Resonating models for the electric power market

Carlo Lucheroni*
Dipartimento di Matematica e Informatica, Università di Camerino, via Madonna delle Carceri 9, 62032 Camerino (MC), Italy
(Received 3 June 2007; published 21 November 2007)

This paper describes the economic phenomenon of *price spiking* in electric power markets and introduces an alternative way to model it. A stochastic FitzHugh-Nagumo dynamics in a special regime is proposed as a basic model for the power market, and an extension of the FitzHugh-Nagumo system is introduced to improve the statistical features of the basic model. Ideas from stochastic and coherence resonance are used to discuss the models.

DOI: 10.1103/PhysRevE.76.056116 PACS number(s): 89.65.−s, 05.40.−a, 05.45.Tp, 05.45.Xt

I. INTRODUCTION

A current financial problem is the explanation and modeling of a disappointing behavior of prices in contemporary electricity markets. After a worldwide change in the electric industry regulation policies, markets have been transformed from regulated and monopolistic to (partially) competitive. Prices have been left free to fluctuate, and fluctuations turned out much wilder than expected, certainly very much wilder than the well-known stock market fluctuations, and very different from them in quality. The most striking aspect of electricity price time series is *spiking*, the phenomenon that is the subject of this paper. No commonly accepted financial model currently exists which can describe spiking in a satisfactory way. Outside finance, econophysics is interested in price dynamics [1,2]. Financial problems are often successfully studied using ideas and models from physics, and physics benefits from this interdisciplinary cross-breeding. Spiking will be discussed in this paper using ideas from the field of stochastic dynamical systems, namely, resonance and the role of noise in resonance, drawing from stochastic and coherence resonance studies [3,4]. Two dynamical systems will be used to study spiking, the well-known FitzHugh-Nagumo system (FNS) [5] and a proposed extension of it. They will be used to implement two models of the electricity market—or *power market* as it is often called [6]—and these models will be compared with real-world data from the Alberta, Canada, power market [7]. A power market will be considered here as an institutional structure that in a certain way can resonate with the consumer electricity demand, which is periodic. In the proposed models, an interplay between nonlinearity and noise controls the resonance patterns, and such resonance patterns can take the form of spikes. Three time scales are involved in the resonance, hidden to some degree in the parameters of the models, and this feature makes the proposed models different in nature from the scale-free models that are often used to study the stock market fluctuations in physics [1,2,8]. The models are continuous in prices and in time, and depend on a single source of noise, which makes them particularly interesting for financial uses and for econometric applications.

The paper is divided into nine sections. Market data are presented and discussed in Sec. II. A brief introduction to the fascinating and multifaceted economics of power markets is given in Sec. III. Current mathematical models of these markets, all of financial origin, are discussed in Sec. IV. Section V reviews two aspects of the well-known FNS on which the proposed models are based. Section VI shows how the FNS can be used as a basic model for the power market and introduces an extension to it. In Sec. VII some tools to extract and compare statistics from data and simulations are reviewed. Section VIII discusses the further possibility for these models to be used even without any regard to electricity demand, and some general problems encountered when using power market data from a physics perspective. Section IX concludes the paper.

II. THE ALBERTA POWER MARKET

Figure 1(a) shows the hourly time series of day-in-advance capped pool prices per MW h quoted by the Canadian Alberta Electric System Operator (AESO) [7], from April 7, 2002 to April 6, 2007. AESO is the independent system operator (ISO) and power exchange for the Alberta

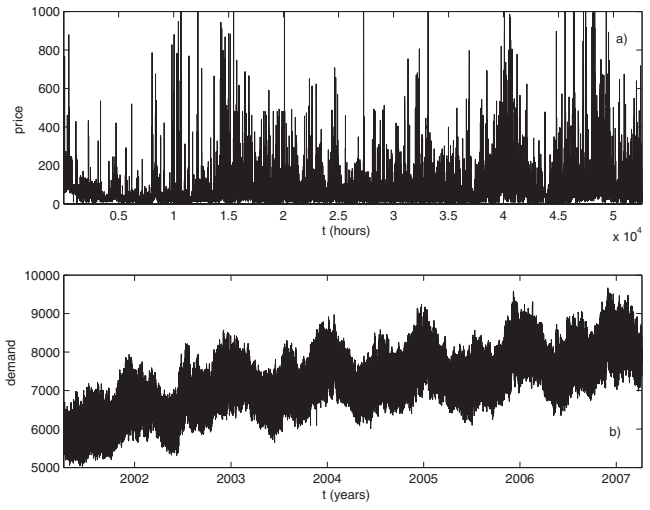

FIG. 1. Prices and demand in the Alberta power market: five years from April 7, 2001 to April 6, 2007. (a) Hourly prices in C$ units vs time in hours counted from 1 a.m. of April 7, 2001; (b) hourly energy demand in MW h; time span as in (a) but expressed in actual years.

*carlo.lucheroni@unicam.it

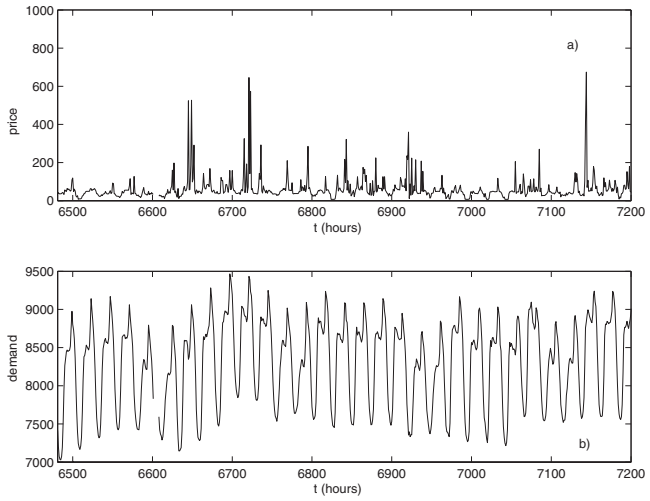

FIG. 2. Prices and demand in the Alberta power market: one month from January 1, 2007 to January 31, 2007, time in hours. (a) Hourly prices in C$ units; (b) hourly energy demand in MW h.

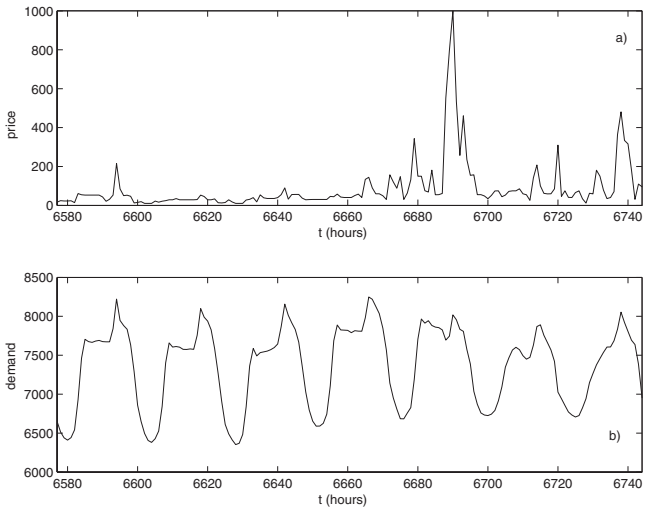

FIG. 3. Prices and demand in the Alberta power market: one week from Monday January 6, 2003 to Sunday January 12, 2003, time in hours. (a) Prices in C$ units; (b) hourly energy demand in MW h.

power grid (AIES), and it is a mandatory pool. The AIES is connected with other major grids, and participating producers burn mostly coal and gas. Pool prices, in Canadian dollars (C$), are hourly averages of system marginal prices (SMPs) determined by AESO every minute, and the cap is 1000 C$ per MW h. AESO operates also a real-time balancing market and other subsidiary and financial markets [9], and posts a default price of 1000 C$ in some very special and rare cases. Figure 1(b) shows the recorded associated demand in MW h.

Price formation in power markets can be very different from price formation in the stock market. Moreover, each power market is different from all other power markets. Electricity price formation relies on so many interacting levels, each of them with many possible variants: the primary fuel markets, the physical electricity transport network, the institutional settings and rules that match demand and supply, the players' decision patterns and interactions, the pure financial level, and the interregional interactions. Electricity price formation ultimately depends on a repeated decisions game with a price as its outcome, played under a number of constraints that vary from market to market. This notwithstanding, and maybe surprisingly, stylized facts more or less common to most power markets can be extracted from price time series coming from very different markets. The existence of stylized facts encourages hopes of framing in a mathematical model the essence of abstract power market behavior.

Figure 1 shows at first sight the signs of the three most important stylized facts about power markets: *A lot of spikes*, which are sudden price jumps, first up and then down; *mean reversion*, typical of interest rates and commodity markets but uncommon in the stock market—in the intermediate and long run, prices tend to return to an average fixed value; and *presence of a trend and seasonality*, with multiple time scales. Figures 2 and 3 explore the two time series at two shorter time scales, one month and one week. Figure 1 shows clear winter-summer 12- and 6-month periodicity (on the back of a constant growth) for the demand, but for prices this periodicity is not obvious. Possibly an increase of spiking activity is associated with demand crests at year changes. Figure 2 shows clear weekly, night-day, and morning-afternoon periodicity for the demand, and reveals a strong 24 h periodicity for the prices. In particular, spikes appear only in a limited period of the day, less than 12 h, on the top of a more regular small-amplitude background. Figure 3 shows working day–weekend periodicity, but most importantly it shows the structure of the spikes in the clearest way. Spikes have a structure and a width, comparable in average with one-quarter of a day. Higher-frequency data for scales shorter than 1 h are not available, and data resolution for an average spike is six data points. Except for strong statistical correlation to daylight hours, and the strong sensation that spikes require some threshold crossing in demand before being fired, Figs. 1–3 give no clue about the mechanism at work when (and right after) a spike is started. Even worse, most of the seasonality's characteristic periods are hidden with regard to the price series and clearly visible only in the associated demand series. Models have to confront these data, but also have to take into account the economics of the underlying phenomenon.

III. ECONOMICS OF THE ALBERTA POWER MARKET AND SOME EXPLANATIONS OF SPIKING

In this section some very basic information about the economics of the Alberta power market is reviewed, mainly for two reasons. First, since a nonlinear dynamics is proposed as a model for such a market, it is important to understand where this nonlinear mechanism could originate. Second, since the model proposed is different from more common stock market models, it is important to understand where the difference between power markets and stock markets comes from, to put the proposed model in the right perspective.

A power market is a market where committed promises of delivery of units of electric energy (usually expressed in MW h), throughout given hours and at given places, are exchanged. Electric energy, which is the source of this market, is a very peculiar commodity [10], with a peculiar economics, which here will be analyzed at four levels: storage and generation, transport network, institutional, and financial.

Electric energy can be stored in a very limited way, so that it has to be produced spot on demand. Its generation costs depend on the intensity of the demand. At standard demand intensity (off-peak time) only less expensive base generators are active, burning steadily nuclear fuel, oil, or coal. In case of a sudden high demand (peak time), more expensive generators, burning gas, are turned on for a short time. Storage difficulty results in difficulty in smoothing offer and demand mismatches.

Transport of electricity from producers to customers is made through the *power network*, or *grid*. Local and nationwide power networks are connected among each other by few links, whereas local topology can be very intricate. In each local network, producers and consumers share the same wires, which have a *limited transport capacity*. Energy transfer from one grid point to another affects immediately, and always all players, in a coherent way, due to the Kirchhoff laws. Coherency is very strict also because the whole network is synchronized to the same frequency. Even small instantaneous imbalances of demand and supply can congest the whole network and disrupt its functioning, not to mention link or equipment failures. Characteristic of such a coherent (and sometimes scale-free [11]) system is the temporal and spatial dynamics of its failures, which sometimes result in blackouts. Another very interesting aspect of power grids is the coevolution in advanced markets of grid development and growth of electricity demand, which can make of a power grid an adaptive, perpetually out-of-equilibrium, critically self-organized system [12].

At the institutional level [6], in the Alberta market few producers match their electricity offer with the demand of many consumers. AESO carries on two distinct functions and for this reason it is called a *pool* [13]. First, it operates the network, taking care of the production and delivery timing and of the transport congestions as an independent system operator. Second, it operates a *power exchange*. AESO aggregates demand and supply following an auction scheme [14], taking into account engineering requirements and computing the SMPs, which are matching prices that also optimize the electricity flow in the grid. For this market, the Canadian regulator imposes a maximum attainable price, called a *cap*. As in most exchanges, to ensure timely delivery to final customers and proper power grid operation, hourly prices for 1 MW h and for each day are auctioned during a day-long session held one day in advance of the day of the physical delivery (*day-in-advance market*). Since producers are few, supply side price formation is the result of competition among oligopolistic producers [15], and it is usually studied with a Cournot model where collusion among suppliers is taken into account as possible and frequent [16]. When achieved, collusion leads to *market power*, which results in prices much higher than perfectly competitive prices (like those in the stock market) and possible restraints on power production (a means to raise prices).

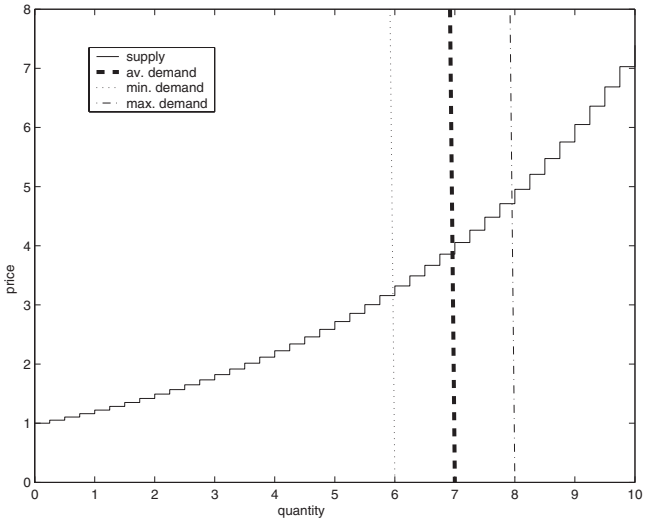

FIG. 4. Generation costs vs quantity produced, in arbitrary units. Supply of goods is offered for a cost price (continuous curve), and demanded quantity (thick dashed line) sets the equilibrium price at its crossing with supply. If demand decreases (dotted line) or increases (slash-dotted line) with respect to some average value, the equilibrium price decreases or increases following demand variations.

At the financial level, bilateral contracts with delivery prices set months in advance can be exchanged inside or outside the market and interact strongly with the exchange price mechanism through expectation formation. Finally, it should always be remembered that a power market is essentially an evolving and adapting social institution, and exchange price time series often embed abrupt rule changes. Not surprisingly, power grids are sometimes considered the most complicated machines man has ever built [17,18], and power markets are just the decentralized and intricate decision schemes designed to operate these machines.

Coming to the explanations of spiking, it should be noted that spikes can be a bad thing. They make free-market energy trading a risky business, and customers exposed to marked prices upset. But since nobody would refrain from turning lights on in the dark before checking electricity quotations, demand is usually considered rigid in the short run, not very sensitive to price levels. Microeconomic theory proposes some explanations, not mutually exclusive, about spike generation. Here only the five most common ones are reported and shortly commented on. *Generation costs* increase sharply with produced quantity, as shown in Fig. 4 (rigidity of demand is represented by the extreme steepness of the demand curves), since after a certain level in required power increasingly costly generators must be turned on in a sequence (and starting itself is costly). Equilibrium prices then cyclically and rigidly follow production costs for a cyclic demand. This explanation alone is not very convincing, since spike shape is far too variable in comparison to just variation in demand. *Fuel costs*, especially for gas [19], are mean reverting, modestly spiking themselves, and very volatile, but they can explain electricity price variations only on scales much longer than a few hours. *Grid congestions* are more

interesting, since they are very frequent and link demand to a nonlinear interplay between the spatial network structure (like the blackouts), the transmission rights market, and strategic behavior of producers and subnetwork owners finalized to raise prices. In the case that a grid can be represented as a scale-free network steadily close to criticality (maybe because of coevolution with demand growth), congestions could have an even stronger impact on price formation at a short time scale. *Correlation* with other spiking grids [20] could be considered, but that would not explain why the other grids spike. Last but not least, the *auction protocol* is an extremely interesting candidate. The power exchange does not compute SMPs as crossings of demand and "objective" supply curves. Producers take strategic decisions when posting quantities and prices, they can read published data series in which competitors hide their countermeasures, and post "strategic" supply curves [21]. When they are few, it is too tempting not to exercise their market power, and even without explicit communication herding behavior can set in. An improvidently designed auction protocol combined with the possibility of easily induced congestions (the power exchange takes into account the ISO constraints when forming prices) can be a powerful mechanism for nonlinear and threshold behavior of prices. Only the triggering instant is not certain, but in this case a noisy environment can play a major role.

IV. CURRENT MATHEMATICS OF POWER MARKETS AND SCALING

Spikes are studied mostly in the frame of standard financial modeling. Continuous time and continuous price stochastic interest rate models are often considered as the starting point for such modeling [10]. As is very common in finance, an additive or multiplicative uncorrelated continuous process is used as a noise source, modulated by a quantity called volatility. An extra Poisson point process can be added, giving up price continuity, and, if continuity is not considered essential, more general Lévy processes can be explored. Jump processes might better mimic power price variability, but they open the door to mean reversion problems because it is hard to have a process strongly revert only after a discontinuous jump, and weakly revert for the rest of the dynamics. A strategy that preserves continuity is that of embedding a continuous model in a Markov M-state regime structure [10,22], where the parameters of the model change in time following an accessory M-state process. Two-state processes are interpreted in terms of system switches between two financial states, maybe turbulent and normal. Other models include stochastic volatility or they are essentially discrete in time. A less standard strategy is to give up dynamical equations while preserving continuity, by using infinitely divisible cascades (IDCs). IDCs [23] can be used to build continuous stochastic processes by defining them starting from their distributional and scaling aspects, and they have multiscaling properties. In finance they were introduced in [24] for the stock market. In [25] such a scheme is used to study data from NordPool, a Scandinavian power exchange. A wavelet analysis is carried out and the results are not conclusive about multifractality, but this kind of analysis is interesting because it is close in spirit to many studies made by physicists about scaling, fat tails, and memory in finance [8]. Figure 5 shows a multiscale wavelet analysis [26] for the Alberta price series. In the diagram, to each hour in the studied series (on the abscissa) and to each scale from a given range (on the ordinate) is associated the modulus of the wavelet coefficient (different gray intensities correspond to different values), giving immediate visual information about the local regularity of the series at a given scale. In this case no special inhomogeneity seems to show up at short or intermediate scales, ruling out multifractality. If scaling is suspected, a direct autocorrelation or spectral analysis of the price series can be made, as shown in Fig. 6, but if nonstationarity is suspected as well, autocorrelation results should

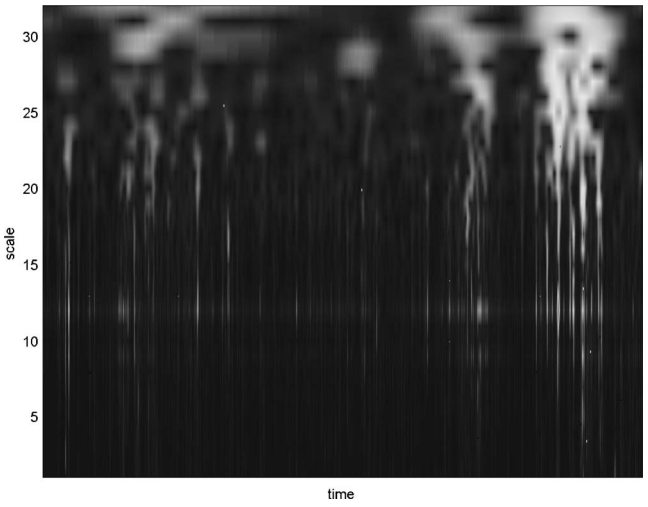

FIG. 5. Multiscale analysis: Morlet continuous wavelet diagram of the Alberta power market price series; on the time axis all hours from April 7, 2001 to April 6, 2007; on the vertical axis the explored scales—the shortest scale at the bottom, in hours.

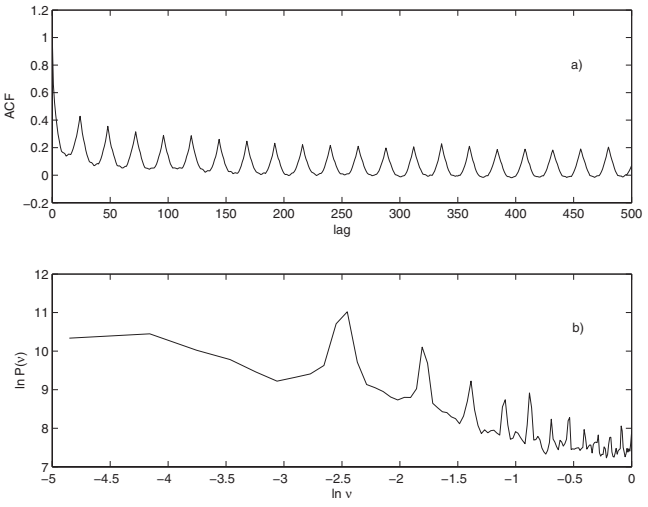

FIG. 6. Alberta power market price autocorrelation function (ACF) and spectrum. (a) ACF vs lags expressed in hours; (b) power spectrum $P(\nu)$ on a log-log scale, where $\nu=1/(\text{lag})$.

be interpreted with care (in this case the wavelet analysis is more suitable than autocorrelation). Figure 6 will be discussed more accurately in Sec. VIII.

V. STOCHASTICALLY RESONATING SPIKING AND THE SOFT ϵ REGIME IN THE FNS

In Sec. VI the FNS will be taken as a basic mathematical model for the power market. In this section, an overview of the stochastic FNS will illustrate why at first sight such a system can seem unsuitable for this purpose, and it will suggest how to find a resonating regime and a choice of parameters that can make the FNS work rather well as a power market model. The focus will be on two features of the FNS stochastic dynamics, namely, the intrinsic presence of three competing time scales and the complex reaction of the FNS to a stochastic periodic forcing.

The FitzHugh-Nagumo model originates from mathematical biology, where it was introduced as a simplified model of the deterministic voltage $x(t)$ and current $y(t)$ dynamics in time t at a chosen spot in the neuron membrane, as a reaction to an external stimulus $f(t)$. In this model, the initially quiescent neuron reacts to a current pulse by firing a spike in voltage. During the spike time and for a stretch of time after the spike the neuron cannot react to other stimuli—this is called its *refractory period*, of time length T_r. Right after the spike the voltage comes back to its previous value, stays there until the end of the refractory period, and comes back to its *quiescent state*. This dynamics is modeled by the two original FitzHugh-Nagumo deterministic coupled nonautonomous ordinary differential equations [5,27] in their stochastic version [4,28,29] as

$$\epsilon \dot{x} = g(x) - y, \tag{1a}$$

$$\dot{y} = \gamma x + b - \beta y - f(t) + \sigma(d)\xi, \tag{1b}$$

where

$$g(x) = \kappa x - \lambda x^3. \tag{2}$$

In Eqs. (1a) and (1b) $\epsilon > 0$, γ, b, $\beta \geq 0$, and $\sigma > 0$ are constants, ξ indicates the derivative of the Wiener process, i.e., a $\delta(t)$-autocorrelated normal process, and $\kappa > 0$, $\lambda > 0$. The standard deviation σ sets the noise intensity and it is commonly defined in terms of a diffusion constant $d > 0$ as $\sigma(d) = \sqrt{2d}$. γ couples the two equations, and for $\gamma = 0$ and $f = 0$ the second equation becomes a stand-alone Ornstein-Uhlenbeck process, which reverts to its mean b on a time scale $1/\beta$. For $\gamma = 0$ Eq. (1a) can be interpreted as describing the unidimensional overdamped motion of a particle in a double-well potential $U(x) = -\int^x g(x')dx'$, with either one or two minima x_{0-} and x_{0+}, local or global. In fact, the number and type of minima are parametrically controlled by y. For $y = 0$, $U(x)$ has two global minima. For $y \neq 0$ it has only one global minimum whereas the other minimum is either only local or disappears. When $\gamma \neq 0$ the two equations can be viewed as subject to different time scales. Taking the scale in which t is measured as a reference, in the first equation time gets multiplied by $1/\epsilon$. The FNS is usually studied in the

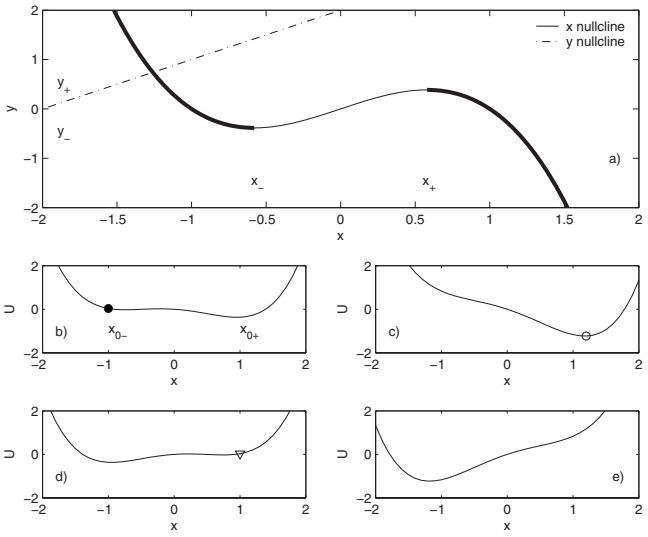

FIG. 7. Nullclines and potential U for $\kappa = 1$, $\lambda = 1$, $b = 2$, $\gamma = 1$, $\beta = 1$, $f = 0$, where the meaning of x_+, x_-, y_+, y_-, x_{0+}, x_{0-}, ●, ○, and ▽ is explained in the text; U, x, and y in arbitrary units. (a) Phase space with x and y nullclines, the left and right branches of the x nullcline being indicated by a thickened line; (b)–(d) potential $U(x)$ for different values of y: $y =$ (b) -0.25, (c) -1, (d) 0.25, and (e) 1.

singular regime where $\epsilon \ll 1$. In this regime every change of x is faster than any change of y. In this case, x is called the *fast variable* and y the *slow variable*.

In a phase plane picture and in the singular regime, it is then significant to draw the x and y *nullclines*, which are the static curves obtained, respectively, for $\dot{x} = 0$ and $\dot{y} = 0$ (and $\sigma = f = 0$). The two nullclines for Eqs. (1a) and (1b) are shown in Fig. 7(a), with a continuous line for the x nullcline and a dot-slashed line for the y nullcline. The x nullcline has three branches, separated along the abscissa by the points x_- and x_+ for which $g'(x) = 0$, which correspond, respectively, to y_- and y_+ along the ordinate. The left and right branches are called the *stable branches*. The branch in between is called the *unstable branch*. For $y \in [y_-, y_+]$ the phase plane region included between the two stable branches can be called the *core region*, where most of the dynamics is confined. The intersection $P_* = \{x_*, y_*\}$ of the two nullclines can set the stable or unstable *fixed points* of the dynamics. Outside the singular regime, for $\epsilon \approx 1$, the fast and slow dichotomy is blurred and the nullclines are more a guide for the eye than an analytic tool.

In the singular regime, a spike can be understood as a feedback cycle in the core region of the phase plane, following the sequence of the three marks (●,○,▽) in Figs. 7(b)–7(d)—where $U(x)$ is depicted for some values of y—and in Figs. 8(a)–8(c). At $y = -0.5 < y_-$ the system state vector $\{x(t), y(t)\}$ leaves the quiescent state corresponding (and close) to the leftmost of the two available minima (●), and quickly rolls from the bottom of the left stable branch where the potential has just lost one of its minima to the bottom part of the right stable branch (○). Then the system waits on the right branch of the x nullcline (▽) until $U(x)$ is lifted up enough to loose its right-hand minimum.

Keeping the parameters of the x nullcline (i.e., κ and λ) fixed, a useful function of the y nullcline is that of control-

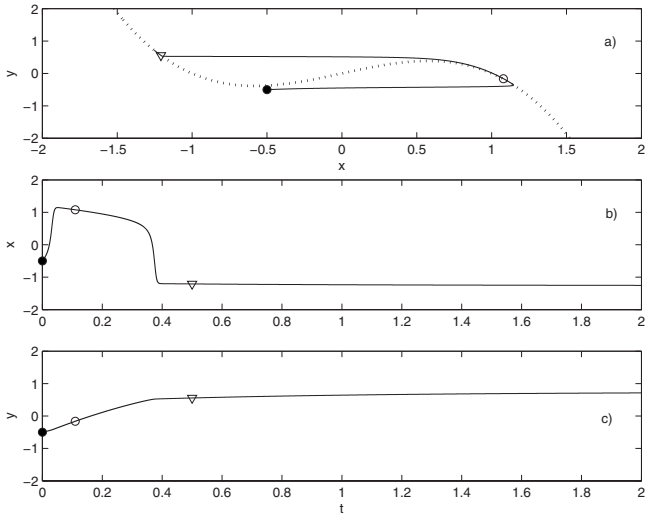

FIG. 8. Spiking mechanism in the absence of noise; ●, ○, and ▽ as explained in the text; x, y, and t in arbitrary units: $\epsilon=0.01$, $\kappa=1$, $\lambda=1$, $b=2$, $\gamma=1$, $\beta=1$, $f=0$, $d=0$. (a) Phase space, (b) $x(t)$, and (c) $y(t)$.

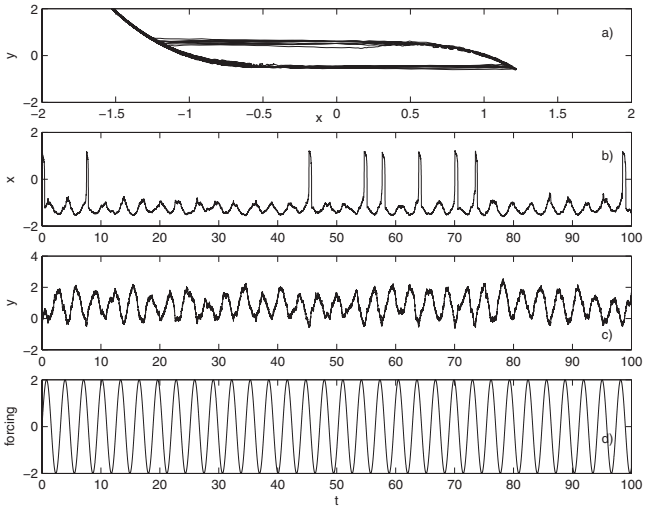

FIG. 9. Stochastic forced FNS in the SRS regime: $\epsilon=0.01$, $\kappa=1$, $\lambda=1$, $b=2$, $\gamma=1$, $\beta=1$, $A=2$, $\omega_f=2$, $d=0.1$, x, y, and t in arbitrary units. (a) Phase space, (b) $x(t)$, (c) $y(t)$, and (d) $f(t)=A\sin(\omega_f t)$.

ling parametrically (as a function of b, for example) whether the FNS operates *suprathreshold* in a *self-sustained oscillation regime* or *subthreshold* in an *excitable dynamics regime*. If the nullclines cross each other only once and on the stable branch, the fixed point is stable and the system is subthreshold. If they cross only once on the unstable branch, the fixed point is unstable and the system is suprathreshold. This means that in this second case the fixed point repels the dynamics, and the system keeps on cycling in the core region, while $x(t)$ traces a *spike train*. In the singular regime, the time shape of a single spike takes the form of two almost vertical ramps—the interbranch rolls are very quick—joined by a flat arch. When the system supports spike trains, all spikes have the same height, which is set once for all by the choice of $g(x)$. This is good for neuron modeling, but not good for power market modeling, where spike heights vary considerably. And this is probably the reason why the FNS has not been used in published finance literature as yet.

Equations (1a) and (1b) can be merged into a single second-order equation for x:

$$\epsilon\ddot{x} = (\kappa - 3\lambda x^2 - \epsilon\beta)\dot{x} - (\gamma - \beta\kappa)x - \beta\lambda x^3 - b + f(t) - \sigma\xi. \tag{3}$$

This shows that the FNS of Eqs. (1a) and (1b) is a nonlinear, damped, forced, stochastically driven oscillator. If a sinusoidal forcing $f(t)=A\sin(\omega_f t)$ $(A>0)$ with frequency ω_f and period $T_f=2\pi/\omega_f$ is applied through f, this oscillator can *resonate* with a broad resonance centered around some ω_i with period $T_i=2\pi/\omega_i$. In Eq. (3) the limit $\epsilon\to 0$ is singular. It has also to be noted that ϵ can be increased toward 1 to slow down the fast variable during a forcing cycle of frequency ω_f. In Eq. (1b) β can be absorbed into t. This is equivalent to saying that time is measured in units of $1/\beta$ and that the time scale ϵ has just to be considered as relative to $1/\beta$. Setting from now on $\beta=1$ means simply that all parameters on the right-hand side (RHS) of Eq. (1b) are taken relative to an implicit β. The forced deterministic FitzHugh-Nagumo oscillator can be studied in its A-ω_f parameter plane, to show the possibility of *phase locking* [30,31]. When the system is set subthreshold by the use of b, for each ω_f there is a region $A\in[0,A_{\max}(b,\omega_f)]$ in which spikes are not fired that is called the *silent region*, where only small oscillations are possible. When A crosses $A_{\max}$ from below, full spikes appear, and no intermediate situation is allowed. For $A>A_{\max}$ special curves called *deterministic* $(m:n)$ tuning curves can be drawn in this space after imposing the condition that for every n forcing periods exactly m spikes are fired. Then, there are at least two time scales characterizing a sinusoidally driven FNS: a T_i due to its internal frequencies or alternatively a T_r due to the basic spike refractory period, and a T_f due to the forcing. T_r (or T_i) is a function of ϵ. In the stochastic case, a third time scale T_n can be attributed to the noise spike activation effects.

If noise and periodic forcing are added to a FNS inside the silent region and close to its edge, with $T_f/2$ close to T_r, A slightly less than $A_{\max}(b,\omega_f)$, and $\epsilon\ll 1$, the system behaves as in Fig. 9. In coincidence with a forcing cycle, sometimes the system fires a spike, and sometimes it does not. The background dynamics is due to small oscillations and noise. Just before a spike, the slow dynamics is reverting to $\gamma x_{0-}+b-A$, heading toward y_-, where the first minimum x_{0-} is lost, but not being able to cross it. Statistically, it can happen that the noise's extra contribution to $y(t)$ pulls the system below y_-, where the spiking mechanism is triggered. This can be seen happening eight times in Fig. 9(c). When this happens, the system is on a so-called *stochastic limit cycle* [4,29]. This means that, even subthreshold, when assisted by the noise, the system is able to behave statistically as suprathreshold. In the presence of noise, the deterministic $(m:n)$ phase-locking condition is replaced by a statistical $(m:n)$ constraint, which requires that on average m spikes are fired every n cycles. For example, the system in Fig. 9

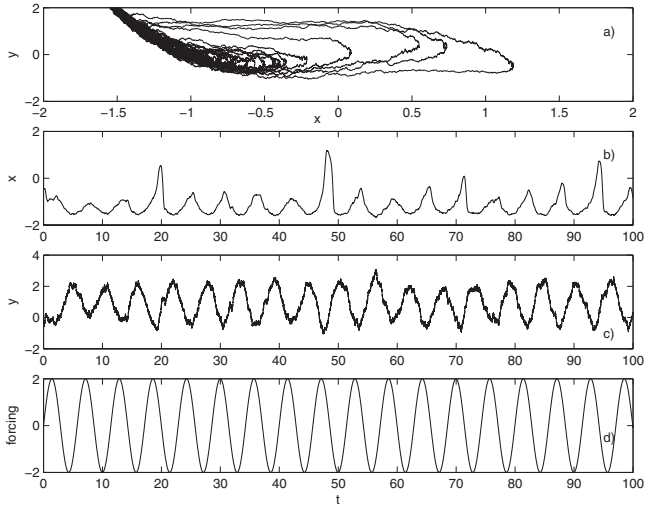

FIG. 10. Stochastic forced FNS in the soft regime: $\epsilon=0.3$, $\kappa=1$, $\lambda=1$, $b=2$, $\gamma=1$, $\beta=1$, $A=2$, $\omega_f=1.1$, $d=0.1$, x, y, and t in arbitrary units. (a) Phase space, (b) $x(t)$, (c) $y(t)$, and (d) $f(t)=A\sin(\omega_f t)$.

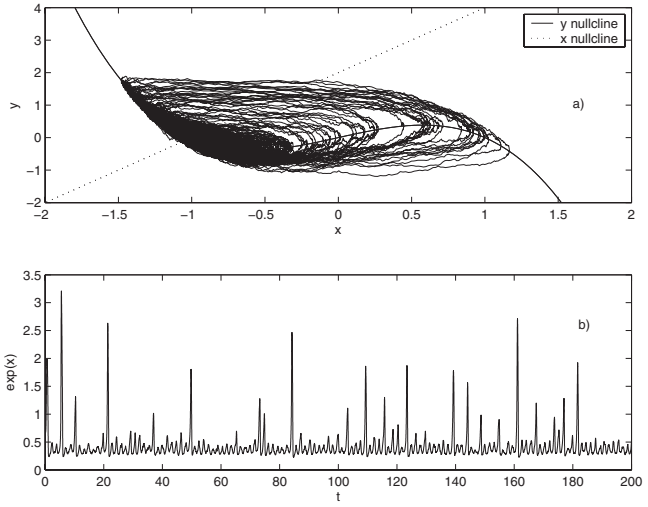

FIG. 11. Simulation of power market prices, basic model (FNS) for $\epsilon=0.15$, $\kappa=1$, $\lambda=1$, $b=2$, $\gamma=2$, $A=2$, $\omega_f=4$, $d=0.1$, $\Delta t=0.02$, x, y, and t in arbitrary units (1 simulated day cycle corresponds to $T_f=2\pi/\omega_f\simeq 1.57$ time units). (a) Phase space and (b) price $\exp[x(t)]$.

could be on a (1:4) stochastic tuning curve. But a statistical tuning curve constraint is not coincident with the regime that is needed here. Under the statistical constraint, absence of spiking in a period can be equilibrated on average by multiple spiking even in the same cycle in another period. A system behaving as in Fig. 9 is selected in a different regime, requiring a deterministic integer number of spikes per cycle at maximum. This regime could be referred to as *stochastically resonating spiking* (SRS). A nonempty set of parametrizations belongs to this regime because of the refractory period T_r mechanism, which for given T_f allows an integer number q of spikes to be accommodated such that $qT_r < T_f/2$. The noise activation time T_n controls the firing–not-firing statistics and the coherence between input $f(t)$ and output $x(t)$, as it competes with T_r and T_f. An effect of the interplay among these three scales can be seen in Fig. 10, where ϵ is raised to 0.3, slowing down the fast variable—and consequently ω_f is decreased, trying to track the resonance. The spiking mechanism is being blurred out, as spikes start losing their identity and gaining variability in height. Looking at the phase space in Fig. 10(a), the $\epsilon\to 1$ effect is clear. The nullcline picture is weakened; the system explores the right nullcline with more difficulty—even though the activation mechanism has not changed—and spends more time inside the inner core region and closer to the fixed point. Increasing ϵ even further toward 1 causes the FNS to abandon the right nullcline as the spiking regime is completely lost. This nonsingular and less analytically appealing regime has not very often been explored in the literature, even though numerical results like those in [4,32] for $\epsilon=0.5$ give a clear indication of what happens there. Then, for financial modeling reasons this soft (in contrast to singular) ϵ regime of the FNS is much more interesting.

A lot of information about the SRS regime is immediately available as a side product of the numerous studies of *stochastic resonance* (SR) [3,4] in the FNS. SR is a property shown by some systems that have an excitable dynamics—along with many others which are not excitable. When driven subthreshold in the silent region for a range of noise intensities d, these systems become particularly reactive to the forcing frequency for a specific value of the noise. Thus, they are often studied in the d-ω_f parameters space. Some of them, including the FNS, show a *doubly stochastic resonance* [33], and become particularly reactive for a specific combination of d and ω_f. It is interesting to note that in Ref. [31] the doubly SR condition is associated with statistical $(m:n)$ phase-locking conditions. With this result of SR in the FNS in mind, it is easy to jump to the conclusion that for a given ω_f there is a range of values for d which *inhibit* the resonance.

VI. TWO MODELS FOR THE POWER MARKET

The FNS of Eqs. (1a) and (1b) with $\beta=1$ is thus proposed as a first basic model for the power market price dynamics. The term $b-f(t)+\sigma\xi(t)$ in Eq. (1b) is identified with the noisy dynamics of the aggregated demand for electricity. Since prices $p(t)$ are positive quantities, $x(t)$ can be interpreted—as is done often in finance [34]—as a *logprice*, so that

$$\ln p(t) = x(t). \tag{4}$$

The price dynamics is attributed mathematically to a nonlinear price-forming mechanism, based on (a) stochastically resonating spiking and (b) a soft ϵ regime, reached where the ϵ and the mean reversion $(1/\beta=1)$ time scales become comparable. In Fig. 11 an instance of the FitzHugh-Nagumo dynamics in this regime is shown, for $f(t)=A\sin(\omega_f t)$ and a choice of the parameters that will be discussed later at the end of Sec. VII. This and all other simulations are done by a stochastic Euler algorithm with time increments Δt shorter than the relevant time scales and with a scaling factor $\sqrt{\Delta t}$

multiplying the Gaussian noise term. Figure 11(a) shows the phase-space dynamics and the nullclines, Fig. 11(b) shows the price process.

The spiking pattern obtained in this way is attractive, since it associates by design the spiking activity with demand crests, the firing events being random and spike heights being nonuniform. The FNS nonlinearity is exploited in almost all the phase-space core region. Since the Alberta price excursions are intrinsically limited in height, due to the price cap, the parameters κ and λ can be used to set this span to the price span observed in the data. This constraint to the span in $x(t)$ is included naturally (but statistically) in the FNS. In this soft ϵ regime, the FNS mean reversion mechanism is more complex than in the $\epsilon \to 0$ case. Since the ϵ and the mean reversion time scales are comparable, $\gamma x_{0+}+b$ cannot be considered a fixed mean reversion target for the y dynamics, since it varies on a scale comparable to 1. After a spike activation, during an $x(t)$ roll in the system potential $U(x)$ toward the right branch's newly formed global minimum x_{0+}, the (formerly slow) y dynamical equation has a changing $\gamma x(t)$ driver that makes y faster and nonlinear in its mean reversion toward x_{0+}. Moreover, the FNS feedback mechanism is such that $U(x_{0+})$ is increased at a velocity comparable to the y mean reversion velocity—or higher, depending on γ. This velocity matching can preclude y from ever reaching x_{0+} as x_{0+} is yet a minimum, and cause y to reverse its motion back toward the stable attractor too soon. In the meantime, the x dynamics cannot reach its full span and cannot fully spike. Spikes tend to lose their identity more quickly as γ gets larger. After a spike is activated because y has crossed y_- downward, y jerks up suddenly and rapidly crosses y_+, then goes back slowly to the attractor P_*. Depending on the interplay between ϵ and γ, a large fluctuation regime mixed with a spiking regime can be attained, which is now good for financial modeling. This soft ϵ regime has its own rigid signature in the distribution of crossings of the phase-space core region, due to the shape of the y nullcline, and in the rounded and widened right-hand edge of the region invaded by the dynamics.

An *extended model* [35] that helps to make the FNS distributional signature more flexible can be obtained by replacing the linear γx term of the FNS with a hyperbolic sine function $\sinh(x)$ (in Ref. [36] another hyperbolic FitzHugh-Nagumo system was proposed, a hyperbolic tangent model studied in the standard way). The system is operated in the SRS and soft ϵ regime. Its equations are

$$\epsilon \dot{x} = \kappa x - \lambda x^3 - y, \tag{5a}$$

$$\dot{y} = \sinh[\gamma(x-c)] + b - \beta y - f(t) + \sigma(d)\xi, \tag{5b}$$

where c is a constant, and the definitions and limitations of the remaining parameters are the same as in the FNS discussed in Sec. V. The choice of the $\sinh(x)$ function is just one among many other possible choices, since what is important now is the mechanism that the shape of the chosen function implements in the phase space. Peculiar to the extended model is the monotonic but twisted shape of its y nullcline. As a hyperbolic function, $\sinh(x)$ explodes exponentially at moderate $|x|$ values, and is close to linearity around 0. Centering it on c, dilating it with γ, and shifting it along y by means of b allows one to place it at any chosen P_* along the x nullcline and to implement a nullcline concavity or convexity in the inner part of the core region. More importantly, in the soft ϵ regime the twist can be used to modify the FNS feedback mechanism and to prevent the system from accessing part of the phase space, due to the extremely large values that $\sinh(x)$ can attain even for moderate values of x. Moreover, the convex parts of the nullcline favor spike formation whereas concave parts disfavor it, modifying the distribution of inner core region crossings. Consider, for example, the following parameter selection: $\epsilon=0.15$, $\kappa=\lambda=\gamma=1$, $b=0$, $\xi=2$, $c=-1$. At $\sigma=0$ and $\omega=4$, a check can be made for a range of values for A, for example from 2 to 10. As A increases, the access to the inner core of the phase space is gradual and there is no abrupt qualitative change between small oscillations ($A=2$) and standard spiking orbits ($A=10$). This behavior is then different from that of the FNS, in which an $A_{\max}$ rigidly separates small oscillations from spiking. In some sense, in the extended case the system can be precluded full access to true excited dynamics, depending on the y nullcline shape and position. The distributional signature of the extended model will be clearly visible in the statistical analysis carried out in Sec. VII.

At first sight the two models have very similar spiking patterns, so that both can be used if only the solution pattern is important. Consider, for example, the extended model with $\epsilon=0.15$, $\kappa=\lambda=\gamma=1$, $b=0$, $\xi=2$, $c=-1$, as before, $d=0.1$, and a multiperiod forcing of the form

$$f(t) = u\left[v \sin\left(\frac{\omega_f}{365} t\right) + \sin(\omega_f t)\right], \tag{6}$$

where $\omega_f=4$, $u=2$, and $v=0.06$. For this choice, $T_f=2\pi/\omega_f \approx 1.57$ corresponds to a day of the market data, and if the simulation time span is chosen as 3500 with $\Delta t=0.035$, approximately 6 years can be simulated. Figure 12 shows a simulation for $d=0.1$. Figure 12(a) shows the price series and Fig. 12(b) shows the forcing with its double periodicity, a very rough approximation of yearly and daily periodicity of the Alberta market. Figure 12(a) has to be compared with Fig. 1(a).

At the large scale, the model is able to spike at the yearly crests of the demand with higher frequency but with almost the same height distribution as at its lowest point, as for the real market. At the intermediate and small scales, as can be seen in Fig. 13 where the time span of 50 time units corresponds to 1 month, the shape of the spikes is smoother than in the real market and only one spike per day is allowed (by design). The impression that the simulation output is smoother than the real price series is partly due to the fact that there are only 24 Alberta market data points per day, whereas in the case of the simulation the mesh can be chosen freely. In the real market each day can support more than one spike. The model can be adapted to this behavior in two ways. Either another subdaily periodicity can be introduced in Eq. (6), or the parameters can be adjusted to admit multiple spiking. Another difference with real data comes from

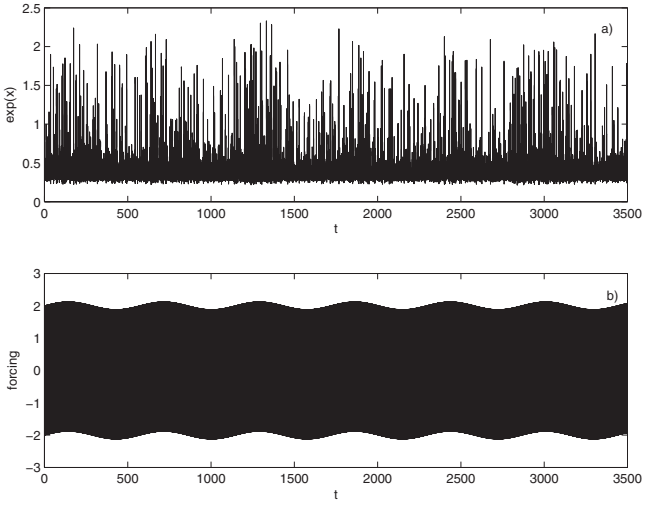

FIG. 12. Simulation of six years of power market prices; extended model for $\epsilon=0.15$, $\kappa=1$, $\lambda=1$, $b=0$, $\gamma=1$, $\xi=2$, $c=-1$, f as in Eq. (6) with $u=2$, $v=0.06$, $\omega_f=4$, $d=0.1$, $\Delta t=0.035$, x, y, and t in arbitrary units (again, 1 simulated day cycle corresponds to $T_f=2\pi/\omega_f\simeq 1.57$ time units). (a) Price process; (b) forcing: the smaller yearly $\omega_f/365$ frequency modulates the much higher daily ω_f frequency, which cannot be resolved in the picture.

the fact that in the simulation each spike occupies a full 12 h time span, and the background small oscillations observed in the real data are lost. Fine tuning of the model can improve this aspect too. In Fig. 14 different parameters, in particular a smaller ω, are used. Figure 14(a) shows the phase space with dynamics and nullclines, in particular the $\sinh(x)$ nullcline. Figure 14(b) shows the price dynamics, where the small oscillations and some multiple spiking are evident. In any case, the main strength of the two models appears to be the smoothness of the height distribution at all scales, certainly not sharp as in the usual FNS singular regime. A main weakness shared by the models is that steady growth in demand

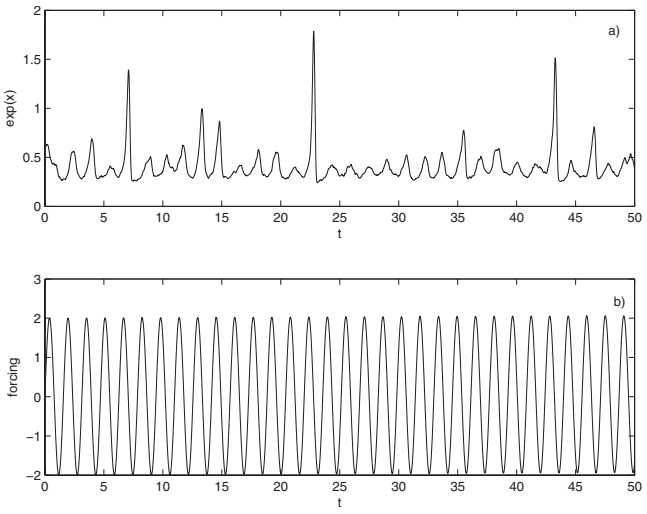

FIG. 13. Simulation of 1 month of power market prices; extended model with the same parameters and units as in Fig. 12. (a) Price process; (b) forcing.

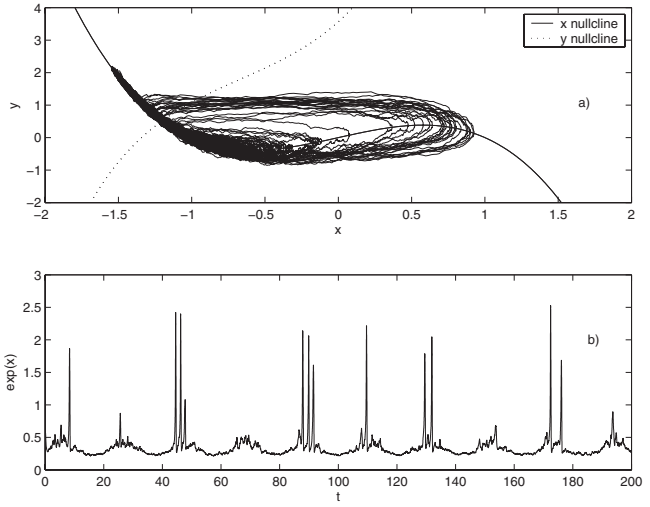

FIG. 14. Simulation of power market prices, extended model, for $\epsilon=0.1$, $\kappa=1$, $\lambda=1$, $b=1.8$, $\gamma=1$, $\xi=2$, $c=-2/3$, $A=2$, $\omega_f=0.3$, $d=0.1$, arbitrary units. (a) Phase space and nullclines; (b) price process.

cannot be implemented in these models for long times. This problem will be considered in Sec. VIII.

VII. SPIKE STATISTICS

The advantage of using the FNS to model the price series of the Alberta power market is double. Not only can a satisfactory spiking pattern be obtained, but also a fundamental statistical property of the price time series is automatically recovered. This can be shown using again suggestions from SR studies. SR in sample data is studied mainly by two means, the *interspike interval* (ISI) distribution [37] and the ensemble-averaged (since these systems are nonstationary) power spectrum. For a signal $x(t_i)$ dependent on a set of N equally spaced times t_i $(i=0,\ldots,N)$ of spacing Δt, an ISI sample frequency histogram is prepared by selecting a reference height h and defining "operative spikes," the contiguous portions of $x(t)$ above h, of number S. The collection of first passage times $\phi_j(h)$ from below to above h for each operative spike j is the set of spike times. The collection of passage times $\rho_j(h)$ from above to below h for each spike j is the set of spike return times. The collection of the differences

$$t_j(h)=\phi_{j+1}-\phi_j \tag{7}$$

of immediately adjacent spike times is the set of interspike widths. The collection of the differences

$$\delta_j(h)=\rho_j-\phi_j \tag{8}$$

is the set of intraspike widths. For any h, the ISI histogram with n on its abscissa is constructed by counting for each integer $n=1,\ldots,N$ the number of interspike widths that match the condition $t_j=n\Delta t$, $j=0,\ldots,S$. An intraspike width (ISW) histogram is constructed by counting in the same way the matchings for the intraspike widths, including $n=0$ in the matching condition. For a FNS operating in the singular regime, all spikes have the same height, so that the shape of

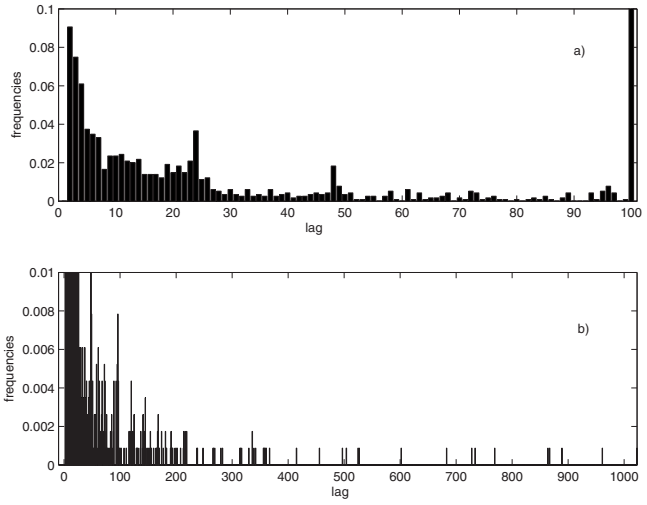

FIG. 15. ISI frequency histogram of the Alberta price series at price reference level 200 C$; on the abscissa time lags in hours. (a) Values on the abscissa are clipped at lag 100, the last bin containing all the remaining frequency weight; (b) the abscissa spans all analyzed lags whereas the ordinate is cut at frequency 0.01.

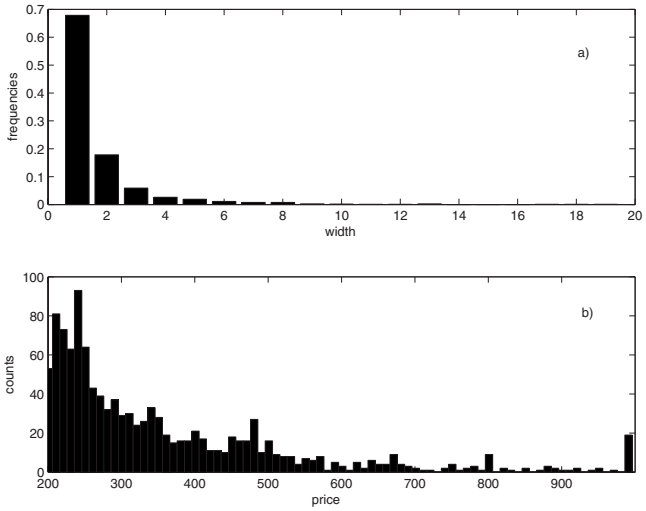

FIG. 16. ISW and SH histograms of the Alberta price series at price reference level 200 C$. (a) ISW frequencies vs time widths in hours; (b) SH counts vs prices in C$.

the ISI histograms is rather insensitive to h—as long as h is far from the background noisy small oscillations. This is certainly not true for time series like the Alberta prices series of Fig. 1(a) or for the output of models that try to explain such series. Information about the height distribution of the spikes can then be obtained by counting the number of spikes which reach a given price (higher than h), in a spike height (SH) histogram. From SR studies, a lot of knowledge is available about ISI properties of the FNS. The main result is that the ISI shape in the SRS regime is a collection of isolated peaks with heights that follow an exponential or a Γ function decline. Since the system cannot fire for half cycle T_f, the corresponding multiples of $T_f/2$ interspike widths are suppressed, whereas multiples of T_f contribute. If the condition of one spike each half cycle is not strictly matched, this picture is somewhat blurred. The exponential (i.e., Poissonian) decay is due to an essential independency of the triggering events.

For a reference height of 200 C$, in Fig. 15 the sample frequencies of interspike widths (each bin contains 1 h and the first bin is at lag 1) are shown for the Alberta market data. Figure 15(a) is a magnification clipped from the abscissa at 100 C$ of Fig. 15(b), which is cut in the ordinate for clarity to show the longest interval counts. The last bin of Fig. 15(a) is there for comparison, since it contains the weight of the clipped-out tail. Figure 15 shows a major feature. In its first part the Alberta ISI distribution seems exponential, but it has clearly a long nonexponential tail in its second part. The exponential seems to be built of peaks that appear—as expected—at the market periodicities, i.e., subday and many-day scales. The highest subday frequencies are under (and include) 3 h (at least two spikes each morning or afternoon) followed by peaks at 6 h (morning-afternoon periodicity), and those at day-night 12 h. After the 24 h peak a long tail begins, and a 48 h periodicity is clearly visible. In Fig. 16(a) the ISW frequency histogram shows that spikes of different widths are present in the Alberta price series, but most of them last only a few hours. Figure 16(b) shows the count of spikes with different heights (each bin has a span of 10 C$). Above an exponential background, some prices seem curiously more frequent, whereas the price cap (and the special events default price) at 1000 C$ shows an expected behavior very different from other prices.

An analogous statistics can be collected for the simulated process of Fig. 12. In this case the model data set can be made larger than the sample of the market data set. Choosing the reference price height at 1 (units are here arbitrary), Fig. 17 shows ISI sample frequencies on a rescaled abscissa, obtained by dividing the simulation times by $T_f/(24\Delta t)$, to express the simulation time in hours. Each bin contains 1 h. The general pattern is similar to the pattern found for the real market, a possible Γ function built by peaks centered on

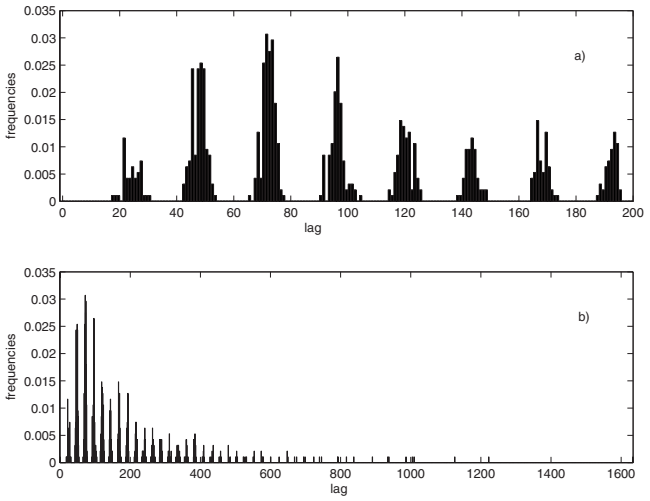

FIG. 17. ISI frequency histogram of the extended model at price reference level 1, lags in units of $T_f/24$ (i.e., simulation hours, see text). (a) The abscissa is cut at lag 200; (b) all analyzed lags.

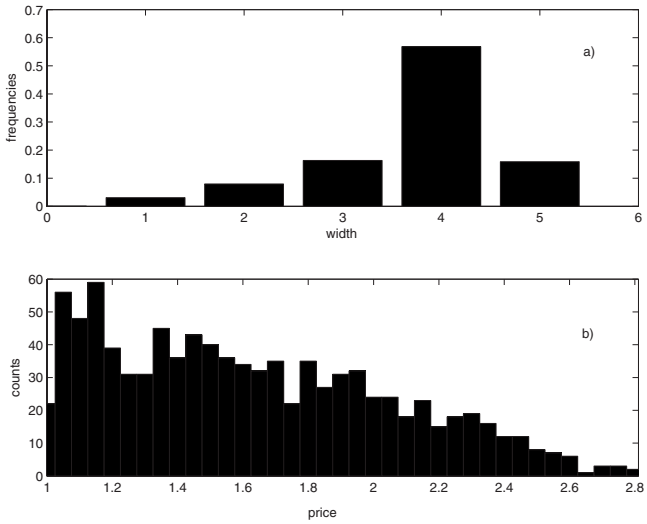

FIG. 18. ISW and SH histograms of the extended model at price reference level 1. (a) ISW frequencies vs time widths in $T_f/24$ time units (i.e., simulation hours); (b) SH counts vs prices (arb. units).

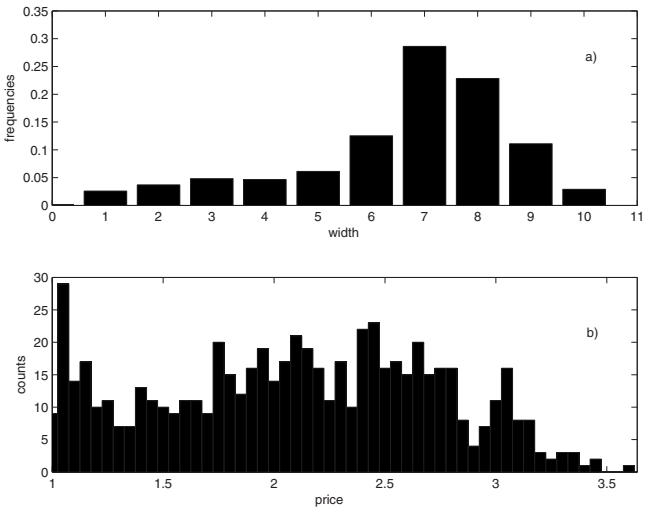

FIG. 19. ISW and SH histograms of the basic model at price reference level 1. (a) ISW frequencies vs time widths in $T_f/24$ time units (i.e., simulation hours); (b) SH counts vs prices (arb. units).

multiples of the 24 h basic frequency, with a long nonexponential tail. The peak of the Γ function depends on the activation time T_n (i.e., on d), and for higher d it can be moved to the left. Figure 18(a) shows the ISW frequency histogram. In the simulation all spikes have a width of 12 h at their base, as can be seen from Fig. 13(a). Moreover, spikes of the extended model are generated by soft excursions in the phase space, in a way that is different from the singular regime. A vertical cut in the phase space intersects some spikes that are fully developed which set the maximum width. Other spikes that cannot develop form the left tail of the distribution. Figure 18(b) shows the SH count histogram, and the effect on the peak heights of the particular position chosen for the y nullcline. Comparison of real-market and simulation statistics corroborates the idea that in the real market a SRS mechanism is at work.

The basic and extended models work essentially in the same way, the advantage of the latter being mainly its finer control of SH profiles, an important advantage for econometric analysis. A comparison between them can be made by considering the analytical form of the y nullcline in the extended case. When the nullcline center $\{c,b\}$ is placed on a P_* of the x nullcline, a first-order Taylor expansion around c,

$$\sinh[\gamma(x-c)] \approx -\gamma c + \gamma x, \tag{9}$$

suggests that it behaves around P_* as the nullcline of a FNS model for which $-\gamma c+b$ replaces b, whereas γ is the same. The parameters used in the simulation of the basic model shown in Fig. 11 are obtained from the parameters used in the simulation of the extended model shown in Fig. 12 by the use of Eq. (9). A difference can be appreciated by comparing their ISW and SH histograms. The data from the basic model shown in Fig. 19 can be compared with the data from the extended model shown in Fig. 18. Both simulations are obtained under the same conditions—except for b of the basic model. The main qualitative difference is the shape of the SH histograms, which is better for the extended model when they are compared with the Alberta data shown in Fig. 16. When accuracy of modeling is important, it can also be considered that in addition to the $\sinh(x)$ function there are other possible functions with an analogous topology, such as monotonic high-order antisymmetric polynomials or the simpler but very interesting $\exp(x)$ function. In this sense, the extended model has to be considered more as the representative of a class of possible fine tuning extensions to the basic model, than a model on its own.

VIII. STATIONARITY, AUTOCORRELATION, AND THE PROBLEM OF SPIKING

Many power market financial data sets come without demand data, and they are sometimes already averaged over 24 h, so that in this last case intraday information is lost together with the 24 h periodicity. Then, even though the nonstationary SRS mechanism cannot be invoked to model them, it is not unreasonable to search for stationary mechanisms that nonetheless try to exploit the resonance idea—making no use of forcing. Another well-studied property of the FNS, *coherence resonance* (CR) [4,28], can again deliver helpful suggestions. As seen in the discussion of Sec. V, in every regime the FNS has an intrinsic resonating capability. When set subthreshold without forcing, noise itself can excite some of its resonating modes and have these modes show up in the FNS dynamic behavior. This is CR, and it is in this aspect a stationary phenomenon. A further and more obvious possibility for a stationary regime is that of exploiting the suprathreshold self-sustained FNS dynamics, associating the daily periodicity with the intrinsic nonlinear oscillator frequency. In both cases, since the model is autonomous, it is more difficult to take into account all the periodicities that real-market data in any case display.

The issue of nonstationarity due to a steady demand growth is more complicated but very interesting. If $f(t)$ has a

component linear in t, the subthreshold system is slowly made more reactive to noise as $f(t)$ grows in time, until the model crosses the threshold and disruptively enters the suprathreshold regime. If some parameter in Eqs. (1a) and (1b) is allowed to change in the meantime, a dynamics can be imagined that holds the system in a given subthreshold SRS regime as f keeps growing, in a sort of out-of-equilibrium steady state. This could happen to some advanced real-world power markets, where demand and network development coevolve. In this coevolution case, adaptive parameters could be difficult to design, but possible to implement. In contrast, looking at the Alberta market data, a certain impression of continuously increasing spiking activity in time might come to mind. In this second case, models based on the FNS could find their limits.

The power spectrum technique was not used in this paper, but the very possibility of its use—besides ISI histograms—is suggestive of a brief consideration. The Alberta demand series and most probably the associated price series are nonstationary. This can be seen—*a posteriori*—from Fig. 6, where the autocorrelation and power spectrum for the Alberta price series are shown. In the absence of demand data, and in the case of preaveraged series, the extracted autocorrelation would show no peaks and it would show a long tail. The power spectrum would show a linear behavior in a log-log scale. This behavior could be interpreted as memory—and not as due to a different mechanism. Even though multifractality can be excluded by some multiscale method, memory would leave open the possibility for non-Brownian monoscaling. Before running the autocorrelation, some form of seasonality could obviously be suspected, and a deseasonalizing procedure would be run on the data before the autocorrelation analysis. If the mechanism of price formation is SRS, where a noisy periodic demand forces a FNS engine to randomly spike, this deseasonalizing would tweak the data, in a further misleading way. Thus, the ISI technique seems to be the most appropriate to compare any proposed model to the market data.

A last issue is linked to the economic fact that spikes are an unwanted phenomenon. The modeling proposed here draws from the FNS knowledge, which has been accumulated in time also because of SR and CR studies. If the main idea of this paper is correct, a market can be seen as an algorithmic machine that is tuned for some reasons (grid properties and auction protocol) to the demand periodicity. In a sense, detuning by change in the noise properties (this is what SR and CR suggest) or in the system parameters could reduce spiking and lessen the problem. For example, availability of demand knowledge is linked to the noise perceived by the suppliers, and day-ahead scheduling in the auction protocols introduces a characteristic frequency. Maybe a change in auction protocols could then disrupt resonance, at least for some time. Microeconomic analysis in this direction could help to soften the spiking problem.

IX. CONCLUSIONS

This paper studied the open financial problem of price spiking in power markets using ideas that come from physics. Spiking in power markets reminds one in a natural way of fluctuations in stochastic dynamical systems. Following this idea, it has been shown that a resonating stochastic FNS dynamics in a soft ϵ regime for the logprices is able to reproduce market spike patterns and their main statistical features. These results suggest the exploration of other similar nonlinear dynamical systems as power market models, and some of them are under current investigation. In the process of adapting the statistical signature of the FNS to real-market features, an extended model has been developed that probably brings in some interesting mathematics—partial suppression of excitability. This is also under current investigation, with its implications for SR and CR.

To conclude, the spirit of this paper is not related to financial econometrics, not even related to a stochastic resonance or a coherence resonance study of the models. The main effort has been that of drawing from existing fields of physics some ideas that could be used to help study an open problem in finance, and in the meantime developing models that could also be used in physics. The approach proposed here for the spiking problem has been different from the more usual search for self-similarity in financial systems—here scales matter, but this is because physics has a lot of tools that can be used when approaching financial problems. On the side of finance, nonlinear continuous time finance is a very recent area of research and it can find in the approach proposed here an interesting spiked drive.

ACKNOWLEDGMENT

The author wishes to thank F. Marchesoni for interesting discussions and stimulating suggestions.

[1] R. N. Mantegna and H. E. Stanley, *Introduction to Econophysics: Correlations and Complexity in Finance* (Cambridge University Press, Cambridge, U.K., 1999).

[2] J. P. Bouchaud and M. Potters, *Theory of Financial Risk and Derivative Pricing: From Statistical Physics to Risk Management* (Cambridge University Press, Cambridge, U.K., 2003).

[3] L. Gammaitoni, P. Hanggi, P. Jung, and F. Marchesoni, Rev. Mod. Phys. **70**, 223 (1998).

[4] B. Lindner, J. Garcia-Ojalvo, A. Neiman, and L. Schimansky-Geier, Phys. Rep. **392**, 321 (2004), and references therein.

[5] R. A. FitzHugh, Biophys. J. **1**, 445 (1961); J. Nagumo, S. Arimoto, and S. Yoshitzawa, Proc. IRE **50**, 2061 (1962).

[6] R. Wilson, Econometrica **70**, 1299 (2002).

[7] All market data used in this paper are available at the AESO site http://www.aeso.ca/. Prices before 2002, even though available, are not taken into consideration here, because they were computed before a major regulatory change.

[8] T. Di Matteo, Quant. Finance **7**, 21 (2007).

[9] C. K. Woo, D. Lloyd, and A. Tishler, Energy Policy **31**, 1103 (2003).
[10] H. Geman, *Commodities and Commodity Derivatives* (Wiley and Sons, New York, 2005).
[11] D. P. Chassin and C. Posse, Physica A **355**, 667 (2005).
[12] B. A. Carreras, D. E. Newman, I. Dobson, and A. B. Poole, IEEE Trans. Circuits Syst., I: Fundam. Theory Appl. **51**, 1733 (2004).
[13] T. Krause, EEH-ETH website.
[14] P. Klemperer, J. Econ. Surv. **13**, 227 (1999).
[15] For an introduction to oligopoly and market power, see A. Mass-Colell, M. Whinston, and J. Green, *Microeconomic Theory*, int. ed. (Oxford University Press, Oxford, 1995), Chap. 12.
[16] S. Borenstein and J. Bushnell, J. Ind. Econ. **47**, 285 (1999).
[17] C. W. Gellings, and K. E. Yeager, Phys. Today, **57** (12), 45 (2004).
[18] E. J. Lerner, Ind. Phys. **9** (5), 8 (2003).
[19] D. Czamanski, P. Dormaar, M. J. Hinich, and A. Serletis, Energy Econ. **29**, 94 (2007).
[20] H. Park, J. W. Mjelde, and D. A. Bessler, Energy Econ. **28**, 81 (2006).
[21] M. H. Rothkopf, Electr. J. **12**, 60 (1999).
[22] C. de Jong, ERIM Report No. ERS-2005–052-F&A, 2005 (unpublished); C. Mari, Physica A **371**, 552 (2006), and references therein.
[23] P. Abry, R. Baraniuk, P. Flandrin, R. Riedi, and D. Veitch, IEEE Signal Process. Mag. **19**, 28 (2002).
[24] B. Mandelbrot, L. Calvet, and A. Fisher (unpublished).
[25] P. Malo, Helsinki School of Economics Report No. W-402, 2006 (unpublished).
[26] The diagram was made with FRACLAB 2.0, available at http://www.irccyn.ec-nantes.fr/hebergement/FracLab
[27] J. P. Keener, *Principles of Applied Mathematics. Transformation and Approximation* (Addison-Wesley, Reading, MA, 1988).
[28] V. Anishchenko, V. V. Astakhov, A. B. Neiman, T. Vadivasova, and L. Schimansky-Geier, *Nonlinear Dynamics of Chaotic and Stochastic Systems* (Springer, Berlin, 2002).
[29] H. Treutlein and K. Schulten, Ber. Bunsenges. Phys. Chem. **89**, 710 (1985).
[30] A. Longtin, Chaos, Solitons Fractals **1**, 1835 (2000).
[31] S. G. Lee and S. H. Kim, J. Korean Phys. Soc. **49**, 31 (2006).
[32] M. Kostur, X. Sailer, and L. Schimansky-Geier, Fluct. Noise Lett. **3**, L155 (2003).
[33] H. E. Plesser and T. Geisel, Phys. Rev. E **59**, 7008 (1999).
[34] T. Bjork, *Arbitrage Theory in Continuous Time* (Oxford University Press, Oxford, 1999); R. Merton, *Continuous-Time Finance* (Blackwell, Oxford, 1992).
[35] The temptation to playfully refer to this model as "FitzHugh-Nagumo Power Market" was great, but curiosly enough a power plant named Fitzhugh Station exists in Ozark, Arkansas, USA. To avoid misunderstandings, the model will be referred to simply as the extended model.
[36] D. Terman and D. L. Wang, Physica D **81**, 148 (1995).
[37] L. Gammaitoni, F. Marchesoni, E. Menichella-Saetta, and S. Santucci, Phys. Rev. Lett. **62**, 349 (1989).