

On a multi-timescale statistical feedback model for volatility fluctuations

Lisa Borland[†] and Jean-Philippe Bouchaud^{+*}

[†] *Evnine-Vaughan Associates, Inc.,
456 Montgomery Street, Suite 800,
San Francisco, CA 94104, USA*

⁺ *Science & Finance, Capital Fund Management,
6-8 Bd Haussmann, 75009 Paris, France, and
Service de Physique de l'État Condensé,
Orme des Merisiers, CEA Saclay,
91191 Gif sur Yvette Cedex, France.*

(Dated: July 10, 2005)

We study, both analytically and numerically, an ARCH-like, multiscale model of volatility, which assumes that the volatility is governed by the observed past price changes over different time scales. With a power-law distribution of time horizons, we obtain a model that captures most stylized facts of financial time series: Student-like distribution of returns with a power-law tail, long-memory of the volatility, slow convergence of the distribution of returns towards the Gaussian distribution, multifractality and anomalous volatility relaxation after shocks. At variance with recent multifractal models that are strictly time reversal invariant, the model also reproduces the time asymmetry of financial time series: past large scale volatility influence future small scale volatility. In order to quantitatively reproduce all empirical observations, the parameters must be chosen such that the model is close to an instability, meaning that (a) the feedback effect is important and substantially increases the volatility, and (b) that the model is intrinsically difficult to calibrate because of the very long range nature of the correlations. By imposing consistency of the model predictions with a large set of different empirical observations, a reasonable range of the parameters value can be determined. The model can easily be generalized to account for jumps, skewness and multiasset correlations.

I. INTRODUCTION

The quest for a faithful mathematical model of price fluctuations has been taunting researchers for more than a century now, starting with Bachelier's random walk model in 1900 [1]. Such an endeavour is important for a bevy of reasons, both from the point of view of (a) fundamental economics (what is the cause of price variations and what information do they reveal?) and (b) of financial engineering, with option pricing, risk control and trading models as obvious applications.

In an ideal world, “the” mathematical model of price changes should be simple enough to allow easy calculations and calibration, yet rich enough to embrace all known stylized facts that the recent access to huge amounts of data has helped establish. It is now widely accepted that price changes reveal (i) fat tails, well described by a power-law decay of the probability distribution for large returns [2, 3, 4], (ii) long range memory in volatility fluctuations or volatility “clustering”, again described by a power-law decay (in time) of the autocorrelation of the volatility [5, 6, 7] and (iii) asymmetric causal correlations between past price changes and future volatilities, often referred to as the “leverage effect” [8] (for reviews, see e.g. [9, 10, 11, 12]). We discuss below other, somewhat related, stylized facts that have been reported in the recent literature, such as multifractal scaling, critical relaxation of the volatility after a shock (the financial analogue of the Omori law for earthquakes), etc. More recently, some statistical *asymmetry* of financial time series under time reversal was pointed out [13] – in other words, financial time series do distinguish past from future. This might appear trivial but constitutes in fact, as we discuss below, a very strong constraint on the family of eligible models for financial time series – for example, Bachelier's random walk model is strictly time reversal symmetric.

Scores of different models have been proposed to improve upon the simple Brownian motion model, which has neither fat tails nor volatility clustering. Lévy processes allow one to superimpose jumps to the Brownian motion, and therefore generate fat tails, but has no volatility clustering [10, 11, 14, 15]. GARCH models or simple stochastic volatility models such as the Heston model allow one to get both fat tails and some sort of volatility clustering, but

*Electronic address: bouchau@spec.saclay.cea.fr

not the long memory observed in the data [16, 17, 18, 19]. Models that mix jumps and stochastic volatility have been investigated [20]. Multifractal stochastic volatility models, initiated by Mandelbrot, Fisher and Calvet [21] and much studied since [22, 23, 24, 25, 26, 27, 28, 29, 30, 31, 32, 33, 34], seem to capture in a parsimonious way a large amount of empirical properties. However, most multifractal models are again strictly time reversal symmetric and lack an intuitive interpretation in terms of agent based trading models [35]. We in fact strongly believe that any serious model of price fluctuations should *in fine* be justified by reasonable behavioral rules and market microstructure effects (see [36, 37, 38, 39, 40] for recent work in that direction). Quite recently, one of us (LB) has proposed, in the context of option pricing, a “statistical feedback” process where the local volatility is large when price moves are deemed rare, leading to a non-linear diffusion equation for the price [41]. This equation can be solved and leads to a Student-Tsallis distribution for price changes at all times [42]. In its original form, however, the model breaks time translation symmetry: there is a well defined starting date and starting price. Although this can be used to price options [41, 43] (in the spirit of the Hull-White model for interest rates [44]), the process has to be modified to be interpreted as a *bona fide* model of returns. Such an extension, and its modification to account for long-range memory, was proposed in [45] and recovers, following a different route, a multiscale GARCH model proposed by Zumbach and Lynch in 2003 [13, 46] (see [6, 47, 48] for earlier work in that direction). Numerical simulations of this model suggest a very rich phenomenology, that seems to account for most stylized facts of financial time series.

The aim of the present paper is to motivate this new model, discuss its relation with previous work, and investigate in full details its statistical properties, both analytically and numerically. We focus in particular on the probability distribution of returns which is the crucial ingredient for option pricing and risk control. Although not an exact result, we find that these distributions can be well fitted by a Student-Tsallis form, with a lag-dependent tail exponent. We reproduce in great details most empirical facts, including the anomalous relaxation of the volatility after a shock, and the past/future asymmetry of the time series. The model can be generalized to include jumps, the leverage effect, and multi-stock correlations. We then discuss the issue of calibration. Within strict econometric standards, calibration is extremely difficult due to the long-memory nature of both the empirical volatility process and the theoretical models that are constructed precisely to capture this long memory. We advocate the idea of ‘soft’ calibration, which in such cases should consist in reproducing semi-quantitatively as many observables as possible. These observables should be chosen to be robust to the details of the model specification, and test different “orthogonal” predictions of the model (these statements will be made clearer in the course of the paper and in Section VI). Consequences for option pricing are briefly discussed, and will be the subject of another paper.

II. SET UP AND MOTIVATION OF THE MODEL

In the following, we will consider a discrete time model, with an elementary time scale equal to τ , for example $\tau = 1$ minute. [A continuous time version of the model will be discussed below]. The price at time $t_i = i\tau$ will be noted p_i . We will conform to the standard of dealing with the log-price $x_i = \ln p_i$ and define returns as $r_i = x_{i+1} - x_i$. [59] The random return is constructed as the product of a time dependent volatility σ_i and a random variable ξ_i of zero mean and unit variance:

$$r_i = \mu\tau + \sigma_i \xi_i \sqrt{\tau}, \quad (1)$$

where μ is the average drift, which we will set to zero in the sequel, meaning that we measure all returns relative to the average drift. The noise ξ_i can a priori have any probability distribution to account for high frequency kurtosis and jumps, but for simplicity we will mostly focus in this paper on the case of a Gaussian noise. However, as we discuss below, the introduction of jumps is needed to faithfully reproduce real price time series.

The seminal insight of ARCH or GARCH models [49] is that the volatility process reflects trading activity and is subordinated to past price changes. Intuitively, the level of activity becomes high when past price changes are, in some sense, anomalous. In the simplest ARCH model, this is expressed as:

$$\sigma_i^2 = \sigma_0^2 \left[1 + g \frac{r_{i-1}^2}{\sigma_0^2 \tau} \right], \quad (2)$$

meaning that the volatility is equal to its ‘base level’ σ_0^2 plus a contribution coming from the last price change. In fact, we have written the feedback term in a way that expresses the comparison between the square of the last return and its expected value, equal to $\sigma_0^2 \tau$. If the last price change was small compared to usual, the volatility today is close to its normal value, whereas in the other limit, the last return is deemed anomalous and leads to a potentially large increase of today’s activity.

An argument motivating Eq. (2) above is as follows. Suppose that some traders open positions (for example, long) at time t_{i-1} , when the price is p_{i-1} . Such trades are often initiated with both a profit objective and a risk limit,

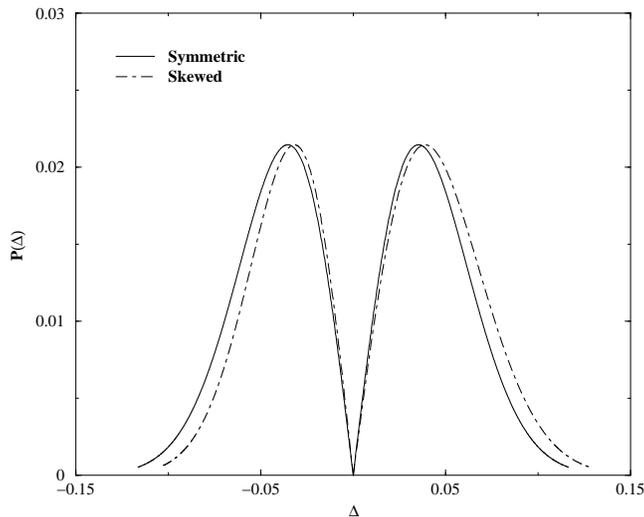


FIG. 1: Schematic shape of the distribution of stop loss/stop gain thresholds around the opening price of the trade. Plain line: symmetric distribution; dotted line: asymmetric distribution, giving rise to the leverage effect discussed in section V.

which would close the position at time t_i if the price has moved up too much (stop gain) or down too much (stop loss). It is very natural to hypothesize that to each opening trade are associated two thresholds, one above, one below p_{i-1} , that trigger a closing trade if exceeded. If many agents open both long and short trades at t_{i-1} , one can expect a quasi-continuous distribution of thresholds at $p_{i-1}(1 + \Delta)$, more or less symmetrically distributed around p_{i-1} , triggering with equal probability sell back or buy back orders. The density $P(\Delta)$ is, in the simplest case, even (but see Section V for the inclusion of the leverage effect) and obviously vanishes at $\Delta = 0$ since nobody opens a trade to close it immediately (see Fig. 1). The width of $P(\Delta)$ is given, in order of magnitude, by $\sigma_0\sqrt{\tau}$ since this gives the natural scale beyond which an event might be deemed anomalous. Hence, a relative change of price r_{i-1} will trigger on the order of: [60]

$$N_i(r_{i-1}) \approx \int_0^{|r_{i-1}|} P(\Delta)d\Delta \quad (3)$$

stop trades. These trades of random sign lead, on the next day, to an increase of the volatility as:

$$\sigma_i^2 = \sigma_0^2 + GN_i/\tau, \quad (4)$$

where G is the average square impact per trade, and σ_0^2 is the volatility due to all other trades. Taking into account that $P(\Delta)$ extends over a range $\sigma_0\sqrt{\tau}$, one finally obtains a general single time scale ARCH model:

$$\sigma_i^2 = \sigma_0^2 \left[1 + \mathcal{G}\left(\frac{|r_{i-1}|}{\sigma_0\sqrt{\tau}}\right) \right], \quad (5)$$

where the function \mathcal{G} depends on the detailed shape of $P(\Delta)$. Taking for simplicity, in accordance with the above discussion,

$$P(\Delta) = P_1 \frac{|\Delta|}{2\beta^2\sigma_0^2\tau} \exp\left(-\frac{\Delta^2}{2\beta^2\sigma_0^2\tau}\right), \quad (6)$$

(where β is a number setting the width of the distribution of thresholds, and P_1 the total number of opened trades) finally leads to:

$$\mathcal{G}(u) = 2g\beta^2 (1 - \exp(-u^2/2\beta^2)), \quad (7)$$

where $g = GP_1/2\beta^2\sigma_0^2\tau$ is the ratio measuring the impact of all stop trades compared to that of all other trades. The simplest ARCH model Eq. (2) corresponds to the limit $u \ll \beta$, that is, neglects saturation effects related to the fact that stop limits are not placed arbitrarily far from the entry point (i.e. β is finite). When this saturation is neglected, $\mathcal{G}(u)$ is simply given by gu^2 (but see below, Fig. 11).

Although the above feedback mechanism is most probably at play in financial markets, a strong limitation of the above model is to consider that all traders have the same time horizon, equal to τ in the above formulation. However, it is well documented that the activity of financial markets is fueled by traders with different time horizons, from a few hours to a few months or even years (see e.g. [13, 48]). Therefore, stop losses or profit objectives are not placed only around the last price p_{i-1} but around possibly all past prices $p_{i-\ell}$, $\ell = 1, 2, \dots$. Correspondingly, the width of the distribution of these thresholds is calibrated to the volatility of the price over the particular trading horizon, i.e. $\sigma_0\sqrt{\ell\tau}$. The generalization of Eq. (5) to this situation therefore reads:

$$\sigma_i^2 = \sigma_0^2 \left[1 + \sum_{\ell=1}^{\infty} \mathcal{G}_\ell \left(\frac{|x_i - x_{i-\ell}|}{\sigma_0\sqrt{\ell\tau}} \right) \right]. \quad (8)$$

Expanding \mathcal{G}_ℓ for small arguments finally leads to the symmetric version of the model studied in the present paper (the inclusion of asymmetry will be discussed in Section V): [61]

$$\sigma_i^2 = \sigma_0^2 \left[1 + \sum_{\ell=1}^{\infty} g_\ell \frac{(x_i - x_{i-\ell})^2}{\sigma_0^2 \ell \tau} \right], \quad (9)$$

with

$$r_i = x_{i+1} - x_i = \sigma_i \xi_i \sqrt{\tau}. \quad (10)$$

The coupling constant g_ℓ is proportional to the number of trades P_ℓ with horizon ℓ . Because traders with a longer horizon have slower trading frequencies and under-react compared to short term traders, it is reasonable to imagine that g_ℓ is a decaying function of ℓ . Both for simplicity and because it allows us to reproduce several stylized empirical facts, we will choose g_ℓ to be an inverse power:

$$g_\ell = g/\ell^\alpha, \quad (11)$$

but other choices are possible. For example, Zumbach and Lynch have presented evidence that g_ℓ has additional peaks on the day, week and month times scales. These authors have proposed a model very close in spirit to Eq. (9), and discussed some of its properties. In fact, Eq. (9) is a special case in the family of quadratic ARCH models, where the volatility is expressed as a general quadratic form of past returns:

$$\sigma_i^2 = \sigma_0^2 + \sum_{j < i, k < i} \mathcal{M}(i; j, k) \frac{r_j r_k}{\tau}, \quad (12)$$

which contains ARCH, GARCH, etc. Our specification insists that only combinations of returns ‘reconstructing’ actual price changes over different time scales occur in the above sum, because they correspond to quantities directly observable to the crowd of traders, which, we argue, strongly influence the trading at time i . Our model corresponds to a particular choice for \mathcal{M} above:

$$\mathcal{M}(i; j, k) = \sum_{\ell=\max(i-j, i-k)}^{\infty} \frac{g_\ell}{\ell}, \quad (13)$$

whereas most ARCH models correspond a certain regression on past *instantaneous* square returns, i.e., to $\mathcal{M}(i; j, k) = K(i-j)\delta_{jk}$, with a certain kernel function K , usually corresponding to an exponential moving average, $K(\ell) = \alpha^\ell$.

With a power-law specification for g_ℓ , and the choice of a Gaussian distribution for the noise term ξ in the definition of returns, our model is fully determined by only four parameters: σ_0 sets the volatility scale, τ sets the shortest time scale over which feedback effects are effective, g measures the strength of these feedback effects and α describes the relative importance of short term traders and long time traders in the feedback process. It may however be that the assumption of a Gaussian noise for ξ is insufficient to account for the high frequency statistics of the returns. In particular, one expects that true ‘jumps’ related to unexpected news are not described in terms of a volatility feedback process. It is easy to extend the model in that direction and choose another distribution for ξ . In the following sections, we will present several analytical and numerical results of the Gaussian version of this model, and compare them to empirically known results. But before doing so, let us give the continuous time formulation of the same model, which can be convenient for some applications, such as option pricing. Introducing the standard Brownian noise dW_t , one may write:

$$dx_t = \sigma_t dW_t, \quad (14)$$

with:

$$\sigma_t^2 = \sigma_0^2 + g\tau^\alpha \int_{-\infty}^t dt' \frac{(x_t - x_{t'})^2}{(t - t' + \tau)^{1+\alpha}}. \quad (15)$$

This model is well defined as soon as $\alpha > 1$, which is the case we will focus on in the sequel.

III. ANALYTICAL RESULTS

A. Unconditional distribution of the volatility

Although our model (Eq. (9)) expresses the volatility as a deterministic function of the past prices and the only source of randomness comes from the noise ξ_i in Eq. (1), the volatility effectively appears as a random variable, and one can ask questions about its distribution, correlations, etc. The simplest question concerns the average value of the volatility, which also coincides, for a stationary process, with the long term volatility of the price. Averaging will always be denoted below with brackets $\langle \dots \rangle$ around the quantity which is averaged. For the average volatility, one has (assuming stationarity) :

$$\langle \sigma^2 \rangle = \langle \sigma_i^2 \rangle = \sigma_0^2 + \sum_{\ell=1}^{\infty} g\ell \frac{\langle (x_i - x_{i-\ell})^2 \rangle}{\ell\tau} = \sigma_0^2 + \left[\sum_{\ell=1}^{\infty} g\ell \right] \langle \sigma^2 \rangle. \quad (16)$$

This equation has a well behaved solution only if:

$$z_2 = \sum_{\ell=1}^{\infty} g\ell < 1, \quad (17)$$

where the above equation defines z_2 , the subscript ‘2’ refers to the fact that we study here the second moment of the volatility. When $z_2 < 1$, the square volatility is amplified by a factor $1/(1 - z_2)$ compared to the initial value σ_0^2 . In the case $z_2 \geq 1$, on the other hand, the process becomes non stationary and the volatility grows without bound as time elapses. It is clear that the condition $z_2 < 1$ can only be met if the sum of $g\ell$ converges, which imposes that the exponent α is larger than one. For $\alpha > 1$, one finds $z_2 = g\zeta(\alpha)$, which delimits a region in the plane g, α where the process is stationary. In the following, we will often assume that α is larger than unity but close to it (which is suggested by empirical data), and use in this limit a continuous approximation for discrete sums. In particular, $\zeta(\alpha) \approx 1/(\alpha - 1)$. We will find below that empirical data on stocks favors values of $z_2 \approx 0.85 - 0.9$, meaning that the square volatility is increased by a factor $\sim 6 - 10$ compared to its initial value σ_0^2 . Therefore feedback effects might be an important cause of the excess volatility in financial markets [50, 51].

In order to compute higher moments of the volatility, one needs in general to know the full temporal correlation of the volatility, that we will establish in the next paragraph. Simplified, approximate calculations can be performed in two extreme cases: (i) no temporal correlations (ii) full temporal correlations. This leads to an equation for $\langle \sigma^4 \rangle$ of the form $(1 - z_4)\langle \sigma^4 \rangle = \text{RHS}$, where the right hand side is finite whenever $\alpha > 1$, and can be computed if necessary (see below). The important discussion concerns the value of z_4 . We will denote $M_k = \mathcal{M}(0; -k, -k)$, which behaves, for large k , as $k^{-\alpha}/(\alpha - 1)$. Using the results established below (see Eq. (26)), one can obtain a lower bound $z_{4,<}$ and an upper bound $z_{4,>}$ on the value of z_4 . If correlations are neglected, one finds:

$$z_4 \geq z_{4,<} = 3g^2 \sum_{k=1}^{\infty} M_k^2. \quad (18)$$

If on the other hand, if correlations are overestimated and taken to be constant in time, one finds an upper bound for z_4 that reads:

$$z_4 \leq z_{4,>} = g^2 \left(\left[\sum_{k=1}^{\infty} M_k \right]^2 + 2 \sum_{k=1}^{\infty} M_k^2 + 4 \sum_{k=1}^{\infty} (k-1) M_k^2 \right). \quad (19)$$

As long as $z_4 < 1$, the fourth moment of σ is finite, but if z_4 reaches unity, it does diverge, leading to an infinite kurtosis for the returns. For $\alpha = 1.15$, we find $z_{4,<} = 0.16 z_2^2$ and $z_{4,>} = 1.44 z_2^2$. This shows that the kurtosis κ is certainly finite for $z_2 < 0.833$; numerical simulations below suggest that κ indeed remains finite beyond that value.

The above argument is easily generalized to higher even moments of σ , leading to an equation $(1 - z_{2n})\langle\sigma^{2n}\rangle = \text{RHS}$ with:

$$z_{2n,<} = (2n - 1)!! g^n \sum_{k=1}^{\infty} M_k^n, \quad (20)$$

and a more cumbersome expression for $z_{2n,>}$. For large n , one finds $z_{2n,<} \sim (2gn/e)^n$, showing that however small the value of g , sufficiently high moments of the volatility are divergent. Since $r_i = \sigma_i \xi_i$, the even moments of the returns are given by:

$$\langle r^{2n} \rangle = (2n - 1)!! \langle \sigma^{2n} \rangle; \quad (21)$$

therefore high moments of returns themselves diverge, suggesting that both the unconditional distribution of volatility and returns have a power-law tail (possibly multiplied by a slow function), with an exponent equal to the order of the last finite moment. We will confirm this prediction numerically in the following section. Remember however that the above discussion is only valid when the noise ξ is Gaussian; if ξ itself has a non zero kurtosis, then its contribution should be taken into account.

B. Temporal correlations of the volatility

A well known stylized fact is that the volatility is a ‘long-memory’ process, which means that the temporal correlations of the square volatility decay as an inverse power of the time lag, $\ell^{-\nu}$, with an exponent ν less than unity. This property turns out to be extremely important because it is at the root of the very slow convergence of the distribution of aggregated returns towards the Gaussian. More precisely, the kurtosis of the return $x_i - x_{i-\ell}$ over scale ℓ , itself decays as $\ell^{-\nu}$ instead of ℓ^{-1} , which is the case when the volatility process has a short memory. Since the empirical value of ν is, for stocks, on the order of $\nu = 0.2 - 0.3$, the slowing down is substantial and essential to explain why long dated options still have a smile.

We therefore turn to the calculation of the correlation function of the volatility, defined as:

$$\mathcal{F}(\ell) = \frac{\langle \sigma_{i+\ell}^2 \sigma_i^2 \rangle}{\langle \sigma^2 \rangle^2} - 1. \quad (22)$$

In the limit $g^2 \ll 1$, one can quite easily perform a perturbative analysis that neglects terms of order g^4 , to get:

$$\mathcal{F}(\ell) = 2g^2 \left[\sum_{0 < k < j} \frac{k^{1-\alpha}}{(\ell + j)^{1+\alpha}} + \sum_{0 < j \leq k} \frac{j^2}{k^{1+\alpha}(\ell + j)^{1+\alpha}} \right]. \quad (23)$$

An analysis of this result for $\ell \gg 1$ finally gives, for $\alpha > 1$ but close enough to unity such that one can use continuous integrals instead of discrete sums:

$$\mathcal{F}(\ell) \sim \frac{4g^2 \Gamma(2 - \alpha) \Gamma(2\alpha - 1)}{\alpha^2 \Gamma(\alpha)} \ell^{2-2\alpha} \equiv \mathcal{F}_\infty \ell^{-\nu}, \quad (24)$$

leading to a kurtosis exponent $\nu = 2\alpha - 2$. The volatility is a long memory process whenever $\nu \leq 1$, i.e. $1 < \alpha \leq 3/2$. Comparison with empirical data, done below, suggests that α is in the range 1.1 – 1.2. The exact equation for $\mathcal{F}(\ell)$, not restricted to small g^2 , can also be written down, although it is more cumbersome. For this calculation, one should note that averages such as $\langle \sigma_i^2 \xi_j^2 \rangle_c$ (where the subscript c denotes a connected average) are non trivial, since the volatility randomness comes entirely from past returns themselves. This contrasts with many stochastic volatility models where the volatility σ_i and the noise ξ_j are often chosen to be independent (unless one wants to model the leverage effect). In the present case, one finds, for $j < i$:

$$\langle \sigma_i^2 \xi_j^2 \rangle_c = g^2 \sum_{k, k' > 0} \mathcal{M}(0; -k, -k') \langle \sigma_{i-k} \sigma_{i-k'} \xi_{i-k} \xi_{i-k'} \xi_j^2 \rangle_c = 2g^2 \langle \sigma^2 \rangle M_{i-j}, \quad (25)$$

Now, the full self-consistent equation for \mathcal{F} reads:

$$\begin{aligned} \mathcal{F}(\ell) = & g^2 [3\mathcal{F}(0) + 2] \sum_{k=1}^{\infty} M_k M_{k+\ell} + 4g^2 \sum_{k > k'=1}^{\infty} M_k M_{k+\ell} [1 + \mathcal{F}(k - k') + 2g^2 M_{k-k'}] \\ & + 2g^2 \sum_{k > k'=1}^{\infty} M_k M_{k'+\ell} [\mathcal{F}(k - k') + 2g^2 M_{k-k'}] + g^2 \sum_{k=1}^{\infty} \sum_{k'=1}^{\ell} M_k M_{k'} [\mathcal{F}(\ell - k' + k) + 2g^2 M_{\ell - k' + k}] \end{aligned} \quad (26)$$

Specializing to $\ell = 0$ leads to the fourth moment of the volatility studied in the above section. The two assumptions made there to obtain a lower and an upper bound correspond to $\mathcal{F}(\ell) = \mathcal{F}(0)\delta_{\ell,0}$ and $\mathcal{F}(\ell) = \mathcal{F}(0)$, respectively. For large ℓ , an asymptotic estimate of the various terms leads to the same decay as that predicted by the above perturbative calculation, i.e., $\mathcal{F}(\ell) \sim \mathcal{F}_\infty \ell^{-\nu}$, with $\nu = 2\alpha - 2$, and a prefactor \mathcal{F}_∞ increased by a factor $1/(1 - z_2^2)$. However, sub-dominant terms also appear, proportional to $\ell^{-2\nu}$, $\ell^{-\alpha}$, etc. The finite ℓ behaviour of $\mathcal{F}(\ell)$ would require to solve the above equation numerically.

From the knowledge of $\mathcal{F}(\ell)$ one can obtain the ℓ dependence of the kurtosis of the returns, following [11]. Again, one should take care of the terms involving $\langle \sigma_i^2 \xi_j^2 \rangle_c$, which, as we discuss below, lead to a new, perhaps unexpected effect. One finds, for the kurtosis of the returns on lag ℓ :

$$\kappa(\ell) = \frac{1}{\ell} \left[\kappa(1) + 6 \sum_{j=1}^{\ell} (1 - \frac{j}{\ell}) [\mathcal{F}(j) + 2g^2 M_j] \right]. \quad (27)$$

For large lags $\ell \gg 1$, one finds, using Eq. (24), and for α close to 1:

$$\kappa(\ell) \sim \frac{3\mathcal{F}_\infty}{(3 - 2\alpha)(2 - \alpha)} \ell^{-\nu}. \quad (28)$$

Therefore, one expects the returns to converge to Gaussian, but only on a very long time scale. Any measure of the distance from a Gaussian – such as the mean absolute moment studied below – will tend to zero very slowly, as $\ell^{-\nu}$, see Figs 4-a, 4-b. If one now studies Eq. (27) for small values of ℓ , say $\ell = 2$, one finds:

$$\kappa(2) - \kappa(1) = 3[\mathcal{F}(1) - \mathcal{F}(0) + 2g^2 M_1] \quad (29)$$

In many models, the last term is absent, and since $\mathcal{F}(1) \leq \mathcal{F}(0)$, one usually finds that the kurtosis of aggregated returns is less than the kurtosis of elementary returns. However, the third term in the above expression suggests that one can observe, in some cases, a kurtosis that first *increases* with lag before decaying to zero. We will see that this is indeed the case in the numerical simulations of our model, although this effect is, again, very sensitive to the assumption that ξ_i is a purely Gaussian noise.

C. Conclusion

The summary of this technical section is that the two major stylized facts (fat tails and volatility long-memory) are present in our model. We have indeed shown that the distribution of returns and of the volatility have power-law like tails, since high moments of these distributions diverge. We have also shown that the temporal correlation of the volatility is decaying as a slow power law. The following sections will be to establish these properties more quantitatively using numerical simulations, and to show that many more stylized facts can be reproduced by the model. Finally, we will turn to the question of calibration and discuss how the model parameters can be chosen to fit empirical data.

IV. NUMERICAL RESULTS

We have established above that the volatility-volatility correlation function, and the kurtosis, decay at long times as $\ell^{-\nu}$ with $\nu = 2(\alpha - 1)$. A large amount of empirical work on financial time series suggest that ν is in the range 0.2 – 0.4 for many different assets. For example, averaging over the 500 largest stocks of the NYSE leads to $\nu \approx 0.25$, while $\nu \approx 0.3$ for the S&P 500 Index [11]. We therefore choose to fix $\alpha = 1.15$ (corresponding to $\nu = 0.30$) in most of the numerical experiments that we have conducted. Other values of α are briefly discussed, in particular in the context of the model calibration. The choice $\alpha = 1.15$, although guided by empirical data, immediately leads to a numerical problem due to its proximity with the critical value $\alpha = 1$ which separates a (theoretically) stationary regime for $\alpha > 1$ from a non stationary regime for $\alpha \leq 1$. The convergence of (say) the average volatility to its asymptotic value is expected to occur at speed $T^{1-\alpha}$, where T is the total length of the time series. For $\alpha = 1.15$, this is extremely slow: even for $T = 10^6 \tau$, one expects corrections of order 10% to the theoretical asymptotic results. For this reason, and also to speed up the numerical calculation of the sum that determines the volatility (Eq. 9), we have truncated the power-law memory kernel g_ℓ beyond $\ell = 5 \cdot 10^4$. The total length of our simulations is usually 10^6 steps, but we discard the first $15 \cdot 10^4$ points of the series before we start measuring any observable. Although this is, again, insufficient to obtain very precise results for such low values of α , we believe that these numerical experiments

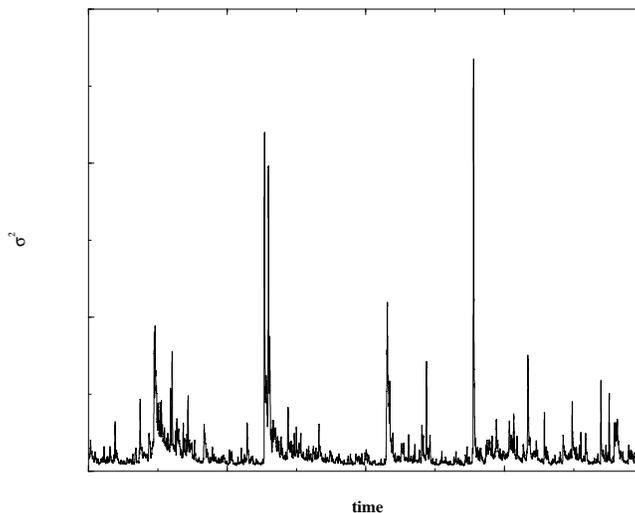


FIG. 2: A typical time series (of length $2 \cdot 10^5$) of the volatility σ^2 , for $z_2 = 0.85$ and $\alpha = 1.15$. We have in fact shown a 300τ moving average of σ_i^2 , aimed at representing the ‘daily’ volatility within our model.

are sufficient to obtain a good estimate of a host of different interesting observables, in any case comparable in quality to the corresponding estimates on real price time series. As will be clear below, we estimate that one day corresponds in our model to $\ell \sim 300$; therefore 10^6 time steps corresponds to 3000 trading days, or twelve years of data. In the following, the base volatility σ_0 is set to $\sigma_0 = 1$, any other value would only change the following results by a trivial multiplicative factor on the returns. We will vary the coupling constant g , which we will in fact express in terms of $z_2 = \sum_{\ell} g_{\ell}$, since we know from the above discussion in section III that it is really z_2 that measures the strength of feedback effects on the volatility. In the limit $z_2 \rightarrow 1$, we know from section III that the volatility will blow up and the process becomes non-stationary for all values of α . Therefore, studying numerically values of z_2 too close to unity will also be difficult (the convergence is now as slow as $[(1 - z_2)T]^{1-\alpha}$), but, ironically, corresponds to the empirical situation. In the following, we restrict our simulations to the range $z_2 \in [0.60, 0.85]$ – smaller values of z_2 lead to a process which is only weakly non Gaussian, whereas larger values of z_2 give rise to a numerically very unstable process, even though in theory the process should still be stationary on extremely long time scales. We will see below that values of z_2 as high as 0.9 might be needed to fit the data, but we have not attempted to simulate the model for such a large value.

Although the issue of calibration will be more deeply discussed in section VI, we will compare in this section our numerical results to empirical data, averaged over a set of 252 US stocks, chosen among the most liquid ones, during a four year time period: 2000-2003.

A. Volatility distribution and volatility correlations

1. Volatility distribution

We first focus on the properties of the ‘true’ volatility σ_i , which we can of course measure numerically but is unobservable directly in practice: only proxies of the volatility, obtained by averaging over several time steps, can be studied. A typical time series of σ^2 is shown in Fig. 2, and reveals apparent shocks and volatility clustering familiar in financial time series. We show in Fig. 3 the histogram of $u = \ln \sigma$ for different values of z_2 . Obviously, since $\sigma > \sigma_0 = 1$, the probability distribution function (pdf) of u is zero when $u \leq 0$. We have found that the pdf $P(u)$ of u can be very accurately fitted by the following form (see Fig. 3):

$$P(u) = Z \exp \left[-\left(\frac{u_0}{u}\right)^{\beta} - \mu u \right] \Theta(u), \quad (30)$$

where $\Theta(u > 0) = 1$ and $\Theta(u < 0) = 0$. We have no detailed justification for this specific functional form for $u \rightarrow 0$. On the other hand, it is easy to show that the exponential tail for large positive u translates into a power-law distribution for σ itself, decaying as $\sigma^{-1-\mu}$, which is indeed expected from our theoretical analysis. Correspondingly,

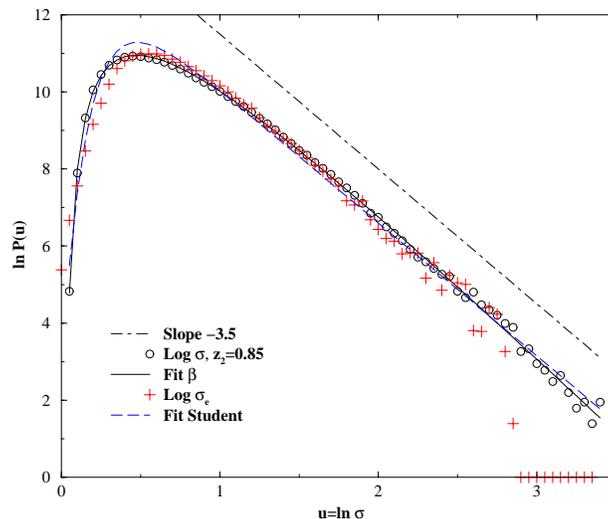


FIG. 3: Histogram of $u = \ln \sigma$ for $z_2 = 0.85$ and $\alpha = 1.15$, and two fits, using Eq. (30) – “Fit β ” – and Eq. (31) – “Fit Student”. We also show the slope $\mu = -3.5$ for comparison with the tail of the distribution of u . The pluses correspond to the histogram of the average volatility over 100 time steps, close to what one would determine empirically from price time series.

z_2	$(1 - z_2)^{-1}$	$\langle \sigma^2 \rangle$	$\mathcal{F}(0)$	μ	β	u_0	$\lambda^2 (\ln)$	$\lambda^2 (\zeta_q)$
0.60	2.50	2.50	0.54	5.70	0.75	0.51	0.45	1.8
0.65	2.85	2.85	0.79	5.27	0.70	0.74	0.57	2.45
0.70	3.33	3.31	1.16	5.02	0.60	1.65	0.69	3.1
0.75	4.0	3.92	1.73	4.72	0.54	3.26	0.85	3.7
0.80	5.0	4.79	2.59	4.42	0.51	5.67	1.03	4.2
0.85	6.66	6.05	3.95	4.03	0.50	6.83	1.24	5.5

TABLE I: Value of different observables and fit parameters for different values of z_2 , at fixed memory kernel ($\alpha = 1.15$, or $\nu = 0.3$). The values of λ^2 must be multiplied by 10^{-2} . Note the 10% discrepancy between the theoretical and empirical value of $\langle \sigma^2 \rangle$ when z_2 reaches 0.85. The value of μ suggests that the kurtosis κ remains finite at least up to $z_2 = 0.85$, in agreement with our theoretical analysis; the numerical value of $\kappa = 3\mathcal{F}(0)$ is found to be ≈ 12 for $z_2 = 0.85$. From this table, one can extrapolate μ to be ≈ 3.6 and $\lambda^2 (\ln)$ to be ≈ 0.015 for $z_2 = 0.9$.

the distribution of returns will also display the same power-law tail. The values of μ that we find using the above fit are summarized in Table I. From Fig. 3, however, we see that the apparent slope of $\ln P(u)$ vs. u in the available range of ‘large’ u values is slightly smaller than the value of μ obtained from a global fit with Eq. (30). For example, for $z_2 = 0.85$, we find $\mu \approx 4$, but the apparent slope is ≈ 3.5 , interestingly closer to the value reported for stocks $\mu \approx 3$ [4]. A slightly larger value of $z_2 = 0.9$ would be in even better agreement with this empirical value of the exponent (see Table I).

We have also tried to fit $P(u)$ assuming an inverse Gamma distribution for the pdf of σ [11, 52], which corresponds to a Student distribution for the returns. In terms of $u = \ln \sigma$, this reads:

$$P(u) = Z' \exp[-Ae^{-Bu} - \mu u], \quad (31)$$

where A and B are parameters, and μ is the power-law tail of the return distribution. Of course, this distribution cannot be exact in the present case since it takes non zero values when $u < 0$. Although it is definitely not as good a fit as Eq. (30), it is quite acceptable, meaning that returns are indeed close to being Student distributed in our model. On the other hand, a log-normal distribution for σ is clearly inadequate to describe our data (it would correspond to a parabola in Fig. 3.)

From Table I, we see that (a) the numerical value of the average volatility is close to its theoretical value up to $z_2 \approx 0.75$, beyond which a systematic underestimation of the true volatility σ^2 is observed, which reaches 10% for $z_2 = 0.85$; (b) the kurtosis $\kappa = 3\mathcal{F}(0)$ increases with z_2 , as expected, and seems to remain finite at least up to $z_2 = 0.85$, beyond which μ appears to drop below 4, signaling a divergence of κ (and correspondingly an even more difficult determination of the statistical properties of the system).

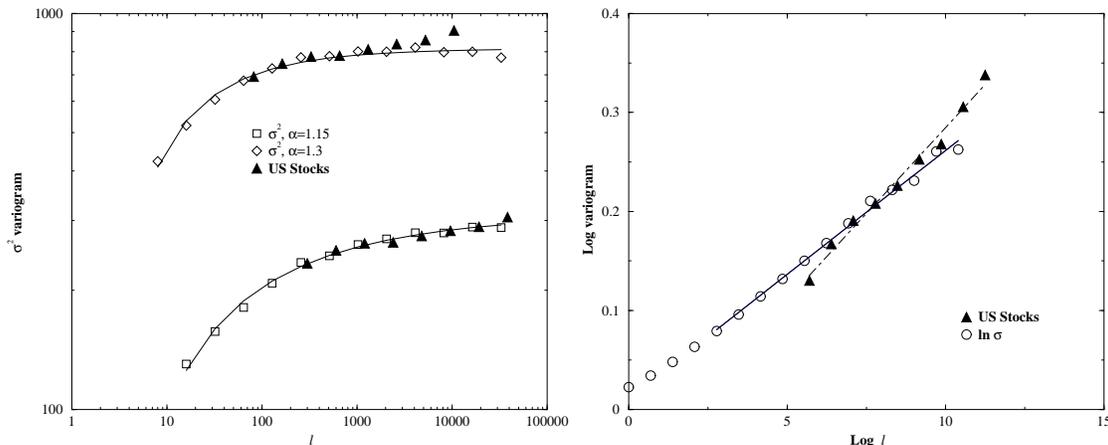


FIG. 4: Left: variogram of the square volatility for $z_2 = 0.85$, $\alpha = 1.15$ and $\alpha = 1.3$ and fit with power-laws with $\nu = 2(\alpha - 1)$. We also show the data for US stocks with one day corresponding to 300τ for $\alpha = 1.15$ and 80τ for $\alpha = 1.3$, for which the agreement is clearly not as good. Right: variogram of the log-volatility for $z_2 = 0.85$, $\alpha = 1.15$, and fit with an affine function of $\ln \ell$, the slope of which yielding (twice) the intermittency parameter λ^2 , here found to be ≈ 0.0125 , whereas the US data suggests a larger value $\lambda^2 \approx 0.0165$ – which would be matched by choosing $z_2 = 0.90$, for which we estimate from Table I $\lambda^2 \approx 0.015$.

2. Volatility correlations

We now turn to the temporal correlations of the volatility. Several characterizations of the “long-memory” property are interesting to consider. Well studied quantities are correlations of different powers of the volatility, or of the logarithm of the volatility. In our model, we of course know exactly the volatility at any instant of time, whereas, as pointed out above, in real conditions one only has access to price changes, from which a (noisy) proxy of the volatility is constructed. We find numerically that the shape of the correlation function can be noticeably different for these two quantities when the noise is large; this observation may be especially important for calibration.

From Table I, one sees that the average volatility is rather ill-determined in the cases most relevant for applications, i.e. z_2 and α both close to unity, variograms should be preferred to correlograms [11]. In other words, we will study the following quantity:

$$V_n(\ell) = \frac{1}{n^2} \langle (\sigma_i^n - \sigma_{i+\ell}^n)^2 \rangle. \quad (32)$$

These variograms are plotted in Figs. 4 a,b for the case $n = 2$ and $n \rightarrow 0$. This last case reproduces, thanks to the $1/n^2$ normalization, the variogram of the logarithm of the volatility which has been much studied in the context of multifractal models (see below). The case $n = 2$ is important because it can be analytically studied, as we did in Section III, and because it is related to the kurtosis of the distribution of returns for different time lags, which determines the smile of option prices. From Eq. (26), one finds that for large ℓ , one should observe:

$$V_2(\ell) \sim 2\mathcal{F}_0 - 2\mathcal{F}_\infty \ell^{-\nu} - 2\mathcal{F}'_\infty \ell^{-2\nu} + \dots \quad (33)$$

When α is close to 1, ν is small and one should a priori be prepared to see corrections to the asymptotic result coming from the $\ell^{-2\nu}$ contribution. Fig. 4-a however shows that, for $z_2 = 0.85$ and $\nu = 0.3$ our numerical result is rather well fitted by the dominant term of Eq. (33). The value of the apparent exponent ν however increases when z_2 decreases (in which case the contribution of the subleading term becomes more important). We also show the US stock data (that corresponds to $\nu \approx 0.25$ [11]), which can be matched quite well with the model. In order to test the sensitivity of $V_2(\ell)$ to the value of α , we also show in Fig. 4-a the case $\alpha = 1.3$, corresponding to $\nu = 0.6$. The agreement with empirical data is clearly not as good, a conclusion confirmed by all other observables we studied.

Another interesting quantity, less noisy than the square volatility, corresponds to $n \rightarrow 0$. As discussed below, this log-variogram appears naturally in the context of multifractal models. The result for $n = 0$ is shown in Fig. 4-b; we see that it can be fitted approximately by the multifractal prediction [30]:

$$V_0(\ell) \sim 2\lambda^2 \ln \frac{\ell}{L}, \quad (\ell \ll L) \quad (34)$$

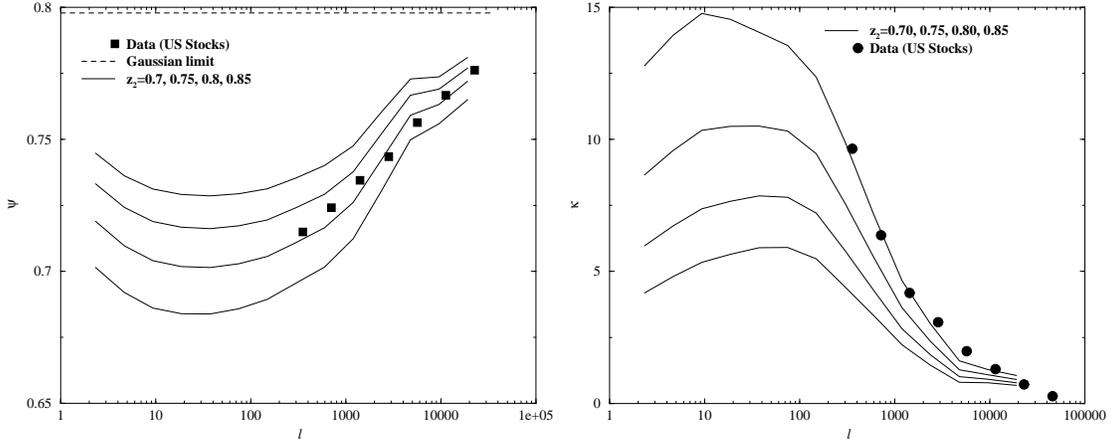


FIG. 5: Evolution of two cumulants of the distribution of returns with the lag ℓ , for different values of z_2 . Left: rescaled mean absolute deviation $\Upsilon(\ell)$. Note that the evolution is non-monotonous as a function of ℓ . Right: Excess kurtosis $\kappa(\ell)$. We have also shown (symbols) the corresponding cumulants for US stocks, where we choose $\ell = 300$ to correspond to one trading day, as in Fig. 4.

where λ^2 is called the intermittency parameter, and $T = L\tau$ is usually called the integral time. The value of λ^2 for different values of z_2 is given in Table I, together with another determination of λ^2 discussed below. For $z_2 = 0.85$, $\alpha = 1.15$, we find $\lambda^2 \approx 0.0125$, whereas our US data gives $\lambda^2 \approx 0.0165$, or $\lambda^2 \approx 0.018$ for the S&P100 Index, given in [53]. This suggests that the optimal value of z_2 might in fact be closer to 0.9, for which we estimate from Table I $\lambda^2 \approx 0.015$. This conclusion is reinforced by the analysis of section VI.

B. Distribution of returns over different time scales

Since the noise variable ξ is Gaussian, one can obtain the distribution of returns on the elementary time scale $\ell = 1$ from the distribution of the instantaneous volatility σ . For example, an inverse Gamma distribution for σ leads to a Student-Tsallis distribution for r . As discussed above, the actual distribution of volatility in our model appears to be slightly different from an inverse Gamma distribution; therefore the distribution of returns in our model will be close to, but different from, a Student distribution. On larger time scales, the distribution progressively becomes Gaussian. However, the convergence is very slow precisely because of the long-memory of the volatility, parameterized by the exponent ν . A way to quantify this convergence is to measure the cumulants of the distribution, for example the excess kurtosis $\kappa(\ell)$, expected from our theoretical analysis to decay as $\ell^{-\nu}$, or the rescaled mean absolute deviation $\Upsilon(\ell)$, defined as:

$$\Upsilon(\ell) = \frac{1}{\sqrt{\langle \sigma^2 \rangle \ell}} \langle |x_{i+\ell} - x_i| \rangle. \quad (35)$$

For a Gaussian distribution, one should find $\Upsilon = \sqrt{2/\pi}$. These quantities are plotted as a function of ℓ in Figs. 5-a,b, for different values of z_2 and for $\nu = 0.3$. An a priori unexpected feature is that non-Gaussian effects actually first *increase* for small ℓ , before decaying back to zero beyond a certain $\ell = \ell^* \approx 50$. The origin of this non monotonicity was discussed in section III and is clearly related to the assumption that the noise ξ_i is Gaussian. Any extra kurtosis coming from unpredictable jumps in the price, not captured by the feedback mechanism of our model, will strongly affect the shape of $\Upsilon(\ell)$ and $\kappa(\ell)$ on short time scales, and remove this non-monotonicity which, to the best of our knowledge, is not observed on empirical data, even on very short time scales. Another possibility is to change the shape of g_ℓ for small ℓ 's.

Of course, the knowledge of κ and Υ is insufficient to fully characterize the whole distribution on different time scales. We have in fact found that a Student-Tsallis distribution with a time dependent number of degrees of freedom is an acceptable fit of this distribution for all values of ℓ . In line the notation of ref. [41], we write this distribution as:

$$P_\ell(\Delta) = \mathcal{N} \frac{\Delta_0^{(3-q)/(q-1)}}{(\Delta_0^2 + (q-1)\Delta^2)^{1/(q-1)}}, \quad \Delta = x_{i+\ell} - x_i \quad (36)$$

with an ℓ dependent parameter $q(\ell)$. In the limit $q \rightarrow 1$, the distribution becomes Gaussian. If the distribution is indeed given by Eq. (36), the relation between Υ and q reads:

$$\Upsilon = \sqrt{\frac{2(\mu-2)}{\pi(\mu-1)^2} \frac{\Gamma(\frac{1+\mu}{2})}{\Gamma(\frac{\mu}{2})}}, \quad \kappa = \frac{6}{\mu-4}, \quad (37)$$

with $\mu \equiv (3-q)/(q-1)$. Using these relations, one can infer, from Fig. 5-a, the value of q that one should use for different times scales in order to get an *approximate* functional form for the distribution of returns. This is useful for option pricing, for example [41].

C. Multifractality

A property related to the systematic change of the distribution of returns with ℓ is *multifractality*, which means that different moments of price changes scale as a power of time, but with different scaling exponents. More precisely, multifractal scaling is the following property:

$$M_n(\ell) = \langle |x_{i+\ell} - x_i|^n \rangle = A_n \ell^{\zeta_n}; \quad \ell \ll L \quad (38)$$

where A_n are constants and ζ_n is an n -dependent exponent. In the monofractal case, where the distribution is the same on all time scales up to a rescaling of the returns, then $\zeta_n = n/2 \zeta_2$. The simplest example is obviously the (geometric) Brownian motion, for which $\zeta_n = n/2$. Any deviation from a linear behaviour of ζ_n is coined multifractality, for which several explicit models were proposed recently [26, 27, 28, 29, 30, 31].

One example is the Bacry-Muzy-Delour (BMD) stochastic volatility model, which makes the following assumptions [30, 31]:

- the log-volatilities $\ln \sigma_i$ are multivariate Gaussian variables (or more generally infinitely divisible [31]).
- the log-volatility variogram is given by Eq. (34)
- the volatilities σ_i are independent from the (Gaussian) noises ξ_i .

From these assumptions, one can compute exactly the moments of the return distribution on different time scales. One finds that these are indeed given by Eq. (38), with:

$$\zeta_n = \frac{n}{2} [1 - \lambda^2 (n-2)], \quad (39)$$

whenever $n < 1/\lambda^2$, beyond which the moments are infinite. (All A_n 's can also be exactly computed [30]). These assumptions and predictions were found to account rather well for some aspects of empirical data.

In the present section, we show that although our model is, strictly speaking, *not* multifractal, many of the multifractal predictions actually hold numerically quite accurately. This means that our model can account very well for apparent multifractal properties of financial time series, and in fact cures some of the deficiencies of standard multifractal models (see below). First, we note that our model is not multifractal since the moments $M_n(\ell)$ can be exactly computed to be *sums* of power-laws with different exponents, and not a unique power-law (see [25] for a related discussion). For example, $M_4(\ell)$ is the sum of ℓ^2 , $\ell^{2-\nu}$, $\ell^{2-2\nu}$, etc., and therefore does not scale as a unique power-law. However, as we show now, the numerical behaviour appears difficult to distinguish from a unique, effective power-law. [62]

We have computed numerically $M_n(\ell)$ for $\ell > \ell^*$, where ℓ^* corresponds to the maximum of $\kappa(\ell)$ or the minimum of $\Upsilon(\ell)$ appearing in Figs. 5-a,b. In this regime, one can neglect the contribution of terms involving $\langle \sigma_i^2 \xi_j^2 \rangle_c$, and $M_n(\ell)$ can be expressed as:

$$M_n(\ell) \approx (n - \frac{1}{2})!! \langle \left(\sum_{i=1}^{\ell} \sigma_i^2 \right)^{n/2} \rangle, \quad (40)$$

which is the quantity that we studied numerically, because it is much less noisy than the direct calculation of moments of returns. The results are shown in Fig. 6-a in a log-log representation, for $z_2 = 0.85$ and $\alpha = 1.15$, from which it is obvious that pure power-laws are indeed excellent fits. From the slope of these lines one obtains the exponents ζ_n , shown for different values of z_2 in Fig 6-b. We note that:

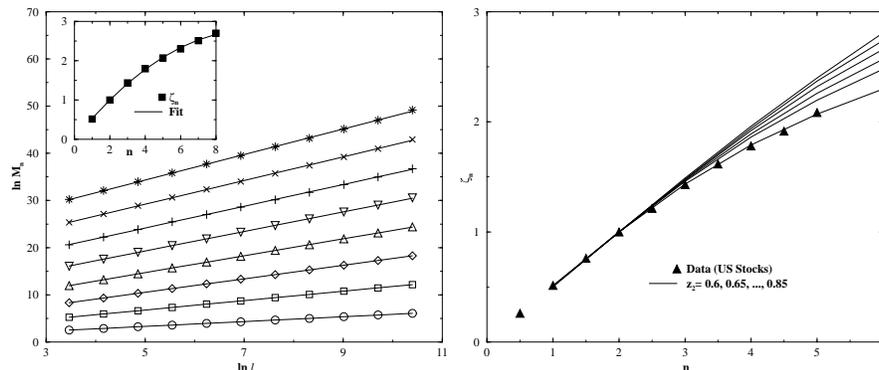


FIG. 6: Left: Evolution of different moments $M_n(\ell)$ of our model with ℓ in a log-log representation, which allows one to extract from the slope of these lines, the exponents ζ_n shown in the inset. Also shown in the inset is the parabolic fit suggested by the BMD model, Eq. (39). Right: The exponents ζ_n as a function of n for different values of the feedback parameter z_2 , and the corresponding results of US stocks (triangles).

- A parabolic fit of ζ_n as a function of n , as in Eq. (39), gives an excellent representation of our data (see Fig 6-a, inset).
- From this parabolic fit, a value of λ^2 can be extracted for different values of z_2 . This value of λ^2 is four times larger than that extracted from the variogram of the log-volatility, in contradiction with the BMD model, where both should be equal (see Table I). However, we note that a similar discrepancy with the BMD model is observed on US stock data as well, but that both observables are fully compatible with our model with the same set of parameters. The discrepancy with the log-normal BMD model is due to the underestimation of the probability of large events in that model.
- The intermittency parameter λ^2 , that gauges the degree of multifractality (i.e. the deviation of ζ_n from a straight line), increases as z_2 increases. The multifractal spectrum extracted from US stock data corresponds to $\lambda^2 \approx 0.055$ and matches quite well our numerical points for $z_2 = 0.85$. Similar values of λ^2 have been reported for other markets as well (see e.g. [24, 30]).

The BMD multifractal model makes other, even more detailed predictions, about the relaxation of volatility after a volatility shock. We now turn to this topic to show that our model can also reproduce these more subtle features.

D. Response to volatility shocks

A question of great importance for option pricing and risk management concerns ‘aftershocks’. It is well known that after a large market move, the volatility remains high for a while. The precise question therefore is: conditioned to a large volatility burst, how fast will the market revert back to normal? This has been addressed both empirically and theoretically, within the context of the BMD model [53]. One finds that after a shock, the volatility reverts to its normal level very slowly, as a power-law of the time ℓ after the shock:

$$\Delta\sigma_{i+\ell} \sim \Delta_0 \ell^{-\theta}, \quad (41)$$

where i is the time of the initial shock, $\Delta\sigma$ the excess volatility over its average value, and Δ_0 the amplitude of the initial shock. For rather large shocks, the empirical data suggests $\theta \approx 1/2$ [53, 54], while θ decreases for smaller shocks. Interestingly, the multifractal BMD model suggests that the exponent θ in fact depends continuously on the amplitude of the initial shock, and decreases from the value $1/2$ as the amplitude of the shock decreases [53]. This prediction was found to be in remarkable agreement with empirical findings, giving strong support to the BMD picture.

We have therefore computed the volatility relaxation process within our model, following the methodology of [53]. We compute the average volatility a time ℓ after the shock, conditioned to a shock of a certain amplitude. The relaxation curves are shown in Fig 7, again in the case $z_2 = 0.85$, $\alpha = 1.15$. We observe that the predictions of the BMD model are again quite accurately verified by our model, which, by the same token, is an alternative candidate to explain empirical results.

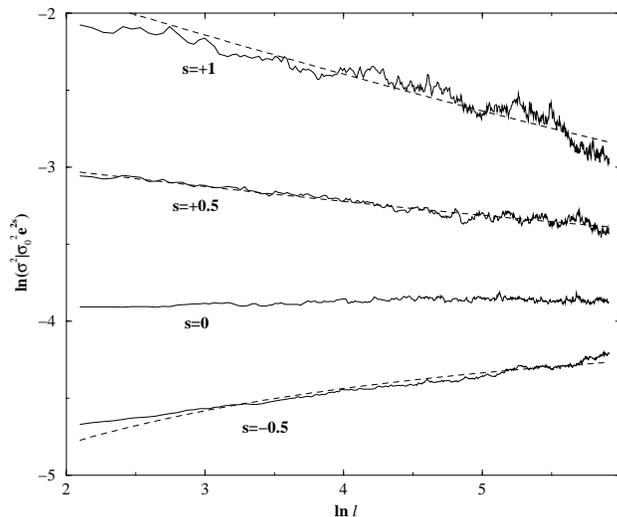


FIG. 7: Evolution of σ^2 conditioned to an initial volatility $\langle \sigma^2 \rangle e^{2s}$, with $s = -1/2, 0, 1/2, 1$ (from bottom to top). The dashed lines are power-law fits, with exponents $\theta(s = -1/2) \approx 0.22$, $\theta(s = 1/2) \approx 0.17$ and $\theta(s = 1) \approx 0.30$, very similar to the results quoted in [53] for the S&P100.

Another characterization of aftershocks inspired from research on earthquakes is the so-called Omori law, which states that the probability of an aftershock larger than a certain threshold occurring a time ℓ after the main shock decays as $1/\ell^p$, with $p \approx 1$. This law was checked for stock markets in [55] on a handful of ‘significant’ crashes. In our model, crashes are self-generated and not related to external news, obviously absent from the model. We show in Fig. 8 the numerically determined Omori law after large endogenous crashes, for which we obtain a significantly different value of $p \approx 0.5$, compatible with the value of θ reported above for large crashes. On the other hand, it is easy to compute the volatility response to an exogenous crash, represented by a large instantaneous jump added ‘by hand’ in the time series. If the amplitude of the jump at time $i = 0$ is J , the volatility after the crash is given by:

$$\langle \sigma_\ell^2 \rangle_J \approx \langle \sigma^2 \rangle + gJ^2 M_\ell. \quad (42)$$

Using $M_\ell \sim \ell^{-\alpha}$, we find that the probability of an aftershock larger than a given threshold also decays as $\ell^{-\alpha}$ for large enough ℓ . Since the value of α is close to unity, an approximate Omori law with $p \approx 1$ will be observed after anomalously large, exogenous crashes in our model. The distinction between endogenous and exogenous crashes, suggested in [53], makes perfect sense in the context of the present model, where endogenous crashes are, in a precise sense, the result of progressive volatility built up, resulting from the ARCH like feedback effect. This volatility built up is in fact related to the non monotonous behaviour of the kurtosis in our model.

E. Time reversal symmetry

A question of general interest is whether financial time series ‘know’ about the arrow of time, i.e. whether it is possible to compute any observable that distinguishes past from future (see [56]). Although the answer of this question would appear, to the layman, to be trivially yes, things turn out to be much more subtle, and of considerable importance. For example, the usual Brownian motion, all Lévy processes and all multifractal models constructed up to now (including Mandelbrot’s cascade, the BMD model or the version studied by Lux in [29]) are strictly invariant under time reversal symmetry (TRS)! Financial data, on the other hand, do reveal non TRS effects. A simple example, on which we will expand in the next section V, is the leverage effect, which is a causal correlation between *past* price changes and *future* volatilities: a drop in price leads to an increased volatility. This effect in turn leads to some (negative) skewness in the distribution of returns (see below).

Here, we want to discuss a distinct effect, recently evidenced by Zumbach and Lynch [13]. In order not to mix this effect with leverage, one can study FX rates between two large currencies, for example Euro vs. Dollar. In this case, any leverage correlation or skewness, if present, is very small. In spite of this, there is a clear time asymmetry in the volatility process: as shown in [13], the correlation between large scale, past volatilities and small scale future volatilities is larger than between small scale, past volatilities and large scale future volatilities. This effect was

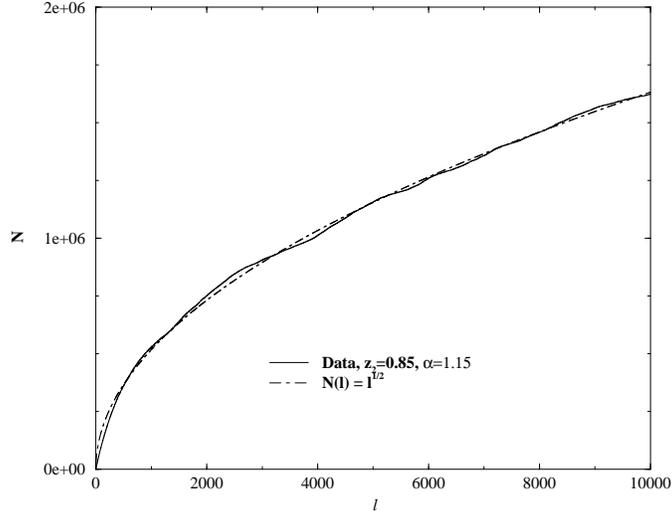


FIG. 8: This Omori plot shows the cumulative number of aftershocks (i.e. returns with an amplitude larger than a certain threshold) following a main shock, and fit with $N(\ell) \sim \sqrt{\ell}$. Main shocks were defined as returns larger than a third of the maximum return observed over the whole time series (of length 850,000), and aftershocks as returns larger than a third of the main shock.

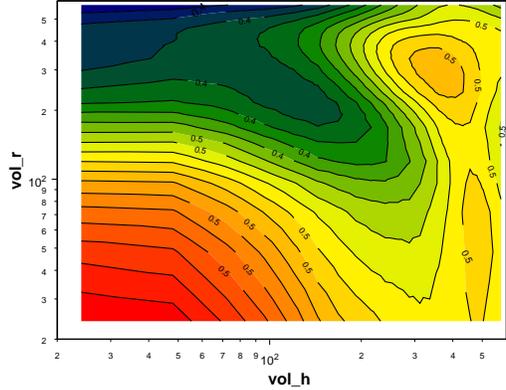


FIG. 9: Zumbach’s mugshot for our model: contour plot of the correlation between past volatility and future volatility, measured on different time scales. For a TRS process, this mug-shot would appear symmetric around the diagonal, whereas empirical data shows, as in this figure, that the region below the diagonal carries more correlation than the region above it.

also noted in [57], but on the example of the S&P 500 index for which the leverage effect is very strong. We have computed this correlation in our model, following the methodology outlined in [13], where the idea of ‘mug-shots’ was introduced to represent graphically such past volatilities/future volatilities correlations. The mug-shot corresponding to our model is shown in Fig. 9. It is clear that our model – almost by construction – captures such a non TRS effect. This was already noted in [13] for a similar model.

We think that the time asymmetry revealed by Zumbach’s mug-shots is extremely important: first, it imposes a theoretical constraint on the eligible models of financial time series that most of them fail to obey. Second, it is a direct proof of the existence of feedback effects in financial markets: the history of past price changes does have a direct impact on the decision and behaviour of traders – in plain contradiction with the efficient markets dogma.

V. THE LEVERAGE EFFECT

Up to now, we have only discussed our feedback model under the assumption of a symmetric reaction of the market participants to price changes. If negative price changes have a larger impact than positive price changes, i.e, if the distribution of thresholds shown in Fig. 1 has some asymmetry, one will observe negative correlations between past price changes and future volatilities (leverage effect) and some skewness in the distribution of returns, totally absent from the above model. The natural way to generalize Eq. (9) to account for such an asymmetry is to write:

$$\sigma_i^2 = \sigma_0^2 \left[1 + \varphi \sum_{\ell=1}^{\infty} g_{\ell} \frac{(x_i - x_{i-\ell})}{\sigma_0 \sqrt{\ell\tau}} + \sum_{\ell=1}^{\infty} g_{\ell} \frac{(x_i - x_{i-\ell})^2}{\sigma_0^2 \ell\tau} \right], \quad (43)$$

where φ measures the strength of the asymmetry. The case $\varphi = 0$ reproduces the model studied above, while $\varphi < 0$ induces a leverage effect. A sufficient condition on φ that ensures that σ^2 always remains positive is to impose that each term of the sum over ℓ contributes positively. Writing as an identity $1 = \sum_{\ell=1}^{\infty} g_{\ell}/z_2$, one obtains:

$$X^2 + \varphi X + \frac{1}{z_2} \geq 0 \quad \forall X, \quad (44)$$

or: $z_2 \varphi^2 < 4$.

The leverage correlation can be defined as:[8] [63]

$$\mathcal{L}(i-j) = \frac{1}{\langle r^2 \rangle^{3/2}} \langle r_i^2 r_j \rangle. \quad (45)$$

This quantity is found empirically to be close to zero for $i \leq j$ and negative for $i > j$. It is not difficult to compute exactly this correlation function in our model, provided the distribution of ξ_i is even.

One finds:

$$\langle r_i^2 r_{i-\ell} \rangle = \langle \sigma^2 \rangle \sigma_0 \tau^{3/2} \varphi \sum_{j=\ell}^{\infty} \frac{g_j}{\sqrt{j}}. \quad (46)$$

One should also note that the average volatility is unchanged by the leverage term φ . Therefore, using $\langle \sigma^2 \rangle = \sigma_0^2/(1-z_2)$, we finally find:

$$\mathcal{L}(\ell) = \varphi g \sqrt{1-z_2} \sum_{j=\ell}^{\infty} \frac{1}{j^{1/2+\alpha}} \sim_{\ell \rightarrow \infty} \ell^{1/2-\alpha}. \quad (47)$$

The decay of the empirical leverage correlation with lag, although noisy, can be fitted by a power-law of exponent close to 0.5, not far from $\alpha - 1/2$ (see Fig. 10). A power-law decay of the leverage correlation was also proposed in the context of the multifractal BMD model in [33, 34].

This quantity is important because it governs the behaviour of the skewness of the return distribution on different time scales. It is indeed easy to show that the normalized skewness of returns on scale ℓ , $\mathcal{S}(\ell)$ is given by [11]:

$$\mathcal{S}(\ell) = \frac{3}{\sqrt{\ell}} \sum_{j=1}^{\ell} \left(1 - \frac{j}{\ell}\right) \mathcal{L}(j) \sim_{\ell \rightarrow \infty} \ell^{1-\alpha} \quad (1 < \alpha < 3/2). \quad (48)$$

From the above expression, we see that even if the return distribution is symmetric on the smallest time scale ($\mathcal{S}(1) = 0$), a negative skewness appears for $\ell > 1$ when $\varphi < 0$, and decays back to zero for very large lags. However, once again, the proximity of the critical line $\alpha = 1$ beyond which the process becomes non stationary, leads to a very slow decay of the skewness. Empirically, for daily returns of individual stocks, one finds $\mathcal{S} \approx -0.1$, corresponding to $\varphi \sim -1$ when $\alpha = 1.15$ and $z_2 = 0.85$.

The skewness of stock indices, on the other hand, is generally much larger (by a factor 10) than that of individual stocks. This is due to an enhanced downside correlation, which should be modeled using the multi-asset model discussed below.

Note that the extra asymmetric term introduced in this section actually contributes also to the volatility-volatility correlation \mathcal{F} computed above and also to the kurtosis. For large ℓ , this extra contribution behaves as $(1-z_2)\varphi^2/\ell^\nu$ and adds to the dominant term computed in section III.

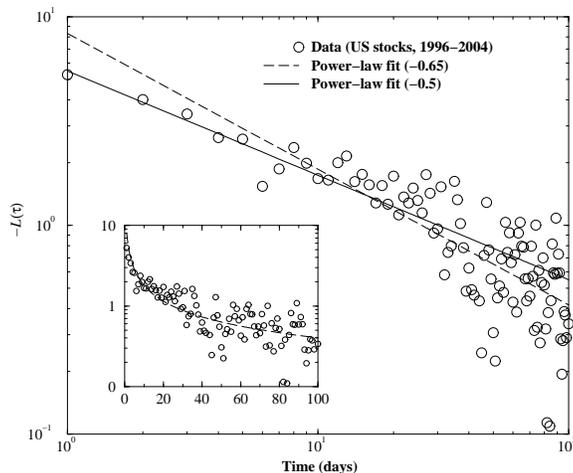


FIG. 10: The leverage correlation \mathcal{L} (with a minus sign) as a function of lag, for US stocks, in a log-log representation. The straight line corresponds to the best power-law fit over the whole range and has slope $-1/2$, whereas the prediction of our model for $\alpha = 1.15$ is a slope of -0.65 (dashed line). Note however the scatter in the empirical data points: the leverage effect for stocks is weak and hard to measure [8]. Inset: same data in a linear representation, with the prediction of our model.

VI. SOFT CALIBRATION WITH REAL TIME SERIES

We have shown in the above sections, using both analytical arguments and numerical simulations, that our model Eq. (9) is able to reproduce semi-quantitatively many of the stylized facts of financial time series that have been reported and studied in the literature. We have in fact shown, in many of the above figures, empirical data that match quite well, at least to the eye, the predictions of the model. What do we mean by ‘semi-quantitatively’? Can one be more quantitative and calibrate, in a standard econometric sense, our model to empirical data?

We believe that our model is interesting precisely because it clearly underlines the limits of such an ambition. The empirical data clearly suggests that any faithful statistical model of financial time series must be somehow close to being non-stationary. This is obvious from the very existence of option markets, which demonstrate the difficulty of measuring and predicting the volatility, even on rather long periods: the at-the-money vol of long-dated options still moves around quite a bit from day to day and there is a persistent smile, symptomatic of a long-memory extending to a few years [11]. We have shown in the above figures that empirical data on stocks seem to favor values of z_2 and α that drive our model very close to instability. This means that even a million step long simulation of our theoretical model for realistic parameters is insufficient to determine the true value of the volatility to better than 5% (see Table I); such an uncertainty affects all the observables of the model. How can one believe that anything more precise than this can be reached on real empirical data? Available time spans are necessarily restricted, true jumps and overnight effects make the returns even more kurtic, true seasonalities (day, week, month, quarters, years) certainly play a role, and non-stationarities (for example, the acceleration of the trading frequency with time) plague any attempt to represent the dynamics of financial markets with fixed values of the parameters on very long time scales. No test, and no model, should aim at more precision than reality itself.

In this situation, we think that the only reasonable strategy is what one could call ‘soft calibration’, in the following sense: instead of focusing on a few observables that one tries to reproduce as accurately as possible to calibrate the model (which is always possible), one should instead find a set of parameters that approximately accounts for *as many* different observations as possible, and cross check the overall consistency of the model. This consistency is more important and more stringent than a perfect fit of an intrinsically elusive target. Calibration in these extreme conditions is an ill-posed problem that, we believe, must be supplemented by intuition on what is important and plausibility.

A. Calibration on stylized facts

How does this work in practice for the model we studied? In the simplest version that we developed, the model has three important parameters: τ , z_2 , α (the value of σ_0 merely sets the scale of the returns, but has no bearing on the structural properties of the model). Accounting for the leverage effect adds one more parameter, φ . But we already know, both from the numerical results that show that our model leads to a non-monotonous kurtosis, and from common sense, that unpredictable jumps must be present and should be factored in through a non-Gaussian noise term ξ , which most probably has itself fat, power-law tails [13, 58]. This adds at least another parameter, which would play an important role in an extended formulation of the model. Neglecting for now this extra complication, our strategy is based on the idea that different observables probe differently the influence of all parameters. This is in fact how we organized the numerical results of section IV:

- The distribution of the volatility or of the returns probes primarily the value of z_2 . The tail exponent μ , and any measure of non-Gaussianity helps restricting the range of acceptable values of z_2 (see Fig. 3 and Table I).
- The temporal correlations of the volatility is primarily sensitive to the value of α , and can be used to limit the acceptable range of this parameter (see Fig. 4-a), whereas the correlation of the log-volatility is most sensitive to the value of z_2 (see Fig. 4-b and Table 1).
- The evolution of the non-Gaussian cumulants κ and Υ is sensitive to z_2 , α , but also to the value of the elementary time scale τ (see Figs. 5-a,b). This has enabled us to fix $\tau \approx 1/300$ day to be consistent with empirical data. Of course this leaves us with the task of curing the unfriendly looking short scale kurtosis, but as mentioned above, this could be easily be dealt with a non-Gaussian noise ξ . This however means that the optimal value of z_2 would be slightly reduced, since part of the kurtosis would already be accounted for.
- The multifractal analysis provides a stringent cross-check of the choice of parameters, since the multifractal spectrum ζ_n is quite sensitive to the value of z_2 (see Figs 6-a,b).
- The consistency of the model can be probed further by analyzing the response of the volatility to shocks of different amplitudes, and studying the Omori plots (see Figs 7, 8). An acceptable description of this rather subtle statistics is, we believe, another useful constraint on the parameter range.
- Interestingly, the leverage correlation is totally decoupled from other observables and can be determined independently from the study of the asymmetry of the distribution of returns and asymmetric volatility correlations, that allow one to fix the parameter φ . [Note however that $\varphi \neq 0$ adds a contribution to the kurtosis of the process].

Following these steps is how we ‘calibrated’ our model on the average behaviour of 252 liquid US stocks in the four-year period 2000-2003. From Figs. 4-6, we see that the value $\alpha = 1.15$ allows one to capture correct time dependence of the volatility correlation and of the evolution of non-Gaussian cumulants, whereas the choice of z_2 in the range 0.85 – 0.90 allows one to capture the correct level of non-Gaussianity and multifractality (the parameter λ^2 appearing in Table I and in Figs 6). These values of z_2 and α allow us to reproduce quite satisfactorily the whole set of observables that we have studied, in particular the Student-like shape of the distribution of returns with a power-law tail index in the right range, and the slow decay of the volatility correlation and of the kurtosis. The choice of the time scale τ is dictated by Figs. 4-5, and is found to be on the order of 1/300th of a day (a few minutes). The value of both z_2 and τ will probably be affected by the inclusion of a non-Gaussian noise ξ – we leave the detailed study of this effect for future work.

B. Volatility prediction

There is another, more direct way to test the consistency of our model, which to some extent avoids the problem of the non-Gaussian nature of the noise ξ (but is still confronted with the intrinsic problems of long memory and non stationarity). The idea is to fix the exponent α and to regress, on empirical data, an estimate of the daily square volatility on the computed feedback strength, defined as:

$$X_i = \sum_{\ell=1}^{\infty} \frac{(x_i - x_{i-\ell})^2}{\ell^{1+\alpha}}. \quad (49)$$

In practice, we have estimated a noisy proxy of square volatility of a stock as $\sigma_i^2 = (H - L + |O - C|)^2 / 4O^2$, where O, H, L, C stand for Open, High, Low, Close. We have computed X_i using the open prices and truncated the sum

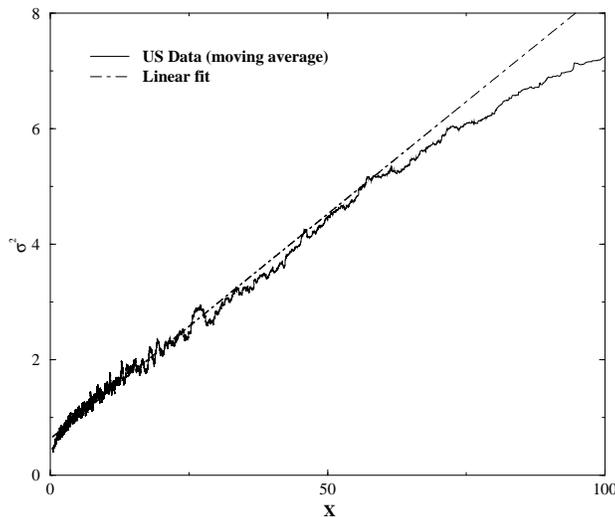


FIG. 11: Scatter plot of σ_i^2 vs. X_i computed daily for 252 US stocks during the four-year period 2000-2003. The coordinates of each point were rescaled by the average square volatility of the stock during that time period. A moving average over 1500 points was performed, unveiling the nearly linear average behaviour of σ_i^2 on X_i , assumed in our model. One can even notice a negative curvature for large X , as suggested by the saturation mechanism we invoked in section II to motivate the model.

over ℓ beyond 500 days, which of course is not very accurate because when α is close to one, the above sum converges only very slowly.

We then plot for all stocks $\sigma_i^2 / \langle \sigma^2 \rangle$ vs. $X_i / \langle X \rangle$. Using our data set, this gives $\sim 400,000$ points; the correlation coefficient between the two sets is found to be 0.285. This value is rather high in view of the roughness of our volatility proxy. The result is shown in Fig. 11, where we have performed a moving average over 1500 points. As one can see the assumption of a linear relation between σ_i^2 and X is rather remarkably borne out, over a rather large range of X_i . From the slope and intercept of the linear relation, we obtain a direct estimate of z_2 , which we find to be ≈ 0.9 , quite close indeed to our previous determination. This direct estimate shows that (a) the basic assumption of the model, that past price changes feedback in the volatility as in Eq. (9), seems to be realistic; and (b) the model is indeed rather close to an instability, with a feedback mechanism that leads to a substantial increase of the volatility.

The direct determination of α using this method is however difficult: one could think of varying α and choosing the value corresponding to the maximal correlation between X_i and σ_i . Unfortunately, the dependence of this correlation coefficient on α is weak and does not allow to extract a meaningful minimum, although one can see that $\alpha = 1.15$ is indeed in the range where the correlation is largest. One could also extend the above method to account for the leverage effect and estimate directly the asymmetry parameter φ .

C. Summary

In summary, we have shown that using a variety of different observables, the range of acceptable values for the parameters of the model can be approximately determined. We have found that using these parameters, all stylized facts can be quantitatively accounted for. Due to the proximity of the unstable regime, however, a very precise determination of optimal parameters seems illusory. On the other hand, the basic assumption of our model, that past price changes feedback in the volatility through Eq. (9), is rather convincingly supported by the results shown in Fig. 11.

VII. GENERALIZATION TO MULTIASSET MODELS

An interesting generalization of the above model concerns the multiasset situation, for example baskets of different stocks with cross-correlations both in the returns and in the volatility. An obvious generalization of our model to this case reads:

$$x_{i+1}^a - x_i^a = r_i^a = \sigma_i^a \zeta_i^a, \quad (50)$$

where i denotes the time index and a labels the stocks. The ξ_i^a 's are characterized by certain correlation matrix $C_{ab} = \langle \xi_i^a \xi_i^b \rangle$ encoding the usual sectorial correlations. For the σ_i^a 's, we write, in full generality:

$$\sigma_i^{a2} = \sigma_0^{a2} \left[1 + \sum_{\ell=1}^{\infty} g_{\ell} \sum_b H^{ab} \frac{(x_i^b - x_{i-\ell}^b)}{\sigma_0^b \sqrt{\ell\tau}} + \sum_{\ell=1}^{\infty} g_{\ell} \sum_b G^{ab} \frac{(x_i^b - x_{i-\ell}^b)^2}{\sigma_0^{b2} \ell\tau} \right]. \quad (51)$$

We leave the investigation of this rich model for future work; thanks to the matrix structure of the feedback effect H and G , one can reproduce a large variety of volatility cross-correlations and leverage effects. Here, we note that the average volatilities obey the following matrix equation:

$$\sum_b \left(\delta_{ab} - \sum_{\ell=1}^{\infty} g_{\ell} G^{ab} \frac{\sigma_0^{a2}}{\sigma_0^{b2}} \right) \langle \sigma^{b2} \rangle = \sigma_0^{a2}, \quad (52)$$

leading to a criterion for the stability of the model, which is that the smallest eigenvalue of the matrix on the left hand side of this equation must remain positive, generalizing the above criterion $1 - z_2 < 1$.

From Eq.(51) one can also estimate the leverage effect for index returns, which can be much enhanced if the matrix H^{ab} has large off diagonal values compared to G^{ab} , meaning that a downward move on any other stock b is perceived as a source of risk for stock a , and triggers extra trades on all other stocks as well.

VIII. CONCLUSION AND PERSPECTIVES

In this work, we have proposed and studied, both analytically and numerically, a multiscale feedback model of volatility. This ARCH-like model (similar to the one studied by Zumbach in [13]) assumes that the volatility is governed by the observed past price changes on different time scales, which, we argue, directly influence the activity of traders. Assuming a power-law distribution of the time horizon of different traders, we obtain a model that captures most stylized facts of financial time series: Student-like distribution of returns with a power-law tail, long-memory of the volatility, slow convergence of the distribution of returns towards the Gaussian distribution, multifractality and anomalous volatility relaxation after shocks. The model, at variance with recent multifractal models that are strictly time reversal invariant, reproduces the time asymmetry of financial time series revealed by Zumbach's mug-shots: past large scale volatility influence future small scale volatility.

The most important conclusion of our work is the following: in order to quantitatively reproduce empirical observations, the parameters must be chosen such that our model is 'doubly' close to an instability, i.e. two parameters are close to values beyond which the process becomes non stationary. This means that (a) the feedback effect is important and substantially increases the volatility, and (b) that the model is intrinsically difficult to calibrate because of the very long range nature of the correlations and the slow convergence of all observables. However, by imposing the consistency of the model predictions with a large set of different empirical observations, a reasonable range of the parameters value can be determined. Furthermore, the adequacy of the basic assumption of our model, i.e. that the instantaneous volatility is directly related to a power-law superposition of past square returns on different time scales, can be directly assessed. The model can easily be generalized to account for jumps (a feature needed to correct an unrealistic non monotonous behaviour of the kurtosis), skewness and multiasset correlations.

The interest of this type of models, compared to (multifractal) stochastic volatility models, is that their fundamental justification, in terms of agent based strategy, is relatively direct and plausible. We believe this is a strong constraint which should guide the construction of any mathematical model of reality. On the other hand, our fundamental assumption, Eq. (9), is in contradiction with the efficient market hypothesis, which asserts that the price past history should have no bearing whatsoever on the behaviour of investors. The large correlation that we find between past price changes and present volatility (see Fig. 11) indicates that this influence is in fact quite strong. This result is, in our view, yet another direct piece of evidence against the efficient market hypothesis, and a clear mechanism leading to excess volatility in financial markets.

Turning to financial engineering applications, such as risk control and option pricing, our model provides a well defined procedure to *filter* the series past price changes, and to compute the probabilities of the different future paths. Similar models have been shown to fare rather well [46, 48]. Once 'softly' calibrated, the model can in principle be used for VaR estimates and option pricing. However, its mathematical complexity does not allow, in general, for explicit analytical solutions and probably one has to resort either to approximate treatments or to numerical, Monte-Carlo methods. The difficulty of long-memory models is that the option price must be computed conditional to the whole past history, which considerably complexifies both analytical solutions and Monte-Carlo methods. In other words, both the option price and the optimal hedge are no longer simple functions of the current price, but *functionals* of the whole price history. Finding operational ways to account for this history dependence seems to us a major challenge, on which we hope to work in the near future.

Acknowledgments

We want to thank D. Farmer and T. Lux for inviting us to write this paper, which was initiated by discussions during the Leiden workshop: “Volatility in financial markets”, October 2004. Discussions with J.F. Muzy and G. Zumbach have been of great help. L.B. also thanks J. Evnine for ongoing discussions and support.

-
- [1] L. Bachelier, *Théorie de la spéculation* (1900), Reprinted by J. Gabay, Editor, Paris 1995.
- [2] T. Lux, *The stable Paretian hypothesis and the frequency of large returns: an examination of major German stocks*, Applied Financial Economics, **6**, 463, (1996).
- [3] F. Longin, *The asymptotic distribution of extreme stock market returns*, Journal of Business, **69** 383 (1996)
- [4] V. Plerou, P. Gopikrishnan, L.A. Amaral, M. Meyer, H.E. Stanley, *Scaling of the distribution of price fluctuations of individual companies*, Phys. Rev. **E60** 6519 (1999); P. Gopikrishnan, V. Plerou, L. A. Amaral, M. Meyer, H. E. Stanley, *Scaling of the distribution of fluctuations of financial market indices*, Phys. Rev. **E 60** 5305 (1999)
- [5] A. Lo, *Long term memory in stock market prices*, Econometrica, **59**, 1279 (1991).
- [6] Z. Ding, C. W. J. Granger and R. F. Engle, *A long memory property of stock market returns and a new model*, *J. Empirical Finance* **1**, 83 (1993).
- [7] Y. Liu, P. Cizeau, M. Meyer, C.-K. Peng, H. E. Stanley, *Correlations in Economic Time Series*, *Physica A* **245** 437 (1997)
- [8] J.P. Bouchaud, A. Matacz, M. Potters, *The leverage effect in financial markets: retarded volatility and market panic* Physical Review Letters, **87**, 228701 (2001)
- [9] D. M. Guillaume, M. M. Dacorogna, R. D. Davé, U. A. Müller, R. B. Olsen and O. V. Pictet, *Finance and Stochastics* **1** 95 (1997); M. Dacorogna, R. Gençay, U. Müller, R. Olsen, and O. Pictet, *An Introduction to High-Frequency Finance* (Academic Press, London, 2001).
- [10] R. Mantegna & H. E. Stanley, *An Introduction to Econophysics*, Cambridge University Press, 1999.
- [11] J.-P. Bouchaud and M. Potters, *Theory of Financial Risks and Derivative Pricing*, Cambridge University Press, 2004.
- [12] T. Lux, *Market Fluctuations: Scaling, Multi-Scaling and Their Possible Origins* in: A. Bunde and H.-J. Schellnhuber, eds., *Theories of Disasters: Scaling Laws Governing Weather, Body and Stock Market Dynamics*. Berlin: Springer (2003).
- [13] G. Zumbach, P. Lynch, *Market heterogeneity and the causal structure of volatility*, Quantitative Finance, **3**, 320, (2003).
- [14] S. I. Boyarchenko, and S. Z. Levendorskii, *Non-gaussian Merton-Black-Scholes Theory*, World Scientific (2002)
- [15] R. Cont, P. Tankov, *Financial modelling with jump processes*, CRC Press, 2004.
- [16] S.L. Heston, *A closed-form solution for options with stochastic volatility with applications to bond and currency options*, Rev. of Fin. Studies, **6**, 327-343, 1993
- [17] J.-P. Fouque, G. Papanicolaou, G. and K. R. Sircar, *Derivatives in financial markets with stochastic volatility*, Cambridge University Press, 2000
- [18] A. A. Dragulescu and V. M. Yakovenko, *Probability distribution of returns in the Heston model with stochastic volatility*, Quantitative Finance **2**, 443-453 (2002).
- [19] J. Perello, J. Masoliver, J.P. Bouchaud, *Multiple time scales in volatility and leverage correlations: a stochastic volatility model*, Appl. Math. Fin. **11**, 1 (2004).
- [20] P. Carr, H. Geman, D. Madan, M. Yor, *Stochastic Volatility for Levy Processes*, Mathematical Finance, **13**, 345 (2003).
- [21] A. Fisher, L. Calvet, B.B. Mandelbrot, *Multifractality of DEM/\$ rates*, Cowles Foundation Discussion Paper 1165.
- [22] S. Ghashghaie, W. Breymann, J. Peinke, P. Talkner, Y. Dodge, *Turbulent cascades in foreign exchange markets* Nature **381** 767 (1996).
- [23] F. Schmitt, D. Schertzer, S. Lovejoy, *Multifractal analysis of Foreign emchange data*, Applied Stochastic Models and Data Analysis, **15** 29 (1999);
- [24] M.-E. Brachet, E. Taffin, J.M. Tchéou, *Scaling transformation and probability distributions for financial time series*, Chaos, Solitons and Fractals, **11** 2343 (2000).
- [25] J.P. Bouchaud, M. Potters, M. Meyer, *Apparent multifractality in financial time series*, Eur. Phys. J. **B 13**, 595 (1999).
- [26] L. Calvet, A. Fisher, *Forecasting multifractal volatility*, Journal of Econometrics, **105**, 27, (2001).
- [27] L. Calvet, A. Fisher, *Multifractality in Asset Returns: Theory and Evidence*, Review of Economics and Statistics **84**, 381-406 (2002).
- [28] T. Lux, *Turbulence in financial markets: the surprising explanatory power of simple cascade models*, Quantitative Finance **1**, 632, (2001)
- [29] T. Lux, *Multi-Fractal Processes as a Model for Financial Returns: Simple Moment and GMM Estimation*, in revision for Journal of Business and Economic Statistics
- [30] J.-F. Muzy, J. Delour, E. Bacry, *Modelling fluctuations of financial time series: from cascade process to stochastic volatility model*, Eur. Phys. J. B **17**, 537-548 (2000); E. Bacry, J. Delour and J.F. Muzy, *Multifractal random walk*, Phys. Rev. E **64**, 026103 (2001).
- [31] J.F. Muzy and E. Bacry, *Multifractal stationary random measures and multifractal random walks with log-infinitely divisible scaling laws*, Phys. Rev. E **66**, 056121 (2002).
- [32] For an inspiring recent account, see B. M. Mandelbrot, R. Hudson, *The (mis)behaviour of markets*, Perseus Books, Cambridge MA (2004)

- [33] B. Pochart and J.P. Bouchaud, *The skewed multifractal random walk with applications to option smiles*, Quantitative Finance, **2**, 303 (2002).
- [34] Z. Eisler, J. Kertesz, *Multifractal model of asset returns with leverage effect*, e-print cond-mat/0403767
- [35] L. Borland, J. P. Bouchaud, J.-F. Muzy, G. Zumbach, *The dynamics of Financial Markets: Mandelbrot's multifractal cascades, and beyond*, Wilmott Magazine, March 2005.
- [36] T. Lux, M. Marchesi, *Scaling and criticality in a stochastic multiagent model*, Nature **397**, 498 (1999); T. Lux, M. Marchesi, *Volatility Clustering in Financial Markets: A Microsimulation of Interacting Agents*, Int. J. Theo. Appl. Fin. **3**, 675 (2000).
- [37] D. Challet, A. Chessa, M. Marsili, Y.C. Zhang, *From Minority Games to real markets*, Quantitative Finance, **1**, 168 (2001) and refs. therein; D. Challet, M. Marsili, Y.C. Zhang, *The Minority Game*, Oxford University Press, 2004.
- [38] W. Brock, C. Hommes, *Econometrica* **65**, 1059 (1997); C. Hommes, *Financial markets as nonlinear adaptive evolutionary systems*, Quantitative Finance, **1**, 149, (2001)
- [39] C. Chiarella, X-Z. He, *Asset price and wealth dynamics under heterogeneous expectations*, Quantitative Finance, **1**, 509, (2001)
- [40] I. Giardina, J.P. Bouchaud, *Bubbles, Crashes and Intermittency in agent based market models*, European Journal of Physics **B 31**, 421 (2003).
- [41] L. Borland, *A Theory of non-Gaussian Option Pricing*, Quantitative Finance **2**, 415-431, (2002).
- [42] Cf. M. Gell-Mann and C. Tsallis, *Non extensive entropies – interdisciplinary applications*, Oxford University Press, NY, 2004.
- [43] L. Borland, J. P. Bouchaud, *A non-Gaussian Option Pricing model with skew*, Quantitative Finance, **4**, 499-514 (2004)
- [44] see, e.g. J. Hull, *Futures, Options and other Financial Derivatives*, Prentice Hall, 2004.
- [45] L. Borland, *A multi-time scale non-Gaussian model of stock returns*, e-print cond-mat/0412526
- [46] G. Zumbach, *Volatility processes and volatility forecast with long memory*, Quantitative Finance **4** 70 (2004).
- [47] R. T. Baillie, T. Bollerslev, H. O. Mikkelsen, *Fractionally integrated GARCH*, J. Econometrics, **31**, 3 (1996); T. Bollerslev, H. O. Mikkelsen, *Modeling and pricing long memory in stock market volatility*, J. Econometrics, **73**, 3 (1996).
- [48] M. Dacorogna, U. Muller, R. B. Olsen, O.Pictet, *Modelling Short-Term Volatility with GARCH and HARCH Models*, Nonlinear Modelling of High Frequency Financial Time Series, 1998
- [49] T. Bollerslev, R. F. Engle, D. B. Nelson, *ARCH models*, in R. F. Engle, D. McFadden, Edts, Handbook of Econometrics, Vol. 4, Elsevier Science, Amsterdam (1994).
- [50] R. J. Shiller, *Do Stock Prices move too much to be justified by subsequent changes in dividends?*, American Economic Review, **71**, 421 (1981); R. J. Shiller, *Irrational Exuberance*, Princeton University Press (2000).
- [51] On feedback effects on the volatility, see also: M. Wyart, J. P. Bouchaud, *Self-referential behaviour, overreaction and conventions in financial markets*, to appear in JEBO.
- [52] S. Miccichè, G. Bonanno, F. Lillo, R. N. Mantegna, *Volatility in financial markets: stochastic models and empirical results*, Physica **A 314**, 756, (2002).
- [53] D. Sornette, Y. Malevergne, J.F. Muzy, *What causes crashes*, Risk Magazine, 67 (Feb. 2003).
- [54] A. G. Zawadowski, J. Kertesz, G. Andor, *Large price changes on small scales*, e-print cond-mat/0401055
- [55] F. Lillo and R. N. Mantegna, *Power law relaxation in a complex system: Omori Law After a Financial Market Crash*, Physical Review E **68**, 016119 (2003)
- [56] Y. Pomeau, *Symétrie des fluctuations dans le renversement du temps*, J. Physique **43** 859 (1985).
- [57] A. Arneodo, J.-F. Muzy, D. Sornette, *Causal cascade in the stock market from the ‘infrared’ to the ‘ultraviolet’*, Eur. Phys. J. **B 2**, 277 (1998)
- [58] see, e.g. J. Doyne Farmer, Laszlo Gillemot, Fabrizio Lillo, Szabolcs Mike, Anindya Sen, *What really causes large price changes?*, e-print cond-mat/0312703.
- [59] Note however that a model based on log-prices has no deep theoretical justifications and might not be the best representation of reality. See [8, 11] for a detailed discussion of this point.
- [60] The following argument is only intended to be qualitative. In reality, the trading range between t_{i-1} and t_i is given by the high minus low of that period, rather than the open to close; all thresholds with that interval will be hit. However, the high-low is of the same order as the open-close r_{i-1} – in fact, for a random walk, the average high-low is exactly twice the average $|r_{i-1}|$, a result due to Bachelier himself!
- [61] It might actually be more consistent, at the level of the theoretical justification of the model, to replace σ_0^2 in the feedback term by a moving average of the past volatility itself, since traders are sensitive to the recent level of volatility. This would make the model more complicated, and we leave this extension for future investigation.
- [62] The marginal case $\alpha = 1$ with a cut-off for $\ell > L$ is quite interesting and might in fact resemble the strictly multifractal BMD model.
- [63] A slightly different normalization was actually used in [8].