronous states, and locking. *Neural Computation*, *12*(1), 43–89.

Ghosh, K., Dixit, P. D., Agozzino, L., & Dill, K. A. (2020). The maximum caliber variational principle for nonequilibria. *Annual Review of Physical Chemistry*, *71*, 213–238.

Goles-Chacc, E., Fogelman-Soulié, F., & Pellegrin, D. (1985). Decreasing energy functions as a tool for studying threshold networks. *Discrete Applied Mathematics*, *12*(3), 261–277.

Goychuk, I., & Goychuk, A. (2015). Stochastic Wilson–Cowan models of neuronal network dynamics with memory and delay. *New Journal of Physics*, *17*(4), 045029.

Greenberg, J. M., & Hastings, S. (1978). Spatial patterns for discrete models of diffusion in excitable media. *SIAM Journal on Applied Mathematics*, *34*(3), 515–523.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. In *Proceedings of the National Academy of Sciences*, *79*(8), 2554–2558.

Hopfield, J. J., & Tank, D. W. (1986). Computing with neural circuits: A model. *Science*, *233*(4764), 625–633.

Hutcheon, B., & Yarom, Y. (2000). Resonance, oscillation and the intrinsic frequency preferences of neurons. *Trends in Neurosciences*, *23*(5), 216–222.

Jia, B., Gu, H., Li, L., & Zhao, X. (2012). Dynamics of period-doubling bifurcation to chaos in the spontaneous neural firing patterns. *Cognitive Neurodynamics*, *6*(1), 89–106.

Jirsa, V. K., & Haken, H. (1997). A derivation of a macroscopic field theory of the brain from the quasi-microscopic neural dynamics. *Physica D: Nonlinear Phenomena*, *99*(4), 503–526.

Keeley, S., Fenton, A. A., & Rinzel, J. (2017). Modeling fast and slow gamma oscillations with interneurons of different subtype. *Journal of Neurophysiology*, *117*(3), 950–965.

Korn, H., & Faure, P. (2003). Is there chaos in the brain? II. Experimental evidence and related models. *Comptes Rendus Biologies*, *326*(9), 787–840.

Llinás, R. R. (1988). The intrinsic electrophysiological properties of mammalian neurons: Insights into central nervous system function. *Science*, *242*(4886), 1654–1664.

Muir, A. (1979). A simple case of the Wilson-Cowan equations. *Biophysical Journal*, *27*(2), 267.

Neves, L. L., & Monteiro, L. H. A. (2016). A linear analysis of coupled Wilson-Cowan neuronal populations. *Computational Intelligence and Neuroscience 2016*, art. 8939218.

Omurtag, A., Knight, B. W., & Sirovich, L. (2000). On the simulation of large populations of neurons. *Journal of Computational Neuroscience*, *8*(1), 51–63.

Panzeri, S., Macke, J. H., Gross, J., & Kayser, C. (2015). Neural population coding: Combining insights from microscopic and mass signals. *Trends in Cognitive Sciences*, *19*(3), 162–172.

Pinto, I. L. D., & Copelli, M. (2019). Oscillations and collective excitability in a model of stochastic neurons under excitatory and inhibitory coupling. *Physical Review E*, *100*(6), 062416.

Pressé, S., Ghosh, K., Lee, J., & Dill, K. A. (2013). Principles of maximum entropy and maximum caliber in statistical physics. *Review of Modern Physics*, *85*(3), 1115.

Rajan, K., & Abbott, L. F. (2006). Eigenvalue spectra of random matrices for neural networks. *Physical Review Letters*, *97*(18), 188104.

Sanchez-Vives, M. V., & McCormick, D. A. (2000). Cellular and network mechanisms of rhythmic recurrent activity in neocortex. *Nature Neuroscience*, *3*(10), 1027–1034.

Sato, Y., Wong, S. M., Iimura, Y., Ochi, A., Doesburg, S. M., & Otsubo, H. (2017). Spatiotemporal changes in regularity of gamma oscillations contribute to focal ictogenesis. *Scientific Reports*, *7*(1), 1–9.

Schneidman, E., Berry, M. J., Segev, R., & Bialek, W. (2006). Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, *440*(7087), 1007–1012.

Shusterman, V., & Troy, W. C. (2008). From baseline to epileptiform activity: A path to synchronized rhythmicity in large-scale neural networks. *Physical Review E*, *77*(6), 061911.

Tkacik, G., Schneidman, E., Berry, I., Michael, J., & Bialek, W. (2006). *Ising models for networks of real neurons.* arXiv:0611072.

Vladimirski, B. B., Tabak, J., O'Donovan, M. J., & Rinzel, J. (2008). Episodic activity in a heterogeneous excitatory network, from spiking neurons to mean field. *Journal of Computational Neuroscience*, *25*(1), 39–63.

Weistuch, C., Agozzino, L., Mujica-Parodi, L. R., & Dill, K. A. (2020). *Inferring a network from dynamical signals at its nodes*. arXiv:2004.02318.

Weistuch, C., Mujica-Parodi, L. R., Amgalan, A., & Dill, K. A. (2020). *Younger, but not older, brains are poised at a critical point of functional activity*. bioRxiv.

Wilson, H. R., & Cowan, J. D. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical Journal*, *12*(1), 1–24.

Wilson, H. R., & Cowan, J. D. (1973). A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. *Kybernetik*, *13*(2), 55–80.