

6

Effects of Noise on Nonlinear Dynamics

André Longtin

6.1 Introduction

The influence of noise on nonlinear dynamical systems is a very important area of research, since all systems, physiological and other, evolve in the presence of noisy driving forces. It is often thought that noise has only a blurring effect on the evolution of dynamical systems. It is true that that can be the case, especially for so-called “observational” or “measurement” noise, as well as for linear systems. However, in nonlinear systems with dynamical noise, i.e., with noise that acts as a driving term in the equations of motion, noise can drastically modify the deterministic dynamics. For example, the hallmark of nonlinear behavior is the bifurcation, which is a qualitative change in the phase space motion when the value of one or more parameter changes. Noise can drastically modify the dynamics of a deterministic dynamical system. It can make the determination of bifurcation points very difficult, even for the simplest bifurcations. Noise can shift bifurcation points or induce behaviors that have no deterministic counterpart, through what are known as noise-induced transitions (Horsthemke and Lefever 1984). The combination of noise and nonlinear dynamics can also produce time series that are easily mistakable for deterministic chaos. This is especially true in the vicinity of bifurcation points, where the noise has its greatest influence.

This chapter considers these issues starting from a basic level of description: the stochastic differential equation. It discusses sources of noise, and shows how noise, or “stochastic processes,” can be coupled to deterministic differential equations. It also discusses analytical tools to deal with stochastic differential equations, as well as simple methods to numerically integrate such equations. It then focuses in greater detail on trying to pinpoint a Hopf bifurcation in a real physiological system, which will lead to the notion of a noise-induced transition. This system is the human pupil light reflex. It has been studied by us and others both experimentally and

theoretically. It is also the subject of Chapter 9, by John Milton; he has been involved in the work presented here, along with Jelte Bos (who carried out some of the experiments at the Free University of Amsterdam) and Michael Mackey. This chapter then considers noise-induced firing in excitable systems, and how noise interacts with deterministic input such as a sine wave to produce various forms of “stochastic phase locking.”

Noise is thought to arise from the action of a large number of variables. In this sense, it is usually understood that noise is high-dimensional. The mathematical analysis of noise involves associating a random variable with the high-dimensional physical process causing the noise. For example, for a cortical cell receiving synaptic inputs from ten thousand other cells, the ongoing synaptic bombardment may be considered as a source of current noise. The firing behavior of this cell may then be adequately described by assuming that the model differential equations governing the excitability of this cell (e.g., Hodgkin–Huxley-type equations) are coupled to a random variable describing the properties of this current noise.

Although we may consider these synaptic inputs as noise, the cell may actually make more sense of it than we can, such as in temporal and spatial coincidences of inputs. Hence, one person’s noise may be another person’s information: It depends ultimately on the phenomena you are trying to understand. This explains in part why there has been such a thrust in the last decades to discover simple low-dimensional deterministic laws (such as chaos) governing observed noisy fluctuations.

From the mathematical standpoint, noise as a random variable is a quantity that fluctuates aperiodically in time. To be a useful quantity to describe the real world, this random variable should have well-defined properties that can be measured experimentally, such as a distribution of values (a density) with a mean and other moments, and a two-point correlation function. Thus, although the variable itself takes on a different set of values every time we look at it or simulate it (i.e., for each of its “realizations”), its statistical and temporal properties remain constant. The validity of these assumptions in a particular experimental setting must be properly assessed, for example by verifying that certain stationarity criteria are satisfied.

One of the difficulties with modeling noise is that in general, we do not have access to the noise variable itself. Rather, we usually have access to a state variable of a system that is perturbed by one or more sources of noise. Thus, one may have to begin with assumptions about the noise and its coupling to the dynamical state variables. The accuracy of these assumptions can later be assessed by looking at the agreement of the predictions of the resulting model with the experimental data.

6.2 Different Kinds of Noise

There is a large literature on the different kinds of noise that can arise in physical and physiological systems. Excellent references on this subject can be found in Gardiner (1985) and Horsthemke and Lefever (1984). An excellent reference for noise at the cellular level is the book by DeFelice (1981). These books provide background material on thermal noise (also known as Johnson–Nyquist noise, fluctuations that are present in any system due to its temperature being higher than absolute zero), shot noise (due to the motion of individual charges), $1/f^\alpha$ noise (one-over- f noise or flicker noise, the physical mechanisms of which are still the topic of whole conferences), Brownian motion, Ornstein–Uhlenbeck colored noise, and the list goes on.

In physiology, an important source of noise consists of conductance fluctuations of ionic channels, due to the (apparently) random times at which they open and close. There are many other sources of noise associated with channels (DeFelice 1981). In a nerve cell, noise from synaptic events can be more important than the intrinsic sources of noise such as conductance fluctuations. Electric currents from neighboring cells or axons are a form of noise that not only affects recordings through measurement noise but also a cell's dynamics. There are fluctuations in the concentrations of ions and other chemicals forming the milieu in which the cells live. These fluctuations may arise on a slower time scale than the other noises mentioned up to now.

The integrated electrical activity of nerve cells produces the electroencephalogram (EEG) pattern with all its wonderful classes of fluctuations. Similarly, neuromuscular systems are complex connected systems of neurons, axons, and muscles, each with its own sources of noise. Fluctuations in muscle contraction strength are dependent to a large extent on the firing patterns of the motorneurons that drive them. All these examples should convince you that the modeling of noise requires a knowledge of the basic physical processes governing the dynamics of the variables we measure.

There is another kind of distinction that must be applied, namely, that between observational, additive, and multiplicative noise (assuming a stationary noise). In the case of observational noise, the dynamical system evolves deterministically, but our measurements on this system are contaminated by noise. For example, suppose a one-dimensional dynamical system is governed by

$$\frac{dx}{dt} = f(x, \mu), \quad (6.1)$$

where μ is a parameter. Then observational noise corresponds to the measurement of $y(t) \equiv x(t) + \xi(t)$, where $\xi(t)$ is the observational noise process. The measurement y , but not the evolution of the system x , is affected by the presence of noise. While this is often an important source of noise with which analyses must contend, and the simplest to deal with mathemati-

cally, it is also the most boring form of noise in a physical system: It does not give rise to any new effects.

One can also have additive sources of noise. These situations are characterized by noise that is independent of the precise state of the system:

$$\frac{dx}{dt} = f(x, \mu) + \xi(t). \quad (6.2)$$

In other words, the noise is simply added to the deterministic part of the dynamics. Finally, one can have multiplicative noise, in which the noise is dependent on the value of one or many state variables. Suppose that f is separable into a deterministic part and a stochastic part that depends on x :

$$\frac{dx}{dt} = h(x) + g(x)\xi(t). \quad (6.3)$$

We then have the situation where the effect of the noise term will depend on the value of the state variable x through $g(x)$. Of course, a given system may have one or more noise sources coupled in one or more of these ways. We will discuss such sources of noise in the context of our experiments on the pupil light reflex and our study of stochastic phase locking.

6.3 The Langevin Equation

Modeling the precise effects of noise on a dynamical system can be very difficult, and can involve a lot of guesswork. However, one can already gain insight into the effect of noise on a system by coupling it additively to **Gaussian white noise**. This noise is a mathematical construct that approximates the properties of many kinds of noise encountered in experimental situations. It is Gaussian distributed with zero mean and autocorrelation $\langle \xi(t)\xi(s) \rangle = 2D\delta(t-s)$, where δ is the Dirac *delta* function. The quantity $D \equiv \sigma^2/2$ is usually referred to as the **intensity of the Gaussian white noise** (the actual intensity with the autocorrelation scaling used in our example is $2D$). Strictly speaking, the variance of this noise is infinite, since it is equal to the autocorrelation at zero lag (i.e. at $t = s$). However, its intensity is finite, and σ times the square root of the time step will be the standard deviation of Gaussian random numbers used to numerically generate such noise (see below).

The Langevin equation refers to the stochastic differential equation obtained by adding Gaussian white noise to a simple first-order linear dynamical system with a single stable fixed point:

$$\frac{dx}{dt} = -\alpha x + \xi(t). \quad (6.4)$$

This is nothing but equation (6.2) with $f(x, \mu) = -\alpha x$, and $\xi(t)$ given by the Gaussian white noise process we have just defined. The stochastic

process $x(t)$ is also known as Ornstein–Uhlenbeck noise with correlation time $1/\alpha$, or as lowpass-filtered Gaussian white noise. A noise process that is not white noise, i.e. that does not have a delta-function autocorrelation, is called “colored noise”. Thus, the exponentially correlated Ornstein–Uhlenbeck noise is a colored noise. The probability density of this process is given by a Gaussian with zero mean and variance $\sigma^2/2\alpha$ (you can verify this using the Fokker–Planck formalism described below). In practical work, it is important to distinguish between this variance, and the intensity D of the Gaussian white noise process used to produce it. It is in fact common to plot various quantities of interest for a stochastic dynamical system as a function of the intensity D of the white noise used in that dynamical system, no matter where it appears in the equations.

The case in which the deterministic part is nonlinear yields a **nonlinear Langevin equation**, which is the usual interesting case in mathematical physiology. One can simulate a nonlinear Langevin equation with deterministic flow $h(x, t, \mu)$ and a coefficient $g(x, t, \mu)$ for the noise process using various stochastic numerical integration methods of different orders of precision (see Kloeden, Platen, and Schurz 1991 for a review). For example, one can use a simple “stochastic” Euler–Maruyama method with fixed time step Δt (stochastic simulations are much safer with fixed step methods). Using the definition of Gaussian white noise $\xi(t)$ as the derivative of the Wiener process $W(t)$ (the Wiener process is also known as “Brownian motion”), this method can be written as

$$x(t + \Delta t) = x(t) + \Delta t h(x, t, \mu) + g(x, t, \mu) \Delta W_n, \quad (6.5)$$

where the $\{\Delta W_n\}$ are “increments” of the Wiener process. These increments can be shown to be independent Gaussian random variables with zero mean and standard deviation given by $\sigma\sqrt{\Delta t}$. Because the Gaussian white noise $\xi(t)$ is a function that is nowhere differentiable, one must use another kind of calculus to deal with it (the so-called stochastic calculus; see Horsthemke and Lefever 1984; Gardiner 1985). One consequence of this fact is the necessity to exercise caution when performing a nonlinear change of variables on stochastic differential equations: The “Stratonovich” stochastic calculus obeys the laws of the usual deterministic calculus, but the “Ito” stochastic calculus does not. One thus has to associate an Ito or Stratonovich interpretation with a stochastic differential equation before performing coordinate changes. Also, for a given stochastic differential equation, the properties of the random process $x(t)$ such as its moments will depend on which calculus is assumed. Fortunately, there is a simple transformation between the Ito and Stratonovich forms of the stochastic differential equation. Further, the properties obtained with both calculi are identical when the noise is additive.

It is also important to associate the chosen calculus with a proper integration method (Kloeden, Platen, and Schurz 1991). For example, an explicit Euler–Maruyama scheme is an Ito method, so numerical results

with this method will agree with any theoretical results obtained from an analysis of the Ito interpretation of the stochastic differential equation, or of its equivalent Stratonovich form. In general, the Stratonovich form is best suited to model “real colored noise” and its effects in the limit of vanishing correlation time, i.e. in the limit where colored noise is allowed to become white after the calculation of measurable quantities.

Another consequence of the stochastic calculus is that in the Euler–Maruyama numerical scheme, the noise term has a magnitude proportional to the square root of the time step, rather than to the time step itself. This makes this method an “order $\frac{1}{2}$ ” method, which converges more slowly than the Euler algorithm for deterministic differential equations. Higher-order methods are also available (Fox, Gatland, Roy, and Vemuri 1988; Mannella and Palleschi 1989; Honeycutt 1992); some are used in the computer exercises associated with this chapter and with the chapter on the pupil light reflex (Chapter 9). Such methods are especially useful for stiff stochastic problems, such as the Hodgkin–Huxley or FitzHugh–Nagumo equations with stochastic forcing, where one usually uses an adaptive method in the noiseless case, but is confined to a fixed step method with noise. Stochastic simulations usually require multiple long runs (“realizations”: see below) to get good averaging, and higher-order methods are useful for that as well.

The Gaussian random numbers ΔW_n are generated in an uncorrelated fashion, for example by using a pseudorandom number generator in combination with the Box–Müller algorithm. Such algorithms must be “seeded,” i.e., provided with an initial condition. They will then output numbers with very small correlations between themselves. A simulation that uses Gaussian numbers that follow one initial seed is called a realization. In certain problems, it is important to repeat the simulations using M different realizations of N points (i.e., M with different seeds). This performs an average of the stochastic differential equation over the distribution of the random variable. It also serves to reduce the variance of various statistical quantities used in a simulation (such as power spectral amplitudes). In the case of very long simulations, it also avoids problems associated with the finite period of the random number generator.

A stochastic simulation yields a different trajectory for each different seed. It is possible also to describe the action of noise from another point of view, that of probability densities. One can study, for example, how an ensemble of initial conditions, characterized by a density, propagates under the action of the stochastic differential equation. One can study also the probability density of measuring the state variable between x and $x + dx$ at a given time. The evolution of this density is governed by a deterministic partial differential equation in the density variable $\rho(x, t)$, known as the Fokker–Planck equation. In one dimension, this equation is

$$\frac{\partial \rho}{\partial t} = \frac{1}{2} \frac{\partial^2 [g^2(x)\rho]}{\partial x^2} - \frac{\partial [h(x)\rho]}{\partial x}. \quad (6.6)$$

Setting the left-hand side to zero and solving the remaining ordinary differential equation yields the asymptotic density for the stochastic differential equation, $\rho^* = \rho(x, \infty)$. This corresponds to the probability density of finding the state variable between x and $x + dx$ once transients have died out, i.e., in the long-time limit. Since noisy perturbations cause transients, this long-time density somehow characterizes not only the deterministic attractors, but also the noise-induced transients around these attractors. For simple systems, it is possible to calculate ρ^* , and sometimes even $\rho(x, t)$ (this can always be done for linear systems, even with time-dependent coefficients). However, for nonlinear problems, in general one can at best approximate ρ^* . For delay differential equations such as the one studied in the next section, the Fokker–Planck formalism breaks down. Nevertheless, it is possible to calculate ρ^* numerically, and even understand some of its properties analytically (Longtin, Milton, Bos, and Mackey 1990; Longtin 1991a; Guillouzic, L’Heureux, and Longtin 1999).

A simple example of a nonlinear Langevin equation is

$$\frac{dx}{dt} = x - x^3 + \xi(t), \quad (6.7)$$

which models the overdamped noise-driven motion of a particle in a bistable potential. The deterministic part of this system has three fixed points, an unstable one at the origin and stable ones at ± 1 . For small noise intensity D , the system spends a long time fluctuating on either side of the origin before making a switch to the other side, as shown in Figure 6.1. Increasing the noise intensity increases the frequency of the switches across the origin. At the same time, the asymptotic probability density broadens around the stable points, and the probability density in a neighborhood of the (unstable) origin increases; the vicinity of the origin is thus stabilized by noise. One can actually calculate this asymptotic density exactly for this system using the Fokker–Planck formalism (try it! the answer is $\rho(x) = C \exp[(x^2 - x^4)/2D]$, where C is a normalization constant). Also, because this is a one-dimensional system with additive noise, the maxima of the density are always located at the same place as for the deterministic case. The maxima are not displaced by noise, and no new maxima are created; in other words, there are no noise-induced states. This is not always the case for multiplicative noise, or for additive or multiplicative noise in higher dimensions.

6.4 Pupil Light Reflex: Deterministic Dynamics

We illustrate the effect of noise on nonlinear dynamics by first considering how noise alters the behavior of a prototypical physiological control system. The pupil light reflex, which is the focus of Chapter 9, is a negative feedback control system that regulates the amount of light falling on the retina. The

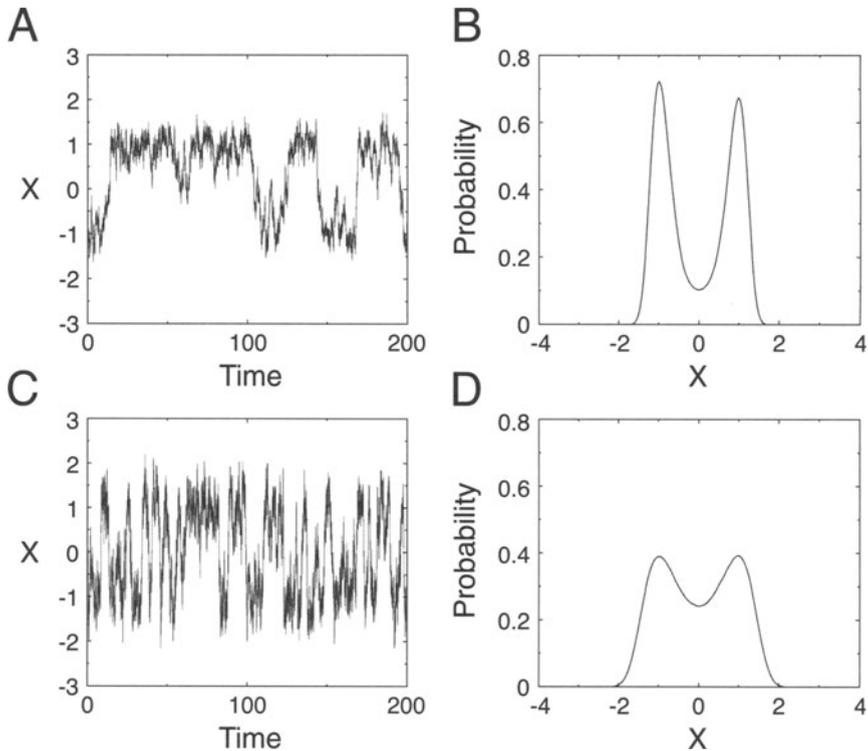


Figure 6.1. Realizations of equation (6.7) at (A) low noise intensity $D = 0.5$ and (C) high noise intensity $D = 1.0$. The corresponding normalized probability densities are shown in (B) and (D) respectively. These densities were obtained from 30 realizations of 400,000 iterates; the integration time step for the stochastic Euler method is 0.005.

pupil is the hole in the middle of the colored part of the eye called the iris. If the ambient light falling on the pupil increases, the reflex response will contract the iris sphincter muscle, thus reducing the area of the pupil and the light flux on the retina. The delay between the variation in light intensity and the variation in pupil area is about 300 msec. A mathematical model for this reflex is developed in Chapter 9. It can be simplified to the following form:

$$\frac{dA}{dt} = -\alpha A + \frac{c}{1 + \left[\frac{A(t-\tau)}{\theta} \right]^n} + k, \quad (6.8)$$

where $A(t)$ is the pupil area, and the second term on the right is a sigmoidal negative feedback function of the area at a time $\tau = 300$ msec in the past. Also, α, θ, c, k are constants, although in fact, c and k fluctuate noisily. The parameter n controls the steepness of the feedback around the fixed point,

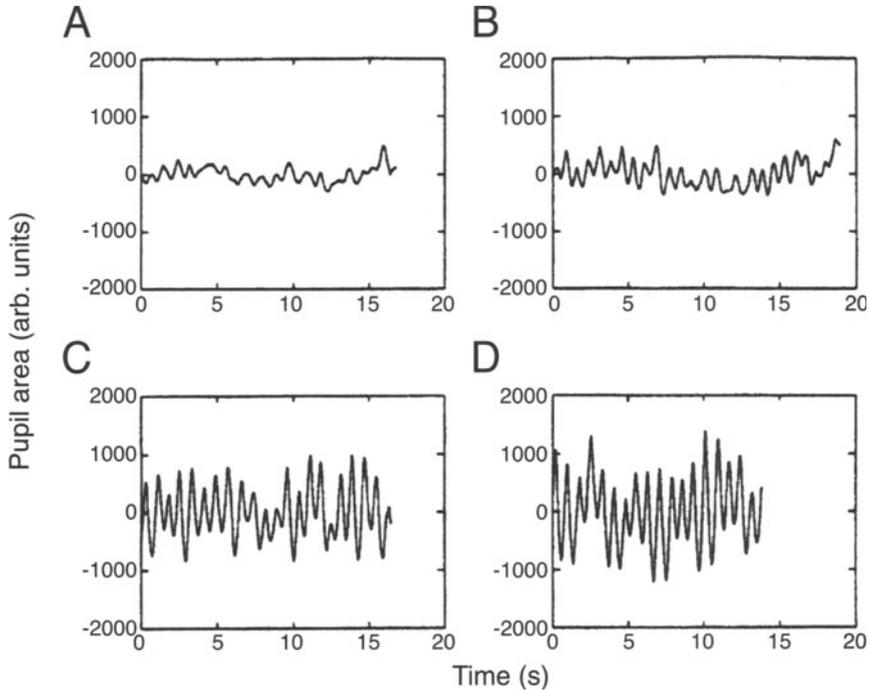


Figure 6.2. Experimental time series of pupil area measured on a human subject for four different values of the feedback gain. Gain values are (A) 1.41, (B) 2.0, (C) 2.82, and (D) 4.0. From Longtin (1991b).

which is proportional to the feedback gain. If n is increased past a certain value n_o , or the delay τ past a critical delay, the single stable fixed point will become unstable, giving rise to a stable limit cycle (supercritical Hopf bifurcation). It is possible to artificially increase the parameter n in an experimental setting involving humans (Longtin, Milton, Bos, and Mackey 1990). We would expect that under normal operating conditions, the value of n is sufficiently low that no periodic oscillations in pupil area are seen. As n is increased, the deterministic dynamics tell us that the amplitude of the oscillation should start increasing proportionally to $\sqrt{n - n_o}$. Experimental data are shown in Figure 6.2, in which the feedback gain, proportional to n , increases from panel A to D.

What is apparent in this system is that noisy oscillations are seen even at the lowest value of the gain; they are seen even below this value (not shown). In fact, it is difficult to pinpoint a qualitative change in the oscillation waveform as the gain increases. Instead, the amplitude of the noisy oscillation simply increases, and in a sigmoidal fashion rather than a square root fashion. This is not what the deterministic model predicts. It is possible that the aperiodic fluctuations arise through more complicated dynamics

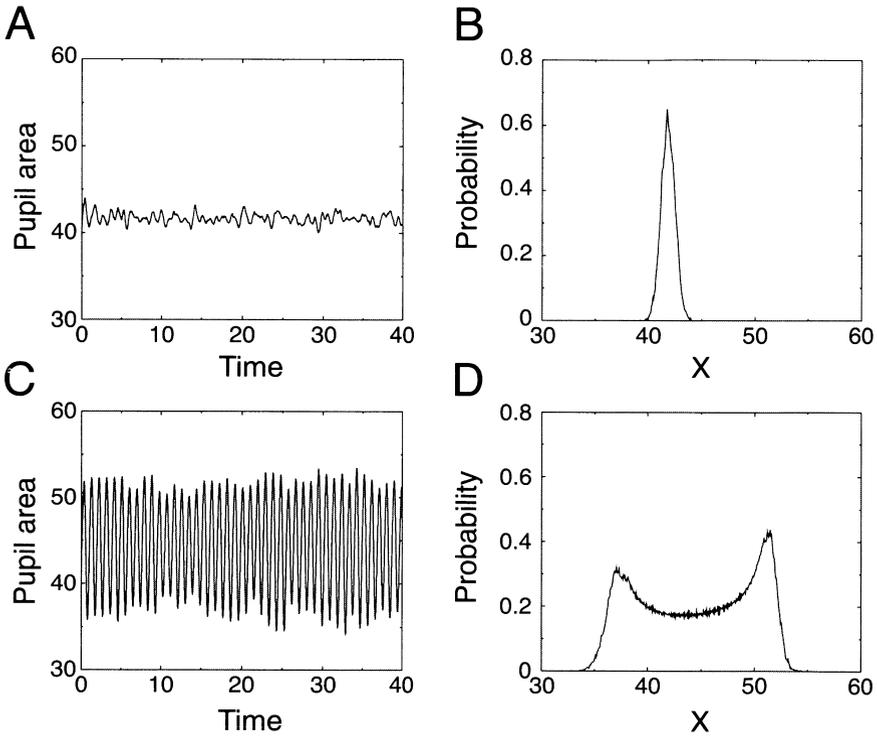


Figure 6.3. Characterization of pupil area fluctuations obtained from numerical simulations of equation (6.8) with multiplicative Gaussian colored noise on the parameter k ; the intensity is $D = 15$, and the noise correlation time is $\alpha^{-1} = 1$. The bifurcation parameter is n ; a Hopf bifurcation occurs (for $D = 0$) at $n = 8.2$. (A) Realization for $n = 4$; the corresponding normalized probability density is shown in (B). (C) and (D) are the same as, respectively, (A) and (B) but for $n = 10$. The densities were computed from 10 realizations, each of duration equal to 400 delays.

(e.g., chaos) in this control system. However, such dynamics are not present in the deterministic model for any combination of parameters and initial conditions. In fact, there are only two globally stable solutions, either a fixed point or a limit cycle.

Another possibility is that noise is present in this reflex, and what we are seeing is the result of noise driving a system in the vicinity of a Hopf bifurcation. This would not be surprising, since the pupil has a well-documented source of fluctuations known as pupillary hippus. It is not known what the precise origin of this noise is, but the following section will show that we can test for certain hypotheses concerning its nature. Incorporating noise into the model can in fact produce fluctuations that vary similarly to those in Figure 6.2 as the feedback gain is increased, as we will now see.

6.5 Pupil Light Reflex: Stochastic Dynamics

We can explain the behaviors seen in Figure 6.2 if noise is incorporated into our model. One can argue, based on the known physiology of this system (see Chapter 9), that noise enters the reflex pathway through the parameters c and k , and causes fluctuations about their mean values \bar{c} and \bar{k} , respectively. In other words, we can suppose that $c = \bar{c} + \eta(t)$, i.e., that the noise is multiplicative, or $k = \bar{k} + \eta(t)$, i.e., that the noise is additive, or both. This noise represents the fluctuating neural activity from many different areas of the brain that connect to the Edinger–Westphal nucleus, the neural system that controls the parasympathetic drive of the iris. It also is meant to include the intrinsic noise at the synapses onto this nucleus and elsewhere in the reflex arc.

Without noise, equation (6.8) undergoes a supercritical Hopf bifurcation as the gain is increased via the parameter n . We have investigated both the additive and multiplicative noise hypotheses by performing stochastic simulations of equation (6.8). The noise was chosen to be Ornstein–Uhlenbeck noise with a correlation time of one second (Longtin, Milton, Bos, and Mackey 1990). Some results are shown in Figure 6.3, where a transition from low amplitude fluctuations to more regular high-amplitude fluctuations is seen as the feedback gain is increased. Results are similar with noise on either c or k . Even before the deterministic bifurcation, oscillations with roughly the same period as the limit cycle that appears at the bifurcation are excited by the noise. Increasing the gain just makes them more prominent: In fact, there is no actual bifurcation when noise is present, only a graded appearance of oscillations.

6.6 Postponement of the Hopf Bifurcation

We now discuss the problem of pinpointing a Hopf bifurcation in the presence of noise. This is a difficult problem not only for the Hopf bifurcation, but for other bifurcations as well (Horsthemke and Lefever 1984). From the time series point of view, noise causes fluctuations on the deterministic solution that exists without noise. However, and this is the more interesting effect, it can also produce noisy versions of behaviors that occur nearby in parameter space for the noiseless system. For example, as we have seen in the previous section, near a Hopf bifurcation the noise will produce a mixture of fixed-point and limit cycle solutions. In the most exciting examples of the effect of noise on nonlinear dynamics, even new behaviors having no deterministic counterpart can be produced.

The problem of pinpointing a bifurcation in the presence of noise arises because there is no obvious qualitative change in dynamics from the time series point of view, in contrast with the deterministic case. The definition

of a bifurcation as a qualitative change in the dynamical behavior when a parameter is varied has to be modified for a noisy system. There is usually more than one way of doing this, depending on which **order parameter** one chooses, i.e., which aspect of the dynamics one focuses on; the location of the bifurcation may also depend on this choice of order parameter.

In the case of the pupil light reflex near a Hopf bifurcation, it is clear from Figure 6.3 that noise causes oscillations even though the deterministic behavior is a fixed point. The noise is simply revealing the behavior beyond the Hopf bifurcation. It is as though noise causes the bifurcation parameter to fluctuate across the deterministic bifurcation. This is a useful way to visualize the effect of noise, but it may be misleading, since the parameter need not fluctuate across the bifurcation point to see a mixture of behaviors below and beyond this point. One can thus say, from the time series point of view, that noise advances the bifurcation point, since (noisy) oscillations are seen where, deterministically, a fixed point should be seen. Further, one can compare features of the noisy oscillation in time with, for example, the same features predicted by a model (see below).

This analysis has its limitations, however, because power spectral (or autocorrelation) measures of the strength of the oscillatory component of the time series do not exhibit a qualitative change as parameters (including noise strength) vary. Rather, for example, the peak in the power spectrum associated with the oscillation simply increases as the underlying deterministic bifurcation is approached or the noise strength is increased. In other words, there is no bifurcation from the spectral point of view. Also, this point of view does not necessarily give a clear picture of the behavior beyond the deterministic bifurcation. For example, can one say that the fixed point, which is unstable beyond the deterministic bifurcation, is stabilized by noise? In other words, does the system spend more time near the fixed point than without noise? This can be an important piece of information about the behavior of a real control system (see also Chapter 8 on cell replication and control).

There are measures that reveal a bifurcation in the noisy system. One measure is based on the computation of invariant densities for the solutions. In other words, let the solution run long enough so that transients have disappeared, and then build a histogram of values of the solution. It is better to repeat this process for many realizations of the stochastic process in order to obtain a smooth histogram.

In the deterministic case, this will produce two qualitatively different densities, depending on whether n is below or above the deterministic Hopf bifurcation point n_o . If it is below, then the asymptotic solution is a fixed point, and the density is a delta function at this fixed point: $\rho^*(x) = \delta(x - x^*)$. If $n > n_o$, the solution is approximately a sine wave, for which

the density is

$$\rho^*(x) = \frac{1}{\pi A \cos[\arcsin(x/A)]}, \quad (6.9)$$

where A is the amplitude of the sine wave.

When the noise intensity is greater than zero, the delta function gets broadened to a Gaussian distribution, and the density for the sine wave gets broadened to a smooth double-humped or bimodal function. Examples are shown in Figure 6.3 for two values of the feedback parameter n . It is possible then to define the bifurcation in the stochastic context by the transition from unimodality to bimodality (Horsthemke and Lefever 1984). The distance between the peaks can serve as an order parameter for this transition (different order parameters can be defined, as in the physics literature on phase transitions). It represents in some sense the mean amplitude of the fluctuations.

We have found that the transition from a unimodal to a bimodal density occurs at a value of n greater than n_o (Longtin, Milton, Bos, and Mackey 1990). In this sense, the bifurcation is postponed by the noise, with the magnitude of the postponement being proportional to the noise intensity. In certain simple cases (although not yet for the delay-differential equation studied here), it is possible to analytically approximate the behavior of the order parameter with noise. This allows one to predict the presence of a postponement, and to relate this postponement to certain model parameters, especially those governing the nonlinear behavior. A postponement does not imply that there are no oscillations if $n < n_p$, where n_p is the extrapolated bifurcation point for the noisy case (it is very time-consuming to numerically determine this point accurately). As we have seen, when there is noise near a Hopf bifurcation, oscillations are present. However, a postponement does imply that if $n > n_o$, *the presence of noise stabilizes the fixed point*. In other words, the system spends more time near the fixed point with noise than without noise. This is why the density for the stochastic differential equation fills in between the two peaks of the deterministic distribution given in equation (6.9).

One can try to pinpoint the bifurcation in the pupil data by computing such densities at different values of the feedback gain. The result is shown in Figure 6.4 for the data used for Figure 6.2. Even for the highest value of gain, there are clearly oscillations, and the distribution still appears unimodal. However, this is not to say that it is unimodal. A problem arises because a large number of simulated data points (two orders of magnitude more than experimentally available) are needed to properly measure the order parameter, i.e., the distance between the two peaks of the probability density. The bifurcation is not seen from the density point of view in this system with limited data sets and large amounts of noise (the higher the noise, the more data points are required). The model does suggest, however, that a postponement can be expected from the density

point of view; in particular, the noisy system spends more time near the fixed point than without noise, even though oscillations occur. Further, the mean, moments, and other features of these densities could be compared to those obtained from time series of similar duration generated by models, in the hope of better understanding the dynamics underlying such noisy experimental systems.

This lack of resolution to pinpoint the Hopf bifurcation motivated us to validate our model using other quantities, such as the mean and relative standard deviation of the amplitude and period fluctuations as gain increases (Longtin, Milton, Bos, and Mackey 1990). That study showed that for noise intensity $D = 15$ and a noise correlation time around one, these quantities have similar values in the experiments and the simulations. This strengthens our belief that stochastic forces are present in this system. Interestingly, our approach of investigating fluctuations across a bifurcation (supercritical Hopf in this case) allows us to amplify the noise in the system, in the sense that it is put under the magnifying glass. This is because noise has a strong effect on the dynamics of a system in the vicinity of a bifurcation point, since there is loss of linear stability at this point (neither the fixed point nor the zero-amplitude oscillation is attracting).

Finally, there is an interesting theoretical aspect to the postponements. We are dealing here with a first-order differential-delay equation. Noise-induced transitions such as postponements are not possible with additive noise in one-dimensional ordinary differential equations (Horsthemke and Lefever 1984). But our numerical results show that in fact, additive noise-induced transitions are possible in a first-order delay-differential equation. The reason behind this is that while the highest-order derivative is one, the delay-differential equation is infinite-dimensional, since it evolves in a functional space (an initial function must be specified). More details on these theoretical aspects of the noise-induced transitions can be found in Longtin (1991a).

6.7 Stochastic Phase Locking

The nervous system has evolved with many sources of noise, acting from the microscopic ion channel scale up to the macroscopic scale of the EEG activity. This is especially true for cells that transduce physical stimuli into neuroelectrical activity, since they are exposed to environmental sources of noise, as well as to intrinsic sources of noise such as ionic channel conductance fluctuations, synaptic fluctuations, and thermal noise. Traditionally, sources of noise in sensory systems, such as the senses of audition and touch, have been perceived as a nuisance. For example, they have been thought to limit our aptitude for detecting or discriminating between stimuli.

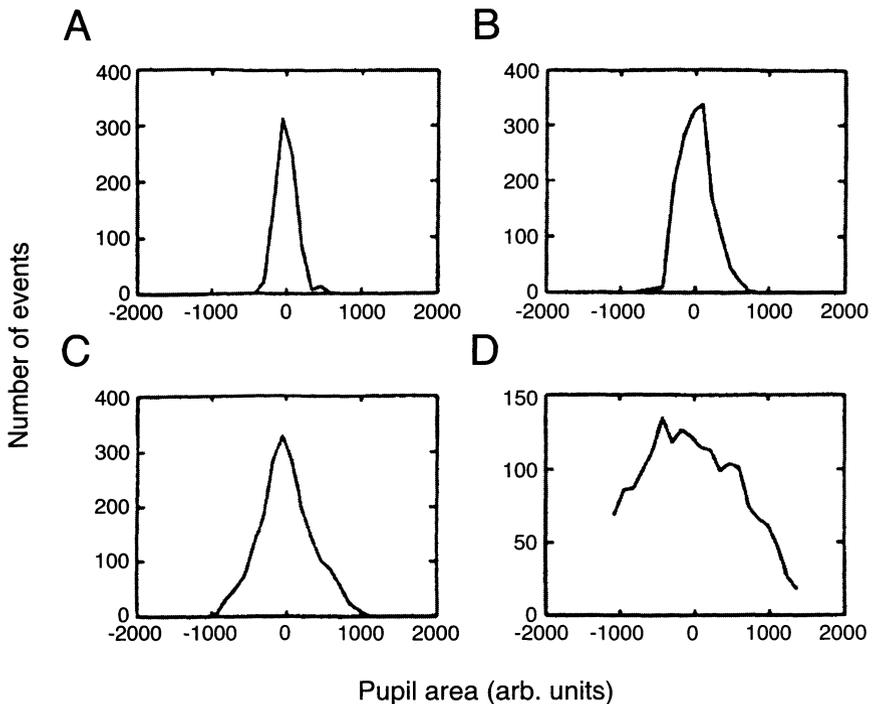


Figure 6.4. Densities corresponding to the time series shown in Figure 6.2 (more data were used than are shown in Figure 6.2). From Longtin (1991b).

In the past decades, there have been studies that revealed a more constructive role for neuronal noise. For example, noise can increase the dynamic range of neurons by linearizing their stimulus–response characteristics (see, e.g., Spekreijse 1969; Knight 1972; Treutlein and Schulten 1985). In other words, noise smoothes out the abrupt increase in mean firing rate that occurs in many neurons as the stimulus intensity increases; this abruptness is a property of the deterministic bifurcation from nonfiring to firing behavior. Noise also makes detection of weak signals possible (see, e.g., Hochmair-Desoyer, Hochmair, Motz, and Rattay 1984). And noise can stabilize systems by postponing bifurcation points, as we saw in the previous section (Horsthemke and Lefever 1984).

In this section, we focus on a special kind of firing behavior exhibited by many kinds of neurons across many different sensory modalities. In general terms, it can be referred to as “stochastic phase locking,” but in more specific terms it is known as “skipping.” An overview of physiological examples of stochastic phase locking in neurons can be found in Segundo, Vibert, Pakdaman, Stiber, and Martinez (1994). Figure 6.5 plots the membrane potential versus time for a model cell driven by a sinusoidal stimulus

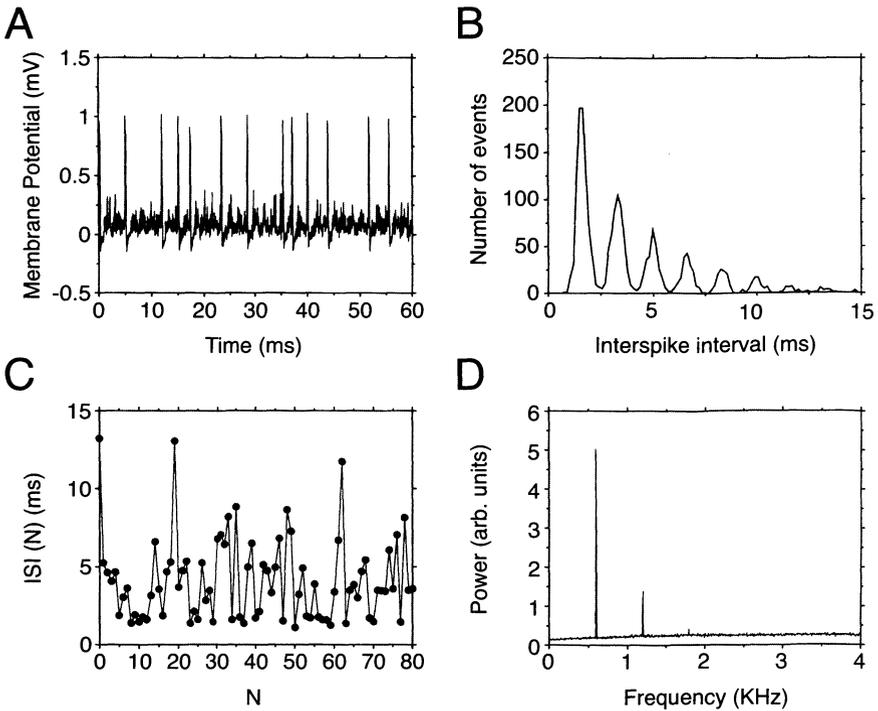


Figure 6.5. Numerical simulation of the FitzHugh–Nagumo equations in the sub-threshold regime, in the presence of noise and sinusoidal forcing. (A) Time series of the membrane potential. (B) Interspike interval histogram obtained from 100 realizations of this stochastic system yielding a total of 2048 intervals. (C) An aperiodic sequence of interspike intervals (ISI). (D) Power spectrum of the spike train, averaged over 100 realizations.

and noise (the model is the FitzHugh–Nagumo equations; see below). The sharp upstrokes superimposed on the “noisy” oscillation are action potentials. The stimulus period here is long in comparison to the action potential duration. The feature of interest is that while the spikes are phase locked to the stimulus, they do not occur at every cycle of the stimulus. Instead, a seemingly random *integer* number of periods of the stimulus are “skipped” between any two successive spikes, thus the term “skipping.” This is shown in the interspike interval histogram in Figure 6.5B. The peaks in this interspike interval histogram line up with the integer multiples of the driving period ($T_o = 1.67$ msec). The lack of periodicity in the firing pattern can be inferred from Figure 6.5C, where the interval value is plotted as a function of the interval number (80 intervals are plotted).

In practice, one usually does not have access to the membrane voltage itself, since the sensory cells or their afferent nerve fibers cannot be impaled by a microelectrode. Instead, one has a sequence of interspike intervals from

which the mechanisms giving rise to signal encoding and skipping must be inferred.

In the rest of this chapter, we describe some mechanisms of skipping in sensory cells, as well as the potential significance of such firing patterns for sensory information processing. We discuss the phenomenology of skipping patterns, and then describe efforts to model these patterns mathematically. We describe the stochastic resonance effect in this context, and discuss its origins. We also discuss skipping patterns in the context of “bursting” firing patterns. We consider the relation of noise-induced firing to linearization by noise. We also show how noise can alter the shape of tuning curves, and end with an outlook onto interesting issues for future research.

6.8 The Phenomenology of Skipping

A firing pattern in which cycles of a stimulus are skipped is a common occurrence in physiology. For example, this behavior underlies $p : m$ phase locking seen in cardiac and other excitable cells, i.e., firing patterns with m responses to p cycles of the stimulus. The main additional properties here are that the phase locking pattern is aperiodic, and remains qualitatively the same as stimulus characteristics are varied. In other words, abrupt changes between patterns with different phase locking ratios are not seen under “skipping” conditions. For example, as the amplitude of the stimulus increases, the skipping pattern remains aperiodic, but there is a higher incidence of short skips rather than long skips.

A characteristic interspike interval histogram for a skipping pattern is shown in Figure 6.5B, and again in Figure 6.6A for in a bursty P-type electroreceptor of a weakly electric fish. The stimulus in this latter case is a 660 Hz oscillatory electric field generated by the fish itself (its “electric organ discharge”). It is modulated by food particles and other objects and fish, and the 660 Hz carrier along with its modulations are read by the receptors in the skin of the fish. This electrosensory system is used for electrolocation and electrocommunication. The interspike interval histogram in Figure 6.6A again consists of a set of peaks located at integer multiples of the driving period. Note from Figure 6.6B that there is no apparent periodicity in the interval sequence. The firing patterns of electroreceptors were first characterized in Scheich, Bullock, and Hamstra Jr (1973). These receptors are known as P-units or “probability coders,” since it is thought that their probability of firing is proportional to, and thus encodes, the instantaneous amplitude of the electric organ discharge. Hence this probability, determined by various parameters including the intensity of noise sources acting in the receptor, is an important part of the neuronal code in this system.

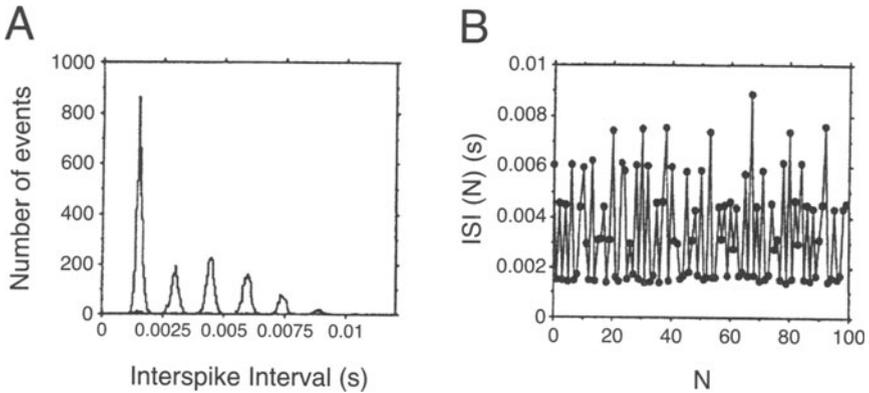


Figure 6.6. Interspike interval histogram and sequence of interspike intervals (ISI) measured from a primary afferent fiber of an electroreceptor of the weakly electric fish *Apteronotus leptorhynchus*. The stimulus frequency is 660 Hz, and is generated by the fish itself. Data provided courtesy of Joe Bastian, U. Oklahoma at Norman.

Another classic example of skipping is found in mammalian auditory fibers. Rose, Brugge, Anderson, and Hind (1967) show that skipping patterns occur at frequencies from below 80 Hz up to 1000 Hz and beyond in a single primary auditory fiber of the squirrel monkey. For all amplitudes, the modes in the interspike interval histogram line up with the integer multiples of the stimulus period, and there is a mode centered on each integer between the first and last visible modes. At low frequencies, multiple firings can occur in the preferred part of the stimulus cycle, and thus a peak corresponding to very short intervals is also seen. At frequencies beyond 1000 Hz, the first peak is usually missing due to the refractory period of the afferent fiber; in other words, the cell cannot recover fast enough to fire spikes one millisecond apart. Nevertheless, phase locking persists as evidenced by the existence of other modes. This is true for electroreceptors as well (Chacron, Longtin, St-Hilaire, and Maler 2000).

The degree of phase locking in all cases is evident from the width of the interspike interval histogram peaks: Sharp peaks correspond to a high degree of phase locking, i.e., to a narrow range of phases of the stimulus cycle during which firing preferentially occurs. As amplitude increases, the multimodal structure of the interspike interval histogram is still present, but the intervals are more concentrated at low values. In other words, the higher the intensity, the lower the (random) integer number of cycles skipped between firings. What is astonishing is that these neurons are highly “tunable” across such a broad range of stimulus parameters, with the modes always lining up with multiples of the driving period. There are many examples of skipping in other neurons, sensory and otherwise, e.g., in mechanoreceptors and thermoreceptors (see Longtin 1995; Segundo, Vibert, Pakdaman,

Stiber, and Martinez 1994; Ivey, Apkarian, and Chialvo 1998 and references therein).

Another important characteristic of skipping is that the positions of the peaks vary smoothly with stimulus frequency, and the envelope of the interspike interval histogram varies smoothly with stimulus amplitude. This is different from the phase-locking patterns governed, for example, by phase-resetting curves leading to an Arnold tongue structure as stimulus amplitude and period are varied. We will see that a plausible mechanism for skipping involves the combination of noise with subthreshold dynamics, although suprathreshold mechanisms exist, as we will see below (Longtin 1998). In fact, we have recently found (Chacron, Longtin, St-Hilaire, and Maler 2000) that suprathreshold periodic forcing of a leaky integrate-and-fire model with voltage and threshold reset can produce patterns close to those seen in electroreceptors of the nonbursty type (interspike interval histogram similar to that in Figure 6.6A, except that the first peak is missing). We focus below on the subthreshold scenario in the context of the FitzHugh–Nagumo equations, with one suprathreshold example as well.

6.9 Mathematical Models of Skipping

The earliest analytical/numerical study of skipping was performed by Gerstein and Mandelbrot (1964). Cat auditory fibers recorded during auditory stimulation with periodic “clicks” of noise (at frequencies less than 100 clicks/sec) showed skipping behavior. Gerstein and Mandelbrot were interested in reproducing experimentally observed spontaneous interspike interval histograms using “random walks to threshold models” of neuron firing activity. In one of their simulations, they were able to reproduce the basic features of the interspike interval histogram in the presence of the clicks by adding a periodically modulated drift term to their random walk model. The essence of these models is that the firing activity is entirely governed by noise plus a constant and/or a periodically modulated drift. A spike is associated with the crossing of a fixed threshold by the random variable.

Since this early study, there have been other efforts aimed at understanding the properties of neurons driven by periodic stimuli and noise. The following examples have been excerpted from the large literature on this subject. French et al. (1972) showed that noise breaks up patterns of phase locking to a periodic signal, and that the mean firing rate is proportional to the amplitude of the signal. Glass et al. (1980) investigated an integrate-and-fire model of neural activity in the presence of periodic forcing and noise. They found unstable zones with no phase locking, as well as quasi-periodic dynamics and firing patterns with stochastic skipped beats. Keener et al. (1981) were able to analytically investigate the dynamics of

phase locking in a leaky integrate-and-fire model without noise. Alexander et al. (1990) studied phase locking phenomena in the FitzHugh–Nagumo model of a neuron in the excitable regime, again without noise. There have also been studies of noise-induced limit cycles in excitable cell models like the Bonhoeffer–van der Pol equations (similar to the FitzHugh–Nagumo model), but in the absence of periodic forcing (Treutlein and Schulten 1985).

The past two decades have seen a revival of stochastic models of neural firing in the context of skipping. Hochmair-Desoyer et al. (1984) have looked at the influence of noise on firing patterns in auditory neurons using models such as the FitzHugh–Nagumo equations, and shown that it can alter the tuning curves (see below). This model generates real action potentials with a refractory period. It also has many other behaviors that are found in real neurons, such as a resonance frequency. It is a suitable model, however, only when an action potential is followed by a hyperpolarizing after-potential, i.e., the voltage goes below the resting potential after the spike, and slowly increases towards it. It is also a good simplified model to study certain neural behaviors qualitatively; better models exist (they are usually more complex) and should be used when quantitative agreement between theory and experiment is sought. The study of the FitzHugh–Nagumo model in Longtin (1993) was motivated by the desire to understand how stochastic resonance could occur in real neurons, as opposed to bistable systems where the concept had been confined. In fact, the FitzHugh–Nagumo model has a cubic nonlinearity, just as does the standard quartic bistable system in equation (6.7); however, it has an extra degree of freedom that serves to reset the system after the threshold for spiking is crossed.

We illustrate here the behavior of the FitzHugh–Nagumo model with simultaneous stimulation by a periodic signal and by noise. The latter can be interpreted as either synaptic noise, or signal noise, or conductance fluctuations (although the precise modeling of such fluctuations is better done with conductance-based models such as Hodgkin–Huxley-type models). The model equations are (Longtin 1993)

$$\epsilon \frac{dv}{dt} = v(v - a)(1 - v) - w + \eta(t), \quad (6.10)$$

$$\frac{dw}{dt} = v - dw - b - r \sin \beta t, \quad (6.11)$$

$$\frac{d\eta}{dt} = -\lambda\eta + \lambda\xi(t). \quad (6.12)$$

The variable v is the fast voltage-like variable, while w is a recovery variable. Also, $\xi(t)$ is a zero-mean Gaussian white additive noise, which is lowpass filtered to produce an Ornstein–Uhlenbeck-type additive noise denoted by η . The autocorrelation of the white noise is $\langle \xi(t)\xi(s) \rangle = 2D\delta(t - s)$; i.e., it is delta-correlated. The parameter D is the intensity of the white noise. This Ornstein–Uhlenbeck noise is Gaussian and exponentially correlated,

with a correlation time (i.e., the $1/e$ time) of $t_c = \lambda^{-1}$. The periodic signal of amplitude r and frequency β is added here to the recovery variable w as in Alexander, Doedel, and Othmer (1990), yielding qualitatively similar dynamics as in the case in which it is added to the voltage equation (after proper adjustment of the amplitude; see Longtin 1993). The periodic forcing should be added to the voltage variable when the period of stimulation is smaller than the refractory period of the action potential.

The parameter regime used to obtain the results in Figure 6.5 can be understood as follows. In the absence of periodic stimulation, one would see a smooth unimodal interspike interval histogram (close to a gamma-type distribution) governed by the interaction of the two-dimensional FitzHugh–Nagumo dynamics with noise. The periodic stimulus thus carves peaks out of this “background” distribution of the interspike interval histogram. If the noise is turned off and the stimulus is turned on, there would be no firings whatsoever. This is a crucial point: The condition for obtaining skipping with the tunability properties described in the previous section is that *the deterministic dynamics must be subthreshold*. This feature can be controlled by the parameter b , which sets the proximity of the resting potential (i.e., the single stable fixed point) to the threshold. In fact, this dynamical system goes through a supercritical Hopf bifurcation at $b_H = 0.35$. It can also be controlled by a constant current that could be added to the left hand side of the first equation.

Figure 6.7 contrasts the subthreshold ($r = 0.2$) and suprathreshold ($r = 0.22$) behavior in the FitzHugh–Nagumo system. In the subthreshold case, the noise is essential for firings to occur: No intervals are obtained when the noise intensity, D , is zero. For $r = 0.22$ and $D = 0$, only one kind of interval is obtained, namely, that corresponding to the period of the deterministic limit cycle. For $D > 0$, the limit cycle is perturbed by the noise, and sometimes comes close to but misses the separatrix: No action potential is generated during one or more cycles of the stimulus. In the subthreshold case, one also sees skipping behavior. At higher noise intensities, the interspike interval histograms hardly differ, and thus we cannot tell from such an interspike interval histogram whether the system is suprathreshold or subthreshold. This distinction can be made by varying the noise level as illustrated in this figure. In the subthreshold case, the mean of the distribution will always move to lower intervals as D increases, although this is not true for the suprathreshold case.

There have also been numerous modeling studies based on noise and sinusoidally forced integrate-and-fire-type models (see, e.g., Shimokawa, Pakdaman, and Sato 1999; Bulsara, Elston, Doering, Lowen, and Lindenberg 1996; and Gammaitoni, Hänggi, Jung, and Marchesoni 1998 and references therein). Other possible generic dynamical behavior might lurk behind this form of phase locking. In fact, the details of the phase locking, and of the physiology of the cells, are important in determining which specific dynamics are at work. Subthreshold chaos might be involved (Kaplan,

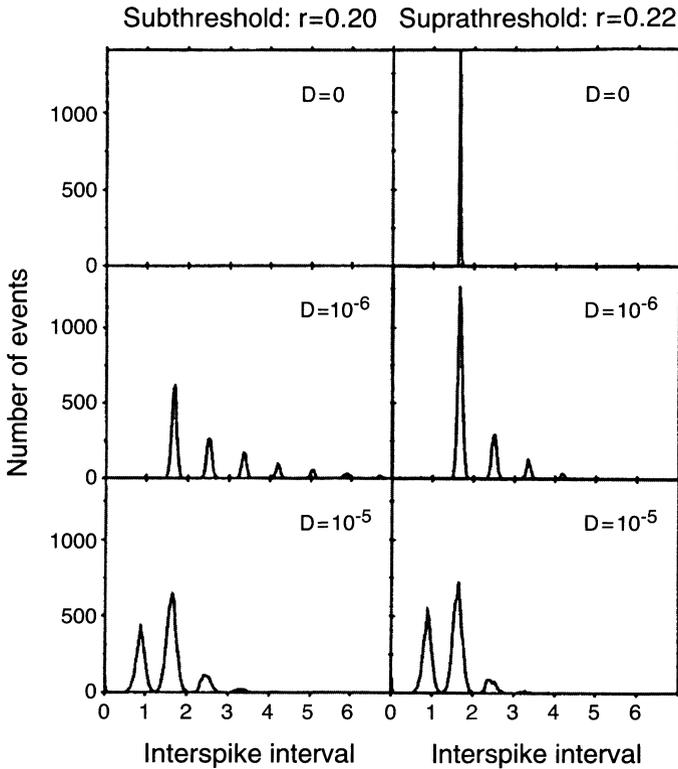


Figure 6.7. Comparison of interspike interval histograms with increasing noise intensity in the subthreshold regime (left panels) and suprathreshold regime (right panels).

Clay, Manning, Glass, Guevara, and Shrier 1996), although probably with noise as well if one seeks smooth interspike interval histograms with symmetric modes and without missing modes between the first and the last (as with the data in Rose, Brugge, Anderson, and Hind 1967; see Longtin 1998). An example of these effects of noise on the multimodal interspike interval histograms produced by the chaos (with pulsatile forcing) is shown in Figure 6.8. In other systems, chaos may be the main player, as in the experimental system (pulsatile stimulation of squid axons) considered in Kaplan et al. (1996).

In Figure 10 of Longtin (1993), a “chaos” hypothesis for skipping was investigated using the FitzHugh–Nagumo model with sinusoidal forcing (instead of pulses as in Kaplan, Clay, Manning, Glass, Guevara, and Shrier 1996). Without noise, this produced subthreshold chaos, as described in Kaplan, Clay, Manning, Glass, Guevara, and Shrier (1996), although clearly, when spikes occur (using some criterion for graded responses) these are “suprathreshold” responses to this “subthreshold chaos”; in other words,

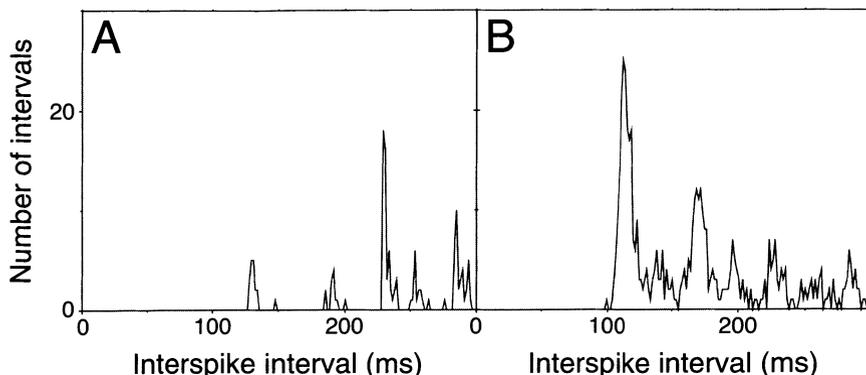


Figure 6.8. Interspike interval histogram from the FitzHugh–Nagumo system $dv/dt = v(v - 0.139)(1 - v) - w + I + \eta(t)$, $dw/dt = 0.008(v - 2.54w)$, where I consists of rectangular pulses of duration 1.0 msec and height 0.28, repeated every 28.5 msec. Each histogram is obtained from one realization of 5×10^7 time steps. The Ornstein–Uhlenbeck noise $\eta(t)$ has a correlation time of 0.001 msec. (A) Noise intensity $D = 0$. (B) $D = 2.5 \times 10^{-5}$.

the chaos could just as well be referred to as “suprathreshold.” In this FitzHugh–Nagumo chaos case, some features of the Rose et al. data could be reproduced, but others not. For example, multimodal histograms were found. But the individual peaks had more “internal” structure than seen in the data (including electroreceptor and mechanoreceptor data); they were not aligned very well with the integer multiples of the driving period; and some peaks were missing, as in the case of pulsatile forcing shown in Figure 6.8. Further, the envelope did not have the characteristic exponential decay (past the second mode) seen in the Rose et al. data (which is what is expected for uncorrelated intervals). Additive dynamical noise on top of this chaos did a better job at reproducing these qualitative features, at least for the parameters explored (Longtin 1998). The modes of the interspike interval histogram were still a bit lopsided, and the envelopes were different from those of the data. Interestingly, solutions that “look chaotic” often end up on periodic orbits after a long while. A bit of noise would probably keep these solutions bouncing around irregularly.

The other reason that some stochastic component may be a necessary ingredient is the smoothness observed in the transitions between interspike interval histograms as stimulation parameters are changed. Changing period or amplitude in the chaotic models leads to sometimes abrupt changes in the multimodal structure (and some peaks just keep on missing). Noise induced firing with deterministically subthreshold dynamics does produce the required smooth transitions in the proper sequence seen in Rose et al. (1967). Note, however, that in certain parameter ranges, it is possible to get multimodal histograms with suprathreshold forcing. This is shown

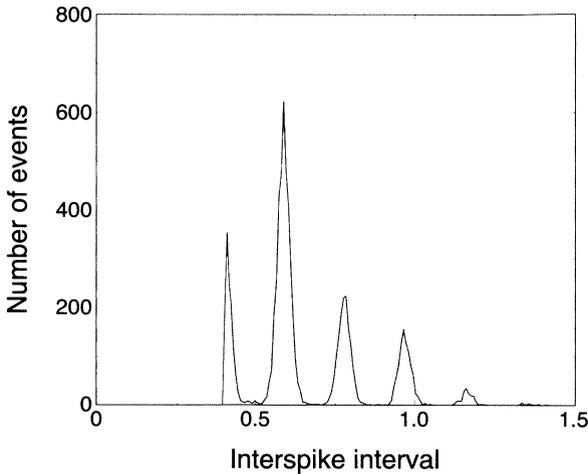


Figure 6.9. Interspike interval histogram from the FitzHugh–Nagumo system with fast sinusoidal forcing $\beta = 32$. Other parameters are $a = 0.5$, $b = 0.15$, $d = 1$, $\epsilon = 0.005$, $I = 0.04$ and $r = 0.06$. The histogram is obtained from 10 realizations of 500,000 time steps.

in Figure 6.9, for which the forcing frequency is high, and the deterministic solution is a periodic 3:1 solution.

All this discussion does not exclude the possibility that chaos alone (e.g., with other parameters or in a more refined model) might give the right picture for this kind of data, or that deterministic chaos or periodic phase locking combined with noise might give it as well. Only good intracellular data can ultimately settle the issue of the origin of the stochastic phase locking, and provide an explanation for the smooth skipping patterns.

6.10 Stochastic Resonance

The notion that skipping neurons in the subthreshold regime rely on noise to fire is interesting from the point of view of signal processing. In order to transmit information about the stimulus (the input) to a neuron in its spike train (the output), noise must be present. Without noise, there are no firings, and with too much noise, we expect to see a very noisy output with again no information (or very little) about the stimulus. Hence, there must be a noise value for which information about the stimulus is optimally transmitted to the output. In other words, starting from zero noise, adding noise will increase the signal-to-noise ratio, and an optimal noise level can be found where the signal-to-noise ratio peaks. This is indeed the case in the FitzHugh–Nagumo model studied above. This effect, in which the signal-to-noise ratio is optimal for some intermediate noise intensity, is

known as **stochastic resonance**. It has been studied for over a decade, usually in bistable systems. It had been studied theoretically in bistable neurons (Bulsara, Jacobs, Zhou, Moss, and Kiss 1991), and predicted to occur in real neurons (Longtin, Bulsara, and Moss 1991). Thereafter, it was studied theoretically in an excitable system (Longtin 1993; Chialvo and Apkarian 1993; Chapeau-Blondeau, Godivier, and Chambet 1996) and a variety of other systems (Gammaitoni, Hänggi, Jung, and Marchesoni 1998), and shown to occur in real systems (see, e.g., Douglass, Wilkens, Pantazelou, and Moss 1993; Levin and Miller 1996). It is one of many constructive roles for noise discovered in recent decades (Astumian and Moss 1998).

This resonance can be studied from the points of view of spectral amplitude at the signal frequency, signal-to-noise ratio, residence-time histograms (i.e., interspike interval histograms), and others as well. In the first case, one computes for a given value of noise intensity D the power spectrum averaged over many spike trains obtained with as many realizations of the noise process. The spectrum is usually in the form of a flat or curved background, on which the harmonics of the small stimulus signal are superimposed (see Figure 6.5D). A dip at low frequencies is often seen, which is due to phase jitter and to the refractory period. A signal-to-noise ratio can then be computed by dividing the height of the fundamental stimulus peak by the noise floor (i.e., the value of the noise background at the frequency of the stimulus). This signal-to-noise ratio can be plotted as a function of D , and the resulting curve will be unimodal, with the maximum corresponding to the stochastic resonance. Alternatively, one can measure the heights of the different peaks in the interspike interval histogram, and plot these heights as a function of D . The different peaks will go through a maximum at different values of D . While there is yet no direct analytical connection between stochastic resonance from these two points of view, it is usually the case that systems exhibiting stochastic resonance from one point of view will also exhibit it from the other.

The power spectrum measures the synchrony between firings and the stimulus. From the point of view of the interspike interval histogram, the measure of synchrony depends not only on the prevalence of intervals at integer multiples of a fundamental interval, but also on the width of the peaks of the interspike interval histogram. As noise increases past the resonance value, these widths increase, with the result that the phase locking is disrupted by the noise, even though there are many firings.

A simple theory of stochastic resonance for excitable systems is being developed. Wiesenfeld et al. (1994) have shown that stochastic resonance will occur in a periodically modulated point process. By redoing the calculation of the classic shot noise effect for the case of a periodic stimulus (the point process is then inhomogeneous, i.e., time-dependent), they have

found an expression for the signal-to-noise ratio (SNR) (in decibels):

$$\text{SNR} = 10 \log_{10} \left[\frac{4I_1^2(z)}{I_0(z)} \exp(-U/D) \right], \quad (6.13)$$

where I_n is the modified Bessel function of order n , $z \equiv rU/D$, and U is a measure of the proximity of the fixed point to the firing threshold (i.e., some kind of activation barrier). For small z , this equation becomes

$$\text{SNR} = 10 \log_{10} \left[\frac{U^2 r^2}{D^2} \exp(-U/D) \right], \quad (6.14)$$

which is almost identical to a well-known result for stochastic resonance in a bistable potential. More recently, analytical techniques have been devised to study stochastic resonance in two-variable systems such as the FitzHugh–Nagumo system (Lindner and Schimansky-Geier 2000).

The shape of the interspike interval histogram, and in particular, its rate of decay, is very sensitive to the stimulus characteristics. This is to be contrasted with the transition from sub- to suprathreshold dynamics in the absence of noise. There are no firings before the stimulus exceeds a threshold amplitude. Once the suprathreshold regime is reached, however, amplitude increases can bring on various $p : m$ locking patterns and even chaos. Noise allows the firing pattern to change smoothly and sensitively over a larger range of stimulus parameters.

The firing patterns of the neuron in the excitable regime are also interesting in the presence of noise only, i.e., without periodic forcing. In fact, such a noisy excitable system can be seen as a stochastic oscillator (Longtin 1993; Pikovsky and Kurths 1997; Longtin and Chialvo 1998; Lee and Kim 1999; Lindner and Schimansky-Geier 2000). The presence of a resonance in the deterministic dynamics will endow this oscillator with a well-defined preferred time between firings; this time scale is closely associated with the period of the limit cycle that arises when the system is biased into its autonomously firing regime. Recently, Pikovsky and Kurths (1997) showed that increasing the noise intensity from zero will lead to enhanced periodicity in the output firing pattern, followed by a decreased periodicity. This has been termed **coherence resonance**, and is related to the induction by noise of the limit cycle that exists in the vicinity of the excitable regime (Wiesenfeld 1985). The effect has also been predicted to occur in bursting neurons (Longtin 1997).

We close this section with a brief discussion of the origin of stochastic resonance in excitable neurons. Various aspects of this question have been discussed in Collins, Chow, and Imhoff 1995a; Collins, Chow, and Imhoff 1995b; Bulsara, Jacobs, Zhou, Moss, and Kiss 1991; Bulsara, Elston, Doering, Lowen, and Lindenberg 1996; Chialvo, Longtin, and Müller-Gerking 1997; Longtin and Chialvo 1998; Neiman, Silchenko, Anishchenko, and Schimansky-Geier 1998; Lee and Kim 1999; Shimokawa, Pakdaman, and Sato 1999; Lindner and Schimansky-Geier 2000. Here we focus on the distri-

bution of the phases at which firings occur, the so-called “cycle histogram.” Figure 6.10 shows cycle histograms for the FitzHugh–Nagumo model with subthreshold parameter settings similar to those used in previous figures, for high (left panels) and low frequency forcing (right panels). The lower panels (low noise) show that the cycle histogram is rectified, with firings occurring only in a restricted range of phases. The associated interspike interval histograms (not shown) are multimodal as a consequence of this phase preference. The rectification is due to the fact that the firing rate for zero forcing is low: When this rate is modulated downward by the signal, the rate goes to zero (and cannot go lower). The rectification for $T = 0.5$ is even stronger, because at higher frequencies, phase locking also occurs (Longtin and Chialvo 1998; Lee and Kim 1999): This is a consequence of the refractory period of the system, responsible for phase locking patterns in the suprathreshold regime, and increasingly important at higher frequencies. At higher noise, the rectification has disappeared: The noise has linearized the cycle histogram. The spectral power of the spike train at the signal frequency is maximal near the noise intensity that produces the “most sinusoidal” cycle histogram (as measured, for example, by a linear correlation coefficient).

This linearization is dependent on noise amplitude only for low frequencies (i.e., for $T > 2$ or so), such as those used in Collins, Chow, and Imhoff 1995a; Collins, Chow, and Imhoff 1995b; Chialvo, Longtin, and Müller-Gerking 1997: The neuron then essentially behaves as a static threshold device. As the frequency increases, linearization requires more noise, due to the increased importance of phase locking. This higher noise also produces an increased spontaneous rate of firing when the signal is turned off. Hence, this rate for the unmodulated system must increase in parallel with the frequency in order for the firings to be maximally synchronized with the stimulus. Also, secondary resonances at lower noise occur for higher frequencies (Longtin and Chialvo 1998) in both the spectra and in the peak heights of the interspike interval histogram, corresponding to the excitation of stochastic subharmonics of the driving force. The noise producing the maximal signal-to-noise ratio is itself minimal for frequencies near the best frequency (i.e., the resonant frequency) of the FitzHugh–Nagumo model (Lee and Kim 1999).

6.11 Noise May Alter the Shape of Tuning Curves

Tuning curves are an important characteristic of neurons and cardiac cells. They describe the sensitivity of these cells to the amplitude and frequency of periodic signals. For each sinusoidal forcing frequency, one determines the minimum amplitude needed to obtain a specific firing pattern, such as 1:1 firing. The frequency–amplitude pairs are then plotted to yield the

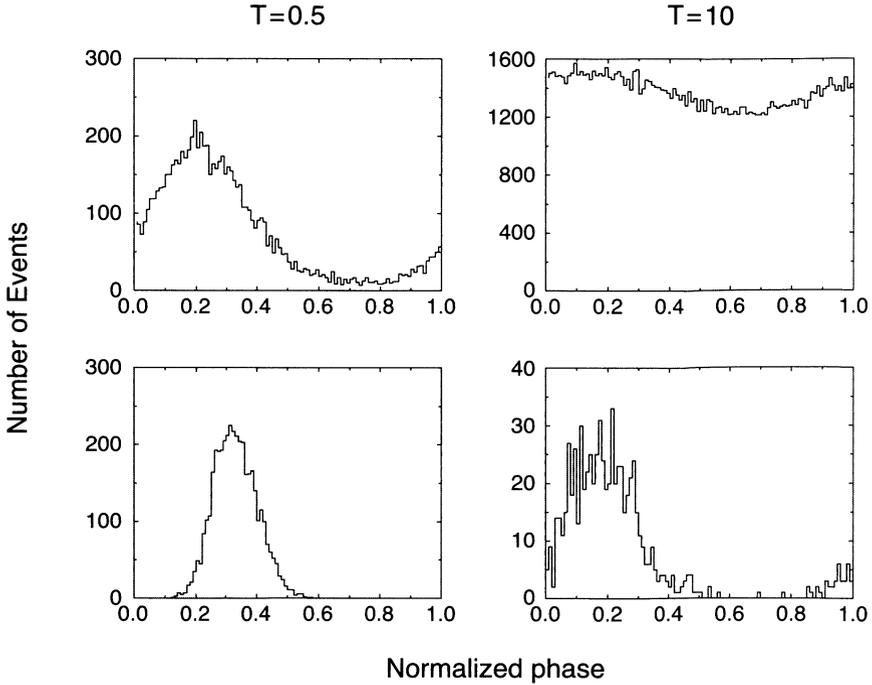


Figure 6.10. Probability of firing as a function of the phase of the sinusoidal forcing (left, $T = 0.5$; right, $T = 10$), obtained by averaging over 50 realizations of 100 cycles. The amplitude of the forcing is 0.01. For the upper panels, noise intensity $D = 8 \times 10^{-6}$, and for the lower ones, $D = 5 \times 10^{-7}$.

tuning curve. We have recently computed the behavior of the 1:1 and Arnold tongues of the excitable FitzHugh–Nagumo model with and without noise (Longtin 2000). Our work was motivated by recent findings (Ivey, Apkarian, and Chialvo 1998) that mechanoreceptor tuning curves can be significantly altered by externally added stimulus noise, and by an earlier numerical study that reported that noise could alter tuning curves (Hochmair-Desoyer, Hochmair, Motz, and Rattay 1984). It was also motivated by the tuning properties of electroreceptors (see, e.g., Scheich, Bullock, and Hamstra Jr 1973), and generally by ongoing research into the mechanisms underlying aperiodic phase locked firing in many excitable cells including cardiac cells.

Figure 6.11 shows the boundary (Arnold tongue) for 1:1 firing for noise intensity $D = 0$. It is V-shaped, highlighting again the resonant aspect of the neuronal dynamics. The minimum threshold occurs for the so-called best frequency which is close to the frequency of autonomous oscillations seen past the Hopf bifurcation in this system. For period $T > 1$, the region below these curves is the subthreshold 1:0 region. For $A > 0$, ratios as

6.12 Thermoreceptors

We now turn our attention in these last two sections to physiological systems that exhibit multimodal interspike interval histograms *in the absence of any known periodic forcing*. The best-known examples are the mammalian cold thermoreceptors and the ampullae of Lorenzini (passive thermal and electroreceptors of certain fish such as sharks). The temperature fluctuates by less than 0.5 degrees Celsius during the course of the measurements modeled in Longtin and Hinzer (1996). The mean firing rate of these receptors is a unimodal function of temperature. Over the lower range of temperatures they transduce, they increase their mean firing rate with temperature, behaving essentially like warm receptors. Over the other half of their range, they decrease their mean firing rate. This higher range includes the normal body temperature, and thus an increase in firing rate signals a decrease in temperature.

This unimodal stimulus–response curve implies that a given firing rate can be associated with two constant temperatures. It has been suggested that the central nervous system resolves this ambiguity by responding to the pattern of the spike train. In fact, at lower temperatures, the firing is of bursting type with a long period between bursts and many spikes per burst (see Figure 6.12). As the temperature increases, the bursting period shortens, and the number of spikes per burst decreases, until there is on average only one spike per burst: This is then a regular repetitive firing, also known as a “beating” pattern. As the temperature increases further, a skipping pattern appears, as cycles of the beating pattern drop out randomly. The basic interval in the skipping pattern is close to the period of the beating pattern. This suggests that there is an intrinsic oscillation in these receptors that underlies all the patterns (see Longtin and Hinzer 1996 and references to Schafer and Braun therein).

Cold receptors are free nerve endings in the skin. The action potentials generated there propagate to the spinal cord and up to the thalamus. An ionic model for the firing activity of mammalian cold receptors has recently been proposed (Longtin and Hinzer 1996). The basic assumption of this model is that cold reception arises by virtue of the thermosensitivity of various ionic currents in a receptor neuron, rather than through a specialized mechanism. Other assumptions, also based on the anatomy and extracellular physiology of these receptors (intracellular recordings are not possible), include the following: (1) The bursting dynamics are of the slow-wave type, i.e., action potentials are not necessary to sustain the slow oscillation that underlies bursting; (2) the temperature modifies the rates of the Hodgkin–Huxley kinetics, with Q_{10} 's of 3; (3) the temperature increases the maximal sodium ($Q_{10} = 1.4$) and potassium ($Q_{10} = 1.1$) conductances; (4) the rate of calcium kinetics increases with temperature ($Q_{10} = 3$); (5) the activity of an electrogenic sodium–potassium pump increases linearly with temperature, producing a hyperpolarization; and (6) there is noise

added to these deterministic dynamics, to account for skipping and for fluctuations in the number of spikes per burst and the interburst period.

Our model is modified from Plant's model (Plant 1981) of slow-wave bursting in the pacemaker cell of the mollusk *Aplysia*. For the sake of simplicity, we have assumed that the precise bursting mechanism is governed by an outward potassium current whose activation depends on the concentration of intracellular calcium. The model with stochastic forcing is governed by the equations

$$\begin{aligned} C_M \frac{dV}{dt} = & G_I m_\infty^3(V) h(V_I - V) + G_x x(V_I - V) \\ & + G_K n^4(V_K - V) + G_{K-Ca} \frac{[Ca]}{0.5 + [Ca]} (V_K - V) \\ & + G_L (V_L - V) + I_p + \eta(t), \end{aligned} \quad (6.15)$$

$$\frac{dh}{dt} = \lambda [h_\infty(V) - h] / \tau_h(V), \quad (6.16)$$

$$\frac{dn}{dt} = \lambda [n_\infty(V) - n] / \tau_n(V), \quad (6.17)$$

$$\frac{dx}{dt} = \lambda [x_\infty(V) - x] / \tau_x, \quad (6.18)$$

$$\frac{d[Ca]}{dt} = \rho [K_c x(V_{[Ca]} - V) - [Ca]], \quad (6.19)$$

$$\frac{d\eta}{dt} = -\frac{\eta}{t_c} + \frac{\xi(t)}{t_c}. \quad (6.20)$$

Here h , n , and x are gating variables, and $[Ca]$ is the intracellular calcium concentration. *There is no periodic input.* The simulations are done for constant temperatures. The precise form of the voltage dependencies of the gating variables and time constants can be found in Plant (1981). The correlation time of the noise $\eta(t)$ (an Ornstein-Uhlenbeck process) was chosen as $\tau_c = 1.0$ msec, so that the noise has a larger bandwidth than the fastest events in the deterministic equations. Our choice of parameters yields action potential durations of 140 msec at $T = 17.8^\circ\text{C}$ down to 20 msec at $T = 40^\circ\text{C}$. The noise, which is intended to represent mainly the effect of conductance fluctuations in the ionic channels and ionic pumps, is made additive on the voltage variable for simplicity.

This model reproduces the basic firing patterns seen at different temperatures (Longtin and Hinzer 1996). The dynamics of our model can be separated into a fast subsystem and a slow subsystem. The slow subsystem oscillates autonomously with a period corresponding to the interval between the beginning of two bursts. When this oscillation sufficiently depolarizes the cell, the fast action potential dynamics become activated. These action potentials "ride" on top of the slow wave. The effect of the temperature is to decrease the period of the slow wave, to speed up the kinetics governing the action potentials, and to hyperpolarize the slow wave (due mostly to

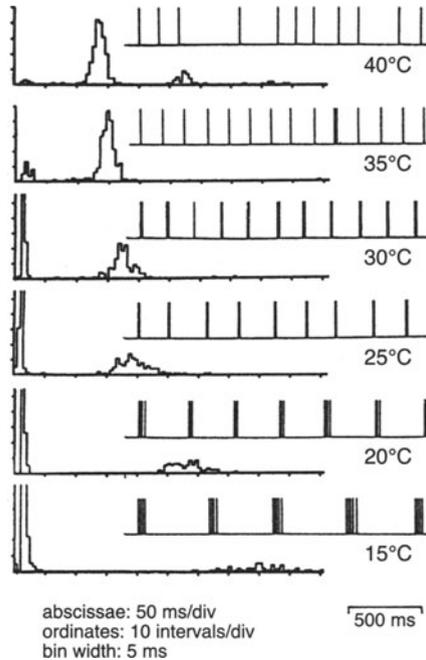


Figure 6.12. Interspike interval histograms and spike trains from cold receptors of the cat lingual nerve at different constant temperatures. From Braun, Schäfer, and Wissing 1990.

the Na–K pump). At higher temperatures, the slow wave is hyperpolarized so much that the threshold for the activation of the fast dynamics is not reached: The slow wave is subthreshold. From this point on, noise is important, since it can induce firings in synchrony with the slow wave. This is similar, then, to the situation seen above for neurons with external stimulation. The main difference is that the cold receptor has a stimulus built into its internal dynamics.

More recently, we have proposed a simplified phase model with periodic and stochastic forcing for the temperature dependence of such bursting activity (Roper, Bressloff, and Longtin 2000). The amplitude is proportional to the slow-wave amplitude, while the frequency is approximately $\sqrt{\lambda\rho}$ in our modified Plant model. This phase model can be analyzed for the boundaries between different solutions with different numbers of spikes per burst, and for the effect of noise on these boundaries.

6.13 Autonomous Stochastic Resonance

Near or in the skipping regime, i.e., at higher temperatures, the degree of phase locking between the spikes and the slow wave, as well as the char-

acteristics of the interspike interval histogram, are again highly sensitive to variations in noise level, or in the period, amplitude, and bias of the slow wave. All these characteristics can change in response to changes in the physicochemical environment of the cold receptor. Thus, in this regime, the cell is sensitive to such changes, due to the presence of noise. In particular, noise appears to be essential for encoding higher temperatures, since without noise, there would be no firings (Braun, Wissing, Schäfer, and Hirsch 1994). The behavior of the interspike interval histogram as a function of the aforementioned parameters is similar to that seen in skipping neurons with external stimulation.

Noise helps express the characteristics of the slow wave in the output spike train. This sensitivity can be characterized by computing averaged power spectra and signal-to-noise ratios. For the slow-wave burster, the spectra are more complicated than in the cases studied above (Longtin 1997). In particular, the background is bumpy, i.e., there are more peaks than the usual harmonics, and the harmonics themselves are broader. This latter feature is expected because the phase of the autonomous oscillation fluctuates. Nevertheless, one finds that there is an increase in the signal-to-noise ratio with noise intensity D , from 0 up to large values. At the point when the signal-to-noise ratio starts to drop again, the numerical simulations, and indeed, the stochastic model itself become doubtful. In fact, at these high noise levels, the action potential waveforms themselves become noisy.

This autonomous stochastic resonance behavior can be contrasted to that studied in Gang, Ditzinger, Ning, and Haken (1994). In that paper, the first claim of autonomous stochastic resonance, the stochastic resonance behavior was characterized in a system right at a saddle-node bifurcation. There, the noise induces a limit cycle that has zero amplitude in the absence of noise. As the noise becomes too large, the coherence of this limit cycle decreases. Thus this behavior is similar to stochastic resonance. However, it requires a saddle-node bifurcation. Our study of the slow-wave burster shows that autonomous stochastic resonance does not directly require a saddle-node bifurcation. Neural dynamics are often sufficiently complex to generate their own autonomous oscillations (through a Hopf bifurcation in the case of our modified Plant model), and can be expected to exhibit stochastic resonance phenomena as the noise intensity varies. A related effect known as coherence resonance has been analyzed in the FitzHugh–Nagumo system (Pikovsky and Kurths 1997). In this latter case, the induction by noise of regularity in the firing pattern can be theoretically linked to the specific statistical dependencies of the escape time to threshold and of the refractory period on noise intensity.

6.14 Conclusions

We have given a brief overview of recent modeling efforts of noise in physiologically relevant dynamical systems, and studied in detail the response of excitable cells to periodic input and noise. We have shown how noise can interact with bifurcations, produce smooth transitions between firing patterns as stimulus parameters are varied, and alter the frequency sensitivity of neurons. In most neural stochastic resonance studies, a neuron was chosen in which the addition of noise to the periodic stimulus could better transduce this stimulus. This approach of adding noise is warranted because the noise level cannot be reduced in any simple way in the system (and certainly not by cooling it down, as we suspect from our study of cold receptors). It is interesting to pursue studies of how internal noise can be changed.

There have been predictions of stochastic resonance in summing networks of neurons (Pantazidou, Moss, and Chialvo 1993), in which the input was aperiodic. There have also been theoretical studies of stochastic resonance in a summing neuron network for slowly varying aperiodic signals, i.e., for situations where the slowest time scale is that of the signal (Collins, Chow, and Imhoff 1995a; Collins, Chow, and Imhoff 1995b; Chialvo, Longtin, and Müller-Gerking 1997). There has also been a recent experimental study of the enhancement by noise of the transduction of broadband aperiodic signals (Levin and Miller 1996); in particular, that study investigated the effect of the amount of overlap between the frequency contents of the signal and of the noise added to the signal. In these latter two studies with aperiodic signals, the resonance is characterized by a maximum, as a function of noise, in the cross-coherence between the input signal and the output spike train, and by related information-theoretic measures such as the transinformation. Further work has been done, using simple neuron models driven by periodic spike trains and Poisson noise, to address situations of periodic and noisy synaptic input (Chapeau-Blondeau, Godivier, and Chambet 1996). Since correlations in spike trains increase the variance of the associated input currents, stochastic resonance has been recently studied from the point of view of increased coherence between presynaptic spike trains (Rudolph and Destexhe 2001). There have been experimental verifications of the stochastic resonance effect away from the sensory periphery, namely, in a cortical neuron (Stacey and Durand 2001). In addition, there is an increasing number of interesting applications of the concept of stochastic resonance to human sensory systems (Collins et al. 1996; Cordo et al. 1996; Morse and Evans 1996). All these studies emphasize the usefulness of the combination of the subthreshold regime and noise, and in some cases such as electroreceptors, of suprathreshold dynamics with noise (Chacron, Longtin, St-Hilaire, and Maler 2000) to transduce biologically relevant signals.

To summarize our stochastic phase locking analysis, we have shown the following:

- Many neurons exhibit skipping in the presence of periodic stimulation.
- Modeling shows that skipping results readily from the combination of subthreshold dynamics and noise. In other words, no deterministic firings can occur; firings can occur only in the presence of noise. It can occur with suprathreshold dynamics, namely, with chaos and/or noise, with some deviations from the skipping picture presented here (such as tunability of the pattern).
- Noise helps in the detection of low-amplitude stimuli through an effect known as “stochastic resonance.” This means that adding noise to these neurons allows them to transduce small stimuli that cannot by themselves make the cell fire. The effect relies on a linearization of the firing probability versus stimulus phase characteristic, which occurs in a stimulus-frequency-dependent fashion except when signals are slower than all neuron time scales. The simplest version of this effect in the presence of periodic forcing can be found in simple threshold-crossing systems when the signal is slow compared to all other time scales.
- Noise can extend the physical range of stimuli that can be encoded.
- Noise can alter the shape of tuning curves, and thus the frequency response characteristics of neurons.
- These results apply qualitatively to systems without external forcing, such as thermoreceptors.

Some important unresolved questions worthy of future investigations are the following:

- Can we identify more precise ionic mechanisms for the skipping patterns, based on data from intracellular recordings?
- To what extent are neurons in successive stages of sensory processing wired to benefit from stochastic resonance?
- What aspects of the skipping pattern determine the firing properties of the neurons they connect to, and why is this code, which combines aspects of a “random carrier” with “precisely timed firings,” so ubiquitous?
- Another important issue to consider in modeling studies is that the axons connecting to receptor cells are driven synaptically, and the synaptic release is governed by the receptor generating potential. In the presence of periodic forcing on the receptor, such an axon can

arguably be seen as driven by both a deterministic and a stochastic component. It will be useful to study how the synaptic noise can produce realistic skipping patterns, and possibly assist signal detection; such investigations will require using more detailed knowledge of synapses and connectivity between neurons.

- More generally, there is still plenty of work to be done, and, probably, effects to be discovered at the interface of stochastic processes and nonlinear physiological dynamics. Among the tasks ahead are the analysis of noise in systems with some form of memory, such as the pupil light reflex, or excitable systems that do not “recover” totally after a firing.

6.15 Computer Exercises: Langevin Equation

Software

There is 1 **Matlab*** program you will use for these exercises:

langevin A **Matlab** script that integrates the Langevin equation, equation (6.4), using the standard Euler–Maruyama method, equation (6.5). The solution is Ornstein–Uhlenbeck noise. The program generates the solution $x(t)$ starting from an initial condition $x(0)$. Time is represented by the vector t , and the solution by the vector x . The solution $x(t)$ can be seen using the command **plot(t,x)**. The histogram of solution can be seen using the command **plot(rhox,rhoy)**.

The following exercises are to be carried out using the **Matlab** program **langevin.m**. The parameters of the equation, i.e., the decay rate α **alpha**, the intensity of the Gaussian white noise D **dnz**, and the initial condition for the state variable **xinit**, are in **langevin.m**. The integration parameters are also in that program. They are the integration time step **dt**, the total number of integration time steps **tot**, the number of times steps to discard as transients **trans**, and the number of sweeps (or realizations) **navgs**. The description of the program parameters is given in comment in the **Matlab** file. These parameters can thus be changed by editing **langevin.m**, and then the numerical integration can be launched from the **Matlab** command window. The program generates the solution $x(t)$ starting from an initial condition $x(0)$. Time is represented by the vector t , and the solution by the vector x . A solution $x(t)$ can thus be plotted at the end of the simulation using the command **plot(t,x)**.

*See Introduction to **Matlab** in Appendix B.

A simulation involving a given set of random numbers is called a “realization.” These random numbers provide the source of noise for the code, and thus mimic the noise in the stochastic differential equation (which, as we have seen, models the noise in the system under study). For our problems, Gaussian-distributed random numbers are needed; they are generated internally by `Matlab` using the function `randn`. A random number generator such as the one built into `Matlab` needs a “seed” value from which it will generate a sequence of independent and identically distributed random numbers. `Matlab` automatically handles this seeding. Since random number generation is done on a computer using a deterministic algorithm, the independence is not perfect, but good enough for our purposes; nevertheless, one should keep this in mind if one uses too many random numbers, since such “pseudorandom” number generators will repeat after a (usually very large) number of iterations.

From a time series point of view, each realization will differ from the other, since it uses different sets of random numbers. However, each realization has the same statistical properties (such as, e.g., moments of probability densities of the state variables and correlation functions). In order to get a good idea of these properties, one typically has to average over many realizations. Each realization can start from the same initial condition for the state variables, or not, depending on which experimental protocol you are trying to simulate, or what kind of theory you are comparing your results to. In some cases, it is also possible to estimate these properties from one very long simulation. But generally, it is best to average over many shorter realizations, each one using a different initial value for the noise variable. This avoids the problem of finite periods for the random number generator, and allows good averaging over the distribution of the noise process. Averaging over multiple realizations also has the advantage of reducing the estimation error of various quantities of interest, such as the amplitude of a peak in the power spectrum. The codes provided here can easily be modified to perform one or many realizations.

The numerical integration scheme used here is a stochastic version of the Euler algorithm for ordinary differential equations, known as the Euler–Maruyama algorithm (6.5).

Exercises

The purpose of these exercises is to study the effect of the noise intensity and of the parameter α on the dynamics of the Langevin equation

$$\frac{dx}{dt} = -\alpha x + \xi(t). \quad (6.21)$$

Note that this equation has, in the absence of noise, only one fixed point, at $x = 0$.

Ex. 6.15-1. Effect of the parameter α on the dynamics.

Run simulations for various values of α and Gaussian white noise intensity D . You should find that the solution looks smoother when α is larger. Also, you should find Gaussian densities for the state variable x ; note that the program actually estimates these densities by unnormalized histograms of the numerically generated solution. Increasing α for a given noise intensity will reduce the variance of the histogram of the x -solution. You should also find that the variance is given by D/α ; you can try to show this by calculating the stationary solution of the associated Fokker–Planck equation.

Ex. 6.15-2. **Effect of noise intensity on the dynamics.** Also, increasing the noise intensity while keeping α constant has the opposite effect. The solution of this equation can be used as a source of “Ornstein–Uhlenbeck” colored noise.

6.16 Computer Exercises: Stochastic Resonance

These exercises use the `Matlab` file `fhnnoise.m` to simulate the FitzHugh–Nagumo excitable system with sinusoidal forcing and colored (Ornstein–Uhlenbeck) additive noise on the voltage equation:

$$\epsilon \frac{dv}{dt} = v(v-a)(1-v) - w + I + r \sin \beta t + \eta(t), \quad (6.22)$$

$$\frac{dw}{dt} = v - dw - b, \quad (6.23)$$

$$\frac{d\eta}{dt} = -\lambda\eta + \lambda\xi(t). \quad (6.24)$$

Here $\lambda = 1/t_{\text{cor}}$; i.e., λ is the inverse of the correlation time of the Ornstein–Uhlenbeck process; we have also chosen a commonly used scaling of the noise term by λ . You can explore the dynamical behaviors for different system, stimulus, and noise parameters. Note that in this program, the periodic forcing is added to the voltage equation. The simulations can be lengthy if smooth interspike interval histograms are desired, so you will have to decide, after a few tries, how long your simulations should be to answer the questions below. An integral (order-1) stochastic Euler–Maruyama algorithm is used here to integrate this system of stochastic differential equations. The time step has to be chosen very small, which limits the integration speed.

Software

There is one `Matlab`[†] program you will use for these exercises:

[†]See Introduction to `Matlab` in Appendix B.

fnnoise A Matlab script (operates in the same way as **langevin.m**) integrates the FitzHugh–Nagumo system equations (6.12) driven by sinusoidal forcing and colored noise. The colored noise is the Ornstein–Uhlenbeck process $\eta(t)$. The program uses the integral Euler–Maruyama algorithm proposed in Fox, Gatland, Roy, and Vemuri (1988). The outputs of the program are the solution $x(t)$ and the interspike interval histogram $\rho(I)$. The solution can be seen using the command **plot(t,v)**. The interspike interval histogram can be seen using the command **plot(rhox,rhoy)**. The sequence of intervals can be plotted using **plot(interv)**.

The description of the program parameters is given in comment in the **Matlab** file. These parameters can thus be changed by editing **fnnoise.m**, and then the numerical integration can be launched from the **Matlab** command window.

The program **fnnoise.m** operates in the same way as **langevin.m**. The main equation parameters you may be interested in changing are the amplitude **amp** and frequency **f** of the sinusoidal forcing; the bias current **ibias**, which brings the system closer to threshold as it increases; and the noise intensity **dnz**. The noise intensity specified in the file refers to the intensity of the Gaussian white noise D , which, when lowpass filtered, gives the colored Ornstein–Uhlenbeck noise. This latter noise is also Gaussian, with an exponentially decaying autocorrelation function, and its variance is given by D/tcor , where **tcor** is the noise correlation time.

The simulation parameters are controlled by the total number of integration time steps **tot**, the number of time steps considered as transients **trans**, and the number of sweeps (or realizations) **navg**. The integration time step is controlled by **dt**, and may have to be made smaller than its reference value 0.0025 for higher sinusoidal forcing frequencies, i.e., for $f > 1$ or so. You can always check that results are accurate with a given time step by checking that they are statistically the same as for a new simulation with smaller time step.

Exercises

- Ex. 6.16-1. **Hopf bifurcation in the absence of noise.** In the absence of noise and periodic forcing, find the Hopf bifurcation point in the FitzHugh–Nagumo model by varying the bias current parameter I . You can also try to calculate this value analytically. This parameter controls the distance between the resting potential and the threshold. Below the Hopf bifurcation, the system is said to be in the subthreshold or excitable regime. Above this bifurcation, it is said to be in the repetitive firing regime.

- Ex. 6.16-2. **Effect of I on firing frequency.** For the same conditions as above, how does the frequency of firing depend on I ? (It changes abruptly near the bifurcation point, but varies little thereafter.)
- Ex. 6.16-3. **Effect of noise intensity on interspike interval histogram (ISIH) with noise.** Compute the interspike interval histogram with noise only (amplitude of sinusoidal forcing is set to zero), and study the behavior of this interspike interval histogram as a function of noise intensity D . (You should find that the maximum of the distribution shifts slightly to smaller interval values, but that the mean shifts over more significantly.)
- Ex. 6.16-4. **Effect of stimulus amplitude on the interspike interval histogram envelope.** Study the behavior of the interspike interval histogram envelope as a function of stimulus amplitude, in the subthreshold regime. Can the second peak be the highest? (The envelope decays more rapidly the higher the amplitude of the sinusoidal forcing. Yes, the second peak can be the highest, especially at higher frequencies or low noise.)
- Ex. 6.16-5. **Effect of stimulus frequency on interspike interval histogram envelope.** Study the behavior of the interspike interval histogram envelope as a function of stimulus frequency $\beta/(2\pi)$, in the subthreshold regime. Is the first peak always the highest? (No; see previous question.)
- Ex. 6.16-6. **Effect of noise intensity on interspike interval histogram envelope.** Study the behavior of the interspike interval histogram as a function of the noise intensity D , in the subthreshold regime. (The noise plays a similar role to the amplitude of the sinusoidal forcing: The higher the noise, the faster the decay of the histogram envelope. However, increasing noise also broadens the peaks in the interspike interval histogram.)
- Ex. 6.16-7. **Subthreshold and suprathreshold regimes.** Compare the behavior of the interspike interval histogram in the subthreshold and suprathreshold regimes. In the subthreshold regime, increasing noise always increases the probability of shorter intervals. In the suprathreshold regime, noise can perturb the limit cycle, producing longer intervals than the cycle period. Hence, in this case, increasing noise does not necessarily decrease the intervals.
- Ex. 6.16-8. **Stochastic Resonance**
Plot the maximum of the peaks (or some area in the interspike interval histogram around this peak, to average out fluctuations) as a function of the noise intensity in the subthreshold regime, and in the suprathreshold regime. You should find that the value of the maximum for the first peak goes through a maximum as a function of D .

Note that, in this regime, no spikes can be generated without noise, and too much noise leads to solutions dominated by noise. Hence, a moderate value of noise causes a predominance of firing intervals around the stimulus period, which is a manifestation of stochastic resonance. You should find that the other peaks also go through maxima. You can also modify the code to compute power spectral densities for the spike trains generated by this stochastic system. For example, you can generate a vector of zeros and ones that resamples the solution $v(t)$ at a lower frequency; a zero represents no spike in the corresponding time bin of this vector, while a one represents one (or more) firing events in that time bin. You can then call the spectral functions (such as the fast Fourier transform) in `Matlab`, and average the results over many realizations to reduce the fluctuations in the spectra. This computation requires more background material and is not pursued here.