

AN INTRODUCTION TO ECONOPHYSICS AND QUANTITATIVE FINANCE

RÉMY CHICHEPORTICHE¹, THIBAUT JAISSON², IACOPO MASTROMATTEO³ AND VINCENT VARGAS⁴

Abstract. This paper gives an account of the talks given by the authors at the 2014 MAS conference in Toulouse. These talks present recent research in the field of econophysics and quantitative finance.

Résumé. Cet article rend compte des exposés donnés par les auteurs durant les journées MAS 2014 à Toulouse. Ces exposés présentent des recherches récentes dans le domaine de l'éconophysique et de la finance quantitative.

Key words or phrases: Econophysics, multifractals, Hawkes processes, market impact.

AMS 2000 subject classifications: 28A80, 60G15, 60G25, 60G57, 60F05, 60F17, 60G55, 62P05 .

1. INTRODUCTION

In this paper, we will try to present recent trends of research related to the field of econophysics and quantitative finance. Therefore, it could be useful to first explain a bit how the field of econophysics differs from classical mathematical finance.

Modern mathematical finance is really born with the founding paper by Black-Scholes on option pricing [10]. The purpose of the pioneering work [10] is to price standard vanilla options, i.e. a contract which gives its holder the right to buy a stock at a future time T at some fixed price K . The main assumption in [10] is to suppose that the stock price process $(S_t)_{t \geq 0}$ is a geometric Brownian motion (GBM) with drift

$$S_t = e^{(\mu - \frac{\sigma^2}{2})t + \sigma B_t}$$

where μ is a drift (supposed to be constant) and σ the average volatility. The main novelty of this work and Merton's work [39] is to combine portfolio theory with option pricing; more precisely, the seller of an option can cover his risk by constituting a dynamic portfolio where he invests his money on a bank account but also on the stock S_t . Though the idea of using a dynamic portfolio in the theory of option pricing is very powerful and general, a few amazing features arise in the specific context of the GBM:

- There is no arbitrage, i.e. the possibility to make riskless money starting with nothing;
- The option price does not depend on the drift μ ;
- With an appropriate portfolio, one can perfectly hedge the risk of an option: the market is said to be complete.

¹ Chaire de Finance Quantitative, École Centrale Paris, 92290 Châtenay-Malabry.

² Centre de Mathématiques Appliquées, CNRS, École Polytechnique, UMR 7641, 91128 Palaiseau.

³ Centre de Mathématiques Appliquées, CNRS, École Polytechnique, UMR 7641, 91128 Palaiseau.

⁴ ENS Ulm, DMA, 45 rue d'Ulm, 75005 Paris.

These features are very desirable for an option seller and make markets look perfectly safe. However, all market participants know that a GBM is not a good model for financial stocks. Indeed, a Brownian motion is Gaussian hence excludes large stock movements in a way which is incompatible with what is observed during a financial crisis for instance. Also, Brownian motion has independent increments whereas it is not hard to convince oneself that stocks have temporal correlations: for instance, it is well known that volatility on markets is highly correlated in time. In particular, the pool of all option prices observed on markets cannot be reproduced at some fixed instant using a GBM model. This led to the introduction of the local volatility model by Dupire in [21]. Though the model looks very counter-intuitive, it can fit any pool of option prices at some fixed instant and it shares exactly the same specific features as the GBM: no arbitrage, completeness, etc... Hence, Dupire's work came as a relief for the community of mathematical finance! So, case closed? Well not really because, although the model fits any pool of option prices at some fixed initial time, the model and the pool of option prices observed on the market diverge dangerously under many realistic scenarios when time passes. This clearly indicates that one should work with realistic models of stock prices: this brings us precisely to the field of econophysics.

In some sense, econophysics can be seen as the endless quest of understanding and modeling stock prices. The pillars of the field are the so-called stylized facts which are the universal statistical features observed on all financial assets. With respect to this problem, one should distinguish two time scales of observation corresponding to low-frequency and high-frequency finance. More precisely, there is some time threshold τ_c around 10 minutes such that prices on time lags below τ_c behave differently than on time lags above τ_c . Therefore, there exist low-frequency and high-frequency stylized facts. When one deals with high frequency finance, one typically studies quantities more general than the price process (which is in fact ill-defined due to the discretization effects at small time scales) like the order book dynamics: in that case, one is interested in the stylized facts of the order book, i.e. its universal statistical features.

Now, typical basic questions in econophysics are

- Can one construct a model which satisfies these stylized facts at all time scales?
- Can one construct a (agent-based) model which explains the high-frequency stylized facts (order book dynamics, price formation process, etc...)?

The next sections will tackle these questions. Each section corresponds to a talk given by one of the authors at the 2014 MAS conference. In particular, the first two sections will study two low-frequency models for the price process: the multifractal random walk (talk by V. Vargas) and the self-reflective dynamics of the stock return fluctuations (talk by R. Chicheportiche). The third section will then present the nearly unstable Hawkes model (talk by T. Jaisson), which reproduces properties of the order flow at both high and low frequencies. The fourth section will finally focus on the problem of market impact (talk by I. Mastromatteo), by illustrating the predictions of a latent liquidity model.

2. FORECASTING VOLATILITY WITH THE MULTIFRACTAL RANDOM WALK (MRW)

V. VARGAS

2.1. Definition of the MRW and stylized facts

The multifractal random walk (MRW) can be seen as a stochastic volatility model, i.e. the price process $(S_t)_{t \geq 0}$ solves the following SDE

$$dS_t/S_t = \sigma_t dW_t \quad (1)$$

where $(W_t)_{t \geq 0}$ is a Brownian motion and the volatility process is independent from W and has the following form

$$\sigma_t = \sigma e^{\lambda \omega_t^T - \lambda^2 E[(\omega_t^T)^2]}$$

where ω^T is a centered Gaussian “process” (independent of W) with covariance

$$E[\omega_s^T \omega_t^T] = \ln^+ \frac{T}{|t-s|}, \quad (2)$$

with $\ln^+(x) = \max(\ln x, 0)$. The problem is that $(\omega_t^T)_{t \in \mathbb{R}}$ lives in the space of distributions (in the sense of Schwartz) but not in the space of functions hence σ_t is ill-defined. To make sense of this expression, first note that if $(\sigma_t)_{t \in \mathbb{R}}$ is a nice process, Brownian scaling gives the following expression

$$S_t = S_0 e^{\int_{[0,t]} \sigma_s dW_s - \frac{1}{2} \int_{[0,t]} \sigma_s^2 ds} = S_0 e^{B_{M[0,t]} - \frac{1}{2} M[0,t]} \quad (3)$$

where B is a Brownian motion and M is the random measure given by the following expression:

$$M[0, t] = \int_0^t \sigma_s^2 ds, \quad t \geq 0.$$

Notice that (3) makes sense in great generality for any (random) measure M (whereas definition (1) makes sense when σ_t is a function, i.e. M is absolutely continuous with respect to the Lebesgue measure). Now, it happens that Kahane’s theory of Gaussian multiplicative chaos [30] enables to make sense of the following formal measure

$$M^T[0, t] = \sigma^2 \int_0^t e^{2\lambda \omega_s^T - 2\lambda^2 E[(\omega_s^T)^2]} ds, \quad t \geq 0,$$

where ω^T is the Gaussian process with covariance (2). M^T is defined by a limit procedure $M^T = \lim_{\tau \rightarrow 0} M_\tau^T$ where

$$M_\tau^T[0, t] = \sigma^2 \int_0^t e^{2\lambda \omega_s^{\tau, T} - 2\lambda^2 E[(\omega_s^{\tau, T})^2]} ds, \quad t \geq 0,$$

and $\omega^{\tau, T}$ is the smooth Gaussian process:

$$E[\omega_s^{\tau, T} \omega_t^{\tau, T}] = \ln^+ \frac{T}{|t-s| + \tau}.$$

The MRW model is then (3) with M^T as choice of M and was first introduced and studied by Bacry, Delour and Muzy in [5] (the idea of looking at a time changed Brownian motion as a model for a stock price goes back to the work of Mandelbrot and Taylor [35]).

One of the interesting features of the MRW model is that it exhibits many of the low-frequency stylized facts observed on financial assets; in particular:

- the volatility fluctuates and is approximately lognormal
- the volatility is highly correlated and exhibits long range correlations.

Note however that the (recentered) model is symmetric since $(B_{M^T[0,t]})_t \stackrel{(Law)}{=} -(B_{M^T[0,t]})_t$ and therefore it does not reproduce the celebrated leverage effect (see [15]). Roughly, the leverage effect is a quantitative measure of the fact that large drawdowns of a stock increases its volatility more than large upward movements. This effect is essentially observed on indexes and large market capitalization stocks. Therefore, the MRW model is essentially used to model currencies where there is usually little leverage effect.

2.2. Forecasting the log-volatility

Since the MRW is a good model for financial assets, it is natural to try to establish prediction formulas for the volatility of the model, namely M^T . In order to use these prediction formulas, one must also calibrate the model on financial data. The value of λ^2 is rather stable in time on an asset and its value (which depends on

the asset) is found to be roughly in the interval $[0.01; 0.08]$. However, the calibration of σ and T to the market is much more delicate. Indeed, we can make the following observation:

Lemma 2.1 (Invariance by integral scale change). *For all $T < T'$:*

$$(M^{T'}[0, t])_{t \in [0, T]} \stackrel{(Law)}{=} e^{2\lambda\Omega_{T'/T} - 2\lambda^2 \ln \frac{T'}{T}} (M^T[0, t])_{t \in [0, T]}$$

where $\Omega_{T'/T}$ is a centered Gaussian variable of variance $\ln \frac{T'}{T}$ that is independent of $(M^T([0, t]))_{t \in [0, T]}$.

As a consequence of the above lemma, if the observation window of M^T is of length less than or equal to T , it is impossible to determine σ and T (ill-posed problem). Typically, when you try to measure the integral scale T on a financial asset, you find it comparable to the scale of your observation window, hence the measure is unreliable. However, one idea to overcome the problem is to let $T \rightarrow \infty$ in order to get rid of the determination of σ and T at the same time!

On the level of the log volatility, recall that ω^T is a Gaussian measure on the space of tempered distributions $\mathcal{S}'(\mathbb{R})$ (in the sense of Schwartz) such that:

$$\begin{aligned} \forall \phi \in \mathcal{S}(\mathbb{R}), \quad E(e^{i \int_{\mathbb{R}} \phi(t)\omega_t^T dt}) &= e^{-1/2 E((\int_{\mathbb{R}} \phi(t)\omega_t^T dt)^2)} \\ &= e^{-1/2 \int \int_{\mathbb{R}^2} \phi(t)\phi(s) \ln^+(T/|t-s|) ds dt}. \end{aligned}$$

Now, we want to let $T \rightarrow \infty$ in the above formula. In order to make sense of the limit, it is rather straightforward to see that ϕ must be such that $\int_{\mathbb{R}} \phi(t) dt = 0$ in which case we get the following convergence

$$e^{-1/2 \int \int_{\mathbb{R}^2} \phi(t)\phi(s) \ln^+(T/|t-s|) ds dt} \xrightarrow{T \rightarrow \infty} e^{-1/2 \int \int_{\mathbb{R}^2} \phi(t)\phi(s) \ln(1/|t-s|) ds dt}.$$

One then naturally introduces ω , a Gaussian measure on the quotient space $\mathcal{S}'(\mathbb{R})/\mathbb{R}$. Now, if one considers $\mathcal{S}_0(\mathbb{R}) = \{\phi \in \mathcal{S}(\mathbb{R}) ; \int_{\mathbb{R}} \phi(t) dt = 0\}$ then:

$$\begin{aligned} \forall \phi \in \mathcal{S}_0(\mathbb{R}), \quad E(e^{i \int_{\mathbb{R}} \phi(t)\omega_t dt}) &= e^{-1/2 E((\int_{\mathbb{R}} \phi(t)\omega_t dt)^2)} \\ &= e^{-1/2 \int \int_{\mathbb{R}^2} \phi(t)\phi(s) \ln(1/|t-s|) ds dt} \\ &= e^{-1/4 \int_{\mathbb{R}} \frac{|\hat{\phi}(\xi)|^2}{|\xi|} d\xi}. \end{aligned}$$

Therefore, ω can be seen as a rigorous definition of the celebrated $1/f$ noise studied in physics. We consider the reproducing kernel Hilbert space of ω :

$$\mathcal{H}^{1/2}(\mathbb{R}) = \{f \in \mathcal{S}'(\mathbb{R})/\mathbb{R} ; \int_{\mathbb{R}} |\xi| |\hat{f}(\xi)|^2 d\xi < \infty\}.$$

Now, one can prove the following prediction formula:

Theorem 2.1 ([36]). *For all $f \in \mathcal{H}^{1/2}(\mathbb{R})$,*

$$\forall t > 0 \quad E[\omega_t | (\omega_s)_{s < 0} = f] = \frac{1}{\pi} \int_{-\infty}^0 \frac{\sqrt{t}}{(t-s)\sqrt{-s}} f(s) ds.$$

In fact, one can use the above theorem to get exact prediction formulas for ω^T . In order to state the result, we introduce the standard Sobolev space:

$$H^{1/2}(\mathbb{R}) = \{f \in \mathcal{S}'(\mathbb{R}) ; \int_{\mathbb{R}} (1 + |\xi|^2)^{1/2} |\hat{f}(\xi)|^2 d\xi < \infty\}.$$

Let L be some observation window and T the integral scale. There is an explicit kernel $K_{L,T}(t, s)$ (see [36] for the exact expression) such that:

Corollary 2.1 ([36]). *For all $f \in H^{1/2}(\mathbb{R})$,*

$$\forall t \in]0, T - 2L[, E[\omega_t^T | (\omega_s^T)_{-2L < s < 0} = f] = \int_{-2L}^0 K_{L,T}(t, s) f(s) ds.$$

Finally, one can discretize the above formulas to get approximate formulas for the discrete model. In that case, one can also transfer the prediction formulas to the volatility itself (in which case, the prediction formulas will also depend on λ^2): see [36] for more details.

3. MODELING STOCK RETURNS FLUCTUATIONS

R. CHICHEPORTICHE

One of the most striking universal stylized facts of financial returns is the volatility clustering effect, which was first reported by Mandelbrot as early as 1963 [34]. He noted that *large changes in the asset price S_t tend to be followed by large changes, of either sign, and small changes tend to be followed by small changes*. The first quantitative description of this effect was the ARCH model proposed by Engle in 1982 [22]. It formalizes Mandelbrot's hunch in the simplest possible way, by postulating that log-returns $r_t \equiv \ln S_t - \ln S_{t-1}$ can be modeled as conditionally Gaussian random variables $r_t = \sigma_t \xi_t$, where the residuals ξ_t 's are IID random variables, of zero mean and variance equal to unity (while many papers take these ξ_t 's to be Gaussian, it is preferable to be agnostic about their univariate distribution. In fact a consensus emerges on that the ξ_t 's themselves have fat-tails: asset returns are *not* conditionally Gaussian and "true jumps" do occur). The volatility σ_t is time dependent and evolves according to:

$$\sigma_t^2 = s^2 + g r_{t-1}^2. \quad (4)$$

In words, this equation means that the (squared) volatility today is equal to a baseline level s^2 , plus a self-exciting term that describes the influence of yesterday's *perceived* volatility r_{t-1}^2 on today's activity, through a feedback parameter g .

It soon became apparent that the above model is far too simple to account for empirical data. For one thing, it is unable to account for the long memory nature of volatility fluctuations. It is also arbitrary in at least two ways:

- First, there is no reason to limit the feedback effect to the previous day only. One can replace r_{t-1}^2 by a weighted moving average of past squared returns, leading to a large family of models such as:

$$\sigma_t^2 = s^2 + \sum_{\tau=1}^{\infty} k(\tau) r_{t-\tau}^2. \quad (5)$$

A slowly (power-law) decaying kernel $k(\tau)$ is able to account for the long memory of volatility.

- Second, there is no reason to single out the day as the only time scale to define the returns: in principle, returns over different time scales could also feedback on the volatility today [11, 12, 33, 41]. For extended time scales *longer* than the day, this leads to another natural extension of the model as:

$$\sigma_t^2 = s^2 + \sum_{\ell} \sum_{\tau=1}^{\infty} g_{\ell}(\tau) R_{t-\tau}^{(\ell)2}, \quad (6)$$

where $R_t^{(\ell)} = \sum_{\tau=1}^{\ell} r_{t-\tau}$ is the cumulative, ℓ day return between $t - \ell$ and t .

The common point to the zoo of generalizations of the initial ARCH model is that the current level of volatility σ_t^2 is expressed as a quadratic form of past realized returns. The most general model of this kind, called QARCH (for Quadratic ARCH), is due to Sentana [43], and reads:

$$\sigma_t^2 = s^2 + \sum_{\tau=1}^{\infty} L(\tau) r_{t-\tau} + \sum_{\tau, \tau'=1}^{\infty} K(\tau, \tau') r_{t-\tau} r_{t-\tau'}. \tag{7}$$

The QARCH can be seen as a general discrete-time model for the dependence of σ_t^2 on all past returns $\{r_{t'}\}_{t' < t}$, truncated to second order. The linear contribution, which involves $L(\tau)$, captures a possible dependence of the volatility on the sign of the past returns. The quadratic contribution contains, through the matrix $K(\tau, \tau')$, all ARCH-like models studied in the literature.

QARCH models only make sense if the expected volatility does not diverge to infinity. The criterion for stability is easy to establish if the ξ 's are IID and of zero mean, and reads:

$$\text{Tr } K < 1. \tag{8}$$

In this case, the volatility is a stationary process such that $\langle \sigma^2 \rangle \equiv \mathbb{E}[\sigma^2] = s^2 / (1 - \text{Tr } K)$: the feedback-induced increase of the volatility only involves the diagonal elements of K . Note also that the leverage kernel $L(\tau)$ does not appear in this equation.

Empirical study and dataset

We have calibrated the above QARCH model using an appropriate methodology presented in detail in [17]. In particular, we have estimated the matrix elements $K(\tau, \tau')$ both under Maximum Likelihood (ML) and Method of Moments (GMM). We have also implemented a constrained model imposing a simple Long-Term structure on K (ML+LT).

We worked with daily stock prices (Open, Close, High, Low) for $N = 280$ names present in the S&P-500 index from Jan. 2000 to Dec. 2009 ($T = 2515$ days), without interruption. The reference price for day t is defined to be the close of that day C_t . The true volatility is of course unobservable; we have replaced it by the Rogers-Satchell estimator [42].

3.1. Large-dimensional effects

The ubiquitous empirical long-memory correlations of the volatility,

$$\langle (r_t^2 - \langle r_{t'}^2 \rangle_{t'}) r_{t-\tau}^2 \rangle_t \propto \tau^{-\beta},$$

with $0 < \beta < 1$, can be reproduced by diagonal¹ power-law models $K(\tau, \tau) = g \tau^{-\alpha} \mathbf{1}_{\{\tau \leq q\}}$ with exponent $\alpha = (3 - \beta) / 2 \in (1, 3/2)$.

For a given α , g must be smaller than a certain $g_c(\alpha, q)$ for $\langle \sigma^2 \rangle$ to be finite according to Eq. (8). Fig. 1 shows the critical frontier $g_c(\alpha, q)$ for $q = 1, 32, 256$ and $q \rightarrow \infty$. The critical frontier in the limit case $q = \infty$ is given by $g_c = 1 / \zeta(\alpha)$, where $\zeta(\alpha)$ is Riemann's zeta function (solid red). Note in particular that the model is always unstable when $\alpha < 1$, i.e. when the memory of past realized volatility decays too slowly. At the other extreme, $q = 1$, the constraint is well known to be $g = k(1) \leq 1$ (solid red). Similar computations (but more complicated, see [17]) lead to the necessary conditions for finite $\langle \sigma^4 \rangle$.

We have found that a fit corresponding to a power-law behavior, $k(\tau) \approx g \tau^{-\alpha}$ for $\tau \ll q_0$, and an exponential decay for larger lags is very accurate. We find the best parameter estimates to be $\alpha \approx 1.11, g \approx 0.081$ and $q_0 \approx 53 \approx 3$ months.

¹ARCH(q), GARCH and FIGARCH models are all diagonal.

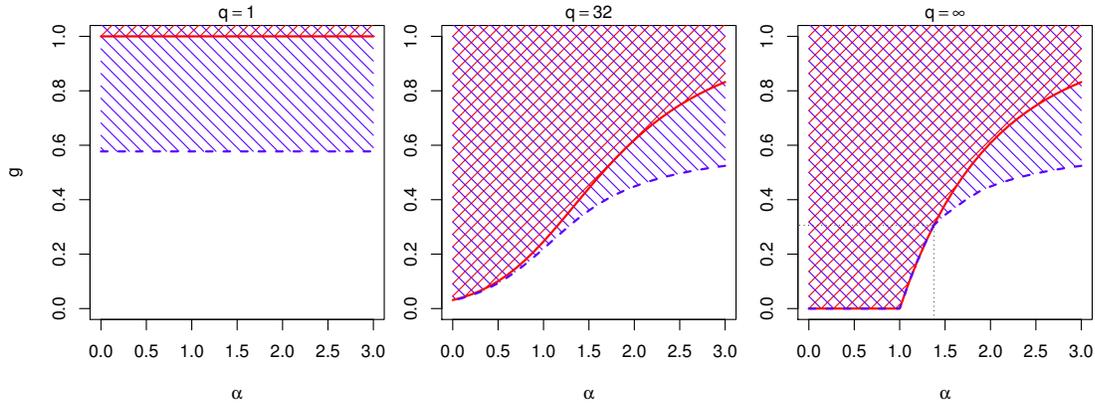


FIGURE 1. Allowed region in the α, g space for $K(\tau, \tau) = g \tau^{-\alpha} 1_{\tau \leq q}$ and $L(\tau) = 0$, according to the finiteness of $\langle \sigma^2 \rangle$ and $\langle \sigma^4 \rangle$. Divergence of $\langle \sigma^2 \rangle$ is depicted by 45° (red) hatching, while divergence of $\langle \sigma^4 \rangle$ is depicted by -45° (blue) hatching. In the wedge between the dashed blue and solid red lines, $\langle \sigma^2 \rangle < \infty$ while $\langle \sigma^4 \rangle$ diverges.

When $q \rightarrow \infty$, the asymptotic value s^2 is equal to only 20% of the observed squared volatility², meaning that volatility feedback effects increase the volatility by a factor ≈ 2.25 . Such a strong amplification of the volatility resonates with Shiller’s “excess volatility puzzle” and gives a quantitative estimate of the role of self-reflexivity and feedback effects in financial markets [19, 44].

3.2. Multi-scale effects

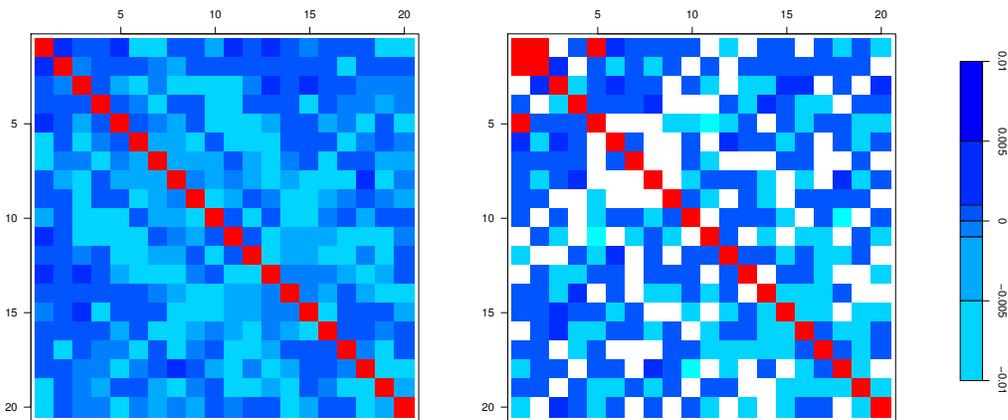


FIGURE 2. Heatmap of the unconstrained model. **Left:** GMM estimation; **Right:** ML estimation (white spots correspond to values smaller than their corresponding error margin). Red spots correspond to values larger than 0.01. Note the negative streaks at large lags, and the significant off-diagonal entry for $\tau = 1, \tau' = 2$ and $\tau = 1, \tau' = 5$.

²When splitting daily returns into close-to-open (night) and open-to-close contributions, we observe an even stronger feedback amplification of the overnight volatility that manifests itself by a nearly vanishing baseline value: most overnight price changes are news-less and purely endogenous [11].

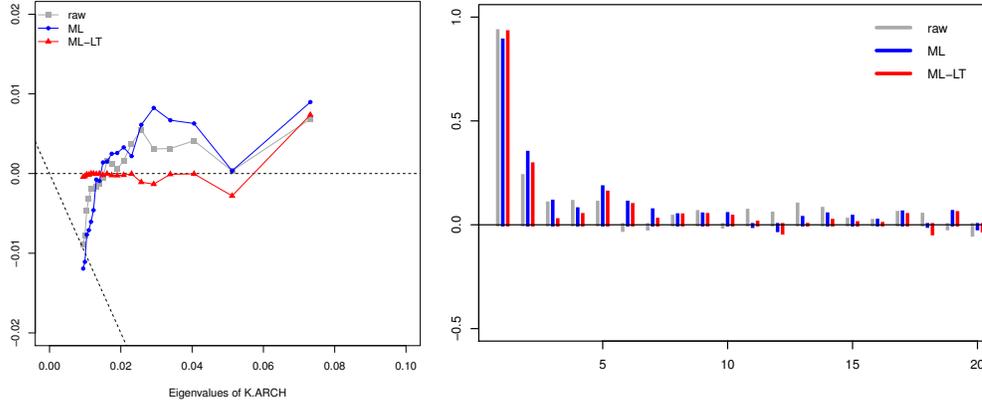


FIGURE 3. Spectral decomposition of the feedback kernel K , for the GMM, ML and ML+LT estimates. **Left:** The difference between the ranked eigenvalues of the estimated kernels and those of K_{ARCH} as a function of the latter (again ranked). The dashed oblique line has slope -1 and separates positive eigenvalues from negative ones. **Right:** Structure of the first eigenvector.

Fig. 2 shows the estimated quadratic kernel K . We find that the diagonal elements $K(\tau, \tau)$ are an order of magnitude larger than the corresponding off-diagonal elements $K(\tau, \tau' \neq \tau)$. This confirms that *daily returns* play a special role in the volatility feedback mechanism, as postulated by ARCH models. Returns on different time scales, while significant, can be treated as a perturbation of the main ARCH effect. The other surprise is that there does not seem to be any obvious structure of the matrix. This means that the fine structure of volatility feedback effects is much more subtle than anticipated.

An illuminating way of interpreting QARCH models is to work in the diagonal basis of the K matrix, where the quadratic term in (7) reads:

$$\sum_{\tau', \tau''=1}^q \left(\sum_n \lambda_n v_n(\tau') v_n(\tau'') \right) r_{t-\tau'} r_{t-\tau''} \equiv \sum_n \lambda_n \widehat{R}_n(t)^2 \tag{9}$$

with (λ_n, v_n) being, respectively, the n -th eigenvalue and eigenvector of K , and

$$\widehat{R}_n(t) = \langle r | v \rangle_t = \sum_{\tau=1}^q v(\tau) r_{t-\tau}, \tag{10}$$

which can be interpreted as an average return over the whole period, with a certain set of weights $v_n(\tau)$. One can therefore say that the square volatility σ_t^2 picks up contributions from various past returns eigenmodes. The modes associated to the largest eigenvalues λ_n are those which have the largest contribution to volatility spikes.

The ARCH(q) model corresponds to a diagonal matrix; in this case the modes are trivially individual daily returns. Another trivial case is when K is of rank one and its spectral decomposition is simply $K(\tau', \tau'') = \lambda v(\tau') v(\tau'')$, where $\lambda = \text{Tr}(K)$ is the only non-null eigenvalue, and $v(\tau) = \sqrt{K(\tau, \tau) / \text{Tr}(K)}$ is the eigenvector associated with this non-degenerate mode.

We show in Fig. 3 the eigenvalues of the estimated K as a function of the eigenvalues of the purely diagonal ARCH model. We see that a) the largest eigenvalue is clearly shifted upwards by the off-diagonal elements; b) the structure of the top eigenvector is non-trivial, and has positive contributions at all lags (up to noise)

indicating a genuine departure from the diagonal ARCH benchmark, for which we would find a single peak at $\tau = 1$; c) the unconstrained estimations — both GMM and ML — lead to 6 very small eigenvalues (perhaps even slightly negative for ML) that all constrained models fail to reproduce.

The positiveness of all eigenvalues was far from granted a priori, because nothing in the calibration procedure imposes the positivity of the matrix K . Although we naively expected that past excitations could only lead to an amplification of future volatility (i.e. that only strictly positive modes should appear in the feedback kernel), we observe that quasi-neutral modes do occur and appear to be significant. This is clearly related to the negative streaks noted above at large lags, but we have no intuitive interpretation for this effect at this stage.

3.3. Conclusion, extensions

We have revisited the QARCH model, which postulates that the volatility today can be expressed as a general quadratic form of the past daily returns r_t . The standard ARCH or GARCH framework is recovered when the quadratic kernel is diagonal, which means that only past squared daily returns feedback on today's volatility. This is a very restrictive *a priori* assumption, and the aim of the present study was to unveil the possible influence of other quadratic combinations of past returns, such as, for example, square weekly returns.

The calibration of these models on US stock returns has revealed several features that we did not anticipate. Quite surprisingly, we find that while the off-diagonal elements of $K(\tau, \tau')$ are significant, they are at least an order of magnitude smaller than the diagonal elements $k(\tau) := K(\tau, \tau)$. The latter are found to decay very slowly with τ , in agreement with previous discussions. Therefore, in a first approximation, the dominant feedback effect comes from the amplitude of *daily returns* only, with minor corrections coming from returns computed on large time spans, at variance with the assumption of the model put forward in [12]. We believe that this finding is unexpected and far from trivial. On the one hand, it is a strong constraint on any attempt to justify the ARCH feedback mechanism from a more fundamental point of view. On the other hand, important corrections to the pure (daily close-to-close) GARCH structure are also found to stem from the interplay of past open-to-close and close-to-open returns, as well as their distinct feedback in the future volatility [11].

4. LIMIT THEOREMS FOR NEARLY UNSTABLE HAWKES PROCESSES

T. JAISSON

4.1. Motivation for the theorem

Because of their tractability and their natural interpretations in term of market quantities, Hawkes processes, see [27, 28] are nowadays widely used in high-frequency finance, see [4, 6]. However, in practice, the statistical estimation results seem to show that very often, only *nearly unstable Hawkes processes* are able to fit the data properly, see [24, 26]. By nearly unstable, we mean that the L^1 norm of their kernel is close to unity. We study in this work such processes for which the stability condition is almost violated. Our main result states that after suitable rescaling, they asymptotically behave like integrated Cox Ingersoll Ross models. Thus, modeling the arrival of trades as nearly unstable Hawkes processes may be a good way to reproduce both their high and low-frequency stylized facts. Indeed, Hawkes processes present the same clustering as order flows at low frequency and our theorem shows that we retrieve some sort of persistence at low frequency.

4.2. Assumptions and asymptotic framework

We consider a sequence of Hawkes point processes $(N_t^T)_{t \geq 0}$ indexed by T . For a given T , (N_t^T) satisfies $N_0^T = 0$ and its intensity process (λ_t^T) is defined for $t \geq 0$ by

$$\lambda_t^T = \mu + \int_0^t \phi^T(t-s) dN_s^T,$$

where μ is a positive real number and ϕ^T a non negative measurable function on \mathbb{R}^+ which satisfies $\|\phi\|_1 < +\infty$.

Let us now give more specific assumptions on the asymptotic behavior of the function ϕ^T .

Assumption 1. For $t \in \mathbb{R}^+$,

$$\phi^T(t) = a_T \phi(t),$$

where $(a_T)_{T \geq 0}$ is a sequence of positive numbers converging to one such that for all T , $a_T < 1$ and ϕ is a non negative measurable function such that

$$\int_0^{+\infty} \phi(s) ds = 1 \text{ and } \int_0^{+\infty} s\phi(s) ds = m < \infty.$$

Moreover, ϕ is differentiable with derivative ϕ' such that $\|\phi'\|_\infty < +\infty$ and $\|\phi'\|_1 < +\infty$.

Thus, the form of the function ϕ^T depends on T so that its shape is fixed but its L^1 norm varies with T . For a given T , this L^1 norm is equal to a_T and so is smaller than one, implying that the stability condition is in force. Note that we do not work in the stationary setting since our process starts at time $t = 0$ and not at $t = -\infty$.

The case where $\|\phi^T\|_1$ is larger than one corresponds to the situation where the stability condition is violated. Since $a_T = \|\phi^T\|_1 < 1$ tends to one, our framework is a way to get close to instability. Therefore we call our processes nearly unstable Hawkes processes. There are of course many other ways to make the L^1 norm of ϕ^T converge to one than the multiplicative manner used here. However, this parametrization is sufficient for applications and very convenient to illustrate the different regimes that can be obtained.

4.3. Observation scales

In our framework, two parameters degenerate at infinity: T and $(1 - a_T)^{-1}$. The relationship between these two sequences will determine the scaling behavior of the sequence of Hawkes processes. Recall that it is shown in [3] that when $\|\phi\|_1$ is fixed and smaller than one, after appropriate scaling, the limit of the sequence of Hawkes processes is deterministic, as it is for example the case for Poisson processes, see [29]. In our setting, if $1 - a_T$ tends “slowly” to zero, we can expect the same result.

On the contrary, if $1 - a_T$ tends too rapidly to zero, the situation is likely to be quite intricate. Indeed, for given T , the Hawkes process may already be very close to instability whereas T is not large enough to reach the asymptotic regime. The last case, which is probably the most interesting one, is the intermediate case, where $1 - a_T$ tends to zero in such a manner that a non deterministic scaling limit is obtained, while not being in the preceding degenerate setting. We largely detail this situation in the next subsection.

4.4. Non degenerate scaling limit for nearly unstable Hawkes processes

Before stating our theorem, we wish to provide intuitions on how it is derived. Let M^T be the martingale process associated to N^T , that is, for $t \geq 0$,

$$M_t^T = N_t^T - \int_0^t \lambda_s^T ds.$$

We also set ψ^T the function defined on \mathbb{R}^+ by

$$\psi^T(t) = \sum_{k=1}^{\infty} (\phi^T)^{*k}(t),$$

where $(\phi^T)^{*1} = \phi^T$ and for $k \geq 2$, $(\phi^T)^{*k}$ denotes the convolution product of $(\phi^T)^{*(k-1)}$ with the function ϕ^T . Note that $\psi^T(t)$ is well defined since $\|\phi^T\|_1 < 1$. In the sequel, it will be convenient to work with another form for the intensity. We have the following result.

$$\lambda_t^T = \mu + \int_0^t \psi^T(t-s)\mu ds + \int_0^t \psi^T(t-s)dM_s^T.$$

Now recall that we observe the process (N_t^T) on $[0, T]$. In order to be able to give a proper limit theorem, where the processes live on the same time interval, we rescale our processes so that they are defined on $[0, 1]$. To do that, we consider for $t \in [0, 1]$

$$\lambda_{tT}^T = \mu + \int_0^{tT} \psi^T(Tt-s)\mu ds + \int_0^{tT} \psi^T(Tt-s)dM_s^T.$$

For the scaling in space, a natural multiplicative factor is $(1 - a_T)$. Indeed, in the stationary case, the expectation of λ_t^T is $\mu/(1 - \|\phi^T\|_1)$. Thus, the order of magnitude of the intensity is $(1 - a_T)^{-1}$. This is why we define

$$C_t^T = \lambda_{tT}^T(1 - a_T). \quad (11)$$

Understanding the asymptotic behavior of C_t^T will be the key to the derivation of a suitable scaling limit for our sequence of normalized processes. We will see that this behavior is closely connected to that of the function ψ^T . About ψ^T , one can first remark that the function defined for $x \geq 0$ by

$$\rho^T(x) = T \frac{\psi^T}{\|\psi^T\|_1}(Tx) \quad (12)$$

is the density of the random variable

$$X^T = \frac{1}{T} \sum_{i=1}^{I^T} X_i,$$

where the (X_i) are iid random variables with density ϕ and I^T is a geometric random variable with parameter $1 - a_T$ ($\forall k > 0, \mathbb{P}[I^T = k] = (1 - a_T)(a_T)^{k-1}$). Now let $z \in \mathbb{R}$. The characteristic function of the random variable X^T , denoted by $\hat{\rho}^T$, satisfies

$$\begin{aligned} \hat{\rho}^T(z) &= \mathbb{E}[e^{izX^T}] = \sum_{k=1}^{\infty} (1 - a_T)(a_T)^{k-1} \mathbb{E}[e^{i\frac{z}{T} \sum_{i=1}^k X_i}] \\ &= \sum_{k=1}^{\infty} (1 - a_T)(a_T)^{k-1} (\hat{\phi}(\frac{z}{T}))^k = \frac{\hat{\phi}(\frac{z}{T})}{1 - \frac{a_T}{1-a_T}(\hat{\phi}(\frac{z}{T}) - 1)}, \end{aligned}$$

where $\hat{\phi}$ denotes the characteristic function of X_1 . Since

$$\int_0^{+\infty} s\phi(s)ds = m < \infty,$$

the function $\hat{\phi}$ is continuously differentiable with first derivative at point zero equal to im . Therefore, using that a_T and $\hat{\phi}(\frac{z}{T})$ both tend to one as T goes to infinity, $\hat{\rho}^T(z)$ is equivalent to

$$\frac{1}{1 - \frac{izm}{T(1-a_T)}}.$$

Thus, we precisely see here that the suitable regime so that we get a non trivial limiting law for X^T is that there exists $\lambda > 0$ such that

$$T(1 - a_T) \xrightarrow{T \rightarrow +\infty} \lambda. \tag{13}$$

When (13) holds, we write $d_0 = m/\lambda$. In fact we have just proved the following result.

Proposition 4.1. *Assume that (13) holds. Under Assumption 1, the sequence of random variables X^T converges in law towards an exponential random variable with parameter $1/d_0$.*

This simple result is of course not new. For example this type of geometric sums of random variable is studied in detail in [31]. Note also that when X_1 is exponentially distributed, X^T is also exponentially distributed, even for a fixed T .

Assume from now on that (13) holds and set $u_T = T(1 - a_T)/\lambda$ (so that u_T goes to one). Proposition 4.1 is particularly important since it gives us the asymptotic behavior of ψ^T in this setting. Indeed, it tells us that

$$\psi^T(Tx) = \rho^T(x) \frac{a_T}{\lambda u_T} \approx \frac{\lambda}{m} e^{-x \frac{\lambda}{m}} \frac{1}{\lambda} = \frac{1}{m} e^{-x \frac{\lambda}{m}}.$$

Let us now come back to the process C_t^T , which can be written

$$C_t^T = (1 - a_T)\mu + \mu \int_0^t u_T \lambda \psi^T(Ts) ds + \int_0^t \sqrt{\lambda} \psi^T(T(t - s)) \sqrt{C_s^T} dB_s^T, \tag{14}$$

with

$$B_t^T = \frac{1}{\sqrt{T}} \sqrt{u_T} \int_0^{tT} \frac{dM_s^T}{\sqrt{\lambda s^T}}. \tag{15}$$

By studying its quadratic variation, we can show that B^T represents a sequence of martingales which converges to a Brownian motion. So, heuristically replacing B^T by a Brownian motion B and $\psi^T(Tx)$ by $\frac{1}{m} e^{-x \frac{\lambda}{m}}$ in (14), we get

$$C_t^\infty = \mu(1 - e^{-t \frac{\lambda}{m}}) + \frac{\sqrt{\lambda}}{m} \int_0^t e^{-(t-s) \frac{\lambda}{m}} \sqrt{C_s^\infty} dB_s.$$

Applying Itô's formula, this gives

$$C_t^\infty = \int_0^t (\mu - C_s^\infty) \frac{\lambda}{m} ds + \frac{\sqrt{\lambda}}{m} \int_0^t \sqrt{C_s^\infty} dB_s,$$

which precisely corresponds to the stochastic differential equation (SDE) satisfied by a CIR process.

We show in [29] that the preceding heuristic can be made rigorous:

Theorem 4.1. *Assume that (13) holds. Under Assumption 1, the sequence of renormalized Hawkes intensities (C_t^T) defined in (11) converges in law, for the Skorohod topology, towards the law of the unique strong solution of the following Cox Ingersoll Ross stochastic differential equation on $[0, 1]$:*

$$X_t = \int_0^t (\mu - X_s) \frac{\lambda}{m} ds + \frac{\sqrt{\lambda}}{m} \int_0^t \sqrt{X_s} dB_s.$$

Furthermore, the sequence of renormalized Hawkes process

$$V_t^T = \frac{1 - a_T}{T} N_{tT}^T$$

converges in law, for the Skorohod topology, towards the process

$$\int_0^t X_s ds, \quad t \in [0, 1].$$

Theorem 4.1 implies that when $\|\phi\|_1$ is close to 1, if the observation time T is suitably chosen (that is of order $1/(1 - \|\phi\|_1)$), a non degenerate behavior (neither explosive, nor deterministic) can be obtained for a rescaled Hawkes process.

5. MARKET IMPACT

I. MASTROMATTEO

Market impact is at the heart of the mechanism through which financial markets process the information contained in an incoming flux of orders, converting it into an efficient price signal [13, 14]. Its relevance is not only theoretical, as for practitioners it represents a cost which needs to be cleverly computed in order to minimize execution costs, justifying the strong interest in this field emerged in the financial industry [1]. Market impact is customarily measured as the average difference among the price prior to the execution of a volume Q of contracts and the one subsequent to the execution. Strong empirical evidence accumulated in recent years shows that it is well-described by a concave function, typically parametrized as

$$\mathcal{I}(Q) = Y \sigma_D \left(\frac{Q}{V_D} \right)^\delta, \quad (16)$$

where σ_D and V_D represent respectively the daily volatility and the daily traded volume, Y is an adimensional factor empirically found to be of order 1 and δ is an exponent estimated in the range $0.4 - 0.7$ [2, 9, 16, 23, 25, 37, 40, 45, 47, 48]. Above formula is quite remarkable from several points of view. First, a concave dependence signals non-additivity of the price response: when executing incrementally a trade of volume $2Q$, the response to the second half of trades is just a fraction of roughly 40% of the response to the first bundle of contracts executed, indicating a non-trivial time dependence of the response. Secondly, it indicates that markets are inherently *fragile*: according to Eq. (16) vanishingly small executed volumes have a disproportionate impact on prices. These findings are in sharp contrast with the traditional economic paradigm according to which traded volumes should impact *linearly* the price [32]. Unexpectedly, these results have been confirmed in a variety of settings: the concavity of impact seems to have little to do with the historical period, the market venue, the maturity and the type of the traded contract. Consequently, the generality of Eq. (16) requires a justification more general than the one provided by ad-hoc models, as it seems to indicate that a more universal mechanism is at play in determining the response of price to traded volumes. In this section, we would like to show how *latent liquidity models* can account for these general findings, providing an implementation-free framework in which the concave shape of the impact function is a consequence of generic assumptions, holding in most type of modern financial markets.

5.1. The latent order book

The notions upon which we would like to build our framework are the one of *dimensionality* (i.e., the fact that prices evolve on a one-dimensional grid) and the one of *market clearing* (buyers and sellers trigger a trade when they agree on a transaction price). We will see that they are sufficient in order to induce a non-trivial shape of the demand and supply curve, characterized by a *V*-shape around the current price that will ultimately justify the concave response of prices to traded volumes. One might naively think that the so-called *order book* provides an empirical proxy for the fundamental demand and supply curve of the market (an *order book* arranges on a regular price grid all the buy and sell orders submitted, and not yet executed, by the market participants). Simple order of magnitude estimations reveal that this intuition is wrong: most of the liquidity

in modern markets is *hidden*, and reveals itself progressively as the trading day goes by. Empirically, one finds that the ratio among the daily traded volume and the one instantaneously available on the order book is a factor of about one hundred [14], suggesting that the notion of liquidity must be captured by a slower and more coarse-grained concept. This is precisely the idea behind latent liquidity models: in this framework, one would like to capture the behavior of market participants at intermediate frequency (say, from a few minutes to some hours) by defining the notion of a slow-evolving *latent order book* in which not only the visible orders are displayed, but also the ones which would reveal themselves if the current price was about to hit their respective price levels.

5.2. A stylized model for the latent order book

The simplest description of the latent order book, proposed in [38], is provided by a model in which particles (*orders*) of two types B and A (respectively *bid*, for buy orders, and *ask*, for sell orders) sit on two opposite sides of a discrete price grid of length L , whose coordinates will be indexed by x , similar to the setting considered in [7, 46]. The mid-point among the rightmost bid-particle and the left-most ask particle will be defined as the price, and will be denoted by x_t . We will suppose that each order can jump either left or right with a Poissonian rate D , and that orders sitting at the same site might undergo a reaction at Poissonian rate λ . The outcome of such reaction is fixed by the following random process depending upon two parameters p and m :

$$A + B \rightarrow \emptyset \quad \text{w. prob.} \quad 1 - p \tag{17}$$

$$A + B \rightarrow B \quad \text{w. prob.} \quad p \frac{1 + m}{2} \tag{18}$$

$$A + B \rightarrow A \quad \text{w. prob.} \quad p \frac{1 - m}{2} . \tag{19}$$

While the case $p = m = 0$ corresponds to the case of ordinary annihilation $A + B \rightarrow \emptyset$, the case $p \neq 0, m \neq 0$ models the presence of a trader which is biasing the flux of particles either left (for $m < 0$) or right (for $m > 0$) by submitting extra orders exactly at the current price. We will ultimately work in the limit $\lambda \rightarrow \infty$ which ensures that the market clearing condition is fulfilled (i.e., that buy and sell orders cannot coexist at the same price level). The presence of a finite number of particles in the system is then ensured by the presence of an external current, adding particles bid particles at the boundary $x = 0$ and ask particles at the boundary $x = L$ at a Poissonian rate J . In continuous-space approximation, such a system is described by the dynamics

$$\frac{\partial \langle b(x, t) \rangle}{\partial t} = D \frac{\partial^2 \langle b(x, t) \rangle}{\partial x^2} - \lambda u_A \langle a(x, t) b(x, t) \rangle \tag{20}$$

$$\frac{\partial \langle a(x, t) \rangle}{\partial t} = D \frac{\partial^2 \langle a(x, t) \rangle}{\partial x^2} - \lambda u_B \langle a(x, t) b(x, t) \rangle , \tag{21}$$

where $a(x, t)$ and $b(x, t)$ are the densities of particles of type A and B, $u_A = 1 - p(\frac{1+m}{2})$ and $u_B = 1 - p(\frac{1-m}{2})$. Even in the symmetric case $u_A = u_B$, the equation above is hard to solve due to the presence of the reaction term containing the product $\langle a(x, t) b(x, t) \rangle$ [8, 18], indeed one can take advantage of the conservation law for the difference of A and B particles by defining the field $\varphi(x, t) = b(x, t) - a(x, t)$. Such field satisfies the equation $\partial_t \varphi = D \partial_{xx}^2 \varphi$, and it is subject to the boundary condition $-D \partial_x \varphi|_{x=0,L} = J$. It can be used in order to retrieve the price by using the relation $\langle \varphi(\langle x_t \rangle, t) \rangle = 0$. In particular, for $u_A = u_B$ the stationary state of the system $\varphi_{st}(x)$ is described by the density profile

$$\varphi_{st}(x) = -(J/D)(x - L/2) , \tag{22}$$

while the stationary price is obviously given by $x_{st} = L/2$. This, in the limit $\lambda \rightarrow \infty$, indicates a V-shaped profile for the original densities of A and B particles around the stationary value of the price, as expected in any latent order book model due to the market clearing condition.

Addressing the issue of market impact requires then analyzing the effect of a perturbation induced on a symmetric, stationary system by a sudden change of $u_A \neq u_B$ occurring at $t = 0$. In order to solve the problem for the price after the perturbation, we can define the field $\psi = u_B b - u_A a$, which again follows the diffusion equation

$$\frac{\partial \langle \psi(x, t) \rangle}{\partial t} = D \frac{\partial^2 \langle \psi(x, t) \rangle}{\partial x^2}, \quad (23)$$

subject to boundary conditions

$$J u_B = -D \frac{\partial \langle \psi(x, t) \rangle}{\partial x} \Big|_{x=0} \quad J u_A = -D \frac{\partial \langle \psi(x, t) \rangle}{\partial x} \Big|_{x=L}. \quad (24)$$

In the limit $\lambda \rightarrow \infty$, the zeros of the field $\psi(x, t)$ coincide with the ones of $\varphi(x, t)$, implying that it is possible to solve the problem for the average price by simply solving Eq. (23).

After some work (see [38]), one finds that the $L \rightarrow \infty$ limit of the price is given by

$$\langle x_t - x_0 \rangle = 2\alpha(u_B/u_A)\sqrt{Dt}, \quad (25)$$

where $\alpha(z)$ is defined by the transcendental equation

$$\alpha(z) \left(\frac{z+1}{z-1} - \operatorname{erf}[\alpha(z)] \right) - \frac{1}{\sqrt{\pi}} e^{-\alpha^2(z)}. \quad (26)$$

This implies that in this limit the evolution of the price is independent of L , and that it follows *exactly* a square-root law. The relation among the impact and the volume executed by the trader can be retrieved by defining the imbalance Q as the extra volume pressure induced by the $m \neq 0$ term. One gets finally

$$\mathcal{I}(Q) = \left(\frac{2\alpha(u_B/u_A)}{\beta^{1/2}(u_B/u_A)} \right) \left(\frac{\langle Q \rangle D}{J} \right)^{1/2} \quad (27)$$

where we used that $\langle Q \rangle = \beta(u_B/u_A)(JT)$ together with

$$\beta(z) = \frac{1}{2z} [(z^2 - 1) - \operatorname{erf}[\alpha(z)](z - 1)^2]. \quad (28)$$

Eq. (27) exactly mirrors the empirical relation (16) with $\delta = 1/2$, $\sigma_D V_D^{-1/2} = (D/J)^{1/2}$ and $Y = 2\alpha\beta^{-1/2}$, thus providing a microscopic interpretation for such a phenomenological formula.

5.3. Conclusions

The simplified setup that we have proposed is able to account for the empirical observation of concave impact, and predicts a square-root shape for the impact function consistently with results previously presented in the literature. These predictions depend on very generic assumptions, namely the existence of a slowly evolving latent order book implementing a one-dimensional notion of price and enforcing the market clearing condition. More sophisticated versions of this model (such as [20, 37, 48]) share its same fundamental properties, and lead to analogous results. This supports the idea that the mechanism that we propose is a universal one, and is thus appropriate in order to account for the stylized facts determining the nature of market impact.

REFERENCES

- [1] R. Almgren and N. Chriss. Optimal execution of portfolio transactions. *Journal of Risk*, 3:5–40, 2001.
- [2] R. Almgren, C. Thum, E. Hauptmann, and H. Li. Direct estimation of equity market impact. *Risk*, 18:5752, 2005.

- [3] E. Bacry, S. Delattre, M. Hoffmann, and J. F. Muzy. Scaling limits for Hawkes processes and application to financial statistics. *arXiv preprint arXiv:1202.0842*, 2012.
- [4] E. Bacry, S. Delattre, M. Hoffmann, and J.-F. Muzy. Modelling microstructure noise with mutually exciting point processes. *Quantitative Finance*, 13(1):65–77, 2013.
- [5] E. Bacry, J. Delour, and J.-F. Muzy. Multifractal random walks. *Physical Review E*, 64:026103–026106, 2001.
- [6] E. Bacry and J.-F. Muzy. Hawkes model for price and trades high-frequency dynamics. *arXiv preprint arXiv:1301.1135*, 2013.
- [7] P. Bak, M. Paczuski, and M. Shubik. Price variations in a stock market with many agents. *Physica A: Statistical Mechanics and its Applications*, 246(3):430–453, 1997.
- [8] E. Ben-Naim and S. Redner. Inhomogeneous two-species annihilation in the steady state. *Journal of Physics A: Mathematical and General*, 25(9):L575, 1992.
- [9] N. Bershova and D. Rakhlin. The non-linear market impact of large trades: Evidence from buy-side order flow. *Quantitative Finance*, 13(11):1759–1778, 2013.
- [10] F. Black and M. Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654, 1973.
- [11] P. Blanc, R. Chicheportiche, and J.-P. Bouchaud. The fine structure of volatility feedback II: Overnight and intra-day effects. *Physica A: Statistical Mechanics and its Applications*, 402:58 – 75, 2014.
- [12] L. Borland and J.-P. Bouchaud. On a multi-timescale statistical feedback model for volatility fluctuations. *The Journal of Investment Strategies*, 1(1):65–104, December 2011.
- [13] J.-P. Bouchaud. Price impact. *Encyclopedia of quantitative finance*, 2010.
- [14] J.-P. Bouchaud, J. D. Farmer, and F. Lillo. How markets slowly digest changes in supply and demand. *Handbook of financial markets: dynamics and evolution*, 1:57, 2009.
- [15] J.-P. Bouchaud, A. Matacz, and M. Potters. Leverage effect in financial markets: The retarded volatility model. *Physical Review Letters*, 87(22):228701, 2001.
- [16] X. Brokmann, J. Kockelkoren, J.-P. Bouchaud, et al. Slow decay of impact in equity markets. *Available at SSRN 2471528*, 2014.
- [17] R. Chicheportiche and J.-P. Bouchaud. The fine-structure of volatility feedback I: Multi-scale self-reflexivity. *Physica A: Statistical Mechanics and its Applications*, 410:174 – 195, 2014.
- [18] S. Cornell and M. Droz. Steady-state reaction-diffusion front scaling for $ma+nb \rightarrow [\text{inert}]$. *Physical review letters*, 70(24):3824, 1993.
- [19] D. M. Cutler, J. M. Poterba, and L. H. Summers. What moves stock prices? *The Journal of Portfolio Management*, 15(3):4–12, 1989.
- [20] J. Donier, J. F. Bonart, I. Mastromatteo, and J.-P. Bouchaud. A fully consistent, minimal model for non-linear market impact. *Minimal Model for Non-Linear Market Impact (November 29, 2014)*, 2014.
- [21] B. Dupire. Theory of rational option pricing. *RISK magazine*, JULY:126–129, 1994.
- [22] R. F. Engle. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica: Journal of the Econometric Society*, pages 987–1007, 1982.
- [23] J. D. Farmer, A. Gerig, F. Lillo, and H. Waelbroeck. How efficiency shapes market impact. *Quantitative Finance*, 13(11):1743–1758, 2013.
- [24] V. Filimonov and D. Sornette. Quantifying reflexivity in financial markets: Toward a prediction of flash crashes. *Physical Review E*, 85(5):056108, 2012.
- [25] C. Gomes and H. Waelbroeck. Is market impact a measure of the information value of trades? market response to liquidity vs. informed metaorders. *Quantitative Finance*, (ahead-of-print):1–21, 2014.
- [26] S. J. Hardiman, N. Bergot, and J.-P. Bouchaud. Critical reflexivity in financial markets: a Hawkes process analysis. *arXiv preprint arXiv:1302.1405*, 2013.
- [27] A. G. Hawkes. Point spectra of some mutually exciting point processes. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 438–443, 1971.
- [28] A. G. Hawkes. Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, 58(1):83–90, 1971.
- [29] T. Jaisson and M. Rosenbaum. Limit theorems for nearly unstable Hawkes processes. *arXiv preprint arXiv:1310.2033*, 2013.
- [30] J.-P. Kahane. Sur le chaos multiplicatif. *Annales scientifiques mathématiques du Québec*, 9:105–150, 1985.
- [31] V. V. Kalashnikov. *Geometric sums: bounds for rare events with applications: risk analysis, reliability, queueing*, volume 413. Springer, 1997.
- [32] A. S. Kyle. Continuous auctions and insider trading. *Econometrica: Journal of the Econometric Society*, pages 1315–1335, 1985.
- [33] P. E. Lynch and G. O. Zumbach. Market heterogeneities and the causal structure of volatility. *Quantitative Finance*, 3(4):320–331, 2003.
- [34] B. Mandelbrot. The variation of certain speculative prices. *The Journal of Business*, 36(4):394–419, 1963.
- [35] B. Mandelbrot and H. M. Taylor. On the distribution of stock price differences. *Operations Research*, 15(6):1057–1062, 1967.
- [36] B. Mandelbrot and H. M. Taylor. On the distribution of stock price differences. *Mathematical Finance*, 22(1):83–108, 2012.
- [37] I. Mastromatteo, B. Toth, and J.-P. Bouchaud. Agent-based models for latent liquidity and concave price impact. *Physical Review E*, 89(4):042805, 2014.

- [38] I. Mastromatteo, B. Tóth, and J.-P. Bouchaud. Anomalous impact in reaction-diffusion financial models. *Physical Review Letters*, 113(26):268701, 2014.
- [39] R. C. Merton. Theory of rational option pricing. *Bell Journal of Economics and Management Science*, 4(1):141–183, 1973.
- [40] E. Moro, J. Vicente, L. G. Moyano, A. Gerig, J. D. Farmer, G. Vaglica, F. Lillo, and R. N. Mantegna. Market impact and trading profile of hidden orders in stock markets. *Physical Review E*, 80(6):066102, 2009.
- [41] U. A. Müller, M. M. Dacorogna, R. D. Davé, R. B. Olsen, O. V. Pictet, and J. E. von Weizsäcker. Volatilities of different time resolutions — analyzing the dynamics of market components. *Journal of Empirical Finance*, 4(2):213–239, 1997.
- [42] L. C. G. Rogers and S. E. Satchell. Estimating variance from high, low and closing prices. *The Annals of Applied Probability*, pages 504–512, 1991.
- [43] E. Sentana. Quadratic ARCH models. *The Review of Economic Studies*, 62(4):639, 1995.
- [44] R. J. Shiller. Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review*, 71(3):421–436, 1981.
- [45] I. Skachkov. Market impact paradoxes. *Wilmott*, 2014(70):71–78, 2014.
- [46] L.-H. Tang and G.-S. Tian. Reaction–diffusion–branching models of stock price fluctuations. *Physica A: Statistical Mechanics and its Applications*, 264(3):543–550, 1999.
- [47] N. Torre. Barra market impact model handbook. *BARRA Inc., Berkeley*, 1997.
- [48] B. Toth, Y. Lempriere, C. Deremble, J. De Lataillade, J. Kockelkoren, and J.-P. Bouchaud. Anomalous price impact and the critical nature of liquidity in financial markets. *Physical Review X*, 1(2):021006, 2011.