ARTICLE **Communicated by Peter Thomas**

A Transition to Sharp Timing in Stochastic Leaky Integrate-and-Fire Neurons Driven by Frozen Noisy Input

Thibaud Taillefumier
ttaillef@princeton.edu
Laboratory of Mathematical Physics, Rockefeller University, New York, NY 10065, U.S.A., and Lewis-Sigler Institute for Integrative Genomics, Princeton University, Princeton, NJ 08544, U.S.A.

Marcelo Magnasco
mgnsclb@rockefeller.edu
Laboratory of Mathematical Physics, Rockefeller University, New York, NY 10065, U.S.A.

The firing activity of intracellularly stimulated neurons in cortical slices has been demonstrated to be profoundly affected by the temporal structure of the injected current (Mainen & Sejnowski, 1995). This suggests that the timing features of the neural response may be controlled as much by its own biophysical characteristics as by how a neuron is wired within a circuit. Modeling studies have shown that the interplay between internal noise and the fluctuations of the driving input controls the reliability and the precision of neuronal spiking (Cecchi et al., 2000; Tiesinga, 2002; Fellous, Rudolph, Destexhe, & Sejnowski, 2003). In order to investigate this interplay, we focus on the stochastic leaky integrate-and-fire neuron and identify the Hölder exponent H of the integrated input as the key mathematical property dictating the regime of firing of a single-unit neuron. We have recently provided numerical evidence (Taillefumier & Magnasco, 2013) for the existence of a phase transition when H becomes less than the statistical Hölder exponent associated with internal gaussian white noise ($H = 1/2$). Here we describe the theoretical and numerical framework devised for the study of a neuron that is periodically driven by frozen noisy inputs with exponent $H > 0$. In doing so, we account for the existence of a transition between two regimes of firing when $H = 1/2$, and we show that spiking times have a continuous density when the Hölder exponent satisfies $H > 1/2$. The transition at $H = 1/2$ formally separates rate codes, for which the neural firing probability varies smoothly, from temporal codes, for which the neuron fires at sharply defined times regardless of the intensity of internal noise.

Color versions of select figures in this article are presented in the online supplement, available at http://www.mitpressjournals.org/doi/suppl/10.1162/NECO_a_00577.

Neural Computation **26, 819–859** (2014)
doi:10.1162/NECO_a_00577

1 Introduction

The dynamics of spike generation within an assembly of neurons underlies neural coding, which is defined as the mapping of behaviorally relevant sensory information onto spatiotemporal spiking patterns (Rieke, Warland, de Ruyter van Steveninck, & Bialek, 1999; Dayan & Abbott, 2001). Accounting for the observed variability of neural responses to identical stimuli, neural noise plays a crucial role in shaping the propagation of information through neural circuits (Faisal, Selen, & Wolpert, 2008). Conceptually, such noise is internal or external (Bressloff, 2010). Internal noise translates the inherent stochasticity of the molecular mechanisms that underpin the electrical activity of neurons and agglomerate the effects of thermal noise and stochastic genetic expression. Sources of internal noise noticeably include random ion channel gating events (White, Rubinstein, & Kay, 2000; Dorval & White, 2005) and random synaptic failure (Dobrunz & Stevens, 1997). External noise is the by-product of the spontaneous activity of surrounding neurons and summarizes the constant bombardment that a neuron undergoes when embedded in an assembly (Fellous, Rudolph, Destexhe, & Sejnowski, 2003). Thus, external noise can be thought of as an activity-dependent perturbation that superimposes itself onto the internal noise, and to first approximation, both noises appear as a nuisance to the faithful transmission of neural information (London, Roth, Beeren, Hausser, & Latham, 2010).

At the single-unit level, the complex interplay between internal and external noise is best exemplified by considering the experiment of Mainen and Sejnowski (1995), which complemented earlier results by Bryant and Segundo (1976). In Mainen and Sejnowski's (1995) study the spiking activity of a neuron is recorded while it is repeatedly injected with a steady current. In response to each stimulation, the neuron emits a train of spikes that tends to start with a fixed delay from the onset of the injection and stabilizes at a fixed given firing rate. Over time, the spiking times gradually desynchronize, as if random perturbations were added independently to each interspike interval. Thus, internal noise appears to be a nuisance resulting from the fluctuation of the cellular machinery, ultimately limiting the ability of spike timing to convey information.

However, if the cell is injected with the same steady input perturbed by a succession of many transient pulses, the neuron no longer fires this coarsely. Instead, it repeatedly produces precise patterns of firing that are closely locked to the time-varying input. In other words, what appears as transient random external fluctuations leads to a reliable precise neural response. To emphasize that, despite being seemingly random, the input is actually a controlled input that drives the neuron's activity; it is referred to as frozen noise (Haas & White, 2002). It is important to realize that due to the stochasticity of neural populations, the overall synaptic input that a neuron integrates is very much akin to such frozen-noise traces.

Mainen and Sejnowski's (1995) experiment shows that the firing pattern of a single unit differs profoundly according to the nature of its input. In particular, it demonstrates that the concept of reliability and precision in the neural response cannot be considered intrinsic to a neuron but has to be seen as message dependent. We recall that the notion of reliability measures the reproducibility of a firing pattern in response to stereotypical stimulations, whereas the notion of spiking precision quantifies the temporal variability of spiking events that encode for equivalent stimuli features (Bair & Koch, 1996; Berry, Warland, & Meister, 1997). It has been shown that spiking precision is a function of the input for virtually all spike-generating mechanisms (Cecchi et al., 2000), even if the reliability of the neural response is left unchanged across trials (Reich, Victor, Knight, Ozaki, & Kaplan, 1997).

It is in the context of this complex message dependence that noise at the population level appears more than a mere nuisance. In principle, different regimes of synchrony within cortical circuits can yield seemingly random, yet precise, spiking patterns (Butts et al., 2007; Stevens & Zador, 1998). If the statistics of these patterns is such that they integrate into a highly fluctuating input, neural integration at the single unit level can propagate temporal precision in the face of internal noise and thus promote temporal coding. This point is consistent with the view that noise can cause synchrony, provided that the perturbations fed in different neurons are correlated (Rosenblum, Pikovsky, & Kurths, 1996; Galán, Fourcaud-Trocmé, Ermentrout, & Urban, 2006).

Envisioned as constructive, the role of noise has been termed *stochastic facilitation* (McDonnell & Ward, 2011). In this context, assuming that noise is relevant to information processing, the experimental perturbation of neural circuits by injection of controlled noise promises to reveal their noise tuning, which will shed light on their encoding strategies at the population level. However, as Poggio and Marr (1977) pointed out, such an inquiry requires a prior solid computational hypothesis, thus identifying the need for a neural stochastic theory that can accommodate controlled perturbation by fluctuating input. Here, we set up a theoretical and numerical framework that will allow us to characterize the behavior of a well-known neuron model in response to the injection of frozen noise input, with the explicit inclusion of internal noise.

1.1. Stochastic Leaky Integrate-and-Fire Model and First-Passage Time. To study the role of noise in shaping neural activity, we give a sample path description of a neuron's stochastic dynamics that preserves the essential mechanistic principle of neural integration. Specifically, we consider a stochastic integrate-and-fire neuron whose membrane potential is modeled as a continuous random process X. Subjected to internal noise, the fluctuating membrane potential X obeys a stochastic differential equation

of the form

$$dX_t = F(X_t, I(t), t)\,dt + \sigma(X_t, I(t), t)dW_t. \tag{1.1}$$

In this equation, dW denotes gaussian white noise as the formal time derivative of the canonical Wiener process W. The response function F drives the deterministic dynamics of the membrane potential, while the diffusion coefficient σ quantifies the internal noise perturbations. The functions F and σ define the integration rules of a neuron through their dependence on the time-varying input I. From a mathematical standpoint, equation 1.1 describes the subthreshold dynamics of the membrane potential as a Markov diffusion process (Stroock & Varadhan, 2006).

We consider that a neuron fires an action potential whenever the noisy integration of the input causes the membrane potential X to exceed a certain voltage threshold l. Thus, spiking events are formally defined as first-passage times,

$$\tau = \inf\left\{t \mid X_t > l, X_0 = r\right\}, \tag{1.2}$$

where we have assumed that the neuron resets to value r after the previous spiking event at time $t = 0$. Although it is simply stated, this spike-generating mechanism is more complex than it seems. Specifically, defined as first-passage times (Redner, 2007), spike occurrences are isolated discrete events that depend on the continuous history of the fluctuating input.

Hereafter, we adopt the stochastic leaky integrate-and-fire (sLIF) model (Ricciardi & Sato, 1988; Lansky & Ditlevsen, 2008; Ricciardi & Sacerdote, 1979; Burkitt, 2006a, 2006b) as the fundamental encoding scheme that transforms neural input into a train of spikes. For simplicity, we restrain our study to the current-based model, where the neuronal input takes the form of an injected current. Thus the subthreshold membrane potential is modeled as an Ornstein-Uhlenbeck (OU) process (Uhlenbeck & Ornstein, 1930) satisfying equation 1.1 with $F(X_t, I(t), t) = -\alpha X_t + I(t)$ and $\sigma(X_t, I(t), t) = \sigma$, where the leak term α is the inverse of the relaxation time of the neuron. The nonlinearity of the model lies entirely in the membrane reset rule that is implemented after reaching the threshold l and triggering a spiking event.

1.2. Hölder Regularity and Phase Transition. In order to motivate our adoption of the sLIF model, We proceed to a thought experiment that idealizes the experiment of Mainen and Sejnowski (1995). This will allow us to gain some insight into the message dependence of neural reliability and precision. We consider in Figure 1 the response of an sLIF neuron to the repeated injection of a steady current with or without a frozen white noise of the same intensity as the internal noise. As opposed to the case of a constant

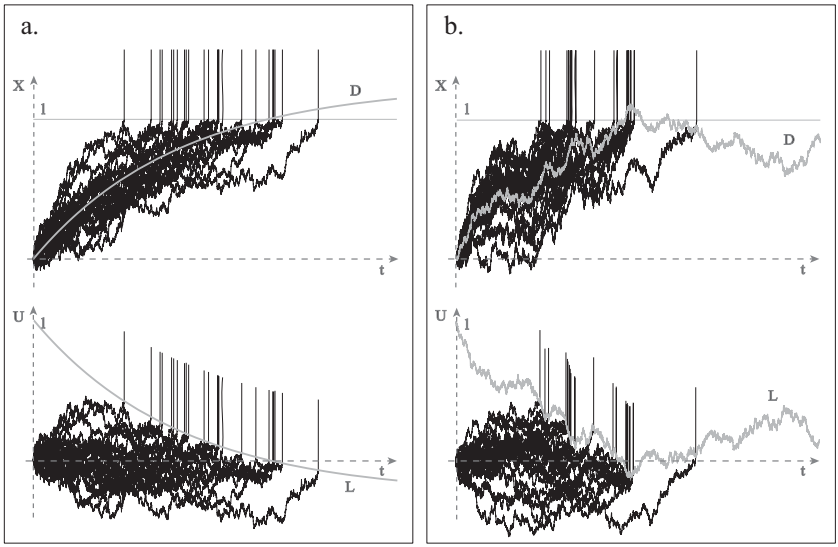

Figure 1: (a) sLIF neuron injected with a steady current. Top: Each black trace represents a membrane potential's trajectory; the gray curve is the drift resulting from the integration of the input current. Bottom: The drift is subtracted from the voltage traces to reveal the canonical OU sample paths. (b) sLIF neuron injected with a steady current and a frozen-noise current. Top: The drift is a frozen Orstein-Uhlenbeck path added to the original exponential rise; spiking times tend to cluster in specific time regions. Bottom: Subtracting the drift shows that the contribution of the frozen noise can be encapsulated in terms of a fluctuating effective barrier.

input, it appears that spikes tend to cluster in well-defined time regions in the presence of frozen noise, with enhanced temporal precision.

For a moment, let us admit that the problem can be formulated in term of an effective barrier, defined as the constant spiking threshold minus the drift resulting from the integration of the input (see Figure 1). It then becomes clear that reliable spiking occurs when the effective barrier intercepts many trajectories that have not yet hit the boundary, whereas precise spiking occurs when the effective barrier falls steeply in a transient fashion. In other words, the spiking regime of the sLIF neuron is dictated by the local regularity of the effective barrier that results from the integration of the frozen noise.

Mathematically, we identify the feature of the effective barrier that controls the firing activity of our externally driven neuron as its Hölder regularity. If we hold ourselves to frozen noise, the degree of regularity of the effective barrier is statistically summarized by the Hölder exponent. This exponent denoted H, also called the roughness exponent, takes values in

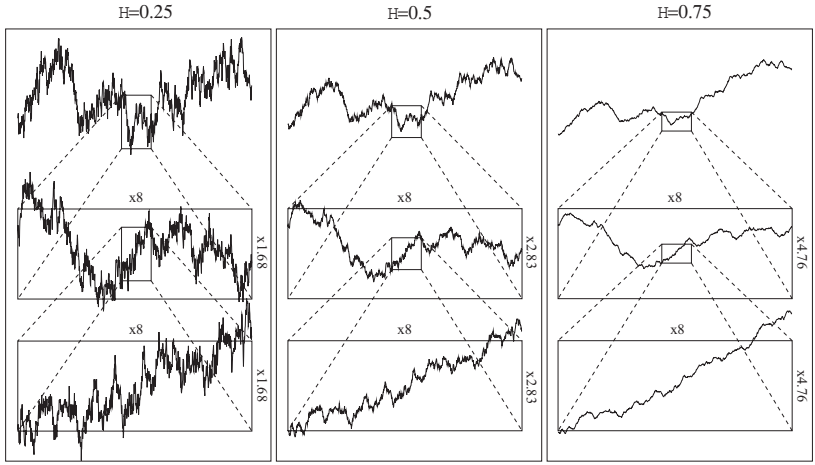

Figure 2: Link between scale invariance and Hölder regularity. We represent the asymptotic behavior of the curve L^H in the limit of vanishing timescale. In the case of H-continuous functions, the Hölder exponent indicates the local scaling exponent of the graph of L^H: if a dilation in time d is accompanied by a dilation in voltage d^H, the statistics of the voltage trace is left unchanged.

$(0, 1)$ and quantifies the strength of the input fluctuation through vanishing scales.[1] For an OU process U, which is asymptotically identical to a Wiener process, it takes the value $H = 1/2$ since the fluctuation scales as $\delta U \sim \sqrt{\delta t}$. Postponing a formal definition of H, it is possible to intuitively capture the meaning of the Hölder exponent for stationary curves through the notion of statistical scale invariance, as shown in Figure 2. Concretely, the Hölder exponent characterizes how infinitesimal fluctuations add up constructively or destructively to generate complex curves, whose statistical profile is typical for a given value H

Consequently, we design our theoretical and numerical framework from the perspective of studying the distribution of spike timing that arises from the injection of frozen noise with different Hölder regularity. In that respect, although the Hölder regularity is a property of the effective barrier, for conciseness, we will call H the Hölder exponent of the frozen noise and will say that the noise is Hölder singular. In a related paper (Taillefumier & Magnasco, 2013), we demonstrate the existence of two regimes of spiking for sLIF neurons that have constant mean firing rate. On one hand, if the frozen injected noise is less singular than the internal neuronal noise ($H > 1/2$), the temporal precision of the driven neuron is classically prescribed by the

[1] For $H = 0$, the curve becomes discontinuous, whereas for $H = 1$, the curve becomes piecewise differentiable.

continuous probability density of spiking times (continuous spiking mode). On the other hand, if the frozen noise is more singular than the internal noise ($H < 1/2$), the neuron exhibits exquisite temporal precision in the sense that the probability density becomes singular, almost everywhere either zero or infinity (singular spiking mode). This observation appears to be generalizable to all current-based integrate-and-fire neurons and becomes the signature of the competition between drift and diffusive components of the dynamics, which shapes the voltage sample path in the vicinity of the threshold. However, our interest here is to set up the theoretical and numerical framework required for the study of this phenomenon through the following case study.

1.3. Study of the Firing Rate of a Periodically Driven Neuron. In the spirit of Mainen and Sejnowski (1995) and following previous work (Fellous et al. 2001; Thomas, Tiesinga, Fellous, & Sejnowski, 2003), we study in depth the sLIF encoding scheme when periodically driven by frozen noisy input. First, under mild hypotheses on the nature of the input, we show that spikes happen as generated through an ergodic Markov chain (Häggström, 2002; Norris, 1998; Stewart, 2009), whose Markov transition kernels are naturally deduced from first-passage distributions. Next, we construct a family of frozen noise currents I_H with Hölder singularity prescribed by a Hölder exponent H satisfying $0 < H < 1$. This family of injected currents is actually constructed by gradually altering a frozen gaussian white noise ($H = 1/2$) that has the same amplitude as the internal noise of the neuron. In dealing with these highly irregular inputs I_H, we resort to Monte Carlo simulations (Metropolis & Ulam, 1949; Robert & Casella, 2004) to infer the instantaneous firing rates by generating extremely long histories of spiking times, with exquisite precision. Bearing in mind that these results are discussed at length in a related paper (Taillefumier & Magnasco, 2013), we then briefly report the effect of varying the noise singularity on the qualitative nature of the distribution of spiking times. Finally, we investigate when the spiking time admits a density function, a property that is readily assumed in most applications but can very well fail to hold true.

2. First-Passage Markov Chain

When an sLIF neuron is driven by an input, it produces a train of spikes that comes under a statistical flavor. If the input is periodic, the train of spikes can be divided such that the spiking events are lined up according to their phases within an input cycle. This procedure generates the raster plot of the neuron's response to the input. Given a binning interval, computing the fraction of time that a spike occurs in each bin yields the normalized peristimulus histogram. Loosely speaking, a peak in the peristimulus histogram indicates that the neuron spikes reliably at a given phase; the sharpness of this peak measures the precision of the spike timing.

To study the relation between driving input and peristimulus histogram, we formalize the mechanism by which an sLIF neuron generates spiking phases. For such a neuron, the train of spikes elicited by a periodic input results from a succession of first-passage problems to an effective barrier. The phases of the spiking events, as opposed to their absolute times, constitute the sample path of a discrete inhomogeneous Markov chain. Then for asymptotically large numbers of input cycles, the peristimulus histogram is defined as the fraction of time the Markov chain spends in the bin phase intervals.

From a mathematical standpoint, the peristimulus histogram that arises from a cyclic input is well defined only if the corresponding Markov chain is ergodic. Indeed, the property of ergodicity ensures that the fraction of time the Markov chain spends in any bin interval is independent of the initial condition of the Markov chain. From a numerical standpoint, ergodicity is also a desired property for the approximate simulation of the Markov chain and, thus, for the estimation of the peristimulus histogram. For a sLIF neuron with smooth input, verifying the ergodic property is straightforward. However, such a verification is not obvious for highly fluctuating input, such as frozen noise inputs.

In this section, we first show that the neural activity of a driven sLIF neuron is advantageously formulated in terms of an effective barrier. Based on this formulation, we define the generation of spiking phases by a driven sLIF neuron as a Markov chain, referred to as the first-passage Markov chain (FPMC). We then establish the ergodic property of the FPMC for effective barriers with a positive Hölder exponent.

2.1. Effective Barrier Formulation. Before defining the FPMC, we show that the spiking activity of a sLIF neuron driven by an external input can be formulated in term of first-passage problems of standard OU processes to a fluctuating barrier. This formulation allows us to clearly separate the contribution of external noise, which is entirely captured by the barrier profile, from the contribution of internal noise, which drives the fluctuations of the OU processes.

In the sLIF model, the subthreshold dynamics of the membrane potential X is described by the inhomogeneous linear stochastic differential equation,

$$dX_t = -\alpha X_t\, dt + \sigma\, dW_t + dC(t), \quad \alpha > 0. \tag{2.1}$$

At the cost of rescaling X and C by σ, we restrain ourselves to the case $\sigma = 1$. In equation 2.1, we write the increment due to the input current $I(t)$ as the infinitesimal variation $I(t)\,dt = dC(t)$ of a load function $C(t)$. We characterize the regularity of the load function by its Hölder exponent H.

When $H > 0$, we say that C is H-continuous. Formally, the Hölder exponent H is defined in $(0, 1)$ as the infinimum of the local Hölder exponents H_t:

$$H = \inf_t H_t < \infty, \quad H_t = \lim_{\delta \to 0+} \sup_{|t-s| \le \delta} \frac{|C(t) - C(s)|}{|t - s|^h}. \tag{2.2}$$

The exponent $H(t)$ quantifies the scaling of the local fluctuations $\Delta_\delta C(t) = C(t + \delta) - C(t)$ with respect to the observation timescale δ. The smaller the Hölder exponent, the larger the fluctuations $\Delta_\delta C$ in the limit of vanishing timescale δ, which justifies considering H as a proxy measure for the degree of singularity of I. As frozen noise, I is generally not bounded in any time intervals and can be represented only as a generalized function. We consider only frozen noise input I whose time integral defines an H-continuous load function C with $0 < H < 1$. In particular, C is nowhere differentiable, and we say that the corresponding frozen noise input I, as a generalized function, is H-singular. Notice that by contrast, if the input current I remains bounded, the corresponding load function C is H-continuous with $H = 1$.

Because the nonlinearity of the sLIF model lies entirely in the spike generation and subsequent reset, we can separately integrate input and noise between spikes. Thus we can transform the first-passage problem, equation 1.2, with constant threshold l and fluctuating input I into a first-passage problem without driving forces to a fluctuating effective barrier. To see this, suppose that when a sLIF neuron is injected by an input current I, it emits a spike labeled by i at time $t_i > 0$. After reaching the spiking threshold l, the neuron's membrane potential X resets to a value $r < l$. To solve equation 2.1 with initial condition $X_{t_i}^+ = r$, we write $X = U^i + l^i$, where we separate the stochastic part U^i, the OU process obtained for $I = 0$, and the deterministic part l^i arising from the integration of the input $I(t) = dC(t)/dt$:

$$U_t^i = r\, e^{-\alpha(t-t_i)} + \int_{t_i}^t e^{-\alpha(t-s)}\, dW_s, \tag{2.3}$$

$$l^i(t) = \int_{t_i}^t e^{-\alpha(t-s)}\, dC(s). \tag{2.4}$$

Observe that the decomposition of $X = U^i + l^i$ is in agreement with the reset rule $X_{t_i^+} = r$ since we have $U_{t_i^+}^i = r$ and $l^i(t_i^+) = 0$. To determine the next spiking time, t_{i+1} can be cast in terms of a first-passage problem for the process U^i with the effective barrier $t \mapsto L^i(t) = l - l^i(t)$:

$$\tau_{i+1} = \inf\{t > t_i \,|\, U_t^i > L^i(t), U_{t_i^+}^i = r\}. \tag{2.5}$$

Therefore, a train of spikes $t_0 < t_1 < \cdots < t_n$ is determined by solving consecutively the first-passage problems, equation 2.5, for independent processes U^i. Note that by definition, the barriers L^i depend on the spike

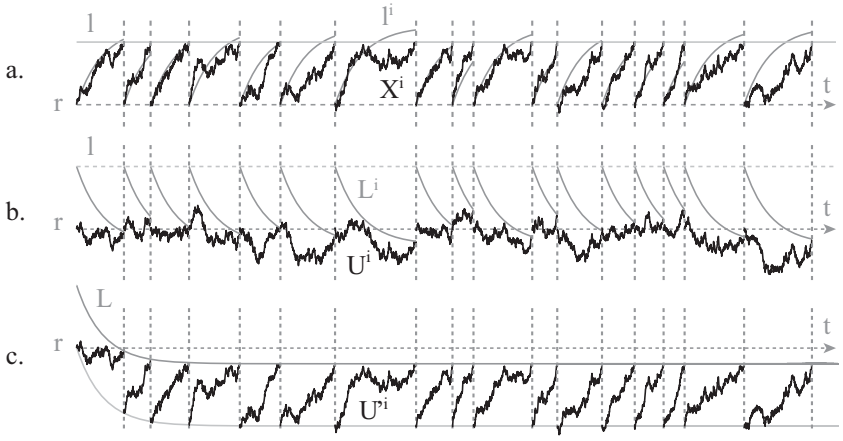

Figure 3: Sequence of first-passage times for a sLIF neuron with constant injected current. (a) In the direct representation, each neuron spikes when its membrane voltage, modeled as an upward-drifted OU process, hits the firing threshold. (b) In the effective representation with a naive reset rule, the deterministic drift is subtracted from the same voltage traces so that the sketched traces are canonical OU trajectories. (c) In the effective representation with a modified reset rule, the effective barrier becomes continuous as a convolution of the injected current. When we simulate a sLIF neuron that is cyclically injected with a current I comprising a frozen noise component, the effective barrier will have a fluctuating profile. Moreover, we will consider the sLIF in steady state, as if many input cycles have already occurred; this implies that the effect of the steady component of the input will be to lower the threshold by a constant value as opposed to the exponential decay observed in panel c.

timing t_i and that due to the reset rule, the barriers L^i may not agree at spiking times: $L^{i-1}(t_i^-) \neq L^i(t_i^+) = l$.

To completely separate the role of noise and the role of input in a sLIF neuron, we must eliminate the dependence of the effective barrier L^i on the spiking time t_i. Fortunately, the linearity of the dynamics, equation 2.1, allows us to recast the successive first-passage problems, equation 2.5, in terms of a sequence of first-passage problems for a single continuous barrier $L = L^0 = l - l^0$ (see Figure 3 for a graphical example). To justify this claim, let us consider the set S_i of white noise realizations ω for which the trajectories $t \mapsto U_t^i(\omega)$ originate from $U_{t_i^+}^i = r$ and first cross L^i after time t:

$$\begin{aligned} S_i &= \left\{ \omega \,\middle|\, \tau_{i+1}(\omega) > t, \tau_i(\omega) = t_i \right\} \\ &= \left\{ \omega \,\middle|\, U_{t_i^+}^i(\omega) = r,\ U_s^i(\omega) < L^i(s), t_i < s < t \right\}. \end{aligned} \tag{2.6}$$

Remembering that $L^i = l - l^i$ for all $i > 0$, we have for $s > t_i$,

$$\left\{\omega \,\middle|\, U_s^i(\omega) < L^i(s)\right\} = \left\{\omega \,\middle|\, U_s^i(\omega) + l^i(s) - l^{i-1}(s) < l - l^{i-1}(s)\right\} \tag{2.7}$$

$$= \left\{\omega \,\middle|\, U_s'^i(\omega) < L^{i-1}(s)\right\}, \tag{2.8}$$

where $U_t'^i = U_t^i + l^i(t) - l^{i-1}(t)$ can be written for $t > t_i$:

$$U_t'^i = e^{-\alpha(t-t_i)}\left(r - \int_{t_{i-1}}^{t_i} e^{-\alpha(t_i-s)} dC(s)\right) + \int_{t_i}^{t} e^{-\alpha(t-s)}\, dW_s. \tag{2.9}$$

In the above expression, we recognize $U_t'^i$ as the solution of equation 2.1 for $dC = 0$, with the new initial condition:

$$U_{t_i^+}'^i = r - \int_{t_{i-1}}^{t_i} e^{-\alpha(t_i-s)} dC(s) = L^{i-1}(t_i) - (l - r). \tag{2.10}$$

Thus, from equation 2.6, the set S_i can be redefined as

$$S_i = \left\{\omega \,\middle|\, U_{t_i^+}'^i(\omega) = L^{i-1}(t_i) - (l - r),\ U_s'^i(\omega) < L^{i-1}(s), t_i < s < t\right\}, \tag{2.11}$$

$$= \left\{\omega \,\middle|\, \tau_{i+1}'(\omega) > t, \tau_i(\omega) = t_i\right\}, \tag{2.12}$$

where

$$\tau_{i+1}' = \inf\{t > t_i \,|\, U_t'^i > L^{i-1}(t), U_{t_i^+}'^i = L^{i-1}(t_i) - (l - r)\}. \tag{2.13}$$

The equality, equation 2.12, implies that, conditional to τ_i, the next spiking time is determined as $\tau_{i+1}' + \tau_i$, where τ_{i+1}' is the first-passage time of a standard OU process U'^i to the effective barrier L^{i-1}, given that $U_{t_i^+}'^i = L^{i-1}(t_i) - (l - r)$. Notice that the effective barrier L^{i-1} intervening in the determination of τ_{i+1}' is now independent of t_i. We can iterate our reasoning for every first-passage problem τ_j given τ_{j-1} with $j \le i$. As a result, given an initial spiking time $t_0 = 0$, a train of spikes $0 < t_1 < \cdots < t_n$ appears as a realization of the sequence of first-passage problems,

$$\tau_{i+1}' = \inf\{t > \tau_i \,|\, U_t'^i > L(t), U_{\tau_i^+}'^i = L(\tau_i) - (l - r)\}, \tag{2.14}$$

where U'^i are standard independent OU processes satisfying the prescribed initial conditions.

By implementing the new reset rules $U'^{i}_{t_i^+} = L(t_i) - (l - r)$, we are able to describe the spiking activity of a driven sLIF neuron in terms of first-passage problems to a single continuous effective barrier L. In this picture, the barrier L fully integrates the input, while the internal noise drives the standard OU processes U'_i . We will adopt this formulation in terms of the effective barrier L to numerically simulate a sLIF neuron that is repeatedly injected with the same current. When we gloss over the treatment of the neuron's initial condition, the introduction of such a formulation greatly simplifies the simulation of consecutive first-passage times. Instead of integrating the input after each reset t_i to form the barriers L^i, we just have to compute the barrier L once and for all, and simulate first passages with the OU processes U'_i and the new reset rule $U'^{i}_{t_i^+} = L(t_i) - (l - r)$.

2.2. Continuous-Time Inhomogeneous Markov Chain. Thanks to the effective-barrier formulation, we can formalize the spike-generating mechanism of a sLIF neuron in the framework of Markov chains. Thus, we consider periodic driving inputs and distinguish spiking events based on their phases, that is, the fraction of the input cycle that has elapsed at the time of the spike emission. The theory of Markov chains is the natural setting to define the notion of peristimulus histograms as a probability measure, which in turn can be approximated by a simple numerical scheme. Because of the fundamental part played by the first-passage problem, we refer to the Markov chain generating spiking phases of sLIF neurons as the first-passage Markov chain (FPMC). Unfortunately, the singular nature of the frozen noise inputs requires to define the FPMC in a rather formal fashion.

Before specifying the FPMC, let us recall that in a typical experiment, the spiking activity of a neuron is recorded in response to repeated presentations of the same stimulus. We idealize this situation by studying the distribution of spiking events when an input, possibly with a frozen noise component, cyclically forces a sLIF neuron with period T. For simplicity, we take the period T as our unit of time, so that $T = 1$. In this way, the phase of a spike occurring at time t is simply given by $\phi = \pi(t) = t - \lfloor t \rfloor$, where $\lfloor t \rfloor$ is the integer part of t. Mathematically speaking, the phase map π identifies the phase space $[0, 1)$ with the circle $\mathbb{S} = \mathbb{R}/\mathbb{Z}$. To avoid discontinuity effects, we consider an input for which the effective barrier already discussed satisfies $L(1) = L(0)$. Reasoning in steady state, this supposes we choose an input current $I(t) = I_c + I_n(t)$ with a constant component I_c and a noisy component $I_n(t)$, giving rise to a load function $C(t) = C_c + C_n(t)$ with

$$\int_0^1 e^{-\alpha(1-s)}\, dC_n(s) = L(1) - L(0) = 0. \tag{2.15}$$

Such currents are easily obtained. Given an arbitrary integrable input I with associated load function C, the new input current $t \mapsto I'(t) = I(t) -$

$\big(L(1) - L(0)\big)\, e^{\alpha(1-t)}$ gives rise to a load function L' with $L'(1) = L'(0)$. The transformation $I \mapsto I'$ is only one of many ways to alter an input current in order to satisfy relation 2.15, and we will use another transformation for the purpose of simulating the spiking activity of a sLIF neuron. Notice that in the steady regime, the only effect of a constant current injection I_c is to lower the effective spiking threshold by $\Delta l = I_c/\alpha$. Then property 2.15 allows us to extend the definition of L on the positive half line $\mathbb{R}^+$ to form a continuous periodic effective barrier by setting $L(t) = L(\pi(t))$, where π is the phase map. Given a periodic barrier L, the successive spiking times of the sLIF neuron form the sequence of random times $\{\tau_i\}$ defined in equation 2.14. Moreover, the periodicity and the continuity of L guarantee an infinite sequence of such spikes.

The spiking times of the sLIF neuron, projected in the phase space as $\{\phi_i = \pi(\tau_i)\}$, define the FPMC. To specify this Markov chain, we must determine its transition kernel by answering the following question: Knowing that the sLIF neuron generates a spike with a given phase, what is the probability law of the next spiking phase? This demands that we first consider the probability of spike timing for the periodic effective boundary L and then to "wrap" that probability on the phase-space $\mathbb{S}$. Specifically, given a spiking time s, say in $[0, 1)$, let us consider the first passage time τ_s of an OU process starting at $U_s = L(s) - (l - r)$ to the barrier L. Because L is a continuous function, it is known that the random variable τ_s admits a continuous nondecreasing cumulative distribution function $F_s : [s, \infty) \to [0, 1]$ (Lehmann, 2002). We then characterize the law of τ_s as the probability measure k_s defined over the intervals $I_{a,b} \subset [s, \infty)$, $s < a < b$ by

$$k_s(I_{a,b}) = F_s(b) - F_s(a). \tag{2.16}$$

Notice that as F is continuous, the measure F does not have Dirac masses, and it does not matter whether the end points of the interval are included. To define the phase measure κ_ψ from the time measure k_s with $\psi = \pi(s)$, we need to recognize the sets of $\mathbb{S}$ that play the same role as the intervals $I_{a,b}$ in defining k_s. As the phase-space $\mathbb{S}$ has the topology of a circle, these sets are easily identified as the arcs $A_{\alpha,\beta}$ running from phase α to phase β . The only added complication is due to the fact that arcs are oriented, meaning that $A_{\alpha,\beta} \neq A_{\beta,\alpha}$. Then for any spiking phase ψ, $0 \leq \psi < 1$, the phase measure k_ψ is entirely specified by

$$\kappa_\psi(A_{(\alpha,\beta)}) = k_\psi\big(\pi^{-1}(A_{(\alpha,\beta)}) \cap [\psi, +\infty)\big), \tag{2.17}$$

where we recall that π is the phase map.

For any time s, the cumulative function F_s is continuous and $F_s(t) \to 1$ at least exponentially when $t \to \infty$. This implies that the cumulative function $\phi \mapsto F_\psi(\phi) = \kappa_\psi(A_{(\psi,\phi)})$ is also continuous for any phase ψ and

the measures κ_ψ have no Dirac masses. However, the measures κ_ψ need not admit a density function, that is, a Lebesgue measurable function $h(\psi)$, such that $\kappa_\psi(d\phi) = h_\psi(\phi)\, d\psi$. Indeed, there exist measures with continuous cumulative functions but no density with respect to Lebesgue measure. An example of such measures arises from uniformly distributing a unit mass over the triadic Cantor set (Mandelbrot, 1982). The resulting measure has no density because it is supported on a set of a zero Lebesgue measure but its cumulative function, the well-known "Devil's staircase," is continuous.

In any case, the collection of measures $\kappa = \{\kappa_\psi\}$ forms a transition kernel κ on the phase-space $\mathbb{S}$, which specifies the probability law of the next spiking phase conditionally to the last spiking phase. Mathematically speaking, given an initial probability measure μ_0 on $\mathbb{S}$, the transition kernel κ defines the FPMC as a continuous-state, discrete-time Markov chain (Häggström, 2002; Norris, 1998; Stewart, 2009). This FPMC, denoted by $\mathcal{T} = \{\mathcal{T}_0, \mathcal{T}_1, \ldots, \mathcal{T}_n, \ldots\}$, induces a probability $\mathcal{P}$ on the phase sequences $\phi_0, \phi_1 \ldots \phi_n$ according to

$$\forall n \in \mathbb{N}, \quad \mathcal{P}(d\phi_n, \ldots, d\phi_0) = \kappa_{\phi_{n-1}}(d\phi_n) \ldots \kappa_{\phi_0}(d\phi_1)\mu_0(d\phi_0). \tag{2.18}$$

2.3. Ergodicity of the Markov Chain. In the context of Markov chains, the peristimulus histogram can be defined as the distribution of spiking phases of a cyclically driven sLIF neuron, that is, as a probability measure μ. For this definition to be unambiguous, we require that the distribution of spiking phases be independent of the initial spiking phase ϕ_0 so that, we associate a unique probability measure μ with each input I. Practically, this ensures that the normalized instantaneous firing rate and the probability of spiking phase coincide. Formally, establishing the identification of the spiking rate with the probability of spiking phases amounts to proving that the Markov chain $\mathcal{T}$ is ergodic. After defining the property of ergodicity, we show that the FPMC $\mathcal{T}$ with H-continuous periodic load function C is indeed ergodic.

Let us first recall that a distribution μ is invariant by the Markov chain $\mathcal{T}$ if it satisfies

$$\mu(d\phi) = \int_0^1 \kappa_\psi(d\phi)\mu(d\psi), \tag{2.19}$$

so that if $\mathcal{T}_n$ is distributed according to μ, so is $\mathcal{T}_{n+1}$. When there exists a unique such measure μ, for any initial distribution μ_0 and any arc set A of spiking phases in $\mathbb{S}$,

$$\lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} \mathbb{1}_A(\mathcal{T}_n) = \mu(A), \quad \mathbb{1}_A(\phi) = \begin{cases} 1 & \text{if } \phi \in A \\ 0 & \text{if } \phi \notin A \end{cases}, \tag{2.20}$$

and the Markov chain is said to be ergodic. Simply stated, the fraction of time that the Markov chain spends in A tends toward the measure of A under μ. Therefore, in the limit of an infinite number of spike observations, the normalized instantaneous firing rate tends to the unique measure μ.

We now show that the Markov chain $\mathcal{T}$ is ergodic for H-continuous barriers.

The phase space of $\mathcal{T}$, identified with the circle $\mathbb{S}$, is compact. As a result, to prove the existence of invariant measures for $\mathcal{T}$, it is enough to show that $\mathcal{T}$ has the strong Feller property (Hernández-Lerma & Lasserre, 2003). The Markov chain $\mathcal{T}$ has the strong Feller property if, for all arc sets $A \subset \mathbb{S}$,

$$\phi_n \to \phi \in \mathbb{S}, \quad \Rightarrow \quad \kappa_{\phi_n}(A) \to \kappa_\phi(A). \tag{2.21}$$

The strong Feller property specifies that if two identical sLIF neurons spike respectively at phases ϕ and ψ, then when ψ asymptotically approaches ϕ, the probability that the first neuron later spikes with a phase in a given phase arc A becomes the same as for the other neuron. In other words, close initial conditions entail similar probability laws for the occurrence of the next spiking events.

To establish the uniqueness of the invariant measure μ, it is enough to show that the Markov chain $\mathcal{T}$ has the irreducibility property (Hernández-Lerma & Lasserre, 2003). The Markov chain $\mathcal{T}$ has the irreducibility property if, for all arc sets $A \subset \mathbb{S}$,

$$\begin{aligned} \exists \phi \in \mathbb{S},\ \kappa_\phi(A) > 0 \quad \Rightarrow \quad & \forall \psi_0 \in \mathbb{S}, \exists n \in \mathbb{N}, \{\psi_n\} \in \mathbb{S}^n, \\ & \kappa_{\psi_n} * \cdots * \kappa_{\psi_1} * \kappa_{\psi_0}(A) > 0, \end{aligned} \tag{2.22}$$

where $*$ denotes the convolution operation. The irreducibility property states that if one spiking phase is achievable for a given starting phase, it is attainable for any starting phase after a finite number of spiking steps. The previous fact follows from an intuitive observation about spiking times: if one trajectory starting at t has a nonzero probability of hitting a barrier in a given time region, we can easily convince ourselves that another trajectory starting at any s prior to t has a nonzero probability of being close to the reset value in t and, from there, to unfold as a trajectory that has been reset in t. When this statement is adapted for spiking phases, a caveat is that because of the wrapping operation, one cannot tell whether a spiking phase is produced by a spiking time that is reached after one or many repetitions of the periodic stimulus. However, this caveat is no obstacle to establish the irreducibility property for spiking phases. Actually, the Markov chain $\mathcal{T}$ has the stronger property that any phase arc can be reached with finite probability after one spiking step: for all arc sets $A \subset \mathbb{S}$,

$$\exists \phi \in \mathbb{S}, \quad \kappa_\phi(A) > 0 \quad \Rightarrow \quad \forall \psi \in \mathbb{S}, \kappa_\psi(A) > 0. \tag{2.23}$$

We deduce more rigorously in the appendix the strong Feller property and the irreducibility property from considerations about the first-passage time problem. Essentially these properties hold for our FPMC for two reasons: the continuity of the barrier, which ensures the continuity of the transition kernels in a broad sense, and the reset rule, which constrains the membrane potential to be reset away from the barrier, thus avoiding pathological situations such as immediate absorption.

3. Monte Carlo Numerical Framework

We have formally described the generation of spiking phases by a sLIF neuron in terms of the ergodic FPMC $\mathcal{T}$. This has enabled us to define the peristimulus histogram (i.e., the normalized instantaneous firing rate), as the unique invariant measure μ of $\mathcal{T}$. We now intend to characterize the effect of the input current's H-singularity on the measure μ. To achieve this goal, we devise an efficient numerical scheme to simulate the production of spiking phases by $\mathcal{T}$. Specifically, we approximate the FPMC $\mathcal{T}$ with another ergodic FPMC $\mathcal{T}^N$ for an effective barrier L_N that is a simpler, approximate version of the effective barrier L.

In practice, we construct a multiresolution sequence of barrier functions $\{L_N\}$ that interpolate L on the dyadic points $D_N = \{k2^{-N} \mid 0 \leq k < 2^N\}$. Such a construction guarantees that L_N converges toward L for increasing N, when the interpolation mesh goes to zero.[2] The approximating sequence $\{L_N\}$ defines a sequence of ergodic FPMCs $\mathcal{T}^N$ with transition kernels κ_ψ^N and invariant measures μ_N. Given an initial spiking phase ψ and for increasingly precise approximations of L, we easily convince ourselves that $\kappa_\psi^N(A)$, the probability that $\mathcal{T}^N$ generates the next spiking phase ϕ in the arc set $A \subset \mathbb{S}$, converges toward the probability $k_\psi(A)$ associated with $\mathcal{T}$. This implies that for increasingly large N, the invariant probability μ_N converges toward the probability μ (Karr, 1975), thus justifying our simple numerical scheme.

In this section, we account for a Monte Carlo method to numerically simulate the FPMC $\mathcal{T}^N$. We first recall the pathwise construction of the linear diffusion processes, which encompass the OU process, as the limit of a finite-dimensional multiresolution approximation scheme. We then construct a finite-dimensional representation of our driving current that is adapted to the sLIF integration scheme while allowing us to control the Hölder regularity of the resulting effective barrier. We finally depict the principles of a probabilistic dichotomic search algorithm (Taillefumier & Magnasco, 2010) that computes efficiently and accurately first-passage times in the case of rough boundaries.

[2] As L is H-continuous with $H > 0$, the convergence $L_N(\phi) \to L(\phi)$ is uniform with respect to spiking phase ϕ and exponential in N.

3.1. Finite-Dimensional Pathwise Representations. In view of specifying simple, approximate barriers L_N, we first recall that a linear diffusion process, such as the OU process U, can be expanded in a well-chosen basis of functions to form a discrete representation. This is generally achieved by writing U as a convergent series of random functions,

$$U = \lim_{N\to\infty} U^N \quad \text{with} \quad U^N = \sum_{n=0}^{N} f_n \cdot \Xi_n, \tag{3.1}$$

where the f_n are deterministic functions and the Ξ_n are independent and identically distributed random variables. The convergence of the series shall be understood as "almost sure" in the sense that although there are realizations of the random coefficients Ξ_n for which the series diverges, the set of such coefficients has probability measure zero.

The Lévy-Cesielski construction (Lévy, 1948) of the Wiener process W on $[0, 1]$ provides an example of discrete representation for a continuous stochastic process. The construction proceeds recursively by defining finite-dimensional processes W_N as linear interpolations of the Wiener process on the dyadic points D_N. In this case, the coefficients Ξ_n are gaussian independent and the elements f_n are called Schauder elements and denoted $s_{n,k}$: $s_{0,0}(t) = t$, and for all $n > 0$,

$$s_{n,k}(t) = \begin{cases} 2^{\frac{n-1}{2}}(t - k2^{-n+1}), & k2^{-n+1} \le t \le (2k+1)2^{-n}, \\ 2^{\frac{n-1}{2}}((k+1)2^{-n+1} - t), & (2k+1)2^{-n} \le t \le (k+1)2^{-n+1}, \\ 0, & \text{otherwise.} \end{cases} \tag{3.2}$$

Interestingly, the Schauder elements are obtained by time-dependent integration of the Haar basis elements, which constitute a complete orthonormal system in $L^2(0, 1)$: $h_{0,0}(t) = 1$ and $h_{n,k}(t)$, where for all $n > 0$,

$$h_{n,k}(t) = \begin{cases} 2^{\frac{n-1}{2}}, & k2^{-n+1} \le t \le (2k+1)2^{-n}, \\ -2^{\frac{n-1}{2}}, & (2k+1)2^{-n} \le t \le (k+1)2^{-n+1}, \\ 0, & \text{otherwise.} \end{cases} \tag{3.3}$$

It is possible to generalize the Lévy-Cesielski construction to OU processes. This requires identifying finite-dimensional interpolations of a process with conditional expectations (Taillefumier & Touboul, 2011). More precisely, the finite-dimensional processes U^N are defined as the expectations of U conditionally to $U_{t\in D_N}$, which may be written as $U^N(t) = \mathbb{E}[U_t \mid U_{k2^{-N}},\ 0 \le k < 2^N]$. Considering these processes leads to the definition of a new Schauder basis of continuous functions $s_{n,k}$ for the OU process.

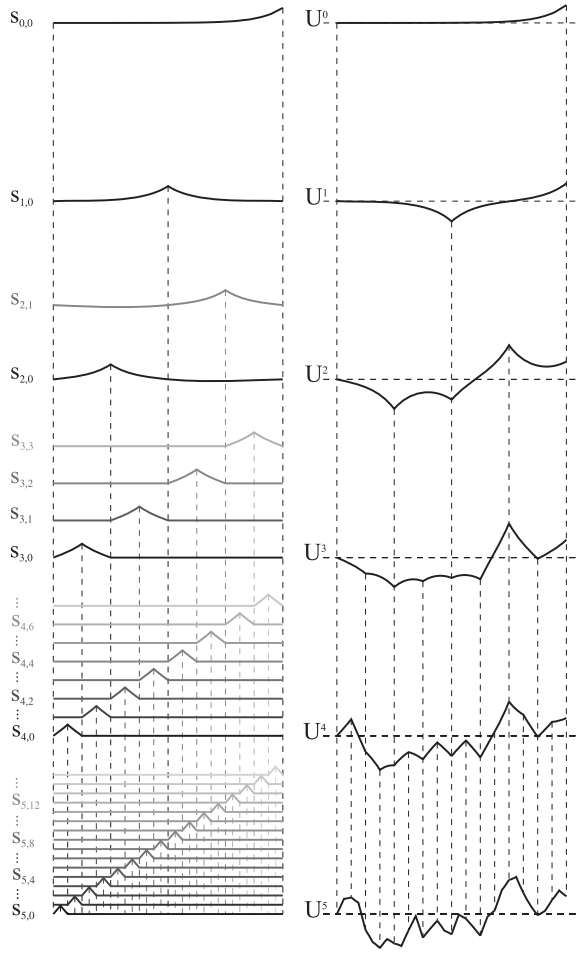

Figure 4: Lévy-Cesielski construction to the OU process. (Left) The elements of the basis $s_{n,k}$ are represented for each rank n with $0 \le n < 6$. (Right) The corresponding conditional averages $U^n = \mathbb{E}\left[U \mid U_{k2^{-n}T},\ 0 \le k < 2^n\right]$ of the OU process are shown for a given noise realization. Notice that the compact supports of $s_{n,k}$ exhibit a dyadic nested structure.

Instead of giving the analytical expression of $s_{n,k}$, we represent the first basis elements in Figure 4 where we also depict the principle of the Lévy-Cesielski construction. Observe that the basis elements $s_{n,k}$, which are compactly supported on $S_{n,k} = [k2^{-n+1}, (k+1)2^{-n+1}]$, quickly tend toward the Wiener process Schauder elements in the limit of vanishing scale.

We can use the above generalization of the Lévy-Cesielski construction to specify finite-dimensional approximations I_N of the input current I such that the corresponding effective barriers L_N interpolate L on the dyadic points D_N. This is made possible by the existence of a basis of functions $h_{n,k}$ that plays the same role for the OU process as the Haar system for the Wiener process. Such Haar-like functions $h_{n,k}$ are derived from the Schauder elements $s_{n,k}$ through $h_{n,k} = s'_{n,k} - \alpha s_{n,k}$, where $s'_{n,k}$ is the time derivative of $s_{n,k}$ and α the leak coefficient. As a result, the Schauder elements $s_{n,k}$ can be recovered by leaky time integration of the Haar-like functions $h_{n,k}$. Importantly, the Haar-like functions $h_{n,k}$ constitute a complete orthonormal basis of $L^2(0,1)$. Then, expanding an input current I on the Haar-like functions $h_{n,k}$ yields finite-dimensional approximations of I that are adapted to the leaky-integration scheme. Actually, for any input current I, the leaky time integration of the finite-dimensional approximation,

$$I_N(t) = \sum_{0 \le n < N} \sum_{0 \le k < 2^{n-1}} h_{n,k}(t) \int_0^1 I(s) h_{n,k}(s)\, ds, \tag{3.4}$$

gives rise to barrier L_N that interpolates L on the dyadic points D_N. Observing that the periodic condition $L_N(0) = L_N(T) = l$ is satisfied by setting the coefficient of $s_{0,0}$ to zero, we propose using the Schauder basis elements $s_{n,k}$ and the Haar-like functions $h_{n,k}$ to construct simple, finite-dimensional approximations to L and I, respectively.

3.2. Family of Hölder Singular Currents. We have seen that the Lévy-Cesielski construction provides a convenient scheme to build simple, finite-dimensional approximations of the current I and its associated barrier L. We now show that the use of the Schauder basis elements allows us to construct I such that the barrier L has the prescribed Hölder regularity H. To see this, we need to realize that the local Hölder regularity of the effective barrier L is entirely captured by the scaling of the coefficients appearing in the expansion of L on the Schauder basis.

For example, suppose the input current is a realization of gaussian white noise $I(t) = \dot{W}(\omega)$. The leaky time integration of the current I gives rise to a barrier L that is the sample path of an OU bridge process, defined as an OU process that is conditioned on its end value. In particular, L is H-continuous of exponent $1/2$. For this reason, we denote an input current with a frozen gaussian white noise component as $I^{1/2}$ and the corresponding effective barrier as $L^{1/2}$. Moreover, we denote the coefficients of $s_{n,k}$ in the expansion of the barrier over the Schauder basis elements by

$$\xi^{1/2}_{n,k} = \int_0^1 I(s) h_{n,k}(s)\, ds. \tag{3.5}$$

By construction, the coefficients $\xi^{1/2}_{n,k}$ are an independent realization of centered gaussian variables with unit variance.

Can we modify the coefficient $\xi^{1/2}_{n,k}$ to build a version of L with arbitrary Hölder exponent H? To positively answer this question, we need to use a well-known result of Daoudi, Lévy Véhel, and Meyer (1998) that rigorously relates the local Hölder exponent of a function to the asymptotic behavior of the coefficients of its decomposition in the Schauder basis. The only complication is that in dealing with the coefficient $\xi^{1/2}_{n,k}$ associated with a frozen gaussian white noise component, we are actually considering realizations of independent gaussian variables. Then, because of the probabilistic nature of $\xi^{1/2}_{n,k}$, statements about Hölder regularity hold only almost surely. However, this is not a relevant limitation for our purpose, and Daoudi et al.'s (1998) results directly entail that for all H, $0 < H < 1$, the barriers L^H defined as

$$L^H(t) = l - \sum_{0\le n} \sum_{0\le k<2^{n-1}} s_{n,k}(t)\cdot \xi^H_{n,k}, \quad \xi^H_{n,k} = 2^{n(H-1/2)}\,\xi_{n,k} \tag{3.6}$$

are precisely almost surely H-continuous (Taillefumier & Touboul, 2011). Moreover, the application $H \mapsto L^H(\xi)$ defined from $(0,1)$ to the set of periodic continuous functions $C(\mathbb{S})$ is a continuous mapping for the L^∞-norm. Therefore, we can continuously (in the L^∞-norm) control the asymptotic Hölder continuity of the effective barrier driving the activity of a sLIF neuron by smoothly changing the coefficient $\xi^H_{n,k}$ used to construct the piecewise approximations L^H_N.

In order to emphasize the effect of varying the Hölder regularity, we adopt a slightly modified version of our barriers L^H by weighting them with a continuous function $H \mapsto c(H)$: $L'^H = c(H)\big(L^H - L^H(0)\big) + L^H(0)$. The function c is chosen so that the newly formed barriers cause the neuron to fire with an overall mean firing rate that remains constant when changing H. We refer to the FPMC associated with the first-passage time to L^H as $\mathcal{T}^H$, and we label all the related quantities, such as the spiking time transition kernel k^H_s or the invariant spiking phase measure μ_H, with an explicit mention of H-dependence.

At this point, it is worth mentioning a minor but conceptually important caveat of our numerical approach. The construction of the Hölder singular currents supposes the choice of a dyadic partition to construct the Schauder basis. Due to the choice of a particular partition, the statistics of the H-continuous limit curve are not stationary in time except when $H = 1/2$. Rigorously stationary H-continuous curves can be constructed as almost sure realizations of the fractional Wiener process. However, we favor our construction because of its computational convenience.

3.3. Probabilistic Dichotomic-Search Algorithm. Using a well-chosen Schauder basis of functions, we are able to construct finite-dimensional

approximations of a barrier L that has been chosen with a prescribed Hölder exponent H. To estimate the invariant measure μ_H of the corresponding FPMC $\mathcal{T}^H$, we must generate first-passage times of an OU process for a fluctuating barrier with varying initial conditions. The law of these first-passage times defines the transition kernels k_s^H, which depends on a particular initial condition prescribed by the timing s of the preceding spike. In that respect, the need for simulating a continuum of transition kernels precludes any attempt to tabulate their probability laws. Besides, probability distributions of passage times are known analytically only for the most trivial situations, such as a Wiener process first crossing affine boundaries. Thus, for our problem, first-passage times must be computed numerically.

When the concern is the sLIF model, numerical integration through the discretization of the Volterra equation is often advocated to directly compute the density of first-passage times (Paninski, Haith, & Szirtes, 2008; Ahmadian, Pillow, & Paninski, 2010). Nevertheless, this otherwise very fast method fails to produce a convergent scheme for nondifferentiable boundaries when the distribution of passage times can become singular. To treat the case of a nondifferentiable barrier, we must estimate the distribution of spiking events by Monte Carlo methods, which consist in numerically simulating many sample paths of the process until it first crosses the boundary. Such an approach traditionally carries both practical and theoretical difficulties, whether focusing on the computational cost of the method or the accuracy of the returned times (Platen, 1999; Honeycutt, 1992; Bouleau & Lépingle, 1994).

To ensure computational efficacy with controlled accuracy, we use a stochastic dichotomic-search algorithm for the first-passage times of an OU process to a rough boundary. We refer to Taillefumier and Magnasco (2010) for a detailed account of the algorithm. Here, we recall just the essential tenets of the method. When first-passage times are being generated by the Monte Carlo method, the computational cost is set by the resolution with which the sample paths of the underlying process are simulated. For the OU process, the Lévy-Cesielski construction offers a recursive scheme that exactly simulates sample paths to any desired level of resolution. Because we are interested only in the first-passage time, the process may be pathwise simulated at high resolution only in regions where a first-passage time is likely, that is, when the process is close to the boundary (Gaines & Lyons, 1997; Römisch & Winkler, 2006). We implement such an adaptive exploration of sample paths through a dichotomic refining procedure that exploits the recursive scheme of the Lévy-Cesielski construction (Taillefumier & Magnasco, 2010).

For being probabilistic in nature, our algorithm can return erroneous estimates for first-passage times. This happens when our method does not refine the sample path of the OU process in a region where a first passage does indeed occur. In such cases, it always produces approximate crossing times that are not first-passage times but later crossings. However, the

probability of occurrences of such errors can be tightly controlled. Specifically, our algorithm is designed not to search for first passages in time intervals where the probability of such crossings is known to be less than a parameter value $\epsilon > 0$, where ϵ can be set very small, typically of order 10^{-20}.

4. Competition Between External and Internal Noise

Using the theoretical and computational framework developed above, we can investigate numerically the peristimulus histogram of a sLIF neuron in response to an input current that comprises a frozen-noise component. In Taillefumier and Magnasco (2013), we exhibit the existence of a transition between a regime where the normalized peristimulus histogram of spiking phases is a continuous density function and a regime where it becomes singular.

The existence of these two regimes is the signature of the competition between the internal noise and the frozen external noise that may result from the integration of background activity. In particular, the Hölder singularity H of the injected frozen noise is the order parameter controlling the transition, which happens precisely when it becomes the same as the Hölder exponent of the internal noise $H = 1/2$.

In this section, we briefly describe the phase transition occurring for input current with Hölder singularity exponent $H = 1/2$. Then we give a heuristic explanation of the phenomenon based on pathwise properties of the subthreshold voltage dynamics. Finally, we rigorously demonstrate that the distribution of spiking phase admits a continuous density when $H > 1/2$.

4.1. Numerical Evidence of a Phase Transition. We simulate the stationary measures μ_H corresponding to a family of H-continuous barriers L_H over a period $T = 1\,\mathrm{s}$. The sLIF model is characterized by the following parameters: the leak constant α (with relaxation time $1/\alpha$), the diffusion coefficient σ, and the threshold value V (the reset is set to $v = 0$). Here, we follow Gerstner and Kistler (2002) and set these parameters to $1/\alpha = 10\,\mathrm{ms}$, $\sigma = 5\,\mathrm{mV}$, and $l = 50\,\mathrm{mV}$. The current I^H includes a suprathreshold steady component that drives the effective barrier close to the mean value of the neuron's voltage and a frozen-noise perturbation (taken with the same amplitude as the intrinsic noise for $H = 1/2$). Additionally, we constrain I_H to be such that the average firing rate of the neuron is held constant at $5\,\mathrm{Hz}$. Thus, the stationary measures μ_H are sampled temporally with the same statistical power. As previously stated, this constraint is satisfied by modulating the amplitude of the injected noise with a continuous function $H \mapsto c(H)$. This treatment increases the amplitude of the noise with larger H (i.e, for smoother barriers), so that over the time course of one period T, the barriers L^H can be seen as sample paths of processes with similar

variance. This approach is numerically correct only as long as the mean spiking interval (here 0.2 s) is larger than the typical fluctuations of the barriers and as long as we operate below the timescale matching regime (Shimokawa, Pakdaman, & Sato, 1999).

In order to better connect our simulations to experimental situations, we represent our results as raster plots. In Figure 5, we draw the raster plots made of 5000 spikes for six different H-continuous barriers, on as many panels (panels a through f). Below each raster plot, we represent the instantaneous firing rate of the sLIF neuron over an input cycle. Above each raster plot, we picture the barrier that has elicited the spiking activity of the sLIF neuron. Observe that if all of the barriers exhibit similar amplitudes of fluctuation, the smoothing of the barriers with larger values of H is accompanied by the formation of a few isolated downward peaks. Accordingly, the larger the Hölder exponent is, the higher the spike reliability is at large scale, as shown by the well-individualized band of firing on the raster plots for $H > 1/2$. In Figures 6 and 7, we represent raster plots in the exact same fashion, except that the panels depict the magnified regions of low spike reliability framed in Figures 5 and 6, respectively. Notice that the barrier is scaled with the same aspect ratio on each panel according to the scaling of an OU process ($H = 1/2$). We deliberately choose this type of scaling to emphasize the "point of view" of a voltage trace exploring the vicinity of the barrier. Zooming in on H-continuous barriers for $H > 1/2$ gradually smoothens their profiles, while the same operation for $H < 1/2$ yields rougher traces. In agreement with this relative smoothing or roughening of the barrier, it is observed that the spike precision, as indicated by a perfect lining up of spiking events on the raster plot, changes abruptly when the barriers have the same roughness as the typical voltage trace ($H = 1/2$).

In general, defining reliability and precision in the neural response requires pooling trains of spikes elicited by stereotypical stimulations into different patterns with well-individualized spiking events (Fellous, Tiesinga, Thomas, & Sejnowski, 2004). The reliability of a spiking event is measured by the fraction of time that the event occurs, while the precision of an individual spiking event is quantified by the variance of its timing conditionally to belonging to a specific pattern (Toups, Fellous, Thomas, Sejnowski, & Tiesinga, 2011, 2012). When considering our numerical results, we think of spiking reliability and spiking precision in a rather different way. On the one hand, if we can identify narrow phase regions that carry most of the spiking phase statistics, we say that the sLIF neuron spikes reliably. On the other hand, we say that the sLIF neuron has great temporal precision if spikes occur only for a restrained set of achievable times within the input cycle. Bearing this in mind, under the constraint of constant average activity, varying the degree of singularity of the injected current causes a sLIF neuron to transition between two regimes: a regime of high spiking reliability but poor temporal precision for weakly singular input and a regime of poor reliability

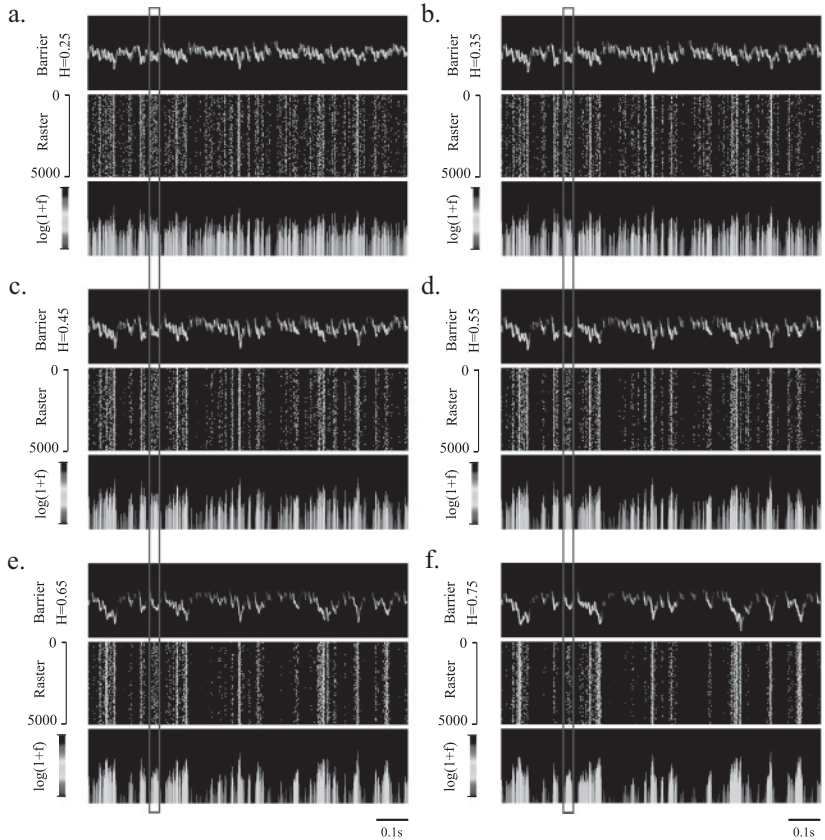

Figure 5: Each panel shows the results of the numerical simulation of a sLIF neuron cyclically driven by an input current I^H of period 1 s. The input current I^H includes a constant steady component I_c and a frozen-noise component I_n^H for different Hölder exponents H. In the top row of a panel, we represent the effective barrier L^H that arises from the integration of I^H. In the middle row of a panel, we represent a raster plot made of 5000 spiking phases. In the bottom row of a panel, we represent the normalized peristimulus histogram in logarithmic scale for 10^8 simulated spikes. The period of the stimulus is divided in 2^{10} phase bins. We denote by f_i, $1 \leq i \leq 2^{10}$ the fraction of spikes that occur in bin i and plot the histogram corresponding to $\{\ln(1 + f_i)\}_i$. Note that $\ln(1 + f_i) = 0$ if no spike happens in bin i. (A color version of this figure is available in the online supplement.)

and exquisite temporal precision for highly singular input. From the mathematical point of view, this transition happens when the effective barrier exhibits the same Hölder exponent as the exponent of the subthreshold voltage trace.

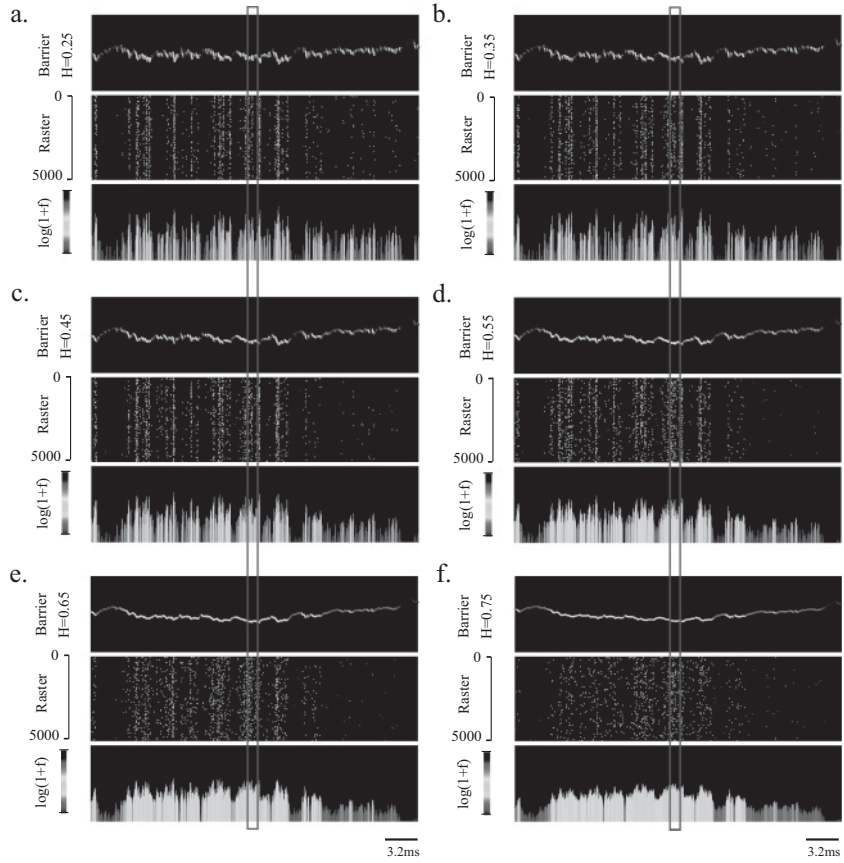

Figure 6: Each panel shows the results of the numerical simulation of Figure 5 in a specific region of the phase space representing a period of 32 ms for different Hölder exponents H. In the top and middle rows of the panels, the effective profile and the raster plots are obtained by zooming in the framed region in Figure 5. The effective barrier is zoomed in according to the scale invariance of the OU process: if time is dilated by d, the voltage is scaled by $\sqrt{d}$. The raster plots show only a fraction of the spikes represented in Figure 5. In the bottom row, the period of 32 ms is again divided in 2^{10} phase bins. We denote by f_i, $1 \le i \le 2^{10}$ the fraction of spikes that occur in bin i and plot the histogram corresponding to $\{\ln(1 + f_i)\}_i$. (A color version of this figure is available in the online supplement.)

4.2. Heuristic Pathwise Interpretation. The transition between the two encoding regimes can be qualitatively explained by pathwise considerations. In that regard, we stress that the effective barrier determines a neuron's spiking pattern through the strength of its fluctuation and also through

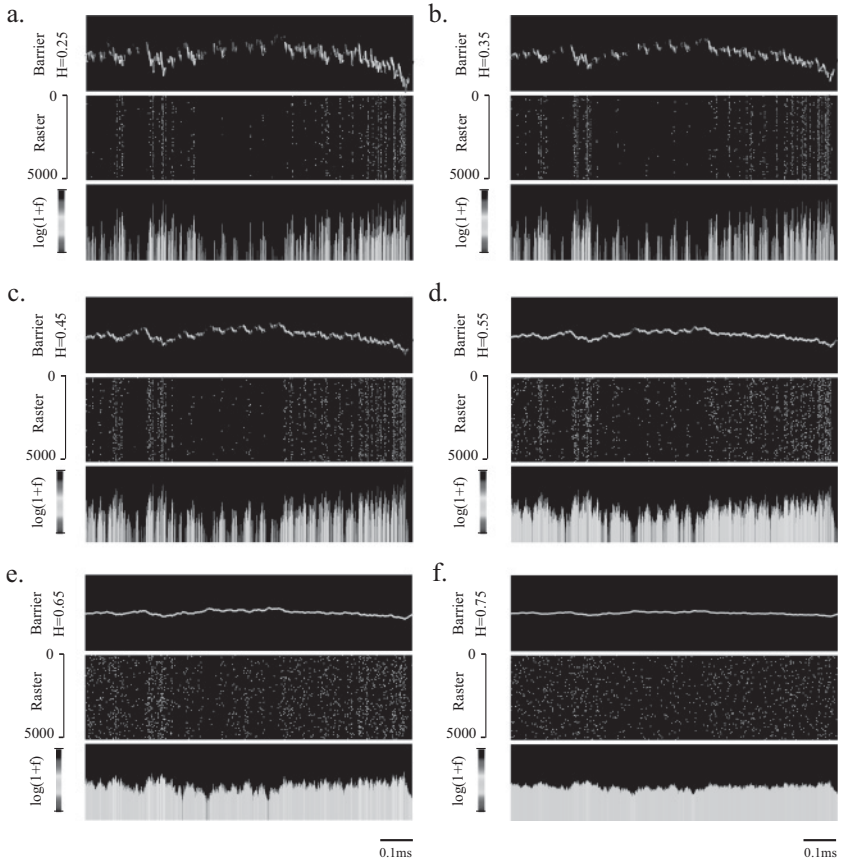

Figure 7: Each panel shows the results of the numerical simulation of Figure 6 in a specific region of the phase space representing a period of 1 ms for different Hölder exponents H. In the top and middle rows of the panels, the effective profile and the raster plots are obtained by zooming in the framed region in Figure 6. The effective barrier is zoomed in according to the scale invariance of the OU process: if time is dilated by d, the voltage is scaled by $\sqrt{d}$. The raster plots show only a fraction of the spikes represented in Figure 6. In the bottom row, the period of 1 ms is again divided in 2^{10} phase bins. We denote by f_i, $1 \leq i \leq 2^{10}$ the fraction of spikes that occur in bin i and plot the histogram corresponding to $\{\ln(1 + f_i)\}_i$. (A color version of this figure is available in the online supplement.)

its Hölder regularity. Actually, if the strength of fluctuations is the dominant effect at a large timescale, the Hölder continuity is always the decisive property that dictates the distribution of spike timing at a small timescale. In a first approximation, a voltage trace that is tasked with hitting a

barrier almost always runs into the accessible downward excursions of the low-resolution depiction of the barrier. Due to the leak term, this process of traveling from one major obstacle to another is essentially memoryless. Yet at a small timescale, the continuous voltage trace bears the memory of the position it just visited. This memory effect plays a crucial part when the trace lies in the vicinity of the barrier.

To realize the crucial part of the above memory effect, we must envision the Hölder exponent as the exponent of the scaling operation that leaves the profile of a curve statistically invariant. This explains why zooming on H-continuous barriers with the scale invariance of the OU process ($H = 1/2$) causes them to smoothen for $H > 1/2$ or roughen for $H < 1/2$. For $H < 1/2$, the voltage trace always runs into an easy-to-access downward fluctuation for $H < 1/2$. Indeed, such obstacles keep presenting themselves at every timescale and become increasingly hard to avoid, since they scale with a lower exponent (higher roughness). Therefore, the voltage trace has to die on the left flank of one of these obstacles and cannot hit the barrier in a sheltered region. For barriers that are everywhere rougher than the typical voltage trace, the barrier's profile is almost surely entirely sheltered, causing the distribution of accessible time to be singular. For $H > 1/2$, the obstacles become easier to avoid as the timescale of observation gets shorter. Therefore, as the voltage trace gets closer to hitting the barrier, it increasingly ignores the obstacles and can reach any part of the curve indiscriminately, giving rise to a continuous density of spiking times.

Understood as a pathwise phenomenon, the transition from a continuous density for first-passage time to a singular distribution is readily generalizable to other types of integrate-and-fire neurons. For instance, if the internal noise is a fractional gaussian noise with Hurst coefficient H, $0 < H < 1$ (Mandelbrot & Ness, 1968), the phase transition of the sLIF neurons precisely happens for H-continuous barriers. We hypothesize that such a behavior, being very general in nature, is the signature of a stochastic integrate-and-fire encoding scheme.

4.3. Integral Equation for the First-Passage Time. Based on heuristic arguments, we hypothesized that a sLIF neuron has a well-defined firing rate for frozennoise components that are less singular than the internal noise. From a theoretical point of view, we wish to show that the spiking phase of a sLIF neuron has a continuous density function for H-singular frozen noise with $H > 1/2$. In order to prove this claim, we first simplify the formulation of our problem and establish an integral equation for the putative density function.

A straightforward simplification of the problem follows from the realization that by ergodicity of the FMPC $\mathcal{T}^H$, if the kernels κ_ϕ^H are continuous for any spiking phase ϕ in $\mathbb{S}$, so is the stationary measure μ^H. Thus, it is enough to show the continuity of the spiking phase kernel κ_ϕ^H, which is itself

implied by the continuity of the spiking time kernels k_s^H defined on $[s, \infty)$ for infinite periodic boundary $t \mapsto L(\lfloor t \rfloor)$. A further simplification follows from the existence of a pathwise mapping of the Ornstein-Uhlenbeck process U with reset initial condition $U_s = L(s) - (l - r)$ to the standard Wiener process W with initial condition $W_s = 0$. This mapping is called the Doob's transformation (Wang & Pötzelberger, 2007) and takes the form

$$\mathcal{D}_s[f](t) = \sqrt{1 + 2\alpha(t-s)}\, f\big(s + \lambda_s(t)\big) - U_s, \tag{4.1}$$

where λ_s is the strictly increasing function

$$\lambda_s(t) = \frac{\ln\big(1 + 2\alpha(t-s)\big)}{2\alpha}. \tag{4.2}$$

Notice that the functional transformation $\mathcal{D}_s$ preserves H-continuity. Incidentally, the kernels k_s^H of the Markov chain corresponding to L^H are equivalently determined as the cumulative distributions of the first-passage times of W with $\mathcal{D}_s[L^H]$ knowing $W_s = 0$. Thus, we need to show only that the first-passage time of a Wiener process for a boundary L^H with $H > 1/2$ admits a density.

Considering a Wiener process for a boundary L^H with $H > 1/2$, we propose to find an integral equation for the putative density of the corresponding first-passage time. Integral equations for the first-passage time naturally arise from probabilistic arguments. Let us consider the event $\{W_t > x\}$ for a continuous barrier L^H satisfying $x > L^H(t)$. Then the first-passage time τ to L^H necessarily occurs before t, and we can write

$$\mathbb{P}(W_t > x) = \mathbb{E}\big[\mathbb{P}(W_t > x \,|\, \tau)\big] = \int_0^t \mathbb{P}(W_t > x \,|\, \tau = s)\, dQ(s), \tag{4.3}$$

where Q denotes the putative first-passage time cumulative function. Because of the Markovian nature of the Wiener process,[3] the probability $\mathbb{P}(W_t > x \,|\, \tau = s)$ is equal to

$$\mathbb{P}\big(W_t > x \,|\, W_s = L^H(s)\big) = \mathbb{P}\big(W_t - W_s > x - L^H(s) \,|\, W_s = L^H(s)\big), \tag{4.4}$$

$$= \mathbb{P}\big(W_{t-s} > x - L^H(s)\big). \tag{4.5}$$

Differentiating equation 4.3 with respect to x, we end up with

$$K\left(\frac{x}{\sqrt{t}}\right) = \int_0^t K\left(\frac{x - L^H(s)}{\sqrt{t-s}}\right) dQ(s), \tag{4.6}$$

[3]This is due to the strong Markov property of the Wiener process.

where K denotes the Heat kernel. Importantly, as long as L^H is H-continuous with $H > 1/2$, we have

$$\lim_{\tau \to t^-} \frac{L^H(t) - L^H(\tau)}{\sqrt{t-\tau}} = 0. \tag{4.7}$$

This warrants taking $x \to L^H(t)$ by superior value in equation 4.6, thus obtaining an integral equation (Park & Paranjape, 1974; Park & Schuurmann, 1976):

$$K\left(\frac{L^H(t)}{\sqrt{t}}\right) = \int_0^t K\left(\frac{L^H(t) - L^H(s)}{\sqrt{t-s}}\right) q(s)\, ds. \tag{4.8}$$

which dates back to original work from Siegert (1951). Notice that we plugged $dQ(s) = q(s)\, ds$ in equation 4.8 to obtain the desired integral equation for the putative density function q of first-passage time.

4.4. Continuity of the Density for Hölder Exponent $H > 1/2$. We are now in a position to prove that as a solution to equation 4.8 the first-passage density q is a continuous function for a barrier L^H with $H > 1/2$. The demonstration of this fact is quite technical and we present only a summary of the main points (see Cannon, 1984, for details).

First, let us observe that the integral equation, 4.8, is of the Volterra type, which comes in two flavors: equations of the first kind and of the second kind (Linz, 1985). To ensure the existence and uniqueness of a solution to the equations of the second kind, we have the following powerful result:

Theorem 1 *(adapted from Courant & Hilbert, 1962; Tricomi, 1985). The linear Volterra equation of the second kind,*

$$g(t) = f(t) + \int_0^t K(t,s) f(s)\, ds, \tag{4.9}$$

where g is a piecewise continuous function, has a unique piecewise continuous solution f for all $t > 0$ if K is bounded on $0 < s < t$ and if there exists a monotone increasing function a with $\lim_{t \to 0} a(t) = 0$, such that for all $0 < s < t$,

$$\int_s^t |K(t,\tau)|\, d\tau \le a(t-s). \tag{4.10}$$

Unfortunately, equation 4.8 is a Volterra equation of the first kind and cannot be dealt with directly. However, for barriers L that are H-continuous, it can be recognized as a linear generalized Abel integral equation, that is,

an equation of the type

$$g(t) = \int_0^t \frac{K(t,s)\,f(s)}{(t-s)^h}\,ds, \tag{4.11}$$

where f is the unknown, g is a continuous function, and K is a continuous kernel for $s \leq t$ and $0 < h < 1$. In our case, $h = 1/2$. Abel integral equations are frequently encountered in physics, and there are methods to prove the existence and uniqueness of a solution by transforming the original equation into a Volterra equation of the second kind. This transformation is made through the use of the Abel integral transform $\mathcal{A}$ and its inverse transform $\mathcal{A}^{-1}$ defined for $h = 1/2$ as

$$\mathcal{A}[f](t) = \int_0^t \frac{f(s)}{\sqrt{t-s}}\,ds, \tag{4.12}$$

$$\mathcal{A}^{-1}[g](t) = \frac{1}{\pi}\frac{d}{dt}\left[\int_0^t \frac{g(s)}{\sqrt{t-s}}\,ds\right]. \tag{4.13}$$

The application of $\mathcal{A}^{-1}$ to equation 4.8 reduces the problem to a Volterra equation of the second kind:

Proposition 1 *(adapted from Cannon, 1984). If L is H-continuous with $H > 1/2$, through the application of the inverse Abel operator, the Volterra equation of the first kind, 4.8, is equivalent to the Volterra equation of the second kind,*

$$\sqrt{2\pi}\,\mathcal{A}^{-1}[g](t) = q(t) + \frac{1}{\pi}\int_0^t H(t,s)q(s)\,ds, \tag{4.14}$$

with the kernel H being defined as

$$H(t,s) = \frac{\partial}{\partial t}\left\{\int_s^t \frac{e^{-\frac{(L(\tau)-L(s))^2}{2(\tau-s)}}}{\sqrt{(t-\tau)(\tau-s)}}\,d\tau\right\}, \tag{4.15}$$

and g denotes the continuous function $g(t) = K(\frac{L(t)}{\sqrt{t}})$.

A careful study shows that the kernel H satisfies the conditions of proposition 1. Thus, the integral equation, 4.14, admits a unique continuous solution, which is the density of the first-passage time to the barrier L.

5. Conclusion

The stochastic linear integrate-and-fire neuron encodes spike timing as the first-passage time of an OU process with a fluctuating barrier (Fourcaud-Trocmé, Hansel, van Vreeswijk, & Brunel, 2003; Sirovich & Knight, 2011). Such an encoding scheme generates a complex message dependence of the firing activity in terms of both spiking reliability, which measures the reproducibility of a firing pattern in response to stereotypical stimulations, and spiking precision, which quantifies the temporal variability of spiking events that encode for equivalent stimuli features (Bair & Koch, 1996; Berry et al., 1997). Because reliability and precision are the two concepts used experimentally to pool spikes in the same meaningful event, it is fundamental to understand the interplay of these two notions in view of elucidating the neural code.

We have revisited the integrate-and-fire model in response to externally imposed, fluctuating neural input. Previous studies that undertook a similar task have narrowed down the concept of timescale matching (Shimokawa et al., 1999), whereby spike reliability and spiking interval precision are maximized if the typical fluctuation of the barrier matches the mean spike interval in the absence of fluctuations (Plesser & Geisel, 1999). The integrate-and-fire model has also been considered in the context stochastic birfurcation theory as a quasi-periodically forced system (Arnold, 1998). From this perspective, firing patterns are likened to attractors for a train of spikes, and reliability is understood as the stability of the attractor in the face of noise perturbation (Tiesinga, 2002). Alternatively, it is also worth mentioning that patterns of spiking times have been shown to converge to random attractors in noiseless integrate-and-fire systems driven by frozen noise (Lajoie, Lin, & Shea-Brown, 2013; Lin, Shea-Brown, & Young, 2009a, 2009b). From a different point of view, it is well documented that the presence of noise in subthreshold stimuli can enhance the ability of a neuron to detect weak signals, giving rise to the concept of stochastic resonance (Rudolph & Destexhe, 2001; Stacey & Durand, 2001).

Here, we clarify how internal and external neuronal noise combine to shape the reliability and precision of the firing activity of the sLIF neuron. We first formulate the most general Markovian framework amenable to simulate the situation of the experiment of Mainen and Sejnowski (1995). Similar to previous work (Shimokawa et al., 2000), we define the succession of spikes as the outcome of an inhomogeneous Markov chain. Identifying the Hölder exponent of the effective barrier as the crucial order parameter, we simulate our Markov chain for injected currents that give rise to various H-continuous effective barriers. Based on our numerical experiments, we briefly discuss the message dependence of the sLIF encoding. We explain the firing regime of a neuron from the interplay of internal noise and external frozen noise: the latter noise gives rise to the statistical profile of the effective

barrier, while the former noise is responsible for the pathwise properties of the subthreshold voltage trace.

Our analysis suggests the existence of two spiking regimes. If the injected frozen noise is less singular than the internal neuronal noise ($H > 1/2$), the firing activity admits a continuous probability density of spiking times, whereas the probability density becomes singular, almost everywhere either zero or infinity, if the frozen noise is more singular than the internal noise ($H < 1/2$). In particular, the nature of the spiking regime depends on the properties of neural noise and is valid for both supra- and subthreshold inputs. In a related paper (Taillefumier & Magnasco, 2013), we explain the nature of the first-passage distribution in the singular spiking mode, while elaborating on the relevance of our analysis to the neural code. This article glosses over these issues, only to focus on the Markovian framework introduced to characterize the firing activity of the sLIF neuron driven by frozen noisy input.

We suggest the exploration of three theoretical avenues in the perspective of widening the scope of our study.

The first direction aims at rigorously characterizing the distribution of spiking events in the singular spiking mode ($H < 1/2$). In particular, it is natural to ask whether the cumulative distribution of the spiking measure is homogeneously Hölder continuous for a given exponent h. Indeed, it can very well be that the spiking measure μ_H admits a nontrivial spectrum of local Hölder exponents h, in which case μ_H can be decomposed into measures μ_h supported by subsets of times for which the cumulative measure has h for local Hölder exponent. Such a measure μ_H can be treated through the multifractal formalism to yield its singularity spectrum, which associates the Hölder exponent with the fractal dimension of μ_h and is characteristic of the statistics of μ_H (Stanley & Meakin, 1988). We predict that the multifractal nature of the spiking measure is closely related to the distribution of the local minima of the effective barrier that are reached or left faster than a typical voltage trajectory (i.e, with a local left or right Hölder exponent $h < 1/2$).

Second, if our study is mainly conceptual, we hope to design numerical methods to bring our analysis closer to experimental relevance. Specifically, we intend to develop an efficient statistical estimator for the local Hölder exponent of the empirical cumulative, tailored to the specifics of the first-passage time problem. If the method would invariably rely on measuring the concentration of the experimental measure in the vicinity of spiking times (Allain & Cloitre, 1991), it should crucially take into account the fundamental asymmetry of the problem: every downward feature of the effective barrier casts a shadow at a later time. As a consequence, left and right Hölder exponents need to be carefully distinguished.

Third, the Hölder singularity of a neuron's afferent input results from the time integration of the activity of the surrounding network of neurons, through the filter of its circuit connectivity. This begs for understanding

which features of the network connectivity promote impeding currents of a given irregularity, whether they arise spontaneously or as a response to some driving input. We envision that the completion of such a program would require describing the stochastic activity of a network of neurons as the evolution of interacting continuous-time random walks. In that respect, we have developed an exact event-driven simulation algorithm to study the precise temporal propagation of firing patterns in networks of stochastic perfect integrate-and-fire neurons (Taillefumier, Touboul, & Magnasco, 2012).

Appendix: Mathematical Proofs

In this appendix, we prove rigorously the two properties that establish the ergodicity of the FPMC: the Feller property and the irreducibility property (Hernández-Lerma & Lasserre, 2003). First, though, we need the following elementary result that justifies approximating the distribution of the first-passage time τ to a barrier L with the law of the first-passage time's τ' corresponding to a barrier L' that is close to L in some sense.

Proposition 2. *Given a real continuous process X and a real x_0, if a sequence of bounded continuous functions $\{L_n\}$ satisfies $L_n(0) > x_0$ and converges uniformly to L on $\mathbb{R}^+$, then $\tau_n = \inf\{t > 0 \mid X_t > L_n(t)\}$ converges in law to $\tau = \inf\{t > 0 \mid X_t > L(t)\}$.*

Proof. Given a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, let $X : \Omega \to C_{x_0}(\mathbb{R}^+)$ be a real continuous process with natural filtration $\mathcal{F}_t$ taking values in the set of continuous functions satisfying $x(0) = x_0$. Note that for any bounded continuous function L in $C_b(\mathbb{R}^+)$ and for every ω in Ω, the first-passage time $\tau(\omega)$ to the barrier L can be seen as a real function defined on $C_b(\mathbb{R}^+)$ by $L \mapsto \tau_L(\omega)$. Moreover, in the complete separable metric space $(C_b(\mathbb{R}^+), |\cdot|_\infty)$, it is a continuous function of its argument. Indeed, given a continuous barrier L and a positive real $\epsilon > 0$, by definition of τ, for any ω, we have $X_t(\omega) < L(t)$ if $t < \tau_L(\omega) - \epsilon$, so that, by continuity of the sample paths, we can define

$$\sup_{0<t<\tau_L(\omega)-\epsilon} (L(t) - X_t(\omega)) = \delta^- > 0. \tag{A.1}$$

In the same way, we can define

$$\sup_{\tau_L(\omega)<t<\tau_L(\omega)+\epsilon} (X_t(\omega) - L(t)) = \delta^+ > 0. \tag{A.2}$$

Setting $\delta = \min(\delta^+, \delta^-)$, for every L' such that $|L' - L|_\infty < \delta$, we have that $|\tau_L(\omega) - \tau_{L'}(\omega)| < \epsilon$.

As a consequence, if L_n converges uniformly toward L when n tends toward infinity, the previous continuity result implies the point-wise sure convergence,

$$\lim_n \tau_n(\omega) = \tau(\omega), \tag{A.3}$$

which entails the convergence in law.

Equipped with the previous elementary fact, we restrain ourselves to the case of the OU process U (which can easily be extended to the time-inhomogeneous case of one-dimensional Gauss-Markov processes) and proceed to prove the Feller property of the FPMC:

Property 1. If the barrier L is homogeneously Hölder continuous, the Markov chain $\mathcal{T}$ is strong Feller:

$$\forall B \in \mathcal{B}(\mathbb{S}), \quad \phi_n \to \phi \in \mathbb{S}, \quad \Rightarrow \quad \kappa_{\phi_n}(B) \to \kappa_\phi(B). \tag{A.4}$$

Proof. Since every open set is a set of continuity, it is enough to show the property for the open arc sets A. We proceed in three steps.

i. Uniform tightness on the real half line: Given the continuous phase map $\pi : [\phi, +\infty) \to \mathbb{R}/\mathbb{Z}$, posit $B = \pi^{-1}(A) \cap [\phi, +\infty)$, which represents the unwrapped version of A in the real half-line $[\phi, +\infty)$. In particular, $\kappa_\phi(A) = k_\phi(B)$, where k_ϕ is defined by equation 2.16. Since the barrier L can be seen as a periodic continuous function on $[\phi, +\infty)$, there exists $M > \sup_{t\in\mathbb{R}^+} L(t)$. Denoting $\tau_\phi = \inf\{t > \phi \,|\, U_t > L(t)\}$ and defining $\tau_\phi^M = \inf\{t > \phi \,|\, U_t > M\}$ for the same OU process U, we have $\tau_\phi \le \tau_\phi^M$. Then $\{\tau_\phi^M < +\infty\} \subset \{\tau_\phi < +\infty\}$ entails $\mathbb{P}(\tau_\phi < +\infty) \ge \mathbb{P}(\tau_\phi^M < +\infty)$. It is well known that $\mathbb{P}(\tau_\phi^M < +\infty) = 1$, that is, an OU process hits a constant barrier in finite time with probability one. As a consequence, $\mathbb{P}(\tau_\phi > N)$ vanishes uniformly in ϕ when N tends to infinity: for any $\epsilon > 0$, there exists $N_\epsilon > 1$, such that for all ϕ in $[0, 1)$, $k_\phi(B \cap [N_\epsilon, \infty)) < \epsilon$.

ii. Convergence in law of the first-passage times: Considering a sequence $\phi_n \to \phi$, define the first-passage time $\tau_n = \inf\{t > \phi_n \,|\, U_t > L(t)\}$ with $U_{\phi_n^+} = L(\phi_n) - (l - r)$, where l is the threshold value and r is the reset value with $r < l$. The shifted first-passage time $\tau_n - (\phi_n - \phi)$ has the same law as $\tau_n' = \inf\{t > \phi \,|\, U_t > L_n(t)\}$ for a shifted barrier $L_n(t) = L\big(t - (\phi - \phi_n)\big)$ with initial condition $U_{\phi^+} = L_n(\phi) - (l - r)$. Since L is homogeneously Hölder continuous for a given exponent $H > 0$, for any $N > 0$, there exists c_N such that

$$\forall \psi, \phi \in [0, N], |\psi - \phi| \le \delta \quad \Rightarrow \quad |L(\psi) - L(\phi)| \le c_N \delta^H. \tag{A.5}$$

By definition of ϕ_n, for any $\delta > 0$, there is $N > 0$ such that for all $n > N$, $|\phi - \phi_n| < \delta$, which implies

$$|L(\psi) - L_n(\psi)| = |L(t) - L\big(\psi - (\phi - \phi_n)\big)| \le c_N |\phi - \phi_n|^H. \tag{A.6}$$

Thus L_n uniformly converges to L on every compact of $\mathbb{R}^+$, which implies that τ_n' converges in law toward τ by proposition 2. Now, since the deterministic quantity $|\tau_n' - \tau_n| = |\phi - \phi_n| \to 0$ when n tends to infinity, the sequence τ_n converges in law toward τ_ϕ (Billingsley, 1999).

iii. Feller property: For ϕ in $[0, 1)$, define the open sets $B' = \pi^{-1}(A) \cap [\phi, N_{\epsilon/4})$, where $N_{\epsilon/4}$ is defined as in step 1. Since for all ϕ in $[0, 1)$, any open set of $[0, N_{\epsilon/4}]$ is a τ_ϕ-continuity set, convergence in law is equivalent to convergent in distribution: there exists $n_\epsilon > 0$ such that

$$\forall n > n_\epsilon, \quad \left|\mathbb{P}\left(\tau_n \in B'\right) - \mathbb{P}\left(\tau \in B'\right)\right| \leq \epsilon/2. \tag{A.7}$$

Recapitulating, we have for all $n > n_\epsilon$:

$$\begin{aligned}
\left|\kappa_{\phi_n}(A) - \kappa_\phi(A)\right| &\leq \left|k_{\phi_n}\left(\pi^{-1}(A) \cap [\phi_n, +\infty)\right)\right. \\
&\qquad \left. - k_\phi\left(\pi^{-1}(A) \cap [\phi, +\infty)\right)\right|, && \text{(A.8)} \\
&\leq \left|k_{\phi_n}(B') - k_\phi(B')\right| && \text{(A.9)} \\
&\qquad + \left|k_{\phi_n}(B \setminus B')\right| + \left|k_\phi(B \setminus B')\right|, && \text{(A.10)} \\
&\leq \epsilon/2 + \epsilon/4 + \epsilon/4 \;=\; \epsilon, && \text{(A.11)}
\end{aligned}$$

proving the Feller continuity property.

We finally demonstrate the property claimed in the main text that implies the irreducibility of the FPMC.

Property 2. The Markov chain $\mathcal{T}$ is irreducible, that is, for all $B \in \mathcal{B}(\mathbb{S})$

$$\text{if} \quad \exists \psi \in \mathbb{S}, \quad \kappa_\psi(B) > 0, \quad \text{then} \quad \forall \phi \in \mathbb{S}, \quad \kappa_\phi(B) > 0. \tag{A.12}$$

Property 2 is very strong stating that if a set can be attained with non zero probability from a given starting time t, it can be reached with positive probability from all other starting times s. This clearly implies the uniqueness of an invariant measure, whose existence stems from the Feller property. Notice that the property of irreducibility holds for measurable sets in $\mathcal{B}(\mathbb{S})$ that can very well be of Lebesgue measure zero in the case of singular kernels κ_ψ.

Proof. *i. Formulation of a sufficient condition:* Suppose that a measurable set $B \in \mathcal{B}(\mathbb{S})$ is such that $\kappa_\psi(B) > 0$, which is equivalent to $k_\psi(\pi^{-1}(B) \cap [\psi, +\infty)) > 0$, where $\pi^{-1}(B) \cap [\psi, +\infty)$ is the unwrapped version of B on $[\psi, +\infty)$. Extending L periodically on $\mathbb{R}$, it is enough to show that if $k_\psi(B) > 0$ with $B \in \mathcal{B}([\psi, +\infty))$, then for all $s < \psi$, we have $k_\phi(B) > 0$. We can further suppose that $\inf B > \psi$. Indeed, if $\inf B = \psi$ and $k_\psi(B) > 0$, we

can always choose $a > 0$ small enough so that $k_\psi(B \cap [\psi + a, +\infty)) > 0$. Consider the OU process U on $[\psi, 1\infty)$ with initial condition $U_\psi = L(\psi) - (l - r)$. There is $a > 0$ such that for $\inf_{0 \le t - \psi \le a} L(t) = l' > U_\psi = L(\psi) - (L - r)$ and thus $\mathbb{P}(\tau < a + \psi) < \mathbb{P}\left(\tau' < a + \psi\right)$ with $\tau = \inf\{t > \psi \,|\, U_t > L(t)\}$ and $\tau' = \inf\{t > \psi \,|\, U_t > l'\}$. As $\lim_{a \to 0} \mathbb{P}\left(\tau' < a + \psi\right) = 0$, denoting $k_\psi(B \cap [\phi, +\infty)) = \delta > 0$, there is $a > 0$ such that $\mathbb{P}\left(\tau' < a + \psi\right) < \delta/2$ so that

$$\begin{aligned} k_\phi(B \cap [a, +\infty)) &= k_\phi(B \cap [\phi, +\infty)) - k_s(B \cap [0, a]) > \delta - \delta/2 \\ &= \delta/2 > 0. \end{aligned} \tag{A.13}$$

ii. Transition kernel of the killed diffusion: Consider the usual OU process U on $[\phi, +\infty)$ with initial condition $U_\phi = L(\phi) - (l - r)$. For each barrier L, we can define a killed Markov process $\bar{U}$ by sending a trajectory of U to an absorbing cemetery state Δ when it first hits the boundary L:

$$\bar{U}_\psi = \begin{cases} U_\psi & \text{if } \psi < \tau \\ \Delta & \text{if } \psi \ge \tau \end{cases}.$$

We can generalize the notion of transition kernel for killed diffusion. If the barrier is regular enough, such a transition kernel $\bar{k}$ is obtained through solving the Fokker-Planck equation with absorbing condition on the barrier. In the case of a general boundary, we can always define the transition kernel as

$$\bar{k}(\phi, x; \psi, y) = k(\phi, x; \psi, y)\big(1 - \Pi(\phi, x; \psi, y)\big) < k(\phi, x; \psi, y), \tag{A.14}$$

where k is the transition kernel for the free process U and $\Pi(\phi, x; \psi, y)$ is the probability that a bridge process between (ϕ, x) and (ψ, y) crosses the barrier

$$\Pi(\phi, x; \psi, y) = \mathbb{P}\big(\sup_{\phi \le v \le \psi} (U_v - L(v)) > 0 \,|\, U_\phi x < L(\phi), U_\psi = y < L(\psi)\big). \tag{A.15}$$

The kernels $\bar{k}$ define a collection of sub-Markov measures $\bar{k}(\phi, x; \psi, \cdot)$ satisfying

$$\begin{aligned} \int_{-\infty}^{L(\psi)} \bar{k}(\phi, x; \psi, y)\, dy = \mathbb{P}(\tau > t) < 1 \quad \text{with} \\ \tau = \inf\{\psi > \phi \,|\, U_\psi > L(\psi)\}. \end{aligned} \tag{A.16}$$

Moreover, for all $y < L(\psi)$, $\bar{k}(\phi, x; \psi, y)$ is positive since $\Pi(\phi, x; \psi, y) < 1$ as long as $x < L(\phi)$ and $y < L(\psi)$.

iii. Strong Markov property: Suppose as in *i* that $B \in \mathcal{B}([\phi, +\infty))$ with $\inf B = a > \phi$ satisfies $k_\phi(B) > 0$, define for all v, $\phi < v < a$, the bounded measurable function $x \mapsto \mu_v(x; B)$ as

$$\mu_v(x; B) = \mathbb{P}(\tau_x \in B) \quad \text{with} \quad \tau_x = \inf\{\psi > v \,|\, U_\psi > L(\psi), U_v = x\}. \tag{A.17}$$

By the strong Markov property, we have

$$\int_{-\infty}^{L(v)} \bar{k}(\phi, L(\phi) - (l - r); v, x)\mu_v(x; B)\, dx = k_\phi(B) > 0. \tag{A.18}$$

Since $\bar{k}(\phi, L(\phi) - (l - r); v, x)$ is positive finite on $x < L(v)$, $\sup_{x<L(v)} \bar{k}(\phi, L(\phi) - (l - r); v, x) = M < \infty$ and we have

$$M \int_{-N}^{L(v)-\epsilon} \mu_v(x; B)\, dx > 0, \tag{A.19}$$

for $N > 0$ large enough and $\epsilon > 0$ small enough. From there, we deduce that there are $\delta > 0$ and a Lebesgue measurable set $D_\delta \subset [-N, L(v) - \epsilon]$ such that $\mu_v(x; B) > 0$ for all $x \in D$. Finally, since we know from *ii* that for all (ϕ, x) with $x < L(\phi)$, $\bar{k}(\phi, x; v, y) > 0$, we conclude that for all $\phi < t$,

$$k_\phi(B) \geq \int_{D_\delta} \bar{k}(\phi, L(\phi) - (l - r); v, x)\mu_v(x; B)\, dx \tag{A.20}$$

$$\geq \delta \int_{D_\delta} \bar{k}(\phi, L(\phi) - (l - r); v, x)\, dx \; > \; 0. \tag{A.21}$$

Acknowledgments

We are grateful to Peter Thomas and an anonymous reviewer for their constructive and helpful suggestions. This work was partially supported by National Science Foundation Grant EF-0929723.

References

Ahmadian, Y., Pillow, J. W., & Paninski, L. (2010). Efficient Markov chain Monte Carlo methods for decoding neural spike trains. *Neural Computation*, *23*(1), 46–96.

Allain, C., & Cloitre, M. (1991). Characterizing the lacunarity of random and deterministic fractal sets. *Phys. Rev. A*, *44*(6), 3552–3558.

Arnold, L. (1998). *Random dynamical systems*. Berlin: Springer-Verlag.

Bair, W., & Koch, C. (1996). Temporal precision of spike trains in extrastriate cortex of the behaving macaque monkey. *Neural Computation*, *8*(6), 1185–1202.

Berry, M. J., Warland, D. K., & Meister, M. (1997). The structure and precision of retinal spike trains. *Proceedings of the National Academy of Sciences*, *94*(10), 5411–5416.

Billingsley, P. (1999). *Convergence of probability measures* (2nd ed.). New York: Wiley.

Bouleau, N., & Lépingle, D. (1994). *Numerical methods for stochastic processes*. New York: Wiley.

Bressloff, P. C. (2010). Metastable states and quasicycles in a stochastic Wilson-Cowan model of neuronal population dynamics. *Physical Review E*, *82*(5), 051903.

Bryant, H., & Segundo, J. (1976). Spike initiation by transmembrane current: A white-noise analysis. *Journal of Physiology*, *260*(2), 279–314.

Burkitt, A. (2006a). A review of the integrate-and-fire neuron model: 1. Homogeneous synaptic input. *Biological Cybernetics*, *95*(1), 1–19.

Burkitt, A. (2006b). A review of the integrate-and-fire neuron model: 2. Inhomogeneous synaptic input and network properties. *Biological Cybernetics*, *95*(2), 97–112.

Butts, D. A., Weng, C., Jin, J., Yeh, C.-I., Lesica, N. A., Alonso, J.-M., & Stanley, G. B. (2007). Temporal precision in the neural code and the timescales of natural vision. *Nature*, *449*(7158), 92–95.

Cannon, J. R. (1984). *The one-dimensional heat equation*. Reading, MA: Addison-Wesley.

Cecchi, G. A., Sigman, M., Alonso, J.-M., Martínez, L., Chialvo, D. R., & Magnasco, M. O. (2000). Noise in neurons is message dependent. *Proceedings of the National Academy of Sciences*, *97*(10), 5557–5561.

Courant, R., & Hilbert, D. (1962). *Methods of mathematical physics. Vol. 2: Partial differential equations*. New York: Wiley.

Daoudi, K., Lévy Véhel, J., & Meyer, Y. (1998). Construction of continuous functions with prescribed local regularity. *Constr. Approx.*, *14*(3), 349–385.

Dayan, P., & Abbott, L. (2001). *Theoretical neuroscience: Computational and mathematical modeling of neural systems*. Cambridge, MA: MIT Press.

Dobrunz, L. E., & Stevens, C. F. (1997). Heterogeneity of release probability, facilitation, and depletion at central synapses. *Neuron*, *18*(6), 995–1008.

Dorval, A. D., & White, J. A. (2005). Channel noise is essential for perithreshold oscillations in entorhinal stellate neurons. *Journal of Neuroscience*, *25*(43), 10025–10028.

Faisal, A. A., Selen, L. P. J., & Wolpert, D. M. (2008). Noise in the nervous system. *Nat. Rev. Neurosci.*, *9*(4), 292–303.

Fellous, J.-M., Houweling, A. R., Modi, R. H., Rao, R., Tiesinga, P., & Sejnowski, T. J. (2001). Frequency dependence of spike timing reliability in cortical pyramidal cells and interneurons. *Journal of Neurophysiology*, *85*(4), 1782–1787.

Fellous, J.-M., Rudolph, M., Destexhe, A., & Sejnowski, T. J. (2003). Synaptic background noise controls the input/output characteristics of single cells in an in vitro model of in vivo activity. *Neuroscience*, *122*(3), 811–829.

Fellous, J.-M., Tiesinga, P.H.E., Thomas, P. J., & Sejnowski, T. J. (2004). Discovering spike patterns in neuronal responses. *Journal of Neuroscience*, *24*(12), 2989–3001.

Fourcaud-Trocmé, N., Hansel, D., van Vreeswijk, C., & Brunel, N. (2003). How spike generation mechanisms determine the neuronal response to fluctuating inputs. *Journal of Neuroscience*, *23*(37), 11628–11640.

Gaines, J. G., & Lyons, T. J. (1997). Variable step size control in the numerical solution of stochastic differential equations. *SIAM J. Appl. Math.*, *57*(5), 1455–1484.

Galán, R. F., Fourcaud-Trocmé, N., Ermentrout, G. B., & Urban, N. N. (2006). Correlation-induced synchronization of oscillations in olfactory bulb neurons. *Journal of Neuroscience*, *26*(14), 3646–3655.

Gerstner, W., & Kistler, W. (2002). *Spiking neuron models: Single neurons, populations, plasticity*. Cambridge: Cambridge University Press.

Haas, J. S., & White, J. A. (2002). Frequency selectivity of layer II stellate cells in the medial entorhinal cortex. *Journal of Neurophysiology*, *88*(5), 2422–2429.

Häggström, O. (2002). *Finite Markov chains and algorithmic applications*. Cambridge: Cambridge University Press.

Hernández-Lerma, O., & Lasserre, J. B. (2003). *Markov chains and invariant probabilities*. Basel: Birkhäuser Verlag.

Honeycutt, R. L. (1992). Stochastic Runge-Kutta algorithms. I. White noise. *Phys. Rev. A*, *45*(2), 600–603.

Karr, A. F. (1975). Weak convergence of a sequence of Markov chains. *Probability Theory and Related Fields*, *33*, 41–48. 10.1007/BF00539859.

Lajoie, G., Lin, K. K., & Shea-Brown, E. (2013). Chaos and reliability in balanced spiking networks with temporal drive. *Physical Review E*, *87*(5), 052901.

Lansky, P., & Ditlevsen, S. (2008). A review of the methods for signal estimation in stochastic diffusion leaky integrate-and-fire neuronal models. *Biol. Cybern.*, *99*(4–l5), 253–262.

Lehmann, A. (2002). Smoothness of first passage time distributions and a new integral equation for the first passage time density of continuous Markov processes. *Adv. Appl. Probab.*, *34*(4), 869–887.

Lévy, P. (1948). *Processus stochastiques et mouvement brownien. Suivi d'une note de M. Loève*. Paris: Gauthier-Villars.

Lin, K., Shea-Brown, E., & Young, L.-S. (2009a). Reliability of coupled oscillators. *Journal of Nonlinear Science*, *19*(5), 497–545.

Lin, K., Shea-Brown, E., & Young, L.-S. (2009b). Spike-time reliability of layered neural oscillator networks. *Journal of Computational Neuroscience*, *27*(1), 135–160.

Linz, P. (1985). *Analytical and numerical methods for Volterra equations*. Philadelphia: Society for Industrial and Applied Mathematics.

London, M., Roth, A., Beeren, L., Hausser, M., & Latham, P. E. (2010). Sensitivity to perturbations in vivo implies high noise and suggests rate coding in cortex. *Nature*, *466*(7302), 123–127.

Mainen, Z., & Sejnowski, T. (1995). Reliability of spike timing in neocortical neurons. *Science*, *268*(5216), 1503–1506.

Mandelbrot, B. B. (1982). *The fractal geometry of nature*. New York: W. H. Freeman.

Mandelbrot, B. B., & Ness, J. W. V. (1968). Fractional Brownian motions, fractional noises and applications. *SIAM Review*, *10*(4), 422–437.

McDonnell, M. D., & Ward, L. M. (2011). The benefits of noise in neural systems: Bridging theory and experiment. *Nat. Rev. Neurosci.*, *12*(7), 415–426.

Metropolis, N., & Ulam, S. (1949). The Monte Carlo method. *J. Amer. Statist. Assoc.*, *44*, 335–341.

Norris, J. R. (1998). *Markov chains*. Cambridge: Cambridge University Press.

Paninski, L., Haith, A., & Szirtes, G. (2008). Integral equation methods for computing likelihoods and their derivatives in the stochastic integrate-and-fire model. *Journal of Computational Neuroscience*, *24*, 69–79. 10.1007/s10827-007-0042-x.

Park, C., & Paranjape, S. R. (1974). Probabilities of Wiener paths crossing differentiable curves. *Pacific J. Math.*, *53*, 579–583.

Park, C., & Schuurmann, F. J. (1976). Evaluations of barrier-crossing probabilities of Wiener paths. *J. Appl. Probability*, *13*(2), 267–275.

Platen, E. (1999). *An introduction to numerical methods for stochastic differential equations* (Research Paper Series 6). Sydney: Quantitative Finance Research Centre, University of Technology.

Plesser, H. E., & Geisel, T. (1999). Markov analysis of stochastic resonance in a periodically driven integrate-and-fire neuron. *Phys. Rev. E*, *59*(6), 7008–7017.

Poggio, T., & Marr, D. (1977). From understanding computation to understanding neural circuitry. *Neurosciences Res. Prog. Bull.*, *15*, 470–488.

Redner, S. (2007). *A guide to first-passage processes*. Cambridge: Cambridge University Press.

Reich, D. S., Victor, J. D., Knight, B. W., Ozaki, T., & Kaplan, E. (1997). Response variability and timing precision of neuronal spike trains in vivo. *Journal of Neurophysiology*, *77*(5), 2836–2841.

Ricciardi, L. M., & Sacerdote, L. (1979). The Ornstein-Uhlenbeck process as a model for neuronal activity. *Biological Cybernetics*, *35*(1), 1–9.

Ricciardi, L. M., & Sato, S. (1988). First-passage-time density and moments of the Ornstein-Uhlenbeck process. *J. Appl. Probab.*, *25*(1), 43–57.

Rieke, F., Warland, D., de Ruyter van Steveninck, R., & Bialek, W. (1999). *Spikes: Exploring the neural code (computational neuroscience)*. Cambridge, MA: MIT Press.

Robert, C. P., & Casella, G. (2004). *Monte Carlo statistical methods* (2nd ed.). New York: Springer-Verlag.

Römisch, W., & Winkler, R. (2006). Stepsize control for mean-square numerical methods for stochastic differential equations with small noise. *SIAM J. Sci. Comput.*, *28*(2), 604–625.

Rosenblum, M., Pikovsky, A., & Kurths, J. (1996). Phase synchronization of chaotic oscillators. *Physical Review Letters*, *76*(11), 1804–1807.

Rudolph, M., & Destexhe, A. (2001). Do neocortical pyramidal neurons display stochastic resonance? *Journal of Computational Neuroscience*, *11*(1), 19–42.

Shimokawa, T., Pakdaman, K., & Sato, S. (1999). Time-scale matching in the response of a leaky integrate-and-fire neuron model to periodic stimulus with additive noise. *Phys. Rev. E*, *59*(3), 3427–3443.

Shimokawa, T., Pakdaman, K., Takahata, T., Tanabe, S., & Sato, S. (2000). A first-passage-time analysis of the periodically forced noisy leaky integrate-and-fire model. *Biological Cybernetics*, *83*, 327–340. 10.1007/s004220000156.

Siegert, A. J. F. (1951). On the first passage time probability problem. *Physical Rev.*, *81*, 617–623.

Sirovich, L., & Knight, B. (2011). Spiking neurons and the first passage problem. *Neural Computation*, *23*(7), 1675–1703.

Stacey, W. C., & Durand, D. M. (2001). Synaptic noise improves detection of subthreshold signals in hippocampal Ca1 neurons. *Journal of Neurophysiology*, *86*(3), 1104–1112.

Stanley, H. E., & Meakin, P. (1988). Multifractal phenomena in physics and chemistry. *Nature*, *335*(6189), 405–409.

Stevens, C. F., & Zador, A. M. (1998). Input synchrony and the irregular firing of cortical neurons. *Nat. Neurosci.*, *1*(3), 210–217.

Stewart, W. J. (2009). *Probability, Markov chains, queues, and simulation*. Princeton, NJ: Princeton University Press.

Stroock, D. W., & Varadhan, S. R. S. (2006). *Multidimensional diffusion processes*. Berlin: Springer-Verlag.

Taillefumier, T., & Magnasco, M. (2010). A fast algorithm for the first-passage times of Gauss-Markov processes with Hölder continuous boundaries. *Journal of Statistical Physics*, *140*(6), 1–27.

Taillefumier, T., & Magnasco, M. (2013). A phase transition in the first passage of a Brownian process through a fluctuating boundary with implications for neural coding. *Proceedings of the National Academy of Sciences*.

Taillefumier, T., & Touboul, J. (2011). Multiresolution Hilbert approach to multidimensional Gauss-Markov processes. *International Journal of Stochastic Analysis*, *2011*.

Taillefumier, T., Touboul, J., & Magnasco, M. (2012). Exact event-driven implementation for recurrent networks of stochastic perfect integrate-and-fire neurons. *Neural Computation*, *24*(12), 3145–3180.

Thomas, P. J., Tiesinga, P. H. E., Fellous, J.-M., & Sejnowski, T. J. (2003). Reliability and bifurcation in neurons driven by multiple sinusoids. *Neurocomputing*, *52–54*, 955–961.

Tiesinga, P. H. E. (2002). Precision and reliability of periodically and quasiperiodically driven integrate-and-fire neurons. *Phys. Rev. E*, *65*, 041913.

Toups, J. V., Fellous, J.-M., Thomas, P. J., Sejnowski, T. J., & Tiesinga, P. H. E. (2011). Finding the event structure of neuronal spike trains. *Neural Computation*, *23*(9), 2169–2208.

Toups, J. V., Fellous, J.-M., Thomas, P. J., Sejnowski, T. J., & Tiesinga, P. H. E. (2012). Multiple spike time patterns occur at bifurcation points of membrane potential dynamics. *PLoS Comput Biol*, *8*(10):e1002615 EP.

Tricomi, F. G. (1985). *Integral equations*. New York: Dover.

Uhlenbeck, G. E., & Ornstein, L. S. (1930). On the theory of the Brownian motion. *Phys. Rev.*, *36*(5), 823–841.

Wang, L., & Pötzelberger, K. (2007). Crossing probabilities for diffusion processes with piecewise continuous boundaries. *Methodol. Comput. Appl. Probab.*, *9*(1), 21–40.

White, J. A., Rubinstein, J. T., & Kay, A. R. (2000). Channel noise in neurons. *Trends in Neurosciences*, *23*(3), 131–137.

Received January 3, 2013; accepted November 13, 2013.