The EWMA Heston model

Léo Parent

The EWMA Heston model

Léo Parent
PRISM Sorbonne
Paris 1 Panthéon-Sorbonne University
leo.parent@etu.univ-paris1.fr

October 5, 2021

Abstract

This paper introduces the exponentially weighted moving average (EWMA) Heston model, a Markovian stochastic volatility model able to capture a wide range of empirical features related to volatility dynamics while being more tractable for simulations than rough volatility models based on fractional processes. After presenting the model and its principal characteristics, our analysis focuses on the use of its associated Euler-discretization scheme as a time-series generator for Monte-Carlo simulations. Using this discretization scheme, and on the basis of S&P500 empirical time series, we show that the EWMA Heston model is overall consistent with market data, making it a credible alternative to other existing stochastic volatility models.

Keywords: Stochastic volatility model, Heston model, quadratic rough Heston model, Zumbach effect, time-reversal asymmetry, volatility distribution, returns distribution.

JEL classification: C15, C32, G10, G14, G17.

1 Introduction

The modeling of asset price dynamics is one of the most important issues in quantitative finance and has resulted in a large body of academic research. It has for long been established that the modeling quality of a price dynamics model depends to a large extent, on its ability to accurately reproduce the empirical features of volatility dynamics. Different approaches have been applied to take account of this.

The ARCH (or autoregressive conditional heteroskedasticity) family of models is one of the most important approaches and includes a wide range of variants able to describe the stylized facts of financial time series. For instance, the famous generalized ARCH or GARCH model (Engle 1982) in which the variance process depends on an exponential moving average of past squared returns provides a good representation of many of the empirical features of financial data, such as tail heaviness, volatility clustering, and feedback effect. More sophisticated extensions such as the EGARCH (Nelson 1991, Bandt and Jones 2006), NGARCH (Lanne and Saikkonen 2005), FIGARCH (Baillie *et al.* 1996, Belkhouja and Boutahary 2011) and QGARCH (Borland and Bouchaud 2005) models, further improved modeling quality by reproducing other features of financial time series including the leverage effect, time-reversal asymmetry, and the feedback effect.

Although some of these models have continuous-time counterparts this family of models in its canonical form adopts a discrete-time approach. However, some financial issues require a continuous-time framework which led to the parallel development of other continuous-time stochastic volatility models.

Among these models, in recent years rough volatility models have emerged as the new standard. These models allow consideration of the empirical roughness of the volatility, and from a more practical point

of view allow remarkable reproduction of the shape of the implied volatility surface. However, although rough volatility models are extremely efficient for capturing empirical phenomena that occur in derivatives markets, their use as a tool for Monte Carlo experiments, and especially in a risk or asset management context, has some limitations. First, these models are founded on non-Markovian and non-semimartingale processes, which makes unbiased simulations based on such models a tricky task (El Euch 2018). In addition to the strictly technical issues, some empirical features of the realized volatility are not entirely captured by this type of model (Blanc *et al.* 2017, Gatheral *et al.* 2020).

The present paper introduces a new stochastic volatility model - the EWMA Heston model (HM) - which aims to address some aspects of these different issues. The aim is to propose a model able to reproduce a broad spectrum of empirical features related to volatility dynamics while being more tractable for simulations than rough volatility models based on fractional stochastic processes. We show that, while being Markovian, the EWMA HM can capture the different dimensions of the Zumbach effect and accurately reproduce the joint dynamics of the asset returns and volatility, as well as the empirical returns and volatility distributions for different time horizons. Unlike most of the academic literature on stochastic volatility models, the present paper deals not with the model's application for derivative pricing issues but focuses instead on the model's use as a tool for Monte Carlo experiments.

The paper is organized as follows. First, we present the EWMA HM and its principal characteristics. Second, we examine a limit-case of the model with interesting properties. Third, we compare the EWMA HM with stochastic volatility models based on quadratic Hawkes processes (Blanc *et al.* 2017, Dandapani *et al.* 2019, Gatheral *et al.* 2020), to highlight their similarities and their differences. We then focus on use of its associated Euler-discretization scheme as a time-series generator for Monte-Carlo simulations. Using this discretization, and based on S&P500 empirical time series, we show that overall the EWMA HM is consistent with market data, making it a credible alternative to other stochastic volatility models.

2 The EWMA Heston model

2.1 Presentation of the model

Let us introduce the EWMA HM defined by the following stochastic differential equations (SDEs) system:

$$\begin{cases} \dfrac{\mathrm{d}S_t}{S_t} &= \mu_t \mathrm{d}t + \sqrt{V_t}\mathrm{d}W_t \\ \mathrm{d}V_t &= \dfrac{1}{\tau_1}\Big(\nu(m_t)^2 - V_t\Big)\mathrm{d}t + \xi\nu(m_t)\sqrt{V_t}\mathrm{d}B_t \\ \mathrm{d}m_t &= \Lambda\mathbf{1}\cdot\left(\dfrac{\mathrm{d}S_t}{S_t} - \eta\mu_t\mathrm{d}t\right) - \Lambda m_t\mathrm{d}t \end{cases} \tag{1}$$

where $\{W\}_{t\in T}$ and $\{B\}_{t\in T}$ are Brownian motions such as $\langle dW, dB\rangle_t = \rho \in [-1:1]$, $\tau_1, \xi \in \mathbb{R}_+$, $\eta \in \{0,1\}$, $\mathbf{1}$ a $d\times 1$ vector of ones, $m_t \in \mathbb{R}^d$, Λ a $d\times d$ diagonal matrix as $\forall j \in \{1,...,d\}, \Lambda_{j,j} \in \mathbb{R}_+$, and with $\nu : \mathbb{R}^d \to \mathbb{R}_+$.

The specificity of this extension of the HM lies in the function ν. Whereas in the standard HM ν^2 is a constant and corresponds to the "long variance" of price, here it is an attraction variance function of the d-dimensional random vector m_t solution of the third SDE of (1)

$$m_t = m_0 e^{-\Lambda t} + \int_0^t e^{-\Lambda(t-u)}\Lambda\left(\frac{\mathrm{d}S_u}{S_u} - \eta\mu_u\mathrm{d}u\right).$$

Therefore, each coordinate of m_t corresponds to an exponential weighted moving average (EWMA) estimator of the price trend adjusted (i.e. $\eta = 1$) or not (i.e. $\eta = 0$) by its deterministic component. Due to the

dependence between m_t and the attraction volatility ν_t, the model is able to capture the leverage effect, even if $\{W\}_{t\in T}$ and $\{B\}_{t\in T}$ are uncorrelated. More importantly, this property makes the EWMA HM structurally adapted to capturing the "strong Zumbach effect" which corresponds to the fact that "conditional dynamics of volatility with respect to the past depend not only on the past volatility trajectory but also on the historical price path" (Gatheral *et al.* 2020 p.3). The properties of the EWMA HM means that it is included in the family of path-dependent volatility models (Guyon 2014). Note also that the genericness of the EWMA HM means that it is able to capture a very wide range of path dependencies through attraction volatility. For instance, exploiting the results in Bochud and Challet (2007) or Abi Jaber (2019), makes it possible to use a specification of Λ and $\nu(.)$ which enables a memory that decays as a power law to be mimicked.

After setting out this general framework, we now focus on an EWMA HM where $\eta = 0$ and m_t depends on a unique EWMA parameter (i.e. m_t is a scalar), defined by the following SDE system:

$$\begin{cases} \dfrac{\mathrm{d}S_t}{S_t} &= \mu_t \mathrm{d}t + \sqrt{V_t}\mathrm{d}W_t \\ \mathrm{d}V_t &= \dfrac{1}{\tau_1}\big({\nu_t}^2 - V_t\big)\mathrm{d}t + \xi\nu_t\sqrt{V_t}\mathrm{d}B_t \\ \mathrm{d}\nu_t &= \left(\dfrac{\psi}{\nu_t - \underline{\nu}} - \nu_t + \underline{\nu} + \alpha\right)\dfrac{\mathrm{d}t}{\tau_2} - \dfrac{\beta}{\tau_2}\dfrac{\mathrm{d}S_t}{S_t} \end{cases} \tag{2}$$

where $V_0, \nu_0, \psi, \tau_1, \underline{\nu}, \beta \in \mathbb{R}_+$ and $\alpha \in \mathbb{R}$. We assume also that $\forall\, t \in T$ the Feller condition is respected[1]:

$$\begin{aligned} \frac{2{\nu_t}^2}{\tau_1} &> \Big(\xi\nu_t\Big)^2 \\ \frac{2}{\tau_1} &> \xi^2. \end{aligned}$$

At first sight, because the attraction volatility is described here by its dynamic through an SDE, membership of this model in the family of EWMA HMs defined above is not obvious. To highlight this and fully understand the concrete effects of the historical price path on the volatility in this EWMA HM, we need to consider the SDE describing attraction volatility. First, notice that if $\beta = 0$, given $\nu_0 > \underline{\nu}$, the dynamic of the attraction volatility is described by an ordinary differential equation (ODE) that converges as follows[2]:

$$\lim_{t\to+\infty} \nu^\star = \frac{\underline{\nu} + \alpha + \sqrt{\alpha^2 + 4\psi}}{2}.$$

Therefore, in this specific case (where $\beta = 0$), the EWMA HM converges toward a standard HM. To consider the dynamics of ν_t in the general case, we set $f\big(\nu_t, t\big) = \nu_t e^{\frac{t}{\tau_2}}$ and apply the Itô lemma:

$$\begin{aligned} \mathrm{d}f\big(\nu_t, t\big) &= \frac{e^{\frac{t}{\tau_2}}}{\tau_2}\nu_t\mathrm{d}t + e^{\frac{t}{\tau_2}}\mathrm{d}\nu_t \\ &= \frac{e^{\frac{t}{\tau_2}}}{\tau_2}\left(\frac{\psi}{\nu_t - \underline{\nu}} + \underline{\nu} + \alpha\right)\mathrm{d}t - \frac{\beta e^{\frac{t}{\tau_2}}}{\tau_2}\frac{\mathrm{d}S_t}{S_t}. \end{aligned}$$

Thus, the solution of the SDE is

$$\nu_t = \nu_0 e^{-\frac{t}{\tau_2}} + \frac{1}{\tau_2}\int_0^t e^{\frac{1}{\tau_2}(u-t)}\left(\frac{\psi}{\nu_u - \underline{\nu}} + \underline{\nu} + \alpha\right)\mathrm{d}u - \frac{\beta}{\tau_2}\int_0^t e^{\frac{1}{\tau_2}(u-t)}\frac{\mathrm{d}S_u}{S_u}.$$

Therefore, we have:

$$\lim_{t\to+\infty} \nu_t = \alpha + \underline{\nu} + \frac{1}{\tau_2}\int_0^t e^{\frac{1}{\tau_2}(u-t)}\frac{\psi}{\nu_u - \underline{\nu}}\mathrm{d}u - \beta\underbrace{\frac{1}{\tau_2}\int_0^t e^{\frac{1}{\tau_2}(u-t)}\frac{\mathrm{d}S_u}{S_u}}_{m_t}.$$

[1] The fact that the Feller condition $\forall\, t \in T$ is respected implies that v stays strictly positive (as in the standard Cox-Ingersoll-Ross (CIR) processes) but has yet to be demonstrated.

[2] See appendix A.

This equation further explains the relationship between the historical price path and attraction volatility. First, we can see that this relationship goes through m_t an EWMA of the past returns of $\{S\}_{t\in T}$ clarifying membership of the model of (1). This EWMA is the solution of the following SDE when $m_0 = 0$

$$\mathrm{d}m_t = \frac{1}{\tau_2}\left(\frac{\mathrm{d}S_t}{S_t} - m_t\mathrm{d}t\right).$$

Further, from a statistical point of view, this process has a theoretical foundation since it can be considered the Kalman filter estimator of an unobservable trend (Harvey 1990, Bruder and Gaussel 2011, Jusselin *et al.* 2017). The sensitivity of the attraction volatility to this trend estimator is defined by the parameter β. Since β is positive, the model is able to capture the leverage effect, even when $\{W\}_{t\in T}$ and $\{B\}_{t\in T}$ are uncorrelated. This aspect of the EWMA HM makes it more appropriate to model stock market dynamics where the empirical leverage effect is stronger than in other markets such as foreign exchanges which have no leverage effects. Another important and related element is the asymmetry between the positive and negative trends in attraction volatility. The fact that the repulsion force approaches infinity when the attraction volatility tends toward $\underline{\nu}$, means attraction volatility cannot fall below $\underline{\nu}$. Therefore, $\underline{\nu}$ constitutes the floor attraction volatility value which means that a long period of a strong positive trend leads to a period where attraction volatility remains around $\underline{\nu}$. Conversely, long periods of a strong negative trend make $\underline{\nu}$ converge toward an affine relationship of m_t. Indeed, a negative trend increases attraction volatility which decreases the repulsion force produced by $\frac{\psi}{\nu_t - \underline{\nu}}\mathrm{d}t$. Consequently, the more negative m_t, the closer $\mathrm{d}\nu_t$ becomes to $-\beta\mathrm{d}m_t$. The speed of the convergence depends on the parameter ψ. All other things being equal, the lower the value of ψ, the greater the rate of convergence toward this situation. Furthermore, similar to the parameter α but related less directly, the value of ψ shows a positive relation to the expected attraction volatility.

These features of the attraction volatility process allow a better specification of the variance process. Based on the above results, the spot variance can be written as:

$$\begin{aligned}
V_t &= V_0 e^{\frac{-t}{\tau_1}} + \frac{1}{\tau_1}\int_0^t e^{\frac{1}{\tau_1}(u-t)}{\nu_u}^2\mathrm{d}u + \xi\int_0^t e^{\frac{1}{\tau_1}(u-t)}\nu_u\sqrt{V_u}\mathrm{d}B_u,\\
\nu_t &= \nu_0 e^{\frac{-t}{\tau_2}} + \frac{1}{\tau_2}\int_0^t e^{\frac{1}{\tau_2}(u-t)}\left(\frac{\psi}{\nu_u - \underline{\nu}} + \underline{\nu} + \alpha\right)\mathrm{d}u - \frac{\beta}{\tau_2}\int_0^t e^{\frac{1}{\tau_2}(u-t)}\frac{\mathrm{d}S_u}{S_u}.
\end{aligned}$$

Also, the variance process can be written such that it more clearly demonstrates the impact of the past trend on the variance:

$$\begin{aligned}
V_t &= V_0 e^{\frac{-t}{\tau_1}} + \frac{1}{\tau_1}\int_0^t e^{\frac{1}{\tau_1}(u-t)}(z_u - \beta m_u)^2\mathrm{d}u + \xi\int_0^t e^{\frac{1}{\tau_1}(u-t)}(z_u - \beta m_u)\sqrt{V_u}\mathrm{d}B_u,\\
z_t &= \nu_0 e^{\frac{-t}{\tau_1}} + \frac{1}{\tau_2}\int_0^t e^{\frac{1}{\tau_2}(u-t)}\left(\frac{\psi}{\nu_u - \underline{\nu}} + \underline{\nu} + \alpha\right)\mathrm{d}u.
\end{aligned}$$

The above expression emphasizes that the strong Zumbach effect is encoded by the model since the volatility is a function of past price trends through m_t. It shows also that the model reproduces the quadratic feedback of the price trends on volatility. Additionally, note that $\{v\}_{t\in T}$ exhibits a very interesting mathematical property related to its dependence on the past. We see that in contrast to rough volatility models (El Euch 2018) the variance process in the EWMA HM is Markovian. This Markovian property makes the model very tractable for Monte-Carlo simulations based on a discretization scheme of the associated SDE system.

2.2 The thresholded version of the model as a limit-case

Following our presentation of the EWMA HM, in the remainder of the article we focus on the case 2 where ψ tends toward zero while remaining strictly positive. In this limit case, attraction volatility is given by:

$$\nu_t = \underline{\nu} + (\alpha - \beta m_t)_+. \tag{3}$$

Thus, the SDE system associated with the EWMA HM can be rewritten as:

$$\begin{cases} \dfrac{\mathrm{d}S_t}{S_t} & = \mu_t \mathrm{d}t + \sqrt{V_t}\mathrm{d}W_t \\ \mathrm{d}V_t & = \dfrac{1}{\tau_1}\Big({\nu_t}^2 - V_t\Big)\mathrm{d}t + \xi\nu_t\sqrt{V_t}\mathrm{d}B_t \\ \mathrm{d}m_t & = \dfrac{1}{\tau_2}\Big(\dfrac{\mathrm{d}S_t}{S_t} - m_t\mathrm{d}t\Big), \end{cases}$$

where ν_t is equal to (3).

Since β is strictly positive, there exists a threshold value of m below which the attraction value of the volatility follows an affine relation with m_t (see figure 1). This critical point is reached when $\alpha - \beta m_t$ is zero. Thus, this value is defined by

$$\bar{m} = \frac{\alpha}{\beta}.$$

Using $\bar{m}$, we can rewrite the attraction volatility as follows:

$$\nu_t = \underline{\nu} + \beta(\bar{m} - m_t)_+.$$

Equivalently, the attraction variance is equal to

$${\nu_t}^2 = \underline{\nu}^2 + \mathbb{1}_{\{m_t<\bar{m}\}}\big(\beta^2(\bar{m} - m_t)^2 + 2\underline{\nu}\beta(\bar{m} - m_t)\big).$$

Thus, when the past trend of the asset measured by an EWMA of its past returns is below the threshold value $\bar{m}$, all variations of this EWMA result in a proportional change in its attraction volatility:

$$\frac{\mathrm{d}\nu_t}{\mathrm{d}m_t} = \begin{cases} 0 & \text{for } m_t \geq \bar{m} \\ -\beta & \text{else.} \end{cases}$$

This result also allows us to deduce the occurrence of a reversal effect if the past market trend is negative. Let us assume that the market portfolio price dynamics follow the SDE system associated with the thresholded version of the EWMA HM, and the trend parameter of the price is given by the following affine relationship:

$$\mu_t = r + \lambda\sqrt{V_t},$$

where r is the free-risk rate. To determine the type of relationship linking the respective dynamics of m and μ, we consider the following simplified case:

$$\lim_{\tau_1,\xi\to 0} V_t = {\nu_t}^2.$$

Under this hypothesis, we have

$$\frac{\mathrm{d}\mu_t}{\mathrm{d}m_t} = \begin{cases} 0 & \text{for } m_t \geq \bar{m} \\ -\beta\lambda & \text{else.} \end{cases}$$

Consequently, in this simplified framework, if $m_t < \bar{m}$, the joint dynamics of m and μ are negatively correlated even if the two Brownian motion $\{B\}_{t\in T}$ and $\{W\}_{t\in T}$ are independent (i.e. $\rho = 0$). More generally, this relation remains true if τ_1 is strictly positive and $\rho \in [-1 : 0]$. In financial terms, this relationship can be viewed as a reversal effect, in the sense that a negative past trend (such as $m_t < \bar{m}$) increases the actual price trend. However, note that unlike the assumption in behaviorist explanations (Dissanaike 1997), this model reversal effect is not a market anomaly. Indeed, the increase in the returns trend of a given asset following a strong fall in its price is due not to the correction of a market overreaction but to a pure market premium effect. In addition, this reversal effect emerges only if $m_t < \bar{m}$ which causes

an asymmetry depending on the type of trend. Thus, positive trends such as $m_t > \bar{m}$ do not produce this type of reversal effect which results from a movement in the attraction volatility.

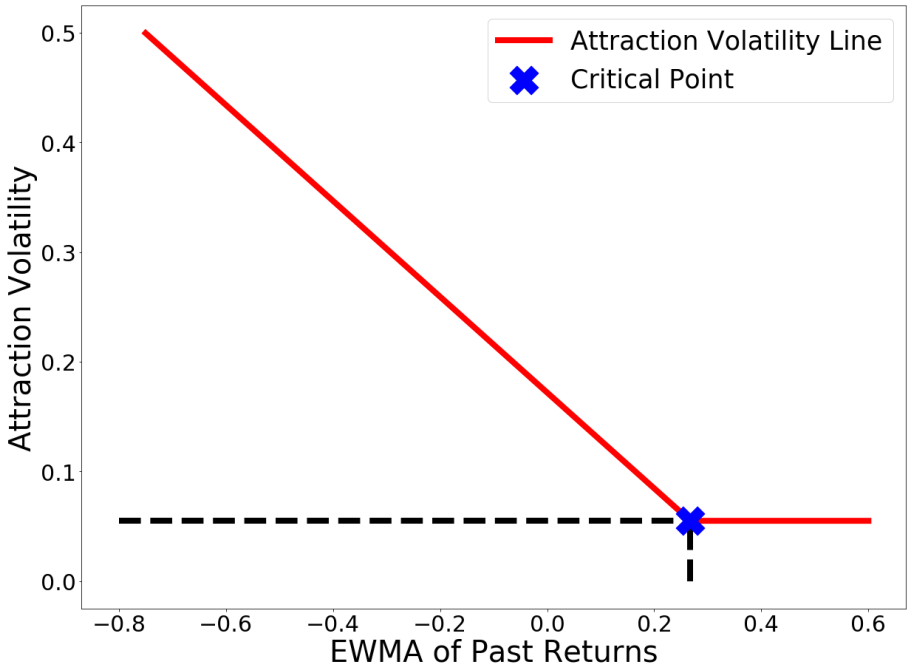

Figure 1: The relationship between the attraction volatility process and the trend parameter m_t in the thresholded EWMA Heston model.

2.3 Comparison with volatility models based on quadratic Hawkes processes

The EWMA HM has some links to the quadratic volatility models introduced by Blanc *et al.* (2017) and developed by Dandapani *et al.* (2019). To highlight the similarities between these models, we focus on the case of the thresholded version of the EWMA HM where $\mu, \xi, \tau_1 = 0$. In this specific case, we have:

$$\begin{cases} \dfrac{\mathrm{d}S_t}{S_t} & = \sqrt{V_t}\mathrm{d}W_t \\ V_t & = \underline{\nu}^2 + \mathbb{1}_{\{m_t < \bar{m}\}}\big(\beta^2(\bar{m} - m_t)^2 + 2\underline{\nu}\beta(\bar{m} - m_t)\big) \\ m_t & = \dfrac{1}{\tau}\displaystyle\int_0^t e^{\frac{1}{\tau}(u-t)}\frac{\mathrm{d}S_u}{S_u}. \end{cases}$$

This system of equations can be compared with the system of equations satisfied by different quadratic volatility models. First, let us focus on the pure quadratic case of the volatility model introduced by Blanc *et al.* (2017), where the kernel function takes the following exponential form $k(t) = \sqrt{\frac{2}{\tau}}e^{-\tau t}$. For simplicity, we describe this model as the pure quadratic volatility model (PQVM). Notice that both the PQVM and the EWMA HM satisfy the following equations:

$$\begin{cases} \dfrac{\mathrm{d}S_t}{S_t} & = \sqrt{V_t}\mathrm{d}W_t \\ V_t & = \underline{\nu}^2 + b_t\big(b_t(\bar{m} - m_t)^2 + 2\sigma(\bar{m} - m_t)\big) \\ m_t & = \dfrac{1}{\tau}\displaystyle\int_0^t e^{\frac{1}{\tau}(u-t)}\frac{\mathrm{d}S_u}{S_u}. \end{cases}$$

with respectively

	Pure Quadratic Volatility Model	EWMA Heston Model
b_t	$\sqrt{\frac{2}{\tau}}$	$\mathbb{1}_{\{m_t<\bar{m}\}}\beta$
$\bar{m}$	0	$\frac{\alpha}{\beta}$
σ	0	$\underline{\nu}$

such as $\tau, \beta, \underline{\nu} \in \mathbb{R}_+, \gamma \in [0:1[$.

Using this formulation, we can see that these two models present several similarities. The most important one is that, in both models, the variance is a quadratic function of an EWMA of past returns. Note however that, in contrast to the PQVM, in the EWMA HM the linear term in m_t is non-zero. This difference stems from the linear sensitivity to m_t of the volatility (if $m_t < \bar{m}$) in the EWMA HM, while in the PQVM, it is the variance that is linearly sensitive to m_t. Another important specificity of the model which constitutes one of the main differences with the various quadratic volatility models discussed in the literature (Blanc *et al.* 2017, Dandapani *et al.* 2019, Gatheral *et al.* 2020) is the switching behavior caused by the indicator function $\mathbb{1}_{\{m_t<\bar{m}\}}$. This is confirmed if we consider the quadratic rough HM (QRHM). In contrast to the PQVM, the QRHM encodes the empirical asymmetry of the feedback effect. However, this asymmetry is not captured in exactly the same way as in the EWMA HM. To highlight this point, note that both the EWMA HM considered and the QRHM described in section 4 in Gatheral et *et al.*'s (2020) article respect the following equations:

$$\begin{cases} \frac{\mathrm{d}S_t}{S_t} &= \sqrt{v_t}\mathrm{d}W_t \\ V_t &= \bar{\theta} + b_t\big(b_t(\bar{m}-m_t)^2 + 2\sigma(\bar{m}-m_t)\big) \\ m_t &= m_0 - \int_0^t f(t-s)m_s\mathrm{d}s + \int_0^t f(t-s)\eta\frac{\mathrm{d}S_u}{S_u}, \end{cases}$$

with respectively:

	Quadratic Rough Heston Model	EWMA Heston Model
b_t	a	$\mathbb{1}_{\{m_t<\bar{m}\}}\beta$
σ	0	p
η	ν	1
$f(u)$	$u^{\alpha-1}\frac{\lambda}{\Gamma(\alpha)}$	τ

such as $\nu, a, \lambda, \tau, p \in \mathbb{R}_+, \alpha \in [0.5:1]$.

Again, one of the main differences between the two models concerns the parameter b_t, and the related switching behavior of the EWMA HM considered. Whereas in the QRHM this parameter is a positive constant, in the EWMA HM it is the product of a constant and an indicator function. This difference implies an alternative way to capture the empirical asymmetry between positive and negative trends in volatility. In the QRHM, this asymmetry is covered only by the parameter $\bar{m}$. In concrete terms, this means that even though positive and negative trends have asymmetric impacts on the volatility, a strong positive trend such as $m_t > \bar{m}$ increases volatility in the QRHM. However, since $m_t > \bar{m}$, the trend measured by $m$has no impact on the volatility in the EWMA HM considered. In addition, while in the QRHM the variance is an affine function of the squared difference between m_t and $\bar{m}$, the variance depends also on the absolute value of this difference. At the same time, unlike in the PQVM, the trend parameter $m$in the QRHM is not an EWMA of past returns. In this case the kernel used to model the dependency between past price trends and spot volatility, is a rough one. This change has several important implications. First, the model is no longer Markovian in terms of its variables (S, V). Second, an important consequence of this choice is that the "memory" of m decays as a power law in the QRHM, while in the EWMA HM the memory decays exponentially. Note however that in both cases, m is highly sensitive to recent returns. Third, the reasons

for the emergence of irregular behavior of the volatility over short time scales are different between the QRHM and the general EWMA HM. In the QRHM, the roughness of the volatility paths stems from the characteristics of its power law kernel whereas in the EWMA HM, this erratic volatility process behavior emerges under certain conditions due to antagonist effects produced on the one hand by the reversal effect of the variance process, and on the other hand by the volatility of this variance.

2.4 Model discretization scheme

As referred to in the introduction, a major strength of the EWMA HM compared to rough volatility models is its ease of use for Monte-Carlo simulations. The simplest way to run discrete-time simulations of the thresholded version of the EWMA HM is to use a Euler-type discretization scheme (Maruyama 1955), analogous to a full truncated Euler scheme for the standard HM (Lord *et al.* 2010). We thus propose the following simple discretization scheme [3]:

$$\begin{aligned}
S_{t+1} &= S_t + S_t\big(\mu_t \Delta t + \sqrt{(V_t)_+}Z_{t+1}\big) \\
m_{t+1} &= m_t + \frac{1}{\tau_2}\left(\frac{S_{t+1}-S_t}{S_t} - m_t\Delta t\right) \\
\nu_{t+1} &= \underline{\nu} + \max(0, \alpha - \beta m_{t+1}) \\
V_{t+1} &= V_t + \frac{1}{\tau_1}\big({\nu_{t+1}}^2 - (V_t)_+\big)\Delta t + \xi\nu_{t+1}\sqrt{(V_t)_+}\big(\rho Z_{t+1} + \sqrt{1-\rho^2}X_{t+1}\big),
\end{aligned}$$

where X_t and Z_t are i.i.d. random variables associated with the Gaussian distribution $\mathcal{N}(0, \Delta t)$.

Regarding the trend component of $\{S\}_{t\in T}$, for the following simulations we assume $\mu_t = \lambda\sqrt{V_t}$. This implies replacing μ_t with $\lambda\sqrt{(V_t)_+}$ in the discretization scheme. In addition, to simulate future potential volatility trajectories for an asset from a given date t, requires m_t to be estimated. The natural estimator of this parameter is given by the following discrete EWMA:

$$m_t = \frac{1}{\tau_2}\sum_{i=1}^{n} e^{-\frac{i\Delta t}{\tau_2}}\frac{S_{t-i\Delta t} - S_{t-(i+1)\Delta t}}{S_{t-(i+1)\Delta t}}.$$

To fix the value of n, we can write that:

$$\frac{1}{\tau_2}\int_{t-n\Delta t}^{t} e^{\frac{1}{\tau_2}(u-t)}\mathrm{d}u = 1 - e^{-\frac{\tau n\Delta t}{\tau_2}}.$$

Consequently, for a given proportion γ of the total weighting of the EWMA (in a continuous framework), we have:

$$\begin{aligned}
\gamma &= 1 - e^{-\frac{\tau n\Delta t}{\tau_2}} \\
n &= -\tau_2\frac{\log(1-\gamma)}{\Delta t}.
\end{aligned}$$

The above equation allows us to fix n, given $\tau_2, \Delta t$ and γ.

[3] The convergence of this discretization scheme is still to be demonstrated.

3 Consistency of the model with market data

To assess the effectiveness of the model for capturing empirical financial phenomena, we use as market data the price and volatility dynamics of the S&P500 for the period July 23, 2001, to July 23, 2021. To estimate the volatility, we use the square root of the realized variance computed from 5-min samples provided by the Oxford-Man Institute of Quantitative Finance[4]. We compare these market data with a set of synthetic data comprised of 1000 simulations for each of 20 years, generated from the Euler scheme of the thresholded version of the EWMA HM, using a time step equal to $\frac{1}{25200}$ expressed in years. To fit the model, we use the ad-hoc estimation procedure presented in appendix B. This restults in the following parameters:

λ	ρ	τ_1	ξ	τ_2	$\underline{\nu}$	α	β
0.5575	-0.465	0.0013	42.95	0.276	0.0595	0.1033	$\sqrt{\frac{2\tau}{\pi}}$

Table 1: The parameters obtained from the calibration procedure.

These parameter values need some explanation. First, let us consider the value taken by λ which is setting the long-term drift. To obtain an order of magnitude for the expected value of the drift (i.e. $\mathbb{E}[\mu_t]$), we can use the product of the mean of the realized volatility of the S&P500 by this parameter λ as a proxy. The mean of the realized volatility is 13.1%, so the value obtained is 7.3% ($0.5575 \times 0.131 \approx 7.3\%$). This rough estimation of $\mathbb{E}[\mu_t]$ is consistent with the annualized empirical daily mean of the returns which is equal to 7.3% (the S&P500 empirical mean of daily returns is equal to 0.028%: thus $(1 + 0.028\%)^{252} \approx 7.3\%$). Also interesting is that the value of τ_1 is extremely low compared to the range of values it takes in the standard HM fitted on market data (Mrázek 2017). Since the unit of τ_1 is in years, τ_1 equal to 0.0013 means that the average duration of the deviation of the variance from this attraction value is of the order of 0.5 days $\big(0.0013 \times 365 \approx 0.5\big)$. In other words, the reversion of the variance process toward the attraction variance is a short term phenomenon. This should understood in the context of the very high value taken by ξ, implying that over very short time scales the variance process is dominated by randomness. The combination of these two effects - high randomness of V_t at the time scale dt and short time reversion toward ${\nu_t}^2$ - results in erratic volatility behavior over the short term, and allows the roughness of empirical volatility paths to be mimicked. Consequently, while in the standard HM the volatility serial correlation is highly dependent on τ_1, here, it depends almost exclusively on the attraction variance ${\nu_t}^2$. As a result and given the value of τ_2, the autocorrelation of the volatility depends on a medium-term price trend of the order of 3.3 months $\big(0.276 \times 12 \approx 3.3\big)$, and not directly on past endogenous volatility movements. Thus, there is a decoupling between the short and the longer-term behavior of the volatility, a property whose importance was emphasized notably by Bennedsen *et al.* (2016).

3.1 The relationship between EWMAs of returns and the volatility

If the EWMA HM is consistent, the market data should exhibit an affine relationship between the EWMA of past returns and realized volatility when the EWMA of past returns is significantly negative. However, if the EWMA of past returns is significantly positive, these two variables should be fairly independent. These relationships are confirmed in figure 2. Another remarkable fact that emerges in relation to the optimal slope coefficient β (optimal in the least square sense) of the regression of the form $\sqrt{V_t} = \underline{\nu} + (\alpha - \beta m_t)_+ + \epsilon_t$ is that for different past EWMAs of the returns (with different values of τ), the optimal slope coefficient is very close to[5]: $\sqrt{\frac{2\tau}{\pi}}$. $\sqrt{2\tau}$ occurs because it corresponds to the inverse of the asymptotic standard deviation of an EWMA with parameter τ of a Brownian motion[6]. The rationale for $\sqrt{\pi}$, is more ambiguous. Although we do not formulate an explicit hypothesis in the present paper, this value could be related directly to the fact that $\sqrt{\pi}$ corresponds to $\frac{1}{\Gamma(0.5)}$, with $\Gamma(.)$ the gamma function. However, the fundamental reason for this relationship has yet to be explained.

[4] The data are available at: https://realized.oxford-man.ox.ac.uk/data.

[5] For other examples of this empirical relationship, see appendix D.

[6] See appendix C.

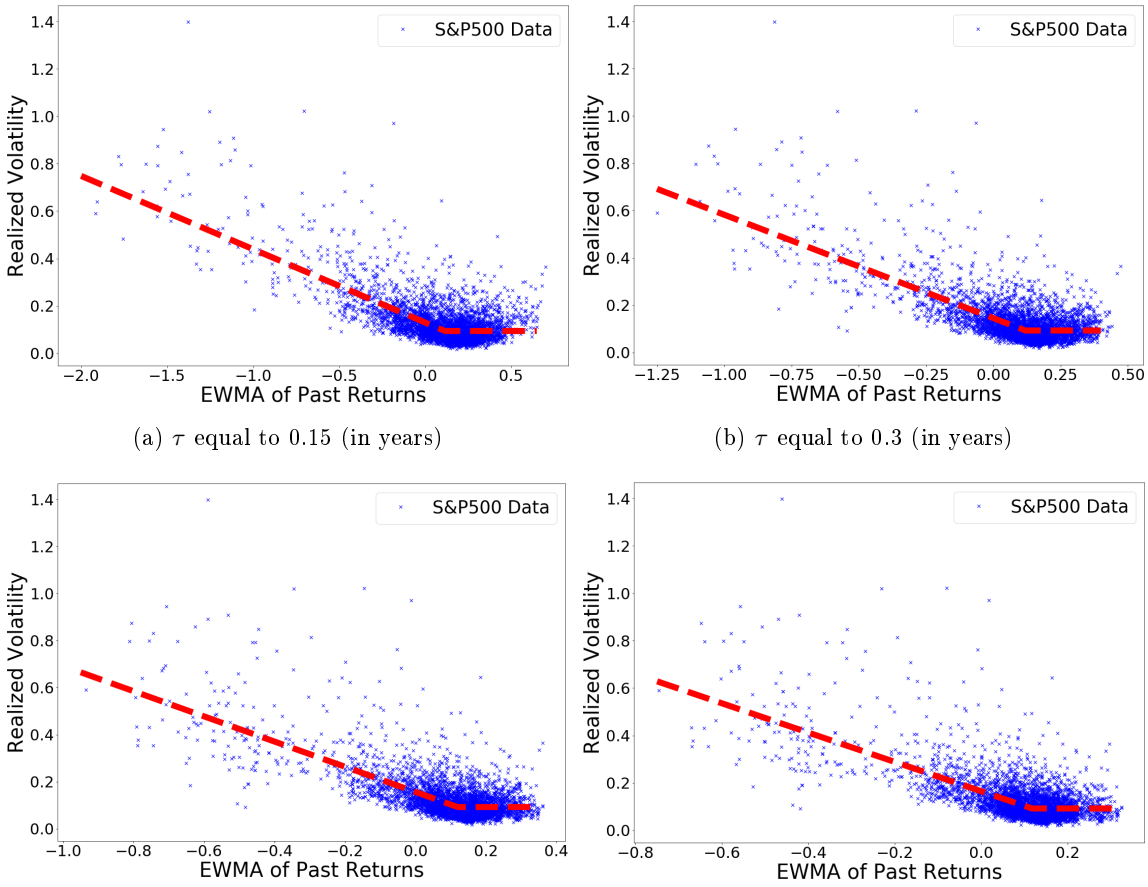

(a) τ equal to 0.15 (in years)

(b) τ equal to 0.3 (in years)

(c) τ equal to 0.45 (in years)

(d) τ equal to 0.6 (in years)

Figure 2: The empirical relationship between the realized volatility of the S&P500 and the EWMAs of its past returns for different values of τ expressed in years. The red lines are regressions of the form $\sqrt{V_t} = \underline{\nu} + (\alpha - \beta m_t)_+$ such as the associated regression lines have all a slope coefficient equal to $\sqrt{\frac{2\tau}{\pi}}$. These regressions are respectively associated with the following R-squared: 0.558,0.547, 0.523, 0.498.

3.2 The dynamics of price and volatility

Let us begin our comparison of the synthetic data produced by the EWMA HM and the S&P500 empirical data, by focusing on the joint price and volatility dynamics. A general observation that emerges from this comparison (see figures 3 and 4) is that the simulated price and volatility trajectory features resemble those of the empirical data. In particular, in both cases, volatility spikes immediately after large negative returns of the index and then decreases gradually. Also, in each case the long positive price trends are periods where volatility is relatively stable and low. If we consider the volatility dynamics more specifically, we see that the model is able also to capture the empirical coexistence of long periods of low volatility with other periods when volatility is very high. Also, we can see that the empirical roughness of the volatility path understood as irregular behavior over short time scales is accurately reproduced by the model[7]. These general observations require completion by a more in-depth comparison between the features characterizing the volatility dynamics produced by the model and the features of the realized volatility. We next focus on two features of the volatility: its autocorrelation function (ACF) and its roughness.

[7]Here the focus is not on the roughness understood according to its rigorous mathematical definition (Bennedsen *et al.* 2016) but on the visual features of the volatility paths.

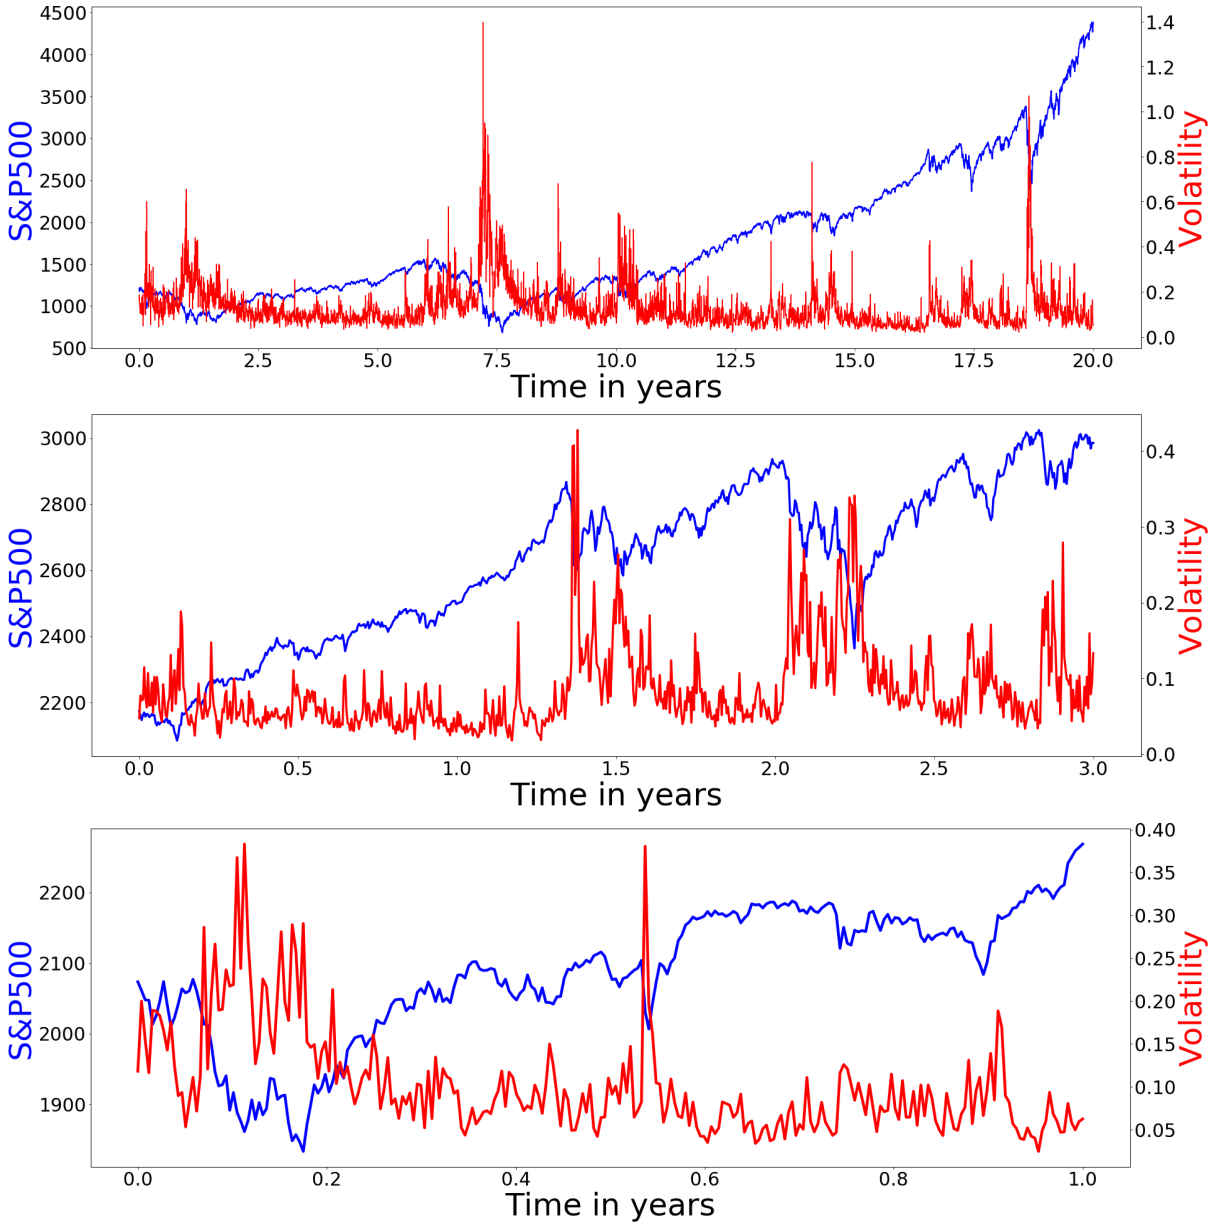

Figure 3: Empirical joint evolution of the S&P500 and its annualized daily realized volatility (5 min samples estimator) over 20, 3 and 1 years.

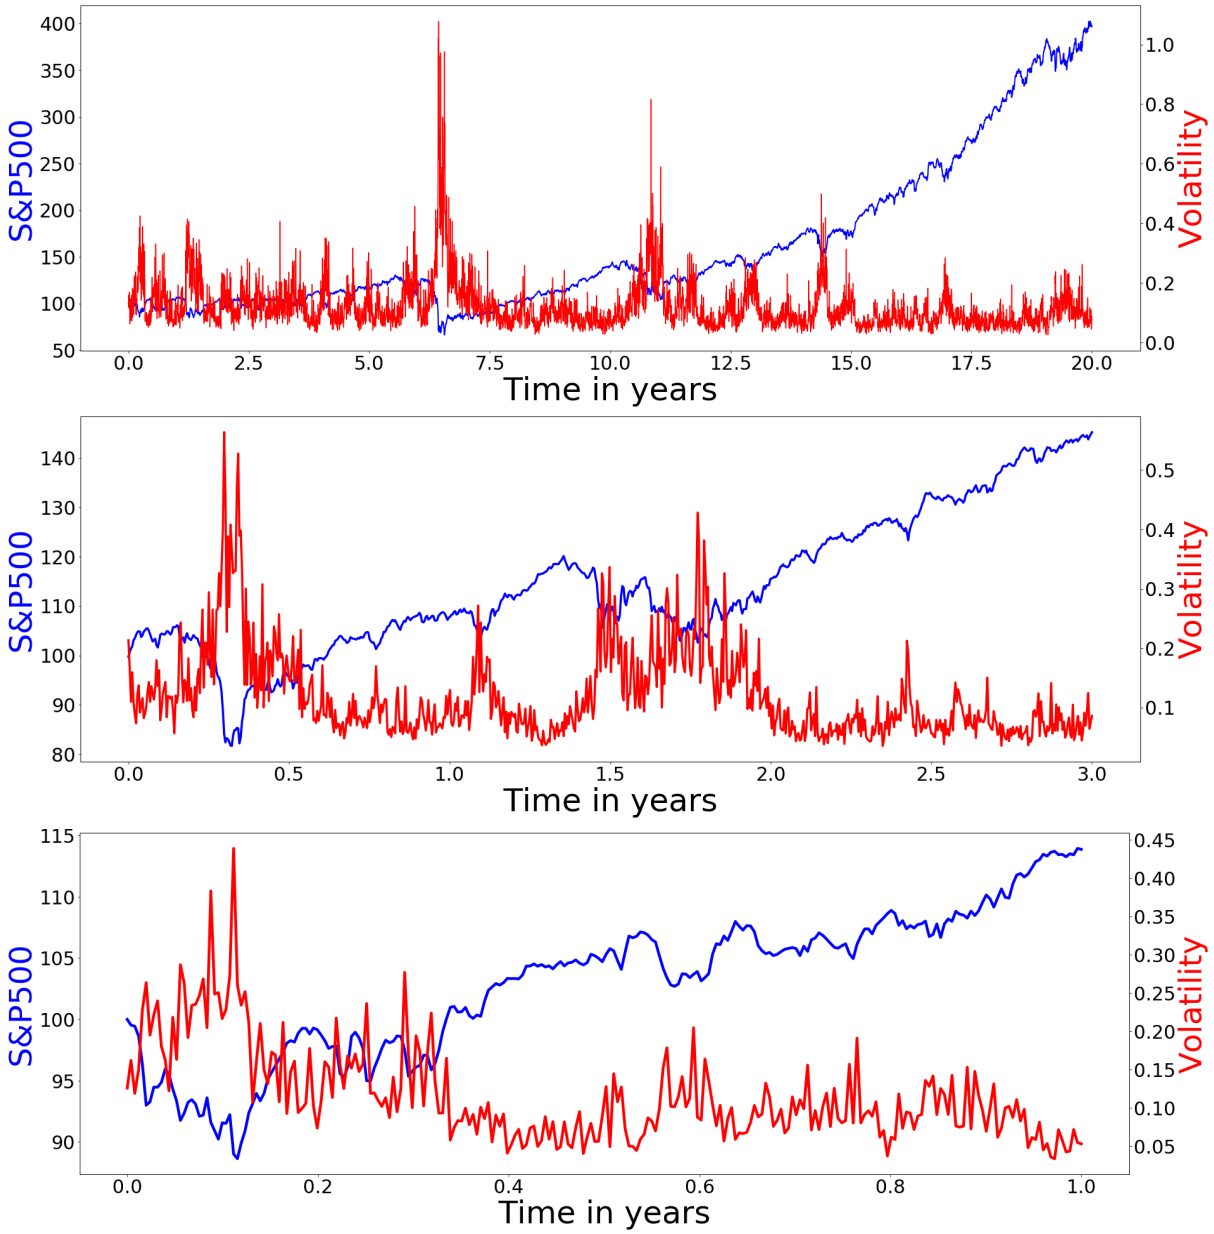

Figure 4: Joint evolution of the S&P500 and its associated volatility over 20, 3 and 1 years simulated from the fitted EWMA Heston model.

3.2.1 The autocorrelation function of the volatility process

The serial correlation of the volatility is an important feature of its dynamics. In figure 5 we observe that the ACF of the simulated volatility paths is generally consistent with the market data. However, the short-medium term (between 1 and 60 trading days) and the longer-term must be analyzed separately. In the case of small lags (less than 60 trading days), the autocorrelations associated with the volatility paths generated by the model are very close to their empirical counterparts. Over this time horizon, not only are all empirical autocorrelations within the model's estimated confidence intervals[8], they are also close to their expected values. In the case of longer lags, the situation changes slightly. The ACF exhibits faster decay in the model compared to the empirical data. However, this difference between the empirical and the model data should be tempered by the fact that the empirical autocorrelations are within the 95% confidence interval of the model. Consequently, it is difficult to be definitive about the capacity of the model to account accurately for the ACF in the real volatility process. However, the academic literature (Ding *et al.* 1993, Cont 2001) tends towards the hypothesis that the model ACF decays too rapidly compared to the market data. This fast decay is due to the fact that the ACF of the volatility is determined primarily by an exponential kernel through the trend process m_t. A possible modification of the thresholded version of the EWMA HM model allowing a volatility process with a more flexible memory structure would involve using a linear combination of several EWMAs of past returns rather than a unique EWMA. In these cases, m_t is replaced by $\bar{m}_t$ defined as follows:

$$\bar{m}_t = \sum_{k=1}^{d} w_k \int_0^t e^{-\tau_k(t-u)} \frac{\mathrm{d}S_u}{S_u}$$

In this framework, the model still belongs to the family of EWMA HMsl introduced previously (1), but no longer corresponds to the specific case of (2).

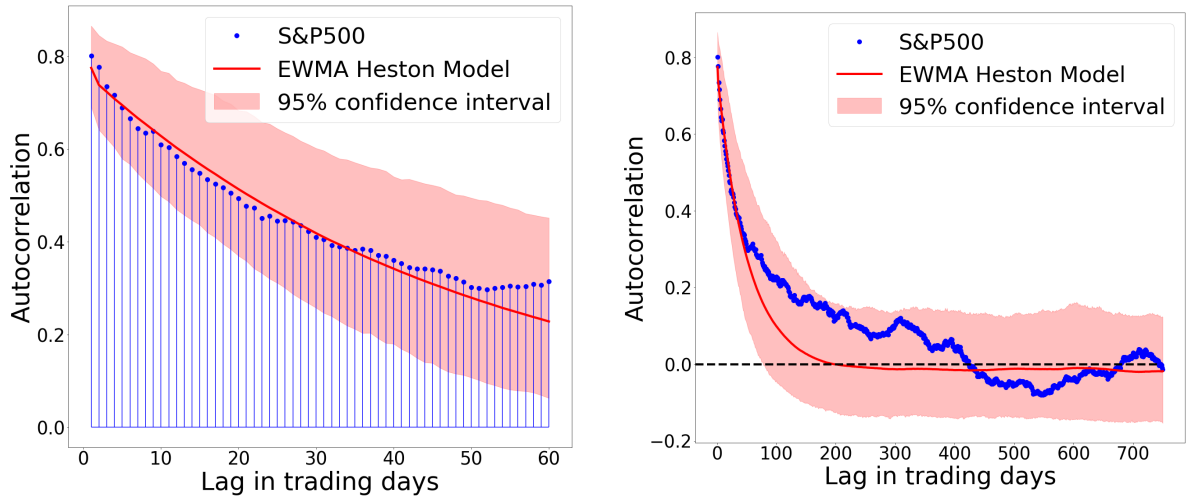

(a) Short/Medium term autocorrelation

(b) Medium/Long term autocorrelation

Figure 5: Autocorrelation of the volatility in the short and longer-term.

3.2.2 The question of volatility roughness

Although visually the EWMA HM would seem to capture the empirical roughness of the volatility paths, this requires deeper investigation. First, note that the EWMA HM is not a rough volatility model. Indeed, the EWMA HM is not based on a fractional process similar to a fractional Brownian motion. Rather, it is a diffusion model whose two sources of randomness are standard Brownian motions which by definition

[8] This confidence interval is estimated using the 2.5% and 97.5% simulation quantiles.

both have a Hurst exponent equal to 0.5. However, Rogers (2019) and more recently Cont and Das (2022) show that Brownian diffusion models may mimic some of the features of rough models of interest for volatility modeling. Cont and Das emphasize also that the empirical roughness of the volatility paths might be caused not by the volatility process but by a microstructure noise. Here, we leave aside consideration of the origins of the empirical rough behavior of realized volatility and adopt a phenomenological approach aimed at evaluating how well the EWMA HM mimics the empirical roughness of the volatility paths. To do this, we measure the roughness of both the empirical and model volatility time series using two different methods.

The first method is the most popular in the academic literature. It consists of estimating a roughness index by running an ordinary least square (OLS) regression of the form:

$$\log\big(\gamma(k\Delta, 2)\big) = c + a \cdot \log(k\Delta) + \epsilon_k, \qquad k = 1, ..., \mathrm{m}$$

where Δ is the time step between observation of the volatility process, m is a bandwidth parameter, and $\gamma(k\Delta, 2)$ is the empirical second order variogram defined as follows:

$$\gamma(k\Delta, 2) = \frac{1}{N-k} \sum_{i=1}^{N-k} \mid \log(\sigma_{(i+k)\Delta}) - \log(\sigma_{\Delta i}) \mid^2,$$

where N is the number of observations in the sample considered. Due to the nature of the data used, sampling in Δ is conducted daily, and not on a logarithmic scale. In order to remove the bias induced by this sampling, we weight each observation k by

$$w_k = \log(k+1) - \log(k).$$

The slope coefficient $\hat{a}$ obtained from the regression is used to estimate a roughness index, denoted $\hat{H}$, equal to $0.5\hat{a}$. In the case of a rough process, $\hat{H}$ can be considered a Hurst exponent estimator (Gatheral 2018). However, in the present case, $\hat{H}$ is treated only as a metric to compare the behaviors of the real and the EWMA HM volatility processes.

The second method used to measure the roughness of time series was proposed by Bennedsen *et al.* (2016). Unlike the first method described, it has the advantage of being able to capture a possible non-affine relationship between $\gamma(k\Delta, 2)$ and $\log(k\Delta)$. The method consists of a non-linear least squares (NLLS) regression of the form

$$\gamma(k\Delta, 2) = c + b(k\Delta)^{2a+1} + \epsilon_t, \qquad k = 1, ..., \mathrm{m}$$

The index $\hat{H}$ is equal to $\hat{a} + 0.5$.

Table 3 reports the roughness index $\hat{H}$ of the S&P500 and simulated volatilities obtained using OLS and NLLS methods with different bandwidth parameters.

Bandwidth parameter	10	20	60	125	250	500	750
$\hat{H}$ of S&P500 data (OLS regression)	0.137	0.136	0.137	0.135	0.13	0.126	0.124
$\hat{H}$ of model data (OLS regression)	0.112	0.121	0.142	0.15	0.149	0.14	0.134
$\hat{H}$ of S&P500 data (NLLS regression)	0.082	0.078	0.116	0.072	0.046	0.091	0.07
$\hat{H}$ of model data (NLLS regression)	0.054	0.126	0.283	0.173	0.042	-0.071	-0.118

Table 2: Roughness index $\hat{H}$ of both S&P500 and simulated volatilities computed from OLS and NLLS methods with different bandwidth parameters.

First, when considering the$\hat{H}$ exponents obtained from both estimation methods, while the realized volatility $\hat{H}$ depends only slightly on the bandwidth parameters used, this parameter has a more significant influence

on the value of $\hat{H}$ for the volatility generated by the model. The reason for this difference becomes clearer when looking at figure 6. In terms of the empirical data, the relationship between $\log\big(\gamma(k\Delta, 2)\big)$ and $\log(k\Delta)$ is captured very well by an OLS model, regardless of the time scale considered. However, the relationship between these two variables is s-shaped in the case of the volatility process generated by the model: it is convex up to about 50 trading days and then turns concave. Therefore, the model does not capture the empirical relationship linking the log expected value of the increments of log-volatility and the log time scale. This is quite important because this empirical relationship is a feature which characterizes volatility behavior regardless of the type of asset considered (Gatheral 2018). The significance of this model limitation requires some qualification insofar as over reasonable time scales the distribution of the log-volatility increments generated by the model remain close to the empirical increments, as shown in figure 7. This is reflected in the fact that the distributions are close to Gaussian for both the empirical and simulated data. Furthermore, the low (but significant) divergence between the variances in the model log-volatility increments and their empirical counterparts over different time scales explains a relatively similar order of magnitude of the OLS estimators of the Hurst coefficient. In contrast, the more important divergence of the exponent $\hat{H}$ obtained using NLLS is explained by the fact that this regression captures the nonlinearity of the relationship linking $\gamma(k\Delta, 2)$ and $\log(k\Delta)$. In addition, it should be noted that $\hat{H}$ takes negative values, which emphasize that in the present context it is not relevant to consider $\hat{H}$ as a Hurst exponent estimator (see Mandlebrot 2003). Nevertheless, in the same way as a volatility process with a long-range autocorrelation, this model limitation could be fixed by replacing m_t with a linear combination of multiple EWMAs of past returns.

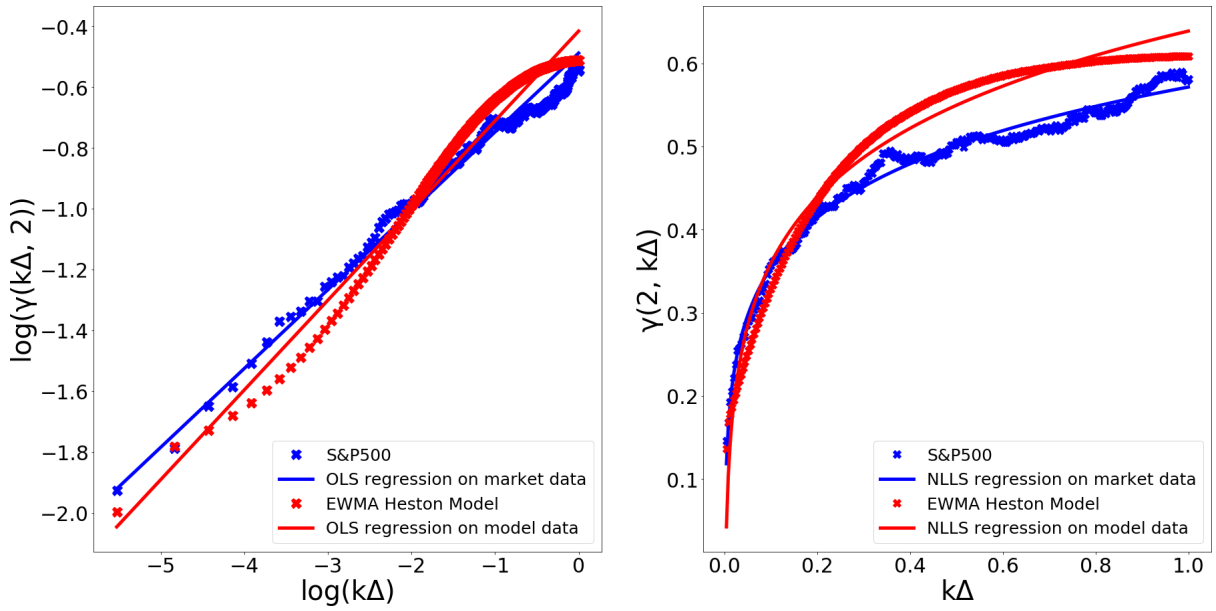

Figure 6: The relationships between $\log\big(\gamma(k\Delta, 2)\big)$ and $\log\big(k\Delta\big)$ (left plot) and between $\gamma(k\Delta, 2)$ and $k\Delta$ (right plot): S&P500 data vs EWMA Heston Model.

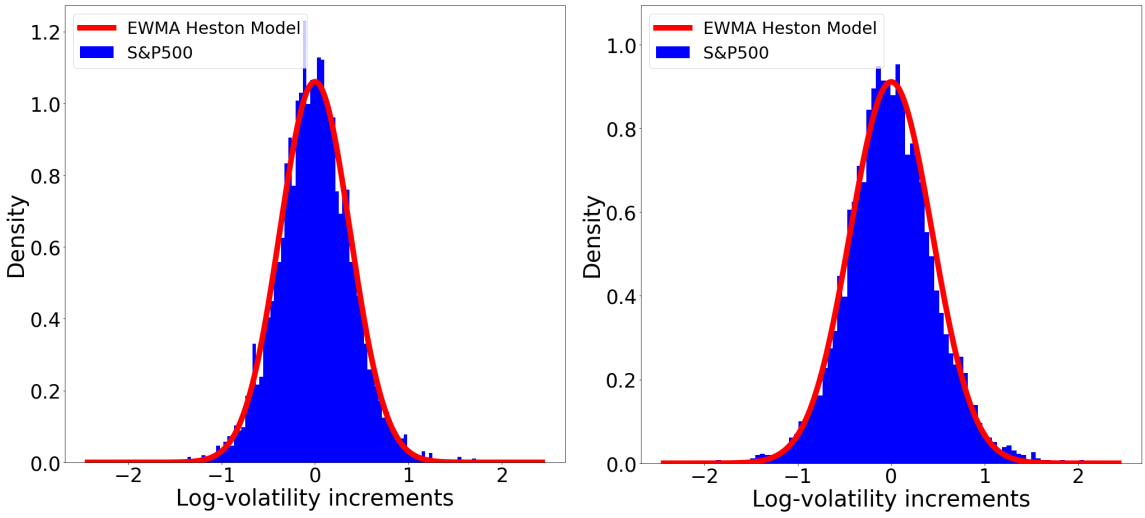

(a) 1 trading day

(b) 5 trading days

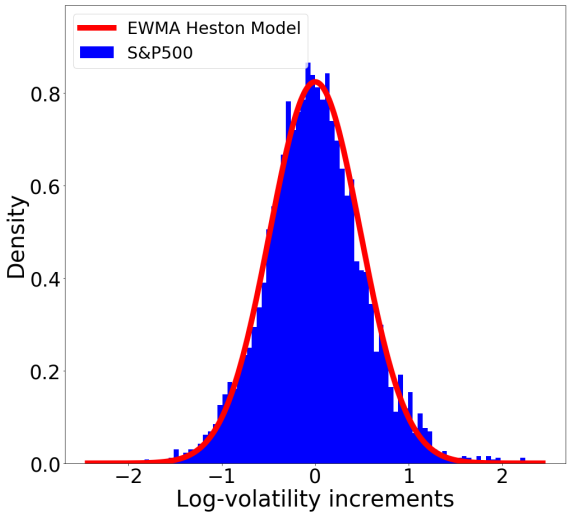

(c) 10 trading days

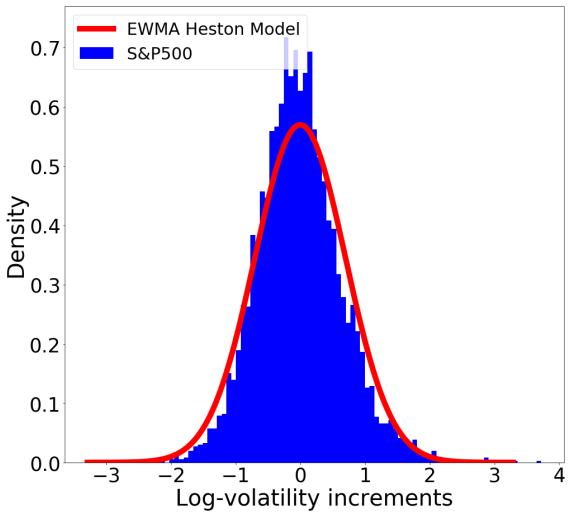

(d) 60 trading days

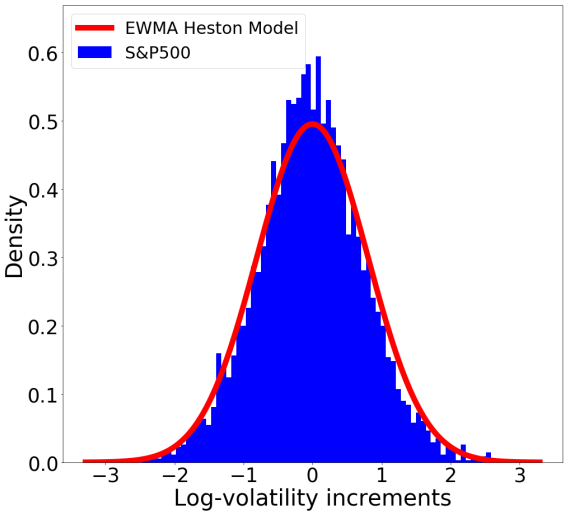

(e) 250 trading days

Figure 7: Log-volatility increments distribution for different time horizons: S&P500 data vs EWMA Heston Model.

3.3 Volatility and log-returns distributions for different time horizons

In what follows we focus on the ability of the model to reproduce the empirical log-returns and average volatility distributions for different time horizons. First, it should be pointed out that assessing the quality of the model on these criteria using standard statistical tools is not straightforward. These probability distributions do not originate from i.i.d. samples but from realizations of a path-dependent process. The practical implications of this are that more often than not, the results of a Kolmogorov-Smirnov test on log-returns or volatility distributions from two different simulations of the fitted model (thus with the same set of parameters) reject the null hypothesis at a 99% confidence level. Consequently, rather than proceeding with this type of statistical test based on the i.i.d. data assumption, it is better to check whether the different moments of the distributions from the market data are in the same range as those emerging from the fitted EWMA HM simulations. To do this, we use the already mentioned synthetic set of 1000 simulations for each 20 years. These simulations allow us to compute the mean as well as the 1^{st} and the 9^{th} deciles of the first four moments of the log-returns distributions and the average volatility distributions for different time horizons. Tables 3 and 4 report these metrics and the corresponding empirical moments.

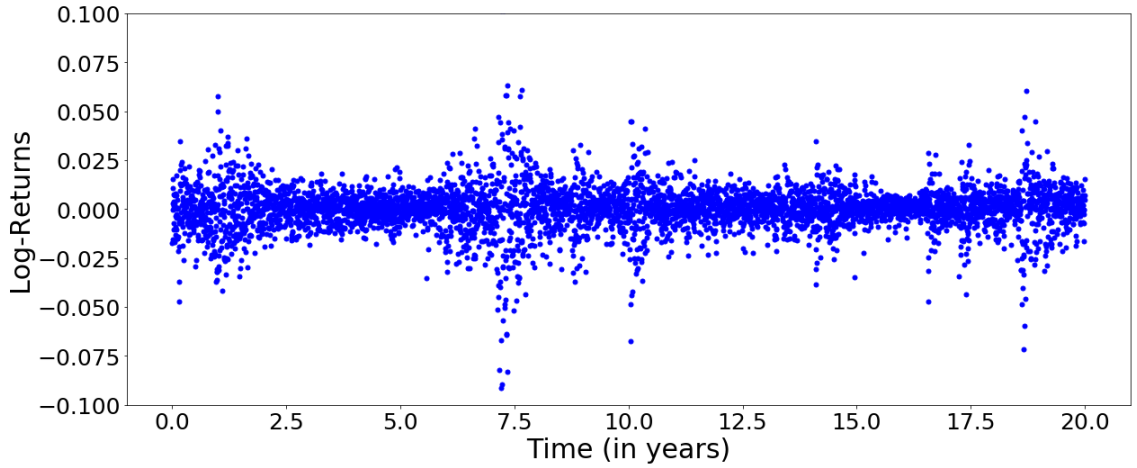

Figure 8: Empirical daily log-returns time-series of the S&P500 between July 23, 2001, to July 23, 2021.

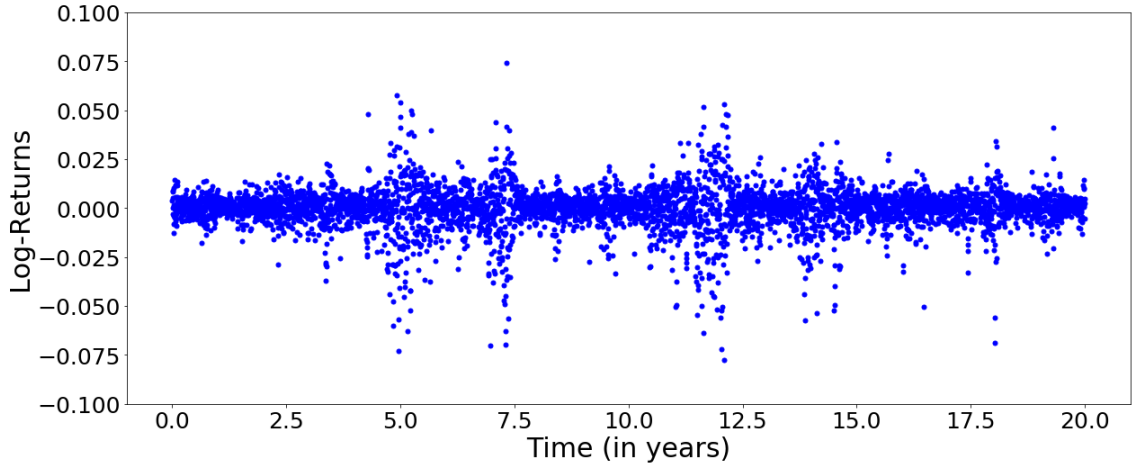

Figure 9: Example of daily log-returns time series from the fitted EWMA Heston model over 20 years.

3.3.1 The log-returns distributions

The log-returns distributions generated by the model are mostly in line with those based on the empirical data. For example, with the exception of the 1 trading day time horizon, the four first moments of the empirical log-returns distributions are between the corresponding 1^{st} and the 9^{th} deciles of the model for all the other time horizons considered [9]. In addition, the relationships between the moments of the log-returns distribution and the time-horizon are generally well reproduced by the model. In particular, as per the market data, in the short-medium terms, the log-returns distribution becomes more and more negatively skewed as the time horizon lengthens. This increasing asymmetry emerges clearly in figure for both the empirical and model distributions. However, the relationship is not monotonic and beyond a certain time-horizon threshold (approximately 65 trading days for the synthetic data), the skewness of the log-returns distribution decreases and in the long run tends to zero. In the case of the fourth moment of the log-returns distributions, the negative convex relationship between the value of the kurtosis and the time-horizon produced by the model is coherent also with the market data. Another characteristic of the empirical log-returns distributions is their heavy tails. Again, the tails of the log-returns generated by the model are mostly consistent with the empirical data, as shown in figure 13. This means that the model is able to capture the possibility of large price swings, a sign of the "wild randomness" of financial markets (Mandelbrot and Hudson 2005). However, despite these model qualities, it should be noted that unlike the log-returns distributions associated with longer time horizons, the model daily log-returns distributions exhibit significant divergence from the empirical distribution (see figure 13). This divergence is demonstrated in the too low standard deviation and a too negative skewness associated with the model distribution compared to the counterpart empirical data. This raises the question of the possibility of a normal source of randomness to accurately capture price fluctuations over short time scales.

3.3.2 The volatility distributions

We now consider the average volatility distributions for different time-horizons (see table 3 and figure 14), and observe a consistency between the model and the market data similar to that observed for the log-returns distributions. Thus, for all the time horizons considered, the first four moments of the average volatility distributions computed using the market data are between the corresponding 1^{st} and the 9^{th} deciles of the model and are close to their expected values. In addition, in the case of both the empirical data and the data generated by the model, the standard deviation, skewness, and kurtosis of the average volatility distribution adopt a decreasing convex relationship with the time horizon considered. More generally, the shapes of the model and empirical distributions are very similar particularly in the case of the daily annualized volatility distribution generated by the model which almost exactly replicates the empirical distribution associated with the S&P500 data. Although the divergence between the empirical and the model average volatility distributions increases slightly for longer time horizons, the goodness of fit remains acceptable. In addition we observe that the distribution of the volatility process generated by the model is very close to both the log-normal and inverse gamma distributions, as shown in figure 10.

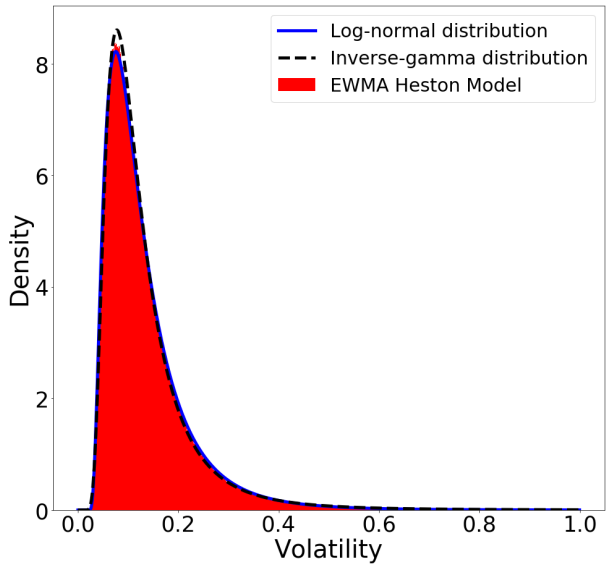

Figure 10: Distribution of the volatility process.

[9] The 1^{st} and 9^{th} deciles can be considered proxies for the lower and upper bounds of the 80% confidence intervals of the moments of the model.

Time horizon	1	5	10	20	60	125	250	500	750
Mean of market data	0.0003	0.0013	0.0026	0.0051	0.0154	0.0321	0.0641	0.1282	0.1923
Avg. mean of sim.	0.0002	0.001	0.0019	00038	0.0114	0.0237	0.0474	0.0949	0.1423
1^{st} decile of mean of sim.	-3e-6	-1e-5	-3e-5	-5e-5	-0.0002	-0.0003	-0.0007	-0.0013	-0.002
9^{th} decile of mean of sim.	0.0004	0.0018	0.0035	0.007	0.0211	0.0439	0.0878	0.1757	0.2635
Std. dev. of market data	0.0113	0.0245	0.0338	0.0476	0.0789	0.1178	0.1673	0.2102	0.2341
Avg. of st.deviation	0.0088	0.0233	0.0333	0.047	0.0791	0.1098	0.1479	0.1974	0.2304
1^{st} decile of std. deviation of sim.	0.0074	0.0193	0.0272	0.0373	0.0581	0.0751	0.0931	0.1168	0.1287
9^{th} decile of std. deviation of sim.	0.0105	0.028	0.0407	0.0581	0.103	0.1502	0.2132	0.3011	0.3631
Skewness of market data	-0.39	-0.97	-1.46	-1.7	-1.49	-1.71	-1.51	-1.52	-1.2
Avg. skewness of sim.	-1.12	-1.19	-1.35	-1.55	-1.77	-1.64	-1.25	-0.79	-0.55
1^{st} decile of skewness of sim.	-1.77	-1.93	-2.22	-2.62	-2.9	-2.72	-2.24	-1.58	-1.24
9^{th} decile of skewness of sim.	-0.63	-0.61	-0.66	-0.78	-0.92	-0.8	-0.48	-0.12	0.08
Kurtosis of market data	9.32	7.77	8.65	8.42	5.25	5.57	3.3	2.66	0.72
Avg. kurtosis of sim.	13.23	10.6	9.87	9.04	7.03	4.81	2.3	0.53	-0.09
1^{st} decile of kurtosis of sim.	5.08	3.51	2.89	2.51	1.69	0.81	-0.1	-0.7	-0.97
9^{th} decile of kurtosis of sim.	24.96	19.96	19.16	17.68	14.49	10.88	6.13	2.33	0.86

Table 3: The first four moments of log-returns distributions for different time horizons in trading days: S&P500 data vs synthetic data generated from the fitted EWMA HM.

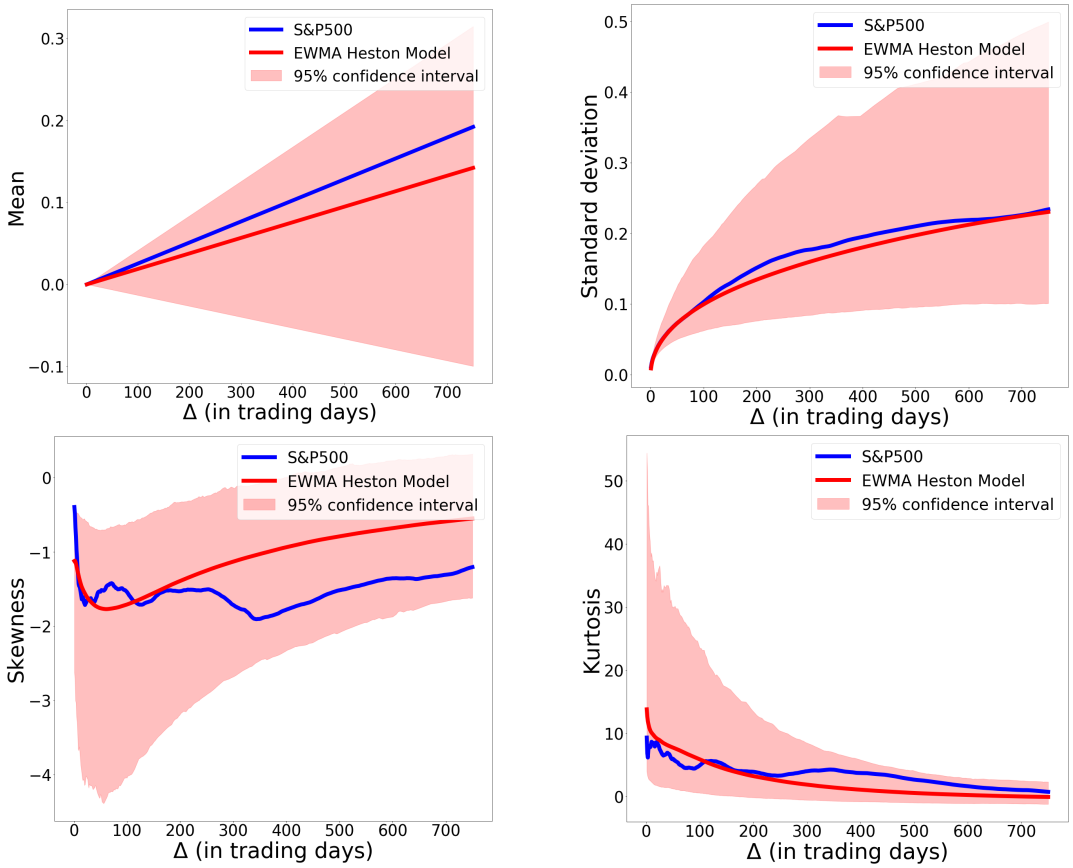

Figure 11: Evolution of the four first moments of the log-returns distributions in the function of the considered time-horizon: S&P500 data vs synthetic data generated from the fitted EWMA HM.

Time horizon	1	5	10	20	60	125	250	500	750
Mean of market data	0.131	0.131	0.131	0.131	0.131	0.131	0.131	0.131	0.131
Avg. mean of sim.	0.129	0.129	0.129	0.129	0.129	0.129	0.129	0.129	0.129
1^{st} decile of mean of sim.	0.115	0.115	0.115	0.115	0.115	0.115	0.115	0.115	0.115
9^{th} decile of mean of sim.	0.145	0.145	0.145	0.145	0.145	0.145	0.145	0.145	0.145
Std. dev. of market data	0.0999	0.0903	0.0869	0.0829	0.0732	0.0651	0.057	0.047	0.0403
Mean of std.dev. of sim.	0.0887	0.0795	0.0771	0.0741	0.0657	0.0559	0.0439	0.0317	0.0253
1^{st} decile of std. dev. of sim.	0.0669	0.0581	0.0557	0.0528	0.0449	0.0366	0.0272	0.0182	0.0139
9^{th} decile of std. dev. of sim.	0.1163	0.1065	0.1038	0.1009	0.0909	0.0796	0.0651	0.0488	0.0402
Skewness of market data	3.42	3.18	3.04	2.88	2.41	2.05	1.55	1.21	0.93
Average skewness of sim.	2.27	2.4	2.32	2.23	1.98	1.67	1.26	0.78	0.54
1^{st} decile of skewness of sim.	1.83	1.53	1.45	1.34	1.12	0.85	0.48	0.07	-0.12
9^{th} decile of skewness of sim.	3.79	3.53	3.49	3.34	3.08	2.72	2.21	1.58	1.24
Kurtosis of market data	19.53	15.47	13.65	11.82	7.97	5.33	2.5	0.7	-0.37
Avg. kurtosis of sim.	12.95	9.64	8.79	7.84	5.75	3.8	1.8	0.27	-0.28
1^{st} decile of kurtosis of sim.	5.1	3	2.49	1.94	1.01	0.14	-0.57	-0.99	-1.18
9^{th} decile of kurtosis of sim.	23.4	18.67	17.78	16.2	12.29	8.78	5.46	2.16	0.64

Table 4: The first four moments of average annualized volatility distributions for different time horizons: market data vs synthetic data generated from the fitted EWMA HM.

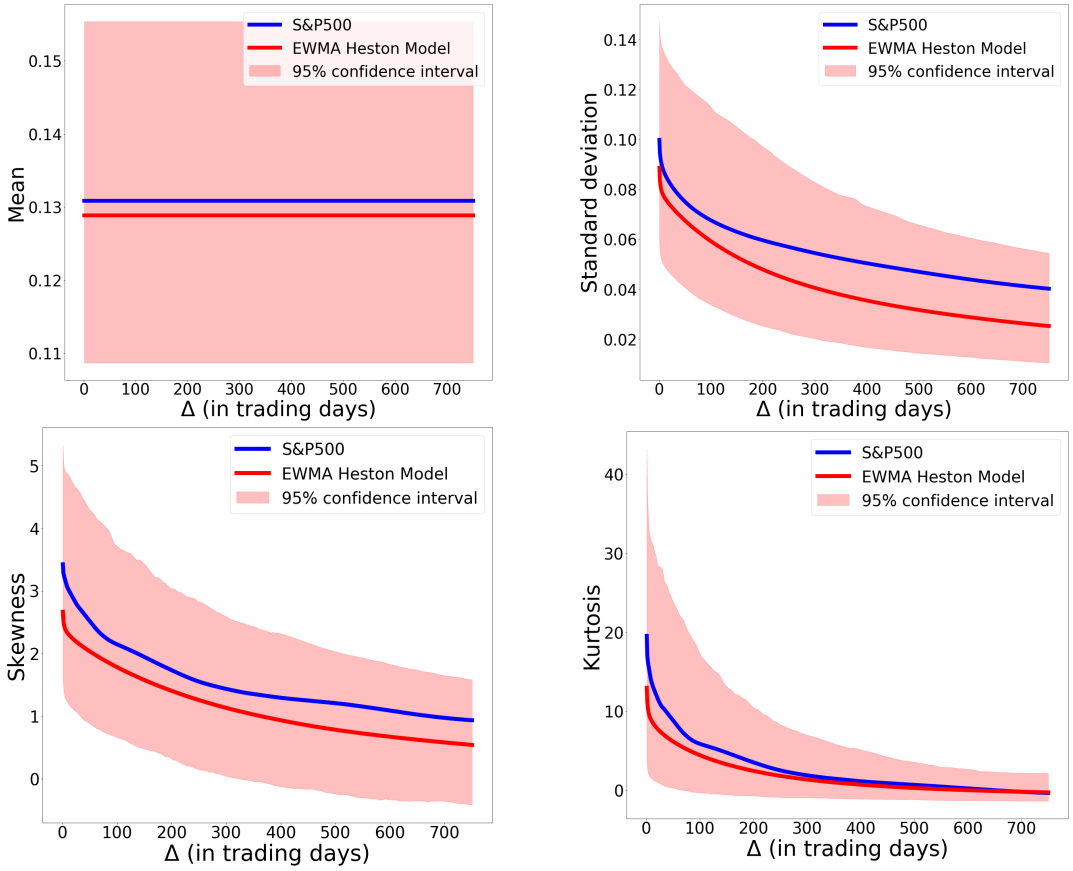

Figure 12: Evolution of the first four moments of average annualized volatility distributions in the function of the considered time-horizon: S&P500 data vs synthetic data generated from the fitted EWMA HM.

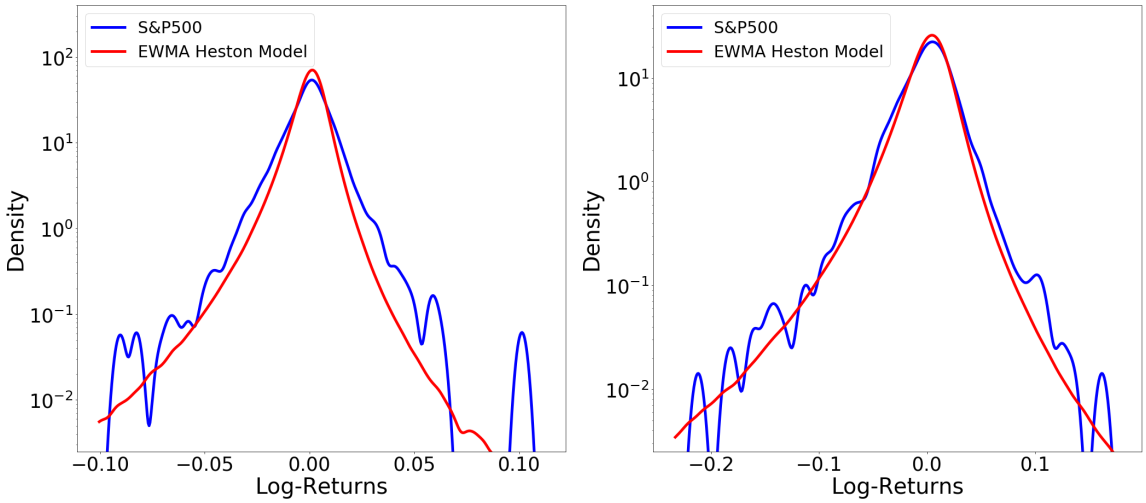

(a) 1 trading day

(b) 5 trading days

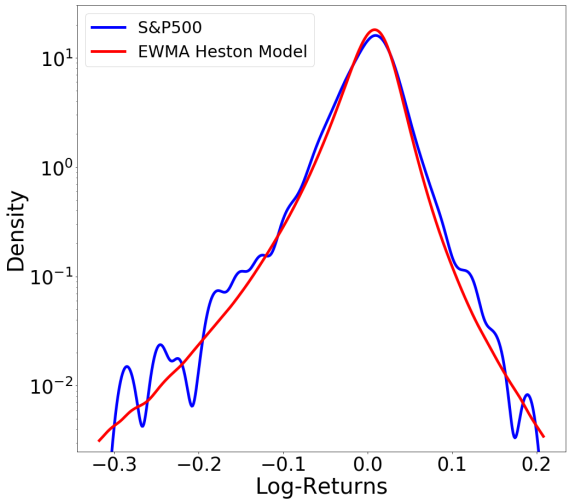

(c) 10 trading days

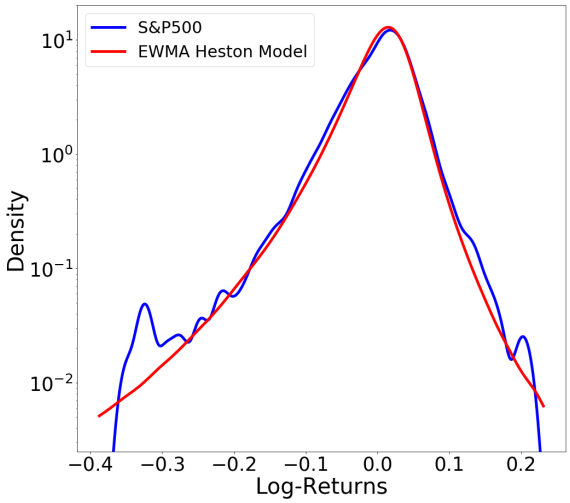

(d) 20 trading days

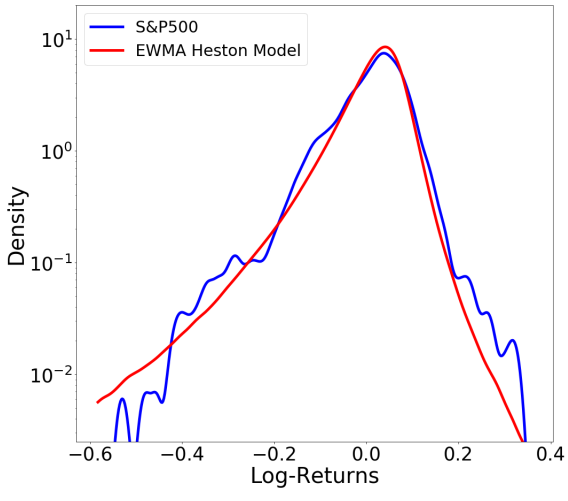

(e) 60 trading days

Figure 13: Log-returns distributions for different time intervals with logarithmic y-axis: S&P500 data vs EWMA Heston Model.

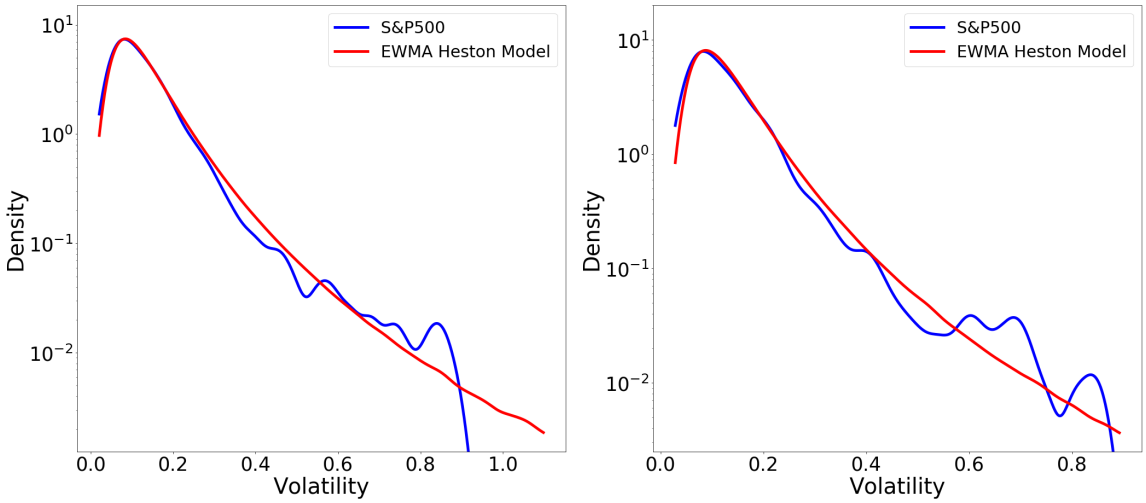

(a) 1 trading day

(b) 5 trading days

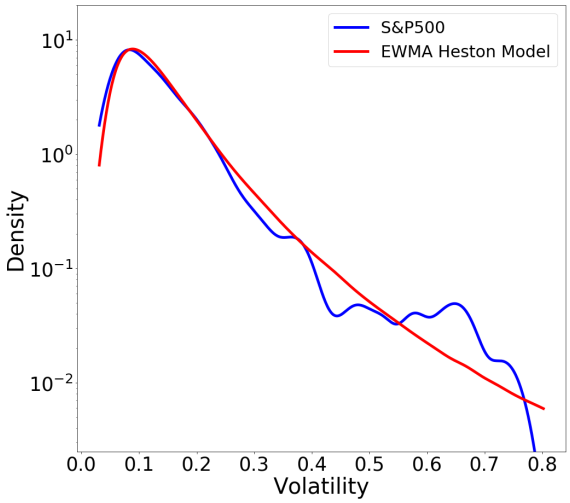

(c) 10 trading days

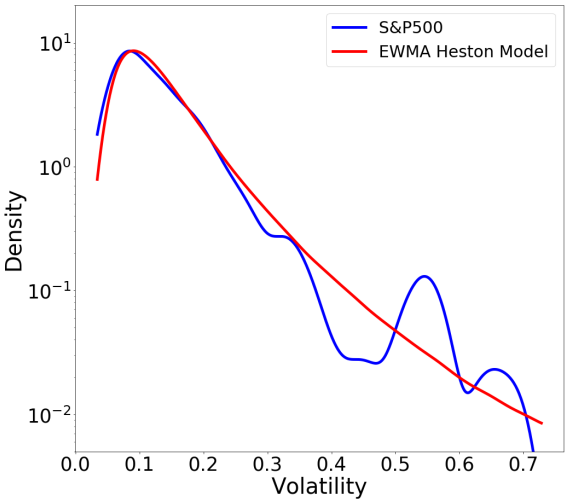

(d) 20 trading days

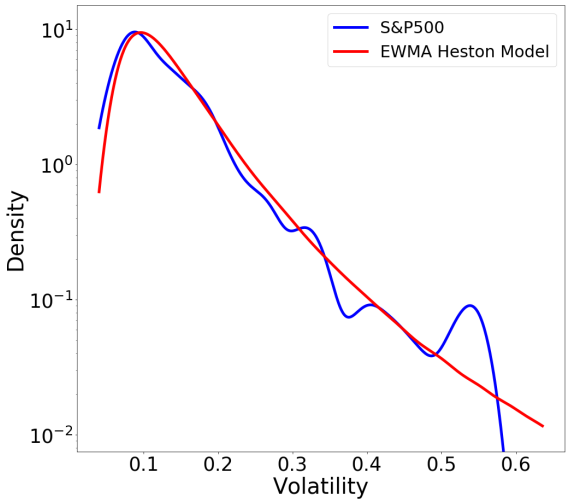

(e) 60 trading days

Figure 14: Annualized average volatility distributions for different time intervals with logarithmic y-axis: S&P500 data vs EWMA Heston Model.

3.4 The time-reversal asymmetry

Finally, we examine the ability of the model to reproduce the empirical time-reversal asymmetry (TRA) of financial time series highlighted Zumbach (Zumbach and Lynch 2001). To assess this aspect of the model, we use the methodology in Blanc *et al.* (2017), and consider the following time asymmetry ratio (TAR) for different lags:

$$\zeta(l) = \frac{\sum_{l'=1}^{l} C(l') - C(-l')}{2\sum_{l'=1}^{l} \max(|C(l')| - |C(-l')|)}$$

where C is the cross-correlation function of the realized volatility and absolute returns

$$C(l) = \frac{\sum_{t=1}^{n} \left(\sigma_t^R - \bar{\sigma}^R\right)\left(|r_{t-l}| - |\bar{r}|\right)}{\sqrt{\sum_{t=1}^{n} \left(|r_t| - |\bar{r}|\right)^2}\sqrt{\sum_{t=1}^{n} \left(\sigma_t^R - \bar{\sigma}^R\right)^2}},$$

where $|r_t|$ and σ_t^R are respectively the absolute return and the realized volatility at time t, $\bar{\sigma}^R$ is the mean of the realized volatility, and $|\bar{r}|$ is the mean of the absolute returns.

Table 5 reports the TARs for both the empirical and synthetic data for different lags:

Lag (in trad. days)	1	2	3	4	5	10	20	60
TAR S&P500 data	0.072	0.047	0.044	0.038	0.032	0.019	0.011	-0.007
Average TAR of sim. data	0.074	0.04	0.026	0.018	0.013	0.004	-0.000	-0.002

Table 5: The TARs for different lags: S&P500 data vs EWMA HM.

The numbers show that the EWMA HM provides a good representation of the empirical TRA. Indeed, as in the case of the S&P500 data, the TAR of the data simulated by the model follows a convex decreasing trend as a function of the lag size. Moreover, the order of magnitude of the TRA generated by the EWMA HM matches the empirical data, and the empirical TAR is within the estimated 95% confidence interval of the model for all considered lags. More broadly, the positive value of the TRA for small lags of these synthetic data confirms that the EWMA HM is able to capture the following aspect of the Zumbach effect: past absolute returns used to forecast future volatilities are more effective than use of past volatilities to forecast future absolute returns. This must be emphasized because it contradicts the idea that "models that use Brownian SDEs are TRS by construction and cannot reproduce this asymmetry" (Blanc *et al.* 2017 p.17). Note also that TAR becomes slightly negative after about 20 trading days for the synthetic data and after about 35 trading days for the empirical data. In the case of our empirical financial time series, we cannot say whether or not this is a structural effect. However, for the data generated by the EWMA HM, the number of lags before the sign of TAR changes depends on the set of parameters used. For instance, for higher values of τ_2, TAR remains positive for longer lags.

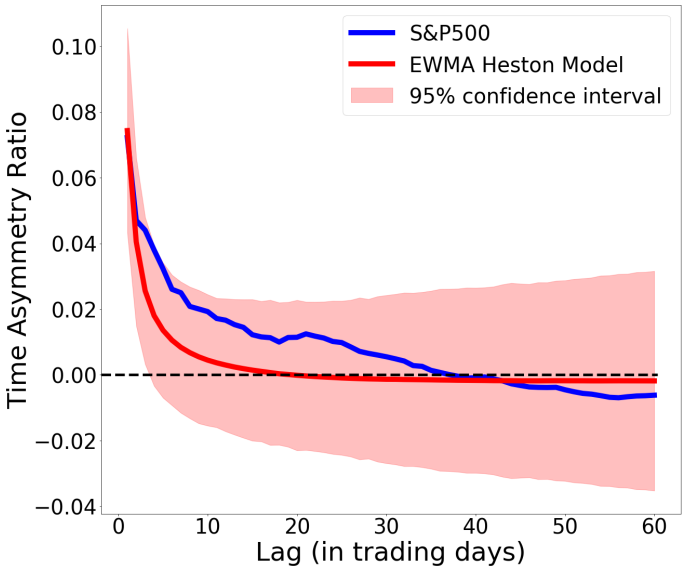

Figure 15: TAR: S&P500 data vs EWMA HM.

4 Conclusion

In this article, we proposed a EWMA HM, a Markovian model which is very tractable for Monte-Carlo simulations and is generally consistent with market data. The model is based on a specific empirical relationship between the EWMA of returns and the realized volatility, which structure makes it naturally adaptable to accounting for a "strong Zumbach effect". We have shown that this model presents some similarities with stochastic volatility models based on quadratic Hawkes processes, models originally conceived to capture this effect (Blanc *et al.* 2017, Dandapani *et al.* 2019, Gatheral *et al.* 2020). Besides these theoretical aspects, we show that based on simulations using a Euler discretization scheme, several aspects of the model are consistent with empirical stock market data. First, it generates realistic evolutions of an asset price combined with its associated volatility, reproducing the empirical irregular behavior of volatility paths occurring over short time scales. We show also that both log-returns and the model volatility distributions are consistent with the empirical results. In particular, the model is able to capture the respective deformations of these distributions in the function of the time horizon. Finally, the simulations show that the model also reproduces the empirical time-reversal asymmetry of financial time series.

However, despite all these strong qualities, the model has some shortcomings. First, in the model serial correlation of the volatility process seems to decay much too quickly compared to the empirical data. Relatedly, the type of relationship linking the expected log-volatility increments and the considered time scale differs between the model generated data and the market data. These two limitations both stem from the exponential memory of the volatility process in the model compared to an empirical volatility process based on a power-law memory (Gatheral *et al.* 2018). However, these shortcomings could be overcome by using a combination of several EWMAs of returns in order to mimic a long-memory property. This potential extension could be conducted using the more general EWMA HM framework and should be the focus of future work.

Another limitation is that we focus only on use of the EWMA HM as a time series generator for Monte Carlo simulations. Its use for pricing issues needs investigation along with deeper study of the model's mathematical properties. We also identified a statistical regularity related to the slope coefficient between the EWMAs of past returns and realized volatility, a statistical regularity which might indicate the microstructural foundation of the EWMA HM. This is also worthy of further investigation and substantive consideration. Thus, our paper offers some directions for further research and a new volatility modeling agenda.

5 Acknowledgments

I would like to thank Jean-Paul Laurent and Thierry Roncalli for useful comments and fruitful discussions.

References

[1] Maruyama G. (1955). Continuous Markov processes and stochastic equations. *Rendiconti del Circolo Matematico di Palermo*, 4:48–90.

[2] Engle R.F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica: Journal of the econometric society*, 987-1007.

[3] De Bondt W. F., & Thaler R. H. (1987). Further Evidence on Investor Overreaction and Stock Market Seasonality. *Journal of Finance*, 42(3), 557-581.

[4] Harvey A.C. (1990). *Forecasting, Structural Time Series Models and the Kalman Filter.*, Institute of Mathematical Statistics.

[5] Nelson D.B. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica: Journal of the Econometric Society*, 347-370.

[6] Heston S.L. (1993). A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options. *The Review of Financial Studies*, 6(2), 327-343.

[7] Ding Z., Granger C.W., and Engle R.F. (1993). A long memory property of stock market returns and a new model. *Journal of empirical finance*, 1(1), 83-106.

[8] Bollerslev T., Engle R.F., and Nelson D.B. (1994). ARCH models. *Handbook of econometrics*, 4, 2959-3038.

[9] Baillie R.T., Bollerslev T., and Mikkelsen, H. O. (1996). Fractionally integrated generalized autoregressive conditional heteroskedasticity. *Journal of econometrics*, 74(1), 3-30.

[10] Dissanaike G. (1997). Do Stock Market Investors Overreact? *Journal of Business Finance & Accounting*, 24(1), 27-50.

[11] Hobson D.G., and Rogers L.C. (1998). Complete Models with Stochastic Volatility. *Mathematical Finance*, 8(1), 27-48.

[12] Zumbach G., and Lynch P. (2001). Heterogeneous Volatility Cascade in Financial Markets. *Physica A: Statistical Mechanics and its Applications*, 298(3-4), 521-529.

[13] Cont R. (2001). Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative finance*, 1(2), 223.

[14] Drăgulescu A.A., and Yakovenko V.M. (2002). Probability Distribution of Returns in the Heston Model with Stochastic Volatility. *Quantitative finance*, 2(6), 443-453.

[15] Mandelbrot B.B. (2003). Heavy tails in finance for independent or multifractal price increments. In *Handbook of heavy tailed distributions in finance* (pp. 1-34). North-Holland.

[16] Lynch P.E., and Zumbach, G.O. (2003). Market heterogeneities and the causal structure of volatility. *Quantitative Finance*, 3(4), 320.

[17] Lanne M., and Saikkonen P. (2005). Non-linear GARCH models for highly persistent volatility. *The Econometrics Journal*, 8(2), 251-276.

[18] Mandelbrot B.B., and Hudson R. L. (2005). *The (mis)behavior of markets: a fractal view of risk, ruin, and reward.* Basic Books.

[19] Borland L., and Bouchaud J. P. (2005). On a multi-timescale statistical feedback model for volatility fluctuations. *arXiv* preprint physics/0507073.

[20] Brandt M.W. and Jones C.S. (2006). Volatility forecasting with range-based EGARCH models. *Journal of Business & Economic Statistics*, 24(4), 470-486.

[21] Brockwell P., Chadraa E., and Lindner A. (2006). Continuous-time GARCH processes. *The Annals of Applied Probability*, 16(2), 790-826.

[22] Bochud T. and Challet D. (2007). Optimal approximations of power laws with exponentials: application to volatility models with long memory. *Quantitative Finance*, 7(6), 585-589.

[23] Albrecher H., Mayer P., Schoutens W., and Tistaert J. (2007). The Little Heston Trap. *Wilmott*, (1), 83-92.

[24] Maller R.A., Müller G. and Szimayer A. (2008). GARCH modelling in continuous time for irregularly spaced time series data. *Bernoulli*, 14(2), 519-542.

[25] Corsi F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics*, 7(2), 174-196.

[26] Zumbach G. (2009). Time Reversal Invariance in Finance. *Quantitative Finance*, 9(5), 505-515.

[27] Zumbach G. (2010), Volatility Conditional on Price Trends. *Quantitative Finance* 10.4 (2010): 431-442.

[28] Lord R., Koekkoek R., and Dijk D.V. (2010). A comparison of biased simulation schemes for stochastic volatility models. *Quantitative Finance*, 10(2), 177-194.

[29] Bruder B., and Gaussel N. (2011). Risk-Return Analysis of Dynamic Investment Strategies. Available at SSRN 2465623.

[30] Belkhouja M., and Boutahary M. (2011). Modeling volatility with time-varying FIGARCH models. *Economic Modelling*, 28(3), 1106-1116.

[31] Chicheportiche R., and Bouchaud J.P. (2014). The fine-structure of volatility feedback I: Multi-scale self-reflexivity. *Physica A: Statistical Mechanics and its Applications*, 410, 174-195.

[32] Guyon J. (2014). Path-dependent volatility. *Risk Magazine*

[33] Bennedsen M., Lunde A., and Pakkanen M.S. (2016). Decoupling the short-and long-term behavior of stochastic volatility. arXiv preprint arXiv:1610.00332.

[34] Blanc P., Donier J., and Bouchaud J.P. (2017). Quadratic Hawkes Processes for Financial Prices. *Quantitative Finance*, 17(2), 171-188.

[35] Goutte S., Ismail A., and Pham H. (2017). Regime-Switching Stochastic Volatility Model: Estimation and Calibration to VIX Options. *Applied Mathematical Finance*, 24(1), 38-75.

[36] Jusselin P., Lezmi E., Malongo H., Masselin C., Roncalli, T., and Dao T.L. (2017). Understanding the Momentum Risk Premium: An In-Depth Journey Through Trend-Following Strategies. Available at SSRN 3042173.

[37] Mrázek M., and Pospíšil J. (2017). Calibration and simulation of Heston model. *Open Mathematics*, 15(1), 679-704.

[38] El Euch O., (2018). *Quantitative Finance under rough volatility*, PhD diss., Sorbonne université.

[39] Gatheral J., Jaisson T., and Rosenbaum M. (2018). Volatility is Rough. *Quantitative Finance*, 18(6), 933-949.

[40] El Euch O., and Rosenbaum M. (2018). Perfect Hedging in Rough Heston Models. *Annals of Applied Probability*, 28(6), 3813-3856.

[41] El Euch O. and Rosenbaum M. (2019). The Characteristic Function of Rough Heston Models. *Mathematical Finance*, 29(1), 3-38.

[42] Abi Jaber E. and El Euch O. (2019). Multifactor approximation of rough volatility models. *SIAM Journal on Financial Mathematics*, 10(2), 309-349.

[43] Abi Jaber E. (2019). Lifting the Heston model. *Quantitative Finance*, 19(12), 1995-2013.

[44] Dandapani A., Jusselin P., and Rosenbaum M. (2019). From Quadratic Hawkes Processes to Super-Heston Rough Volatility Models with Zumbach Effect. *arXiv* preprint arXiv:1907.06151, 2019.

[45] Rogers L.C.G. (2019). Things we think we know. Preprint, available at *https://www. skokholm. co. uk/lcgr/downloadable-papers.*

[46] Gatheral J., Jusselin P., and Rosenbaum M. (2020). The Quadratic Rough Heston Model and the Joint S&P500/VIX Smile Calibration Problem. *arXiv* preprint arXiv:2001.01789.

[47] Cont R., and Das P. (2022). Rough volatility: fact or artefact?. *arXiv* preprint arXiv:2203.13820.

Appendix A Convergence of the attraction volatility when β is 0

Let us consider the following differential equation

$$\mathrm{d}\nu_t = \left(\frac{\psi}{\nu_t - \underline{\nu}} - \nu_t + \underline{\nu} + \alpha\right)\frac{\mathrm{d}t}{\tau_2}$$

such as $\psi, \tau_2, \underline{\nu} \in \mathbb{R}_+$, $\alpha \in \mathbb{R}$, and $\nu > \underline{\nu}$.

If this differentiel equation converges as $\lim_{t\to+\infty} \nu_t = \nu^\star < +\infty$, thus $\dfrac{\mathrm{d}\nu^\star}{\mathrm{d}t} = 0$. Let us start by noticing the following equalities:

$$\begin{aligned}
0 &= \frac{\mathrm{d}\nu^\star}{\mathrm{d}t} \\
0 &= \frac{\psi}{\nu_t - \underline{\nu}} - \nu_t + \underline{\nu} + \alpha \\
0 &= \nu_t^2 - \big(2\underline{\nu} + \alpha\big)\nu_t + \underline{\nu} + \alpha\underline{\nu} - \psi
\end{aligned}$$

Therefore:

$$\nu^\star = \frac{2\underline{\nu} + \alpha \pm \sqrt{\alpha^2 + 4\psi}}{2}$$

Moreover, because $\nu_0 > \underline{\nu}$, the following inequalities must be respected:

$$\begin{aligned}
2\underline{\nu} &< 2\underline{\nu} + \alpha \pm \sqrt{\alpha^2 + 4\psi} \\
-\alpha &< \pm\sqrt{\alpha^2 + 4\psi}
\end{aligned}$$

Since $\psi, \tau \in \mathbb{R}_+$, therefore $(\pm) = +$. It follows:

$$\nu^\star = \frac{2\underline{\nu} + \alpha + \sqrt{\alpha^2 + 4\psi}}{2}$$

Moreover, $\dfrac{\mathrm{d}\nu^\star}{\mathrm{d}t}$ is a continuous function positive on $\left]\underline{\nu} : \nu^\star\right]$ and negative on $\left[\nu^\star : +\infty\right[$. Consequently:

$$\lim_{t\to+\infty} \nu_t = \frac{2\underline{\nu} + \alpha + \sqrt{\alpha^2 + 4\psi}}{2}$$

It follows that if β is 0, the EWMA HM degenrates into the following standard HM:

$$\begin{cases}
\dfrac{\mathrm{d}S_t}{S_t} &= \mu_t \mathrm{d}t + \sqrt{V_t}\mathrm{d}W_t \\
\mathrm{d}V_t &= \dfrac{1}{\tau_1}\big(\nu_t^2 - V_t\big)\mathrm{d}t + \bar{\xi}\sqrt{V_t}\mathrm{d}B_t,
\end{cases}$$

with

$$\begin{cases}
\theta &= \dfrac{\left(2\underline{\nu} + \alpha + \sqrt{\alpha^2 + 4\psi}\right)^2}{4} \\
\bar{\xi} &= \xi\nu_t.
\end{cases}$$

Appendix B The model calibration procedure

To obtain numerical results and enable comparison with empirical data, we need to fit the model from H, the daily empirical time-series of the S&P500 and its realized volatility from July 23, 2001, to July 23 2021. This involves identifying the following parameters

$$\theta_{\mathcal{G}} = (\lambda, \rho, \tau_1, \xi, \tau_2, \bar{\nu}, \alpha, \beta).$$

Due to the relative complexity of the model, obtaining these parameters using standard statistical methods such as maximum likelihood estimation is tricky. For this reason, to calibrate the model we use an ad-hoc procedure involving a neural network (NN), with NN denoted $\mathcal{P}$. This NN, which is used to provide an estimator of the vector of the parameters, takes as its input a vector encoding the index path and its volatility. In other words, the market data matrix H passes through an encoder denoted $\mathcal{E}$ which transforms it into an input vector E, and the elements of this vector then constitute the input layer of $\mathcal{P}$ which passes through the network to provide an estimate of $\theta_{\mathcal{G}}$. In summary:

$$\mathcal{P}\Big(\underbrace{\mathcal{E}(\mathrm{H})}_{\mathrm{E}}; \theta_{\mathcal{P}}\Big) = \widehat{\theta}_{\mathcal{G}}$$

with $\theta_{\mathcal{P}}$ the parameters of $\mathcal{P}$, and $\widehat{\theta}_{\mathcal{G}}$ the estimator of $\theta_{\mathcal{G}}$.

More specifically, the encoder $\mathcal{E}$ extracts from the time-series the following features:

- the 4 first moments of the log-returns for the following lags in trading days: 1, 5, 20, 60, 125,
- the 4 first moments of the average volatility over the following time intervals in trading days: 1, 5, 20, 60, 125,
- the serial corellation for the following lags expressed in trading days: 1, 2, 3, 4, 5, 10, 20, 60, 125, 250.

Thus, the dimension of E is thus equal to 1×30.

Obviously, the NN responsible for providing the model parameters derived from the encoded market data requires to be calibrated from a learning set. For this purpose, we randomly generated 20 000 parameter vectors such that each parameter was obtained using a uniform distribution within the bounds defined in table 6:

	λ	ρ	τ_1	ξ	τ_2	$\bar{\nu}$	α	β
Lower bound	0.2	-1	0.0008	$0.05\nu\sqrt{\frac{2}{\tau_1}}$	0.15	0.01	0.01	$0.9\sqrt{\frac{2\tau}{\pi}}$
Upper bound	0.5	0	2	$\nu\sqrt{2\frac{2}{\tau_1}}$	1	0.1	0.15	$1.1\sqrt{\frac{2\tau}{\pi}}$

Table 6: The respective bounds of the uniform distributions used to generate the training set of parameter vectors.

The bounds of ξ allow that all generated vectors of parameters respect the Feller condition.

With each of the 20 000 vectors of parameters obtained, we generate a sample path of the index and its volatility from $\mathcal{G}$. We thus pass through the encoder $\mathcal{E}$ each of the 20 000 synthetic time-series, in order to obtain the following set $\left\{E^{(i)}\right\}_{1\leq i\leq 20000}$. This step provides we have our learning set $\left\{\left(E^{(i)}, \theta_{\mathcal{G}}^{(i)}\right)\right\}_{1\leq i\leq 20000}$ necessary to fit $\mathcal{P}$. The loss function used for this purpose is:

$$\mathcal{L}\left(\theta_{\mathcal{G}}^{(i)}, \widehat{\theta}_{\mathcal{G}}^{(i)}\right) = \left\| w \odot \left(\theta_{\mathcal{G}}^{(i)} - \widehat{\theta}_{\mathcal{G}}^{(i)}\right)\right\|_2^2 ,$$

where w is a weighting vector, such that the value of the k^{th} coordinate is equal to the inverse of the standard deviation of $\left\{\left(\theta_{\mathcal{G}}^{(i)}\right)_k\right\}_{1\leq i\leq 20000}$.

Given that, $\mathcal{P}$ is fitted by solving the following optimization problem:

$$\theta_{\mathcal{P}}^{\star} = \underset{\theta_{\mathcal{P}}\in\Theta_{\mathcal{P}}}{\arg\min} \sum_{i=1}^{20000} \left\| w \odot \left(\theta_{\mathcal{G}}^{(i)} - \mathcal{P}\big(E_i;\theta_{\mathcal{P}}\big)\right)\right\|_2^2$$

The NN being fitted, we calibrate finally $\theta_{\mathcal{G}}$ from market data:

$$\mathcal{P}\Big(\mathcal{E}(\mathrm{H});\theta_{\mathcal{P}}^{\star}\Big)\Big) = \widehat{\theta}_{\mathcal{G}}$$

We obtain the following parameters[10]:

λ	ρ	τ_1	ξ	τ_2	$\underline{\nu}$	α	β
0.5575	-0.465	0.0013	42.95	0.276	0.0595	0.1033	$\sqrt{\frac{2\tau_2}{\pi}}$

Table 7: The parameters obtained from the calibration procedure.

Appendix C Standard deviation of an EWMA of a Brownian motion

The asymptotic variance of the EWMA of a Brownian motion is given by the following equalities:

$$\begin{aligned}\mathrm{Var}\left(\frac{1}{\tau}\int_{-\infty}^{t} e^{\frac{1}{\tau}(u-t)}\mathrm{d}W_u\right) &= \frac{1}{\tau^2}\int_{-\infty}^{t} e^{\frac{2}{\tau}(u-t)}\mathrm{d}u \\ &= \frac{1}{\tau^2}\left[\frac{\tau}{2}e^{\frac{2}{\tau}(u-t)}\right]_{-\infty}^{t} \\ &= \frac{1}{2\tau}\end{aligned}$$

It follows that the asymptotic standard deviation of the EWMA of a Brownian motion is equal to $\sqrt{\frac{1}{2\tau}}$.

[10] The exact value of β obtained from the neural network was equal to 0.422. This result is approximated in table 7 by $\sqrt{\frac{2\tau_2}{\pi}}$.

Appendix D Complementary results

The type of relationship between the EWMA of past returns and the volatility described by the EWMA HM is not unique to the S&P500 and can apply also to Nasdad and Euro Stoxx 50 data[11]. Also, the optimal slope coefficient β is very close to $\sqrt{\frac{2\tau}{\pi}}$ as in the S&P500. The presence of this specific value coefficient in different market data sets would seem to suggest a fundamental reason underlying this statistical regularity. Thus, it is reasonable to consider that the EWMA HM might have microstructural foundations based on the interactions of the market participants at the high-frequency scale.

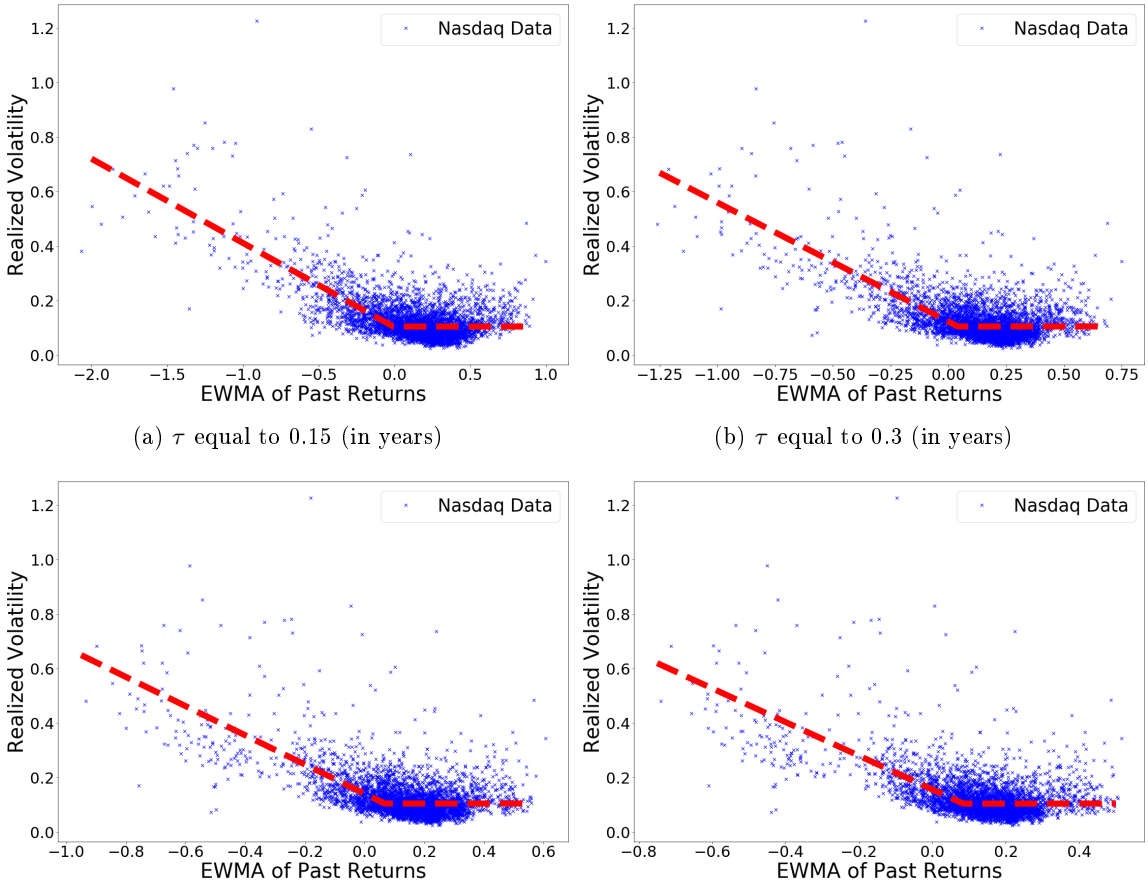

(a) τ equal to 0.15 (in years)

(b) τ equal to 0.3 (in years)

(c) τ equal to 0.45 (in years)

(d) τ equal to 0.6 (in years)

Figure 16: The empirical relationship between the realized volatility of the S&P500 and the EWMAs of its past returns for different values of τ expressed in years. The red lines are regressions of the form $\sqrt{V_t} = \underline{\nu} + (\alpha - \beta m_t)_+$ such as the associated regression lines have all a slope coefficient equal to $\sqrt{\frac{2\tau}{\pi}}$. These regressions are respectively associated with the following R-squared: 0.476, 0.434, 0.405, 0.384.

[11] As in the S&P500 example, the data sets cover the period July 23 2001 to July 23 2021, and as an estimator we use volatility, and the square root of the realized variance computed from 5-min samples provided by the Oxford-Man Institute of Quantitative Finance.

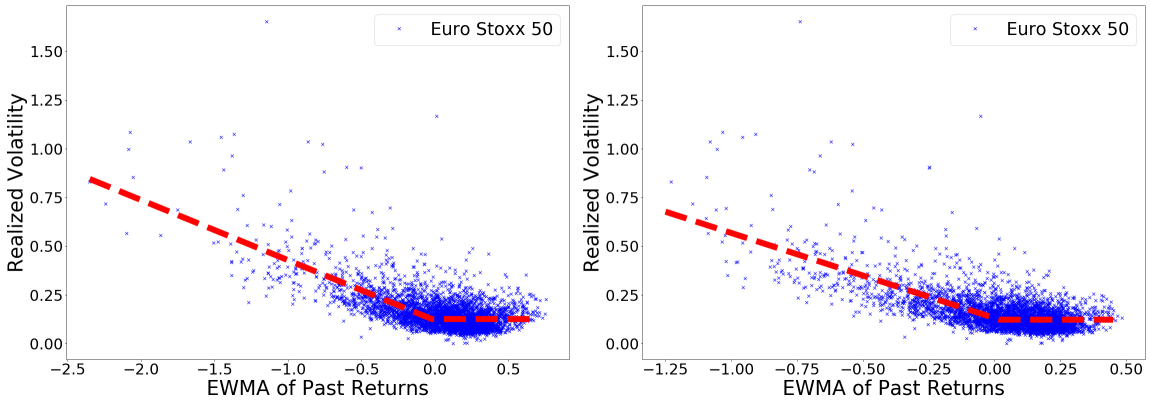

(a) τ equal to 0.15 (in years)

(b) τ equal to 0.3 (in years)

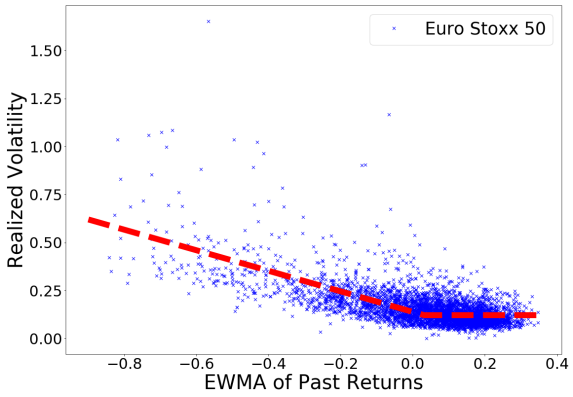

(c) τ equal to 0.45 (in years)

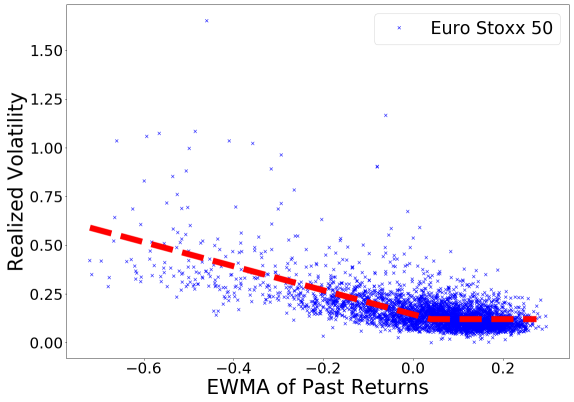

(d) τ equal to 0.6 (in years)

Figure 17: The empirical relationship between the realized volatility of the S&P500 and the EWMAs of its past returns for different values of τ expressed in years. The red lines are regressions of the form $\sqrt{V_t} = \underline{\nu} + (\alpha - \beta m_t)_+$ such as the associated regression lines have all a slope coefficient equal to $\sqrt{\frac{2\tau}{\pi}}$. These regressions are respectively associated with the following R-squared: 0.476, 0.466, 0.445, 0.427.