Investigating Approaches to Modeling Rough Path-Dependent Volatility: Insights and Implications

Léo Parent

Investigating Approaches to Modeling Rough Path-Dependent Volatility: Insights and Implications

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

September 15, 2023

Abstract

This article investigates different approaches to modeling rough path-dependent volatility (RPDV) and their consistency with market data. Accordingly, we proceed with a comparative study of various RPDV model specifications based on the realized volatility data of five stock indices. We then discuss the implications of these results in terms of modeling choices, notably regarding whether the volatility process needs to depend on a historical volatility factor and how to integrate the exogenous component of volatility. Then, through numerical simulations, we examine the properties of volatility dynamics generated by a specific RPDV model. Based on these experiments, we demonstrate that, on the one hand, this type of modeling can replicate most of the properties that characterize empirical volatility dynamics. On the other hand, we show that the apparent positive volatility feedback observed in market data can be explained, at least partially, by the effect of the exogenous component of volatility.

Keywords: Volatility modeling, path-dependent volatility, rough volatility.

JEL classification: C12, C15, C25, C32, C35, G10, G12, G13, G14, G17.

1 Introduction

The rough behavior of volatility and the dependence of volatility on the historical price path are two major characteristics of financial time-series widely documented in the academic literature ([6], [16], [14]). Due to their significance, many volatility models aimed at jointly capturing these two phenomena have emerged in recent years ([4], [7], [10]). Furthermore, a substantial portion of them are included in the family of RPDV models introduced in the eponymous article *Rough Path-Dependent Volatility Models* ([13]), among which the well-known rough Heston model ([5]) and the quadratic rough Heston model ([7]) are included. More broadly, this class of models constitutes a general framework that structurally encompasses a wide range of modeling approaches for capturing the rough behavior and the path-dependence of volatility.

If the abundance of approaches included in this framework can be an advantage, for instance in reducing model risk by employing various modeling approaches for volatility dynamics, it nevertheless raises an important question: what are the right ways to model rough path-dependent volatility? This article aims to provide insights into answering this question. Of course, given the magnitude of the task, this article does not claim to provide a comprehensive answer to this question but rather considers the consistency of certain modeling choices included in the framework of RPDV models with empirical data.

The paper is organized as follows. Section 2 involves a comparative study of various RPDV model specifications based on the historical volatility of 5 stock indices. In this section, we analyze the respective performance of these different specifications in explaining realized volatility and discuss the implications of these results in

terms of modeling choices. The methodology used here is largely inspired by that employed in [10]. Section 3 examines, through numerical experiments, the properties of volatility dynamics generated from a parsimonious form of RPDV model, which includes an exogenous source of randomness but no historical volatility factor. Through these experiments, we demonstrate that, on one hand, this type of model can replicate most of the properties that characterize empirical volatility dynamics. On the other hand, we show that the apparent positive volatility feedback observed in market data can be explained, at least partially, by the effect of the exogenous component of volatility.

2 Comparison of model specifications

The objective of this section is to assess the explanatory power of different RPDV model specifications on historical volatility. To achieve this, we adopt a comparative approach inspired by Guyon and Lekeufack ([10]), where we evaluate the performance of discrete-time versions of RPDV models in explaining spot realized volatility. The dataset used in this study consists of historical data from five stock market indices: the S&P500, Nasdaq, FTSE, DAX, and Euro Stoxx 50, covering the period from January 2000 to June 2021.

2.1 The class of rough path-dependent volatility models

All the models considered in this article belong to the family of RPDV models introduced in [13], defined by the following system of equations:

$$\begin{cases} \dfrac{\mathrm{d}P_t}{P_t} & = \mu_t \mathrm{d}t + \sigma_t \mathrm{d}B_t, \\ (\sigma_t)^p & = \beta_0 + \beta_1^{(+)} \left(m_{1,t} - \bar{m}_1\right)_+^{a_1} + \beta_1^{(-)} \left(\bar{m}_1 - m_{1,t}\right)_+^{a_1} + \beta_2 (m_{2,t})^{p/a_2}, \\ m_{1,t} & = \displaystyle\int_{I_t} (t-u)^{-\alpha_1} \frac{\mathrm{d}P_u}{P_u} + \kappa_1 \int_{I_t} (t-u)^{-\alpha_1} \left(\theta_{1,u} - m_{1,u}\right) \mathrm{d}u, \\ m_{2,t} & = \displaystyle\int_{I_t} (t-u)^{-\alpha_2} (\sigma_u)^{a_2} \mathrm{d}u + \kappa_2 \int_{I_t} (t-u)^{-\alpha_2} \left(\theta_{2,u} - m_{2,u}\right) \mathrm{d}u, \end{cases} \tag{1}$$

with B a Brownian motion, I_t the time interval of integration with an upper bound less than or equal to t, where $p, a_1, a_2 \in \{1, 2\}$, and such as θ_j are $\mathcal{F}_t$-measurable processes.

The first equation of the system, which is quite standard in finance, describes the dynamics of the asset price P. The second equation of the system defines the volatility process σ through a very general expression. Specifically, when the exponent p in the second equation is set to 1, the right-hand side of the equation defines volatility; when p is set to 2, the right-hand side defines the variance process. In fact, this choice of the exponent's value, p, implies two different ways of modeling the volatility process that coexist in the academic literature. Regardless of the specification of p, in this family of models, volatility depends on two stochastic processes: m_1 and m_2. The process m_1 corresponds to a past price trend variable and can be understood as a type of weighted moving average of past returns. The variable m_2 capture the asset price activity regardless of its trend or a historical volatility factor. It can be viewed as a moving average of the volatility process either $k = 1$, or of the variance process σ^2 if $k = 2$. In addition, if θ_1 and θ_2 are deterministic functions, the volatility process is a function of a unique source of randomness B: the Brownian motion driving the asset price. In this case, the volatility is then purely path-dependent. Conversely, the θ_j functions can also be stochastic processes linked to their own source of randomness. The model is then only partially path-dependent.

Thus, thanks to its very general form, the family of RPDV models defined by 1 offers a wide variety of approaches to jointly model the rough behavior and the path-dependence of volatility.

2.2 The tested RPDV models

2.2.1 The building blocks of models

To enhance the clarity and understanding of the different RPDV models tested in this section, we begin by breaking down their components into variables that serve as the elementary building blocks of the tested models. Additionally, we provide a discrete approximation of these variables, enabling their computation from empirical time series[1]. These elements are reported in table 1.

	Expression	**Approximated expression**
$\mathbf{R_{1,t}}$	$\int_{-\infty}^{t-\epsilon}(t-u)^{-\alpha_1}\frac{\mathrm{d}P_u}{P_u}$	$\sum_{i=1}^{p}\Big(t-0.5(t_i+t_{i-1})\Big)^{-\alpha_1}\frac{P_{t_i}-P_{t_{i-1}}}{P_{t_{i-1}}}$
$\mathbf{\Theta_{1,t}}$	$\int_{-\infty}^{t-\epsilon}(t-u)^{-\alpha_1}\theta_u\mathrm{d}u$	$\sum_{i=1}^{p}\big(t-0.5(t_{i-1}+t_i)\big)^{-\alpha_1}\theta_{t_{i-1}}(t_i-t_{i-1})$
$\mathbf{S_{1,t}}$	$\int_{-\infty}^{t-\epsilon}(t-u)^{-\alpha_1}\sigma_u\mathrm{d}u$	$\sum_{i=1}^{p}\big(t-0.5(t_{i-1}+t_i)\big)^{-\alpha_1}\sigma_{t_{i-1}}(t_i-t_{i-1})$
$\mathbf{S_{2,t}}$	$\int_{-\infty}^{t-\epsilon}(t-u)^{-\alpha_2}\sigma_u\mathrm{d}u$	$\sum_{i=1}^{p}\big(t-0.5(t_{i-1}+t_i)\big)^{-\alpha_2}\sigma_{t_{i-1}}(t_i-t_{i-1})$
$\mathbf{V_{2,t}}$	$\int_{-\infty}^{t-\epsilon}(t-u)^{-\alpha_2}(\sigma_u)^2\mathrm{d}u$	$\sum_{i=1}^{p}\big(t-0.5(t_{i-1}+t_i)\big)^{-\alpha_2}(\sigma_{t_{i-1}})^2(t_i-t_{i-1})$

Table 1: The building blocks of models and their numerical quadratures.

Several important remarks should be made here. First, the elementary building blocks defined in table 1 are integrals over $]-\infty : t-\epsilon]$. Here, ϵ is a small positive constant close to zero, representing a latency in the transmission of the dynamics of the explanatory variables to the volatility level. However, the reason for choosing an upper bound of $t-\epsilon$ instead of t is primarily technical. In the case where $\epsilon = 0$, the values of α are constrained to be less than 0.5 to avoid inducing divergence in the volatility process. By choosing a strictly positive value for ϵ, this constraint can be relaxed. In this regard, it can be seen as an alternative to the shifted power-law.

Another important point to explain concerns the choice of approximating the integrals that constitute the elementary building blocks of the tested models. In practice, these integrals are truncated to a length of 1260 trading days (approximately 5 years) and then estimated using a Riemann method. This type of approximation is employed, for instance, by Gatheral to approximate the integral defining the Rough Fractional Stochastic Volatility ([6]) model[2]. If all the different building blocks of the models are approximated using Riemann sums, the quadrature method used to approximate $R_{1,t}$ differs from the method used for the other variables. Specifically, the midpoint method is employed for $R_{1,t}$, while the left rectangle method is used for the remaining variables. If the choice of the left rectangle method for approximating Θ_t, $R_{1,t}$, $S_{1,t}$, $S_{2,t}$, and $V_{2,t}$ may seem peculiar at first glance, it is, in fact, driven by the objective of avoiding any tautological or quasi-tautological formulations. Thus, contrary to the midpoint method, the chosen numeral quadrature ensures that the estimation of σ_t does not depend on itself or a proxy of itself. This difference in the approximation method also implies that only the building block R_{1,t_i} requires information known after $t-1$: the price at time t, P_t. On the other hand, all other discretized variables depend solely on elements known at time $t-1$.

[1] Refer to the details in appendix A.

[2] For more details, refer to the section "Implement variance forecast in Python" of the article "Rough Volatility with Python," accessible at `https://tpq.io/p/rough_volatility_with_python.html`.

2.2.2 The tested models

Given the previously mentioned flexibility of the RPDV model, conducting a comprehensive study of all possible specifications is impractical. Instead, our focus lies in examining the main RPDV models discussed in academic literature, along with other interesting specifications, including those addressed in [13]. Therefore, we consider the model specifications reported in table 2.

	Model	**Free paramaters**
M.1	$\sqrt{b_0 + b_1 R_{1,t}}$	$b_0, \alpha_1 \in \mathbb{R}_+, \quad b_1 \in \mathbb{R}$
M.2	$b_0 + b_1 R_{1,t}$	$b_0, \alpha_1 \in \mathbb{R}_+, \quad b_1 \in \mathbb{R}$
M.3	$\sqrt{b_0 + b_1 \big(R_{1,t} - \bar{R}_1\big)^2}$	$b_0, \alpha_1 \in \mathbb{R}_+, \quad b_1, \bar{R}_1 \in \mathbb{R}$
M.4	$b_0 + b_1 \big\lvert R_{1,t} - \bar{R}_1 \big\rvert$	$b_0, \alpha_1 \in \mathbb{R}_+, \quad b_1, \bar{R}_1 \in \mathbb{R}$
M.5	$b_0 + b_1 R_{1,t} + b_2 \sqrt{V_{2,t}}$	$b_0, \alpha_1, \alpha_2 \in \mathbb{R}_+, \quad b_1, b_2 \in \mathbb{R}$
M.6	$b_0 + b_1 R_{1,t} + b_2 S_{2,t}$	$b_0, \alpha_1, \alpha_2 \in \mathbb{R}_+, \quad b_1, b_2 \in \mathbb{R}$
M.7	$b_0 + b_1 R_{1,t} + b_1 S_{1,t} + b_3 \Theta_{1,t}$	$b_0, \alpha_1 \in \mathbb{R}_+, \quad b_1, b_2, b_3 \in \mathbb{R}$

Table 2: The building blocks of models and their numeral quadratures.

M.1, M.3, and M.5 are existing models in the volatility literature. M.1 can be viewed as a version of the RH model (El Euch 2018, 2019) in which the asset and the variance process share the same Brownian motion. M.3 corresponds to the QRH model (Gatheral *et al.* 2020), and M.5 to a version of the PDV model proposed by Guyon and Lekeufack ([10]). M.2 is a simplified version of model M.7. M.4 can be seen as the counterpart of the QRH model for $p = 1$. Similarly, M.6 is the counterpart of M.5 for $a_2 = 1$. This allows for a comparison of two different approaches to capturing the autoregressivity of the volatility process.

M.7 corresponds to a simplified version of the model presented in section 3 of [13]. Different assumptions about the specification of θ, which were detailed in [13], are considered here. The first one, denoted as **M.7.1**, corresponds to the assumption where:

$$\theta_{1,t} = 0.$$

The two specifications of θ considered fall within the framework presented in section 3.2 of [13]. More precisely, for **M.7.2**, we use:

$$\theta_{1,t} = \sqrt{\int_{-\infty}^{t} e^{\gamma(u-t)} (\sigma_u)^2 \mathrm{d}u}$$

For **M.7.3**, the specification is:

$$\theta_{1,t} = \Big| \int_{-\infty}^{t} e^{\gamma(u-t)} \frac{\mathrm{d}P_u}{P_u} \Big|.$$

It should be noted that, because no assumption is directly made about the drift of the asset, no specific constraint is imposed on the coefficients b_1, b_2, and b_3 in this case.

Furthermore, models are trained using data from 2000 to the end of 2014, minimizing the mean squared

error. The test set, on the other hand, covers the period from 2015 to 2021. This data segmentation allows for evaluating the generalization capability of the different specifications by assessing their performance not only on the training data (in-sample) but also on the out-of-sample data from the test set.

2.3 Comparative regression results

2.3.1 Model performance evaluation

The results of the regressions obtained for the models defined in section 2.2.2 are presented in table 3. A first general observation is that models incorporating a historical volatility process (represented by the building blocks S_i and V_i) outperform significantly the models solely based on the price trend process R_1. This finding, which had already been highlighted by Guyon and Lekeufack, is reflected in the much better performance of models M.5, M.6, and M.7 compared to models M.1, M.2, M.3, and M.4, both in-sample and out-of-sample.

		S&P500		**Nasdaq**		**FTSE**		**DAX**		**Stoxx 50**	
		Train	Test	Train	Test	Train	Test	Train	Test	Train	Test
M.1	r^2	0.379	0.136	0.439	0.053	0.346	-0.15	0.406	-0.47	0.348	0.145
	RMSE	0.092	0.08	0.068	0.08	0.087	0.094	0.082	0.088	0.09	0.083
M.2	r^2	0.596	0.442	0.613	0.141	0.467	0.105	0.547	0.009	0.494	0.365
	RMSE	0.074	0.065	0.057	0.076	0.078	0.083	0.071	0.073	0.079	0.072
M.3	r^2	0.6	0.543	0.638	0.377	0.5	0.22	0.56	0.122	0.484	0.417
	RMSE	0.074	0.058	0.055	0.065	0.076	0.077	0.07	0.068	0.08	0.069
M.4	r^2	0.605	0.517	0.617	0.245	0.483	0.181	0.555	0.088	0.494	0.369
	RMSE	0.073	0.06	0.057	0.071	0.077	0.079	0.071	0.07	0.079	0.071
M.5	r^2	0.765	0.736	0.789	0.731	0.69	0.545	0.777	0.67	0.7	0.614
	RMSE	0.056	0.044	0.042	0.042	0.06	0.059	0.05	0.042	0.061	0.056
M.6	r^2	0.766	0.733	0.791	0.724	0.691	0.56	0.775	0.669	0.7	0.621
	RMSE	0.056	0.045	0.042	0.043	0.06	0.058	0.05	0.042	0.061	0.055
M.7.1	r^2	0.757	0.728	0.786	0.729	0.681	0.541	0.77	0.663	0.692	0.608
	RMSE	0.057	0.045	0.042	0.043	0.061	0.059	0.051	0.042	0.062	0.056
M.7.2	r^2	0.767	0.739	0.79	0.734	0.696	0.556	0.774	0.675	0.704	0.62
	RMSE	0.056	0.044	0.042	0.042	0.059	0.058	0.05	0.042	0.061	0.055
M.7.3	r^2	0.765	0.737	0.79	0.736	0.689	0.552	0.776	0.671	0.697	0.623
	RMSE	0.057	0.044	0.042	0.042	0.06	0.058	0.05	0.042	0.061	0.055

Table 3: r^2 and RMSE associated with various RPDV regression models for different stock indices.

If we delve deeper into the results, the performance of models M.1, M.2, M.3, and M.4 itself varies significantly. It is clear that the explanatory power of M.1, especially out-sample, is extremely weak and inferior to M.2, M.3, and M.4. The comparison is particularly interesting with M.2, as it corresponds to the counterpart of M.1 with $p = 1$ ($p = 2$ for M.1). Hence, the respective results obtained by these two models suggests that a

modeling approach where the instantaneous volatility is linearly sensitive to a price trend process (here R_1) is more preferable than a modeling approach where the instantaneous variance is linearly sensitive to such a process. About the relative underperformance of M.2 out-of-sample compared to M.3 and M.4, this one can be explained by the specificity of the test period considered. Indeed, this period includes the Covid-19 crisis of March 2020, which led to a significant drop in stock indices followed by a strong rebound. While this rebound resulted in a sharp increase in the variable R_1 and thus a decrease in the volatility estimated by the model (due to $b_1 < 0$), the realized volatilities remained higher than their median level. Although this result may initially seem to favor a QRH-type modeling when a historical volatility factor is not directly included, it will be demonstrated in the following sections that this apparent evidence can be misleading. Besides these results, it is worth to note that for these 4 models (M.1-M.4) and all indices, the optimal value of α_1 obtained from the regression falls within the range of 0.4 to 0.5, aligning with the standard assumption in rough volatility literature where $\alpha_1 \in [0, 0.5]$ and $\epsilon = 0$.

On the contrary, for models incorporating a historical volatility factor (M.5-M.7) , the selected value of the exponent α_1 is greater than 0.5: between 0.65 and 0.8 for models M.5 and M.6, and in the vicinity of 1.3 for models M.7. This higher value of α_1 results in a more significant weighting of the recent past in the trend process R_1 in practice. As for the exponent α that determines the kernel of the historical volatility factor for models M5-M.7 (α_2 in the case of M5 and M.6, and α_1 for model M.7), the selected value is between 1.25 and 1.35. In the median case $\alpha = 1.3$, this implies that the periods $t-1$, $t-2$, and $t-3$ contribute approximately 46%, 11%, and 6% of the total weighting, respectively, considering a period of 2520 trading days. Furthermore, the cumulative contribution of the last 20 days amounts to 82% of the total weighting. Furthermore, for all these models (M.5-M.7), volatility is primarily explained by historical volatility factors, and the contribution of the R_1 factor in the explained volatility is relatively minor. Consequently, spot volatility becomes predominantly explained by an autoregressive relationship that heavily weighs its recent past.

It is also interesting to note that the performances of models M.5-M.7 are fairly homogeneous. Thus, the similar performances of models M.5 and M.6 suggest that the definition of the historical volatility factor either as a moving average of past volatility ($p = 1$) or as the square root of a moving average of past variance ($p = 2$) has little impact on the explanatory power of the model. Similarly, M.7.1 only slightly underperforms M.5 and M.6, which suggests that differentiating the exponent α for the historical volatility factor and the price trend factor is relatively marginal, at least within this static regression framework. In a similar manner, the addition of a factor θ (models M.7.2 and M.7.3) provides a non-zero but very small improvement in performance.

2.3.2 Residuals and the missing randomness question

To delve deeper into the analysis of the regression results reported in section 2.3.1, it is interesting to examine the residuals of these regressions. Firstly, it should be noted that the residuals of the regressions exhibit significant disparities in structure depending on the model under consideration. A notable manifestation of these disparities concerns the autocorrelation function of these residuals. Here again, a clear separation appears between the models incorporating a historical volatility factor and the other models. For the former, the autocorrelation is weak, and the Durbin-Watson statistic is either very close to or greater than 2, indicating insignificant autocorrelation. On the contrary, the residuals of the models not incorporating a historical volatility factor exhibit long and significant autocorrelation, which can explain the relatively poor out-of-sample performance of these models ([9]). The figure 1 illustrating the respective autocorrelation function (ACF) of the residuals for models M.2 and M.7.1 on the training data of the S&P500 provides an illustration of this disparity.

In addition to its usefulness in explaining the performance of different models, this distinction between models M.1-M.4 and M.5-M.7 seems to be tailored to different assumptions regarding the integration of exogenous sources of randomness. Firstly, it is worth noting that the incorporation of such an exogenous randomness component appears necessary for all the considered models. In other word, volatility appears to be only partially path-dependent. Indeed, standard deviations of log-returns of realized volatilities are significantly higher than the standard deviations of log-returns of volatilities simulated from discrete models using the

parameters obtained through regressions. As an example, for the case of the S&P500, the empirical standard deviation of daily log-returns of volatility is 0.347. However, this standard deviation is approximately 0.2 and 0.11 for the simulations carried out using models M.2 and M.7.1, respectively (see figure 2a)[3]. The phenomenon of missing randomness in the tested models is thus highly significant. Moreover, this effect is even more pronounced for models incorporating a historical volatility factor.

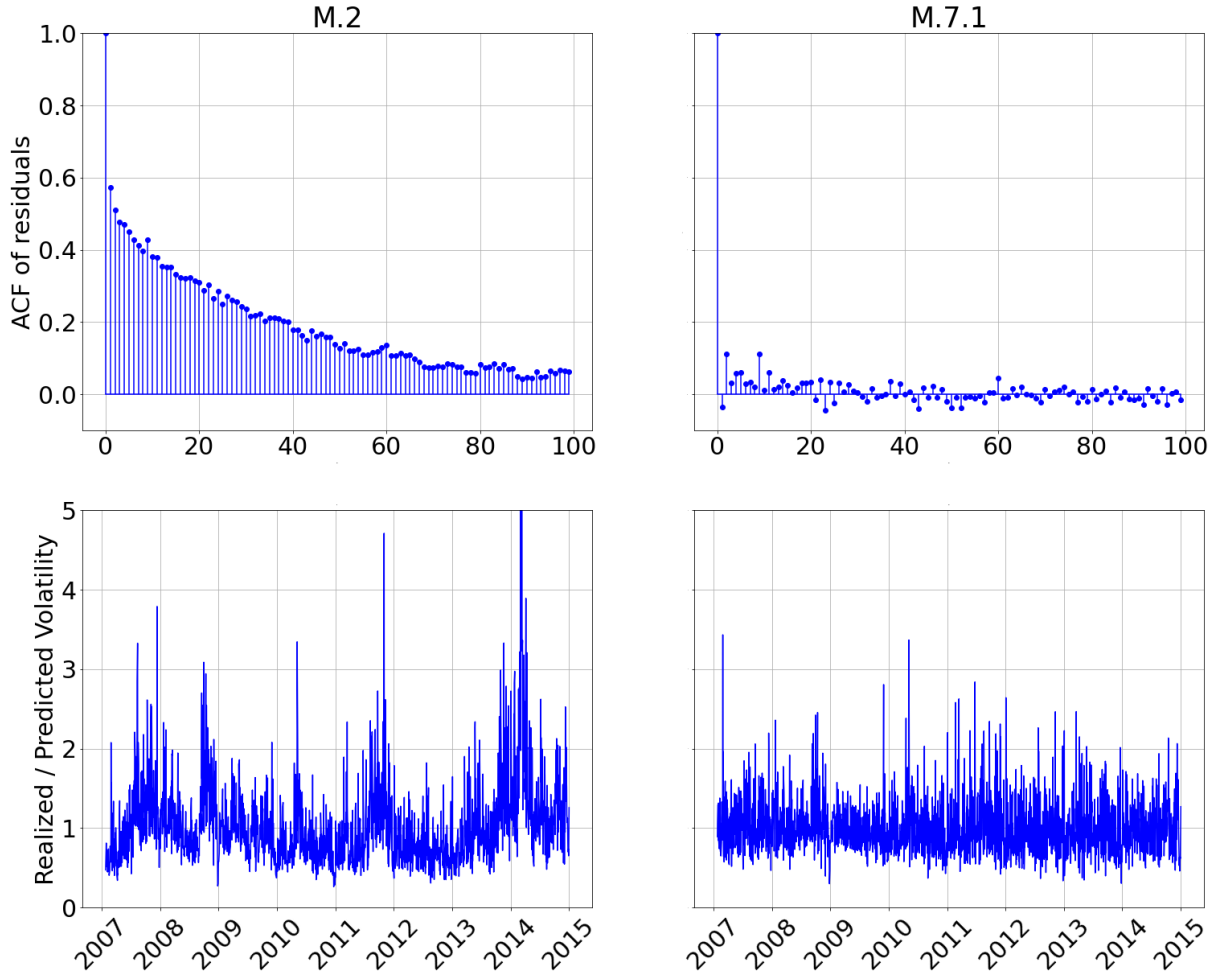

Figure 1: Comparison of residuals from models M.2 and M.7.1 in-the-sample. The ACF graphs show that the autocorrelation of M.2 model residuals is strong and persistent, while that of the M.7.1 model is lowly significant at all lags. This difference is clearly noticeable in in the graphs displaying the residual ratio of realized volatility on predicted volatility for these two models.

For these latter ones (M.5-M.7), a modeling approach like the one proposed by Guyon and Lekeufack ([10]) seems particularly coherent. This involves defining the volatility expression as the product (in the sense of multiplication) of a purely path-dependent model with a stochastic process, where the random component is independent of the asset's price Brownian motion. The volatility then takes the form:

$$\sigma_t = \sigma(R_{1,t}, ...)Y_t \tag{2}$$

where $\sigma(.)$ is the purely-path dependent model and Y is a stochastic process responsible for capturing the exogenous component of volatility. Guyon and Lekeufack suggest modeling Y using either an Ornstein-Uhlenbeck (OU) process or an exponential-OU process. This second option appears particularly suitable when

[3]The simulations are performed using the ARPDV model introduced in the appendix B.

considering the properties of the empirical ratios of realized volatility to predicted volatility for models M.5-M.7. Indeed, these ratios approximately follow a log-normal distribution $\mathcal{LN}(-0.5s^2, s^2)$, where s corresponds to the daily standard deviation of the logarithm of this quantity. Moreover, the asymptotic distribution of an exponential-OU process is a log-normal distribution. These factors, combined with the fact that the autocorrelation of the residuals in models M.5-M.7 is very low, even at a horizon of 1 day, suggest an average duration of the deviation of the process Y_t from 1 of less than 1 trading day. Consequently, the mean reversion parameter must be extremely high, implying that exogenous random frictions constitute an intraday phenomenon. In the case of discrete modeling with a time step equal to or greater than 1 trading day, one can directly use a log-normal distribution $\mathcal{LN}(-0.5s^2, s^2)$, where the realizations are independent from one trading day to the next. Figure 2b illustrates the impact of integrating an exogenous randomness component into the M.7.1 model, employed as a simulation device. While the purely path-dependent version of the model generated volatility dynamics considerably divergent from their empirical counterparts, this modified version, wherein $Y_t = 1 \ \forall \ t$ is replaced with $Y_t \sim \mathcal{LN}(-0.5s^2, s^2)$, produces much more realistic trajectories. Specifically, the standard deviation of daily log-returns of both historical and simulated volatility trajectories becomes approximately equal. More generally, this modification enables the model M.7.1 to accurately replicate the main characteristics of empirical volatility behavior.

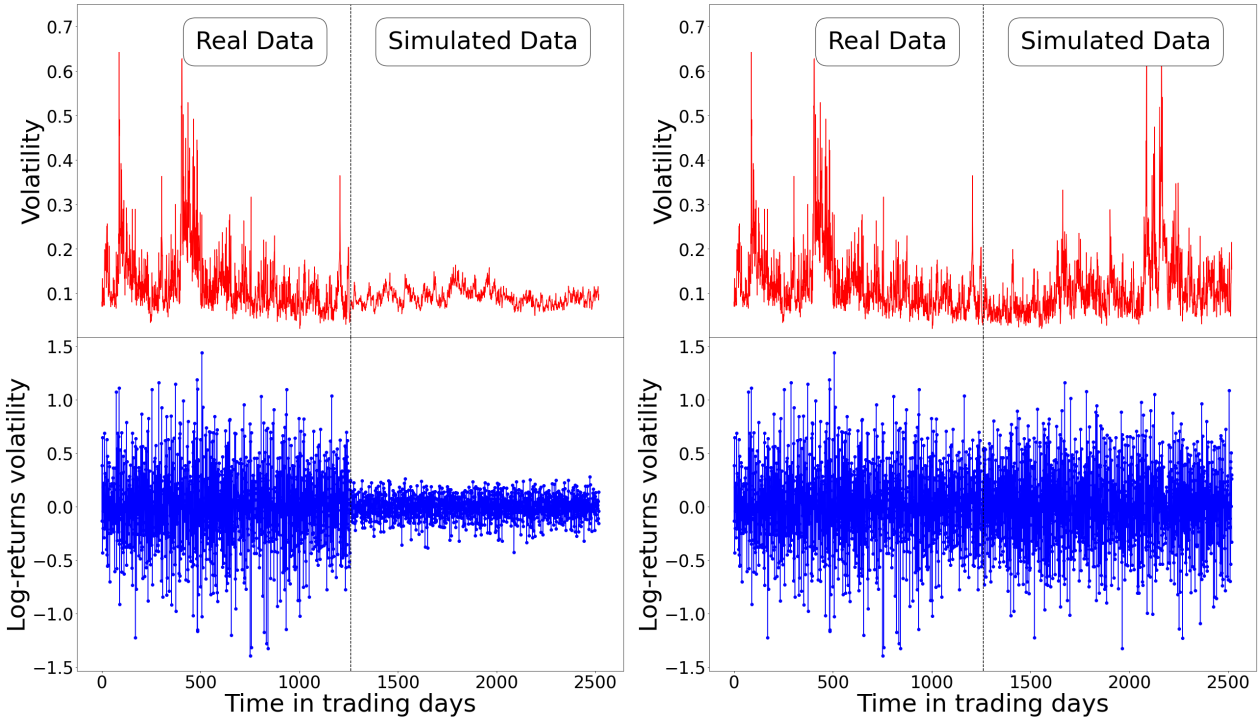

(a) Simulation without exogeneous randomness.

(b) Simulation with exogeneous randomness.

Figure 2: Comparison of simulated data from model M.7.1 calibrated on historical data (S&P500), without incorporating an exogenous source of randomness (figure 2a), and with the incorporation of an exogenous source of randomness (figure 2b). In the first case, the model produces volatility trajectories with characteristics that are very different from those of the historical volatility. In the second case, the model generates realistic volatility trajectories that capture most of the features characterizing historical volatility dynamics.

If we now turn our attention to models M.1-M.4, a question must be posed right away: given the poor performance of these models in explaining spot volatility compared to models M.5-M.7, are they truly relevant for modeling volatility, even when incorporating an exogenous source of randomness? For their part, Guyon and Lekeufack interpret these results as evidence "of the importance of including (an) historical volatility factor". While this observation is coherent, it is, however, not the only possible rational explanation. Indeed, the poor performance of models M.1-M.4 could be attributed, at least partially, to the neglect of exogenous

random frictions. However, it does not seem coherent to integrate these in the same manner as for models M.5-M.7. Indeed, for models M.1-M.4, the residuals of the regressions exhibit long memory, and the behavior of the realized volatility to predicted volatility ratio is too erratic to be modeled by OU or exponential OU processes (see figure 1). Therefore, for these models, it appears more coherent to incorporate exogenous randomness through θ functions, using a modeling approach of the following form:

$$\theta_{j,t}\mathrm{d}t = \bar{\theta}_{j,t}\mathrm{d}t + \sigma_t \mathrm{d}Z_t, \tag{3}$$

where $\bar{\theta}_j$ is a deterministic process and Z is a Brownian motion independent of B. In other words, using the terminology of Guyon and Lekeufack, it makes more sense to favor a "stochastic volatility" type modeling ([10]). Under this approach, the volatility dynamics partly depend on a process that is not observable from the asset price data. A phenomenon then may arise, that can be designated as "spurious volatility feedback", denoting an apparent structural dependence of volatility on its past trajectory, generated by a temporal dependence of volatility on an unobserved third variable. In this case, volatility is not inherently dependent on its past values; however, the fact that both current and past volatilities depend on the same variable creates the illusion of intrinsic autoregressiveness. This phenomenon could potentially explain the superior performance of models M.5-M.7, even though volatility is not structurally dependent on a "historical volatility factor". The difficulty in verifying this hypothesis through tools like the Granger causality test ([8]) lies in the fact that, if they exist, the exogenous random frictions giving rise to this phenomenon are generated by one or more hidden variables. Building on this observation, section 3 indirectly assesses the plausibility of the "spurious volatility feedback" through numerical experiments.

3 Numerical experiments

In this section, we delve into the behavior of the volatility process generated by a RPDV model considered in section 3 of [13], through a series of numerical experiments. The objective is twofold: first, to assess the capability of this type of model in capturing the empirical characteristics of volatility dynamics, and second, to evaluate the credibility of the "spurious volatility feedback" hypothesis introduced in section 2.3.2.

3.1 Experimental setup

3.1.1 The model used

The numerical experiments in this section are conducted using the following RPDV model:

$$\begin{cases} \dfrac{\mathrm{d}P_t}{P_t} &= \sigma_t(\lambda \mathrm{d}t + \mathrm{d}B_t) \\ \sigma_t &= \beta \displaystyle\int_{-\infty}^{t} (t-u)^{-\alpha} \left((\bar{\sigma} - \sigma_u)\mathrm{d}u + \sigma_u \mathrm{d}Z_u\right), \end{cases} \tag{4}$$

and with $\langle \mathrm{d}B_t, \mathrm{d}Z_t \rangle = \rho \mathrm{d}t$.

The idea here is to have a simple version of the model considered in section 3 of [13], in order to observe whether such a modeling is able to replicate the main properties characterizing volatility trajectories, without delving into second-order properties. In the same spirit, the parameters used for the numerical experiments in this section are as follows:

α	ρ	λ	$\bar{\sigma}$	β
0.425	$\sqrt{0.5}$	0.4	0.15	0.35

Table 4: The parameters used for simulations.

The value of α is set to the value of α_1 obtained in section 2 for model M.2 calibrated on the S&P500 data (rounded to the nearest 0.005). The value of the correlation coefficient $\rho = \sqrt{0.5}$ (thus $\nu = 1$) implies an equal weighting between the Brownian motion B associated with the SDE of the price and the Brownian motion W modeling the exogenous random shocks, which the volatility is also dependent on. This choice is also consistent with the magnitude of missing randomness in the M.2 models considered in section 2. Furthermore, the values taken by $\bar{\sigma}$ and λ imply that when volatility equals its attraction value ($\sigma_t = \bar{\sigma}$), the drift of the asset price is equal to 6% ($0.15 \times 0.4 = 0.06$).

3.1.2 Simulation methodology

The simulations are conducted using an Euler scheme for the Markovian approximation of model 4, a discretization method commonly employed in the literature (Rosenbaum and Zhang 2022, Parent 2023). This scheme is presented in the appendix D. The chosen time step Δt is equal to $\frac{1}{252\times 78}$ expressed in years. This choice is motivated by the fact that our proxy for daily realized volatility corresponds to the square root of the sum of squares of 78 5-minute returns. In order to obtain a similar proxy, the same methodology for calculating daily volatility is used for synthetic data. As for the state variables of the model, those are all initialized to $\frac{\bar{\sigma}}{\sum_{i=1}^{n} w_i}$, and a 20-year diffusion of the process is performed to minimize this initialization bias. The state variables obtained after this phase are then extracted to carry out the simulation of the volatility process (and the asset price process) that are considered. This initialization phase is repeated for each simulation, constituting the synthetic dataset, ensuring that the simulations begin with diverse state variables distributed in alignment with the model's structure. Following this procedure, 1000 time-series spanning 20 years each are generated. These time-series form the synthetic dataset, the characteristics of which are analyzed in the following section.

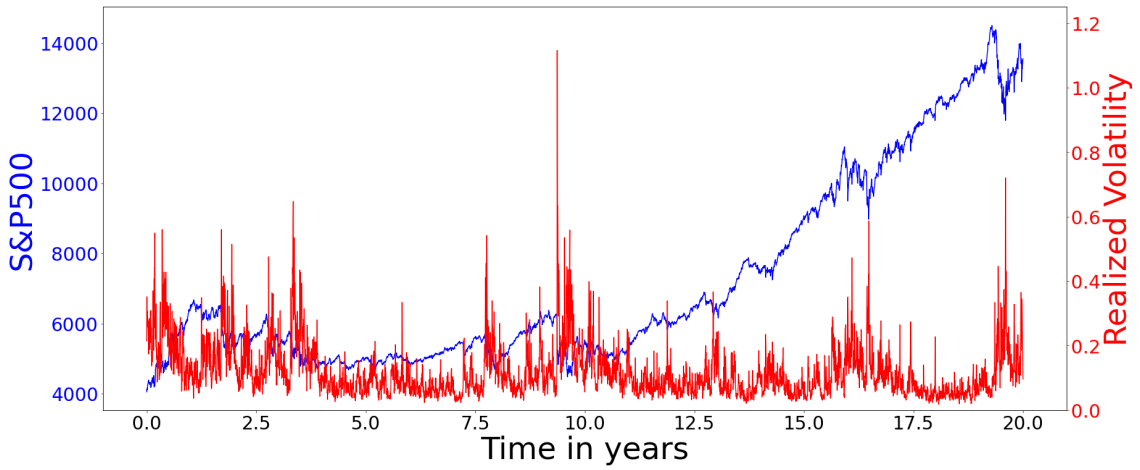

Figure 3: Example of joint dynamics between the S&P500 and its volatility generated by the model 4.

3.2 Results of numerical experiments

3.2.1 Comparison with market data

The purpose of this section is to compare certain key properties of the simulated volatility trajectories generated from model 4 (section 3.1.1) with the properties of actual realized volatility data. Since the parameter choice is tailored to the S&P500 data, the comparison of synthetic data is specifically made here with the real data of this index.

If we begin by examining the unconditional distribution of volatility, the synthetic data are generally highly consistent with the empirical data. For instance, as shown in figure 4, the distribution of daily volatility closely resembles the empirical daily volatility of the S&P500. In a more objective assessment, table 5 demonstrates the consistency of the distributions of volatilities generated by the model at various time scales

with their empirical counterparts by comparing their respective moments. Thus, for all considered time horizons, the first four moments of the empirical distributions of the S&P500 are within the first and the ninth decile of the model distributions. In addition, for both the empirical data and the data generated by the model, the standard deviation, skewness, and kurtosis of the volatility distribution follow a decreasing convex relationship with the time horizon considered. Similarly, the autocorrelation of the volatility process generated by the model is fairly close to the empirical autocorrelation of S&P500 volatility (see figure 4). Hence, all empirical autocorrelations fall within the model's estimated 90% confidence intervals, and, the model effectively reproduces the long memory of volatility observed in financial time series, characterized by a slower decay of autocorrelation compared to exponential decay.

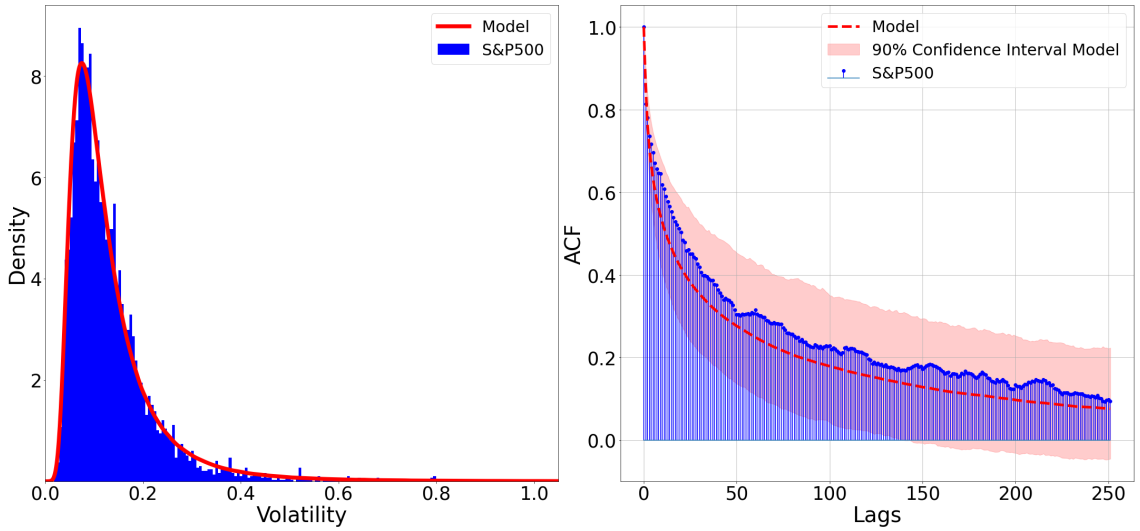

Figure 4: The left plot compares the distribution of daily volatility generated by the model 4 with the empirical distribution of daily volatility of the S&P500. The right plot compares their respective autocorrelation functions.

Time horizon	1	5	21	63	126	252
Mean of market data	0.134	0.134	0.134	0.134	0.134	0.134
Avg. mean of sim.	0.135	0.135	0.135	0.135	0.135	0.135
1^{st} decile of mean of sim.	0.105	0.105	0.105	0.105	0.105	0.105
9^{th} decile of mean of sim.	0.168	0.168	0.168	0.168	0.168	0.168
Std. dev. of market data	0.098	0.088	0.081	0.071	0.064	0.056
Avg. std.deviation of sim.	0.097	0.088	0.078	0.068	0.06	0.051
1^{st} decile of std. deviation of sim.	0.063	0.056	0.047	0.039	0.033	0.027
9^{th} decile of std. deviation of sim.	0.137	0.128	0.114	0.099	0.092	0.081
Skewness of market data	3.357	3.05	2.756	2.268	1.935	1.414
Avg. skewness of sim.	2.982	2.663	2.224	1.781	1.458	1.115
1^{st} decile of skewness of sim.	2.075	1.753	1.347	0.955	0.681	0.39
9^{th} decile of skewness of sim.	4.299	3.876	3.291	2.787	2.381	1.988
Kurtosis of market data	19.598	14.881	11.437	7.573	5.249	2.35
Avg. kurtosis of sim.	17.594	13.184	8.383	4.817	2.886	1.343
1^{st} decile of kurtosis of sim.	6.839	4.553	2.227	0.602	-0.154	-0.681
9^{th} decile of kurtosis of sim.	34.748	26.037	16.32	10.713	6.845	4.134

Table 5: The first four moments of the volatility distributions for different time horizons in trading days: S&P500 data vs synthetic data generated from the model 4.

More broadly, most of the phenomena characterizing the joint dynamics of price and volatility are well

reproduced by the model. As shown in figure 3, the combined phenomenon of volatility spikes leading to sharp price drops akin to near-jumps is captured by the model, just as the presence of relatively long periods of volatility coinciding with positive price trends. Furthermore, certain phenomena that are unattainable from a purely path-dependent model without incorporating a historical volatility factor are replicated here. Among these, for instance, are the "rallies" following a price drop that rapidly propels the asset price to a level significantly higher than its previous value ([11]). This is illustrated in in figure 5 . The phenomenon is as follows: the decrease in the Brownian motion B initially leads to an increase in the volatility level and a decrease in the asset price. This dual movement is even stronger if W also experiences a significant increase. When the dynamics of B becomes positive, the asset price starts to increase again. The sensitivity of the price dynamics to B is directly related to volatility. As a result, there are two polar cases. In the first case, W follows a downward trajectory: the rise in W combined with the decrease in B leads to a rapid drop in volatility, resulting in low exposure to the upward dynamics of B. Consequently, the rebound in the asset price is very moderate. In the second case, which corresponds to the scenario illustrated in figure 5, W follows an upward trajectory: the rise in B, which exerts downward pressure on the volatility level, is slowed down by the opposing effect of the increase in W. The volatility remains at a high level, allowing the asset price to be highly exposed to the upward dynamics of B, leading to a very strong rebound in the price.

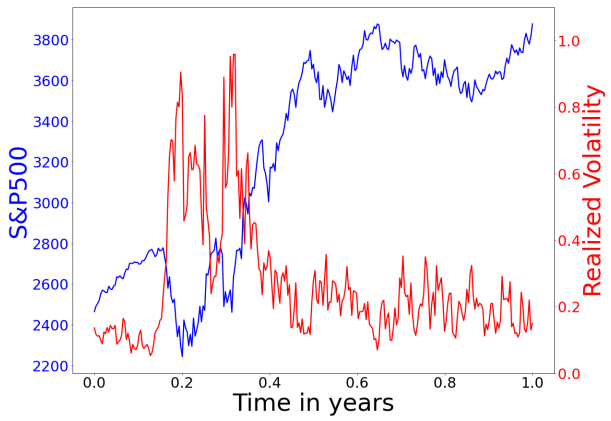

Figure 5: Example of a rally phenomenon following an increase in volatility, generated by the model 4.

This example illustrates another crucial point implied by this modeling framework: the rate of volatility decrease following a spike is highly dependent on the dynamics of the exogenous volatility component. Thus, the phenomenon of slow volatility decay after such a spike is contingent within the considered modeling framework. However, empirical volatility data suggests that this phenomenon is more systematic, making the time reversal asymmetry of volatility "clear to the naked eye" ([10]). While many properties are well captured by the model 4, this last point seems to advocate for a modeling approach that incorporates a historical volatility factor. This question is further explored in the following section.

3.2.2 Regressions on synthetic data from different RPDV models

In this section, we perform the same regressions as in section 2, but this time using the entire set of synthetic time series generated from the model 4. The objective is twofold: firstly, to determine whether the properties of the synthetic data yield regression results similar to those obtained from market data; secondly, to assess the plausibility of the "spurious volatility feedback" hypothesis stated in section 2.3.2.

First and formemost, the results presented in table 6 exhibit a consistent pattern with the outcomes derived from market data outlined in table 3. Thus, on one hand, the performance ranking of the models remains preserved, and on the other hand, the magnitude of the metrics for these models aligns with the results obtained from empirical data.

Another noteworthy element, which might seem surprising, the average performance gap between models M.1-M.4 and models M.5-M.7 is slightly more pronounced for synthetic data than for empirical data. Hence, the outperformance of models incorporating a historical volatility factor compared to models without such a factor is even more evident here, even though the synthetic data considered are generated from a model that does not include such a factor. This element thus seems to lend credence to the hypothesis of spurious volatility feedback, or at the very least, does not allow us to reject this hypothesis. In other words, the outperformance of models incorporating a historical volatility factor does not necessarily imply a structural

volatility feedback effect, but could also be a result of exogenous frictions that are not directly measurable. In this hypothesis, the inclusion of a historical volatility factor helps account for the impact of these frictions on volatility, but it does not reveal a causal relationship between historical volatility and spot volatility. However, this hypothesis, in its pure form, meaning a complete absence of positive volatility feedback effects, appears challenging to align with the excessive symmetry between volatility increase and decrease movements generated by the model, as discussed in section 3.2.1.

		Median		Q1		Q3		D1		D9	
		Train	Test	Train	Test	Train	Test	Train	Test	Train	Test
M.1	r^2	0.35	0.105	0.291	-0.15	0.419	0.245	0.229	-0.52	0.479	0.351
	RMSE	0.063	0.067	0.052	0.056	0.076	0.087	0.045	0.05	0.096	0.116
M.2	r^2	0.406	0.174	0.328	-0.1	0.488	0.323	0.233	-0.57	0.576	0.437
	RMSE	0.059	0.065	0.049	0.053	0.073	0.084	0.041	0.046	0.091	0.114
M.3	r^2	0.469	0.281	0.419	0.11	0.539	0.378	0.374	-0.19	0.599	0.481
	RMSE	0.056	0.059	0.047	0.049	0.068	0.078	0.04	0.043	0.085	0.108
M.4	r^2	0.455	0.286	0.407	0.093	0.531	0.368	0.357	-0.21	0.596	0.448
	RMSE	0.056	0.06	0.047	0.05	0.068	0.079	0.04	0.044	0.085	0.109
M.5	r^2	0.82	0.78	0.795	0.748	0.849	0.814	0.776	0.728	0.877	0.85
	RMSE	0.032	0.032	0.028	0.028	0.039	0.04	0.024	0.024	0.047	0.049
M.6	r^2	0.821	0.782	0.796	0.749	0.85	0.816	0.778	0.728	0.877	0.85
	RMSE	0.032	0.032	0.028	0.028	0.039	0.04	0.024	0.024	0.047	0.049
M.7.1	r^2	0.821	0.78	0.796	0.749	0.85	0.816	0.777	0.728	0.876	0.85
	RMSE	0.032	0.032	0.028	0.028	0.039	0.04	0.024	0.024	0.047	0.048
M.7.2	r^2	0.821	0.78	0.796	0.749	0.852	0.816	0.778	0.729	0.877	0.85
	RMSE	0.032	0.032	0.028	0.028	0.039	0.04	0.024	0.024	0.047	0.049
M.7.3	r^2	0.822	0.78	0.797	0.749	0.851	0.816	0.778	0.728	0.877	0.851
	RMSE	0.032	0.032	0.028	0.028	0.039	0.04	0.024	0.024	0.047	0.049

Table 6: r^2 and RMSE associated with various RPDV regression models applied to synthetic data generated from model 4.

A less stringent assumption, better suited to reconcile the various phenomena under consideration, is that positive volatility feedback does exist, but its magnitude appears greater than its actual scale due to the presence of exogenous randomness. In this hypothesis, the historical volatility factor obtained through regressions performed on empirical data (as those presented in section 2 or in Guyon and Lekeufack ([10]) encompasses two heterogeneous components: a genuine positive volatility feedback effect and a component capturing the impact on spot volatility of past exogenous random frictions. This hypothesis thus justifies a modeling approach that, on one hand, includes a historical volatility factor, and on the other hand, incorporates the exogenous randomness component in the same manner as its endogenous counterpart through the process θ_1. Therefore, unlike a modeling approach of the form 2, which aims to explain volatility as much

as possible through endogenous factors, the exogenous source of randomness of volatility is, in this type of modeling, no longer treated as a residual but as a component of the same nature as the source of randomness associated with the price dynamics. If these two modeling forms constitute competing approaches, they are not necessarily incompatible and could represent two valid ways of describing volatility dynamics from two different angles. Nevertheless, this perspective implies the existence of a mathematical equality relationship between these two modeling forms, allowing for a dialectical resolution of this apparent contradiction.

In a more pragmatic perspective of financial engineering, both approaches have their own weaknesses and strengths. For models incorporating a historical volatility factor in the form of 2, they hold the significant advantage of making the purely path-dependent component dependent solely on observable or quasi-observable variables - returns and past volatility - and of a single source of endogenous randomness. This property offers undeniable advantages in terms of tractability, ease of computation, and simulation, as listed in section 2 of the article by Guyon and Lekeufack, when the model is purely path-dependent (Y is a constant). However, this comparative advantage diminishes when incorporating an exogenous source of randomness through Y, which appears to be necessary based on empirical data (see section 2.3.2). On the other hand, the stochastic volatility modeling approach allows for the incorporation of an exogenous component of random frictions without the need for an additional process Y. However, this more natural way of integrating random frictions within the volatility process can lead to greater complexity in terms of calibration. Consequently, in a financial engineering context, the specific problem at hand may lead to a preference for either one of these two approaches.

4 Conclusion

This article allowed to highlight a set of elements for outlining the conditions of a coherent modeling of rough path-dependent volatility that is consistent with empirical realized volatility data.

The section 2 assessed various forms of RPDV models based on their respective ability to explain the realized volatility of five stock indices. The results obtained highlighted several significant findings. Firstly, as previously demonstrated by Guyon and Lekeufack ([10]), purely path-dependent models incorporating a historical volatility factor significantly outperform purely path-dependent models solely based on a price trend process. In addition, it has been shown that these two categories of models differ from each other in the structure of their respective residuals. On one hand, the models incorporating a historical volatility factor exhibit little to no significant autocorrelation in their residuals. On the other hand, the models that do not incorporate such factors show a very strong and persistent autocorrelation in their residuals. This result has been paralleled with the fact that all the models calibrated on market data - with or without a historical volatility factor - implied a too low randomness of volatility in comparison to empirical volatility randomness. It has therefore been demonstrated that these correlated results suggest that the mode of integrating an exogenous randomness source, necessary regardless of the considered model, should be differentiated between models with or without a historical volatility factor. In the first case, it is coherent to model volatility as the product of a purely path-dependent model (integrating a historical volatility factor) with a stochastic process whose random component is independent of the asset's price Brownian motion, process that mean-reverts very fast. This corresponds to the approach of Guyon and Lekeufack. For models that do not incorporate a historical volatility factor, considering the autocorrelation structure of there associated residuals, it appears much more coherent to integrate the exogenous randomness source as a component of the same nature as the endogenous source of randomness associated with the price dynamics. It corresponds to the standard approach of stochastic volatility models. Moreover, this approach to integrating the exogenous component of volatility has led to the formulation of the hypothesis of spurious volatility feedback, which suggests that the apparent structural dependence of volatility on its past trajectory is actually caused by a temporal dependence of volatility on an unobservable source of randomness, rather than by the presence of a true structural volatility feedback. The credibility of this hypothesis has been investigated in section 3.

More generally, section 3 has presented the results of various numerical experiments conducted using a RPDV

model considered in section 3 of [13], which includes an exogenous source of randomness but no historical volatility factor. First of all, it was shown that most of the phenomena that characterize the empirical dynamics of volatility can be reproduced from this type of model which does not explicitly include a volatility feedback effect. Furthermore, it has been demonstrated that by replicating the same comparison carried out in section 2, this time using synthetic data generated from the considered model, the obtained results were closely aligned with those obtained using empirical data. In particular, similar to the assessments conducted on market data, regression models incorporating a historical volatility factor greatly outperform those that do not include such a factor, even though the data considered in this case are generated by a model that lacks a historical volatility process. These results thus seem to lend credence to the hypothesis of spurious volatility feedback, or at least, do not allow its rejection. However, it has been pointed out that these findings should be tempered by the observation that the volatility increase and decrease movements generated by the model are slightly too symmetrical compared to empirical data. The connection of these two findings has thus led to favoring the hypothesis that there is indeed a positive feedback effect of volatility to some extent, but the magnitude of this effect may be overestimated by the estimation method used due to the presence of exogenous randomness. This hypothesis thus justifies a volatility modeling approach that, on one hand, includes a historical volatility factor, and on the other hand, incorporates the exogenous randomness component in the same manner as standard stochastic volatility models. However, this does not delegitimize the primarily path-dependent approach as adopted in Guyon and Lekeufack ([10]), which treats the exogenous part of volatility as a residual. Indeed, it has been emphasized that, on one hand, these two ways of modeling volatility are not necessarily incompatible and, on the other hand, they each present distinct advantages that position them as complementary approaches in a financial engineering perspective.

References

[1] Andersen T.G., Bollerslev T., Diebold F.X., and Labys P. (2003). Modeling and forecasting realized volatility. *Econometrica*, 71(2), 579-625.

[2] Abi Jaber E., Cuchiero C., Larsson M., and Pulido S. (2021). A weak solution theory for stochastic Volterra equations of convolution type. *The Annals of Applied Probability*, 31(6), 2924-2952.

[3] Abi Jaber E., Larsson M., and Pulido S. (2019). Affine volterra processes.

[4] 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.

[5] El Euch O., and Rosenbaum M. (2019). The characteristic function of rough Heston models. *Mathematical Finance*, 29(1), 3-38.

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

[7] 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.

[8] Granger C.W. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: journal of the Econometric Society*, 424-438.

[9] Granger C.W., and Newbold, P. (1974). Spurious regressions in econometrics. *Journal of econometrics*, 2(2), 111-120.

[10] Guyon J., and Lekeufack J. (2023). Volatility is (mostly) path-dependent. *Quantitative Finance*, 23(9), 1221-1258.

[11] Lillo F., and Mantegna R.N. (2000). Symmetry alteration of ensemble return distribution in crash and rally days of financial markets. *The European Physical Journal B-Condensed Matter and Complex Systems*, 15, 603-606.

[12] Mejía Vega C.A. (2018). Calibration of the exponential OrnsteinUhlenbeck process when spot prices are visible through the maximum log-likelihood method. Example with gold prices. *Advances in Difference Equations*, 2018, 1-14.

[13] Parent L. (2022). Rough Path-Dependent Volatility Models. Available at SSRN 4270481.

[14] Rosenbaum M., Zhang J. (2022). On the universality of the volatility formation process: when machine learning and rough volatility agree. *arXiv* preprint arXiv:2206.14114.

[15] Schöbel R., and Zhu J. (1999). Stochastic volatility with an OrnsteinUhlenbeck process: an extension. *Review of Finance*, 3(1), 23-46.

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

Appendix A Numerical quadratures

We are aiming to approximate the different integrals that make up the various elementary building blocks of the tested RPDV models. The first integral we are seeking to approximate is as follows:

$$\int_{-\infty}^{t-\epsilon} (t-u)^{-\alpha} \frac{\mathrm{d}P_u}{P_u} \tag{5}$$

First, using the midpoint method, we have the following approximation:

$$\int_{t_{i-1}}^{t_i} (t-u)^{-\alpha} \frac{\mathrm{d}P_u}{P_u} \approx \big(t - 0.5(t_{i-1}+t_i)\big)^{-\alpha} \frac{P_{t_i} - P_{t_{i-1}}}{P_{t_{i-1}}}.$$

Therefore, because ϵ is assumed to be negligible ($\epsilon \approx 0$), we can approximate the integral 5 as follows:

$$\sum_{t_i \leq t} \big(t - 0.5(t_{i-1}+t_i)\big)^{-\alpha} \frac{P_{t_i} - P_{t_{i-1}}}{P_{t_{i-1}}}.$$

The other integrals that we are seeking to approximate are of the form:

$$\int_{-\infty}^{t-\epsilon} (t-u)^{-\alpha} f(u) \mathrm{d}u. \tag{6}$$

Using the midpoint rectangle rule, we have the following approximation:

$$\int_{t_{i-1}}^{t_i} (t-u)^{-\alpha} f(u) \mathrm{d}u \approx \big(t - 0.5(t_{i-1}+t_i)\big)^{-\alpha} f(t_{i-1})(t_i - t_{i-1}).$$

Consequently, we can approximate the integral 6 as follows:

$$\sum_{t_i \leq t} \big(t - 0.5(t_{i-1}+t_i)\big)^{-\alpha} f(t_{i-1})(t_i - t_{i-1}).$$

Appendix B The ARPDV model as a simulation device

The purpose of this appendix is to introduce a discrete version of the model M.6 (and consequently of M.2 and M.7 when $\theta = 0$) formulated as an extended autoregressive model, enabling the direct utilization of parameters obtained from the regression method discussed in section 2.2. For simplification purposes, we assume here that $\mu_t = 0$, which is equivalent to assuming that the asset price is a martingale. Furthermore, in order to incorporate the exogenous component of volatility, we assume here that volatility follows a relation of the form 2, with Y being an exponential OU process (refer to Appendix C). Based on this premise, we propose the following discrete model, which we refer to as the *Autoregressive Rough Path-Dependent Volatility* (ARPDV) model:

$$r_t = \sigma_{t-1}\sqrt{\Delta t} X_t,$$

$$P_t = P_{t-1}(1 + r_t)$$

$$\sigma_t = \left(b_0 + b_1 \sum_{i=1}^{l} w_{1,i} r_{t+1-i} + b_2 \sum_{i=1}^{l} w_{2,i} \sigma_{t-i}\right) \exp\left(-0.5s^2 + sZ_t\right), \qquad w_{j,i} = \Big(t - (i+0.5)\Delta t\Big)^{-\alpha_j},$$

where X_t and Z_t are i.i.d realizations from a standard normal distribution. In the case of $s = 0$, the model is purely path-dependent. As soon as $s \neq 0$, the path-dependent component is multiplied by a log-normal distribution with an expected value of 1.

Remark. Interestingly, when we impose the constraint $\alpha_1 = \alpha_2$ (model M.7), the optimal parameters (in the least squares sense) for the various indices are such that $b_0 \approx 0$ and $b_2 \sum_{i=1}^{l} w_{2,i} \approx 1$. Given that, it makes sense to use the following 3-parameter ARPDV model:

$$r_t = \sigma_{t-1}\sqrt{\Delta t} X_t,$$

$$P_t = P_{t-1}(1 + r_t)$$

$$\sigma_t = \left(b_1 \sum_{i=1}^{l} w_i r_{t+1-i} + \sum_{i=1}^{l} w_i \sigma_{t-i} \right) \exp\left(-0.5s^2 + sZ_t\right), \qquad w_i = \frac{\left(t - (i+0.5)\Delta t\right)^{-\alpha}}{\sum_{k=1}^{l} \left(t - (k+0.5)\Delta t\right)^{-\alpha}}.$$

Thus, as shown in figure 6, even though very parsimonious, this model enables the generation of realistic joint dynamics of price and volatility.

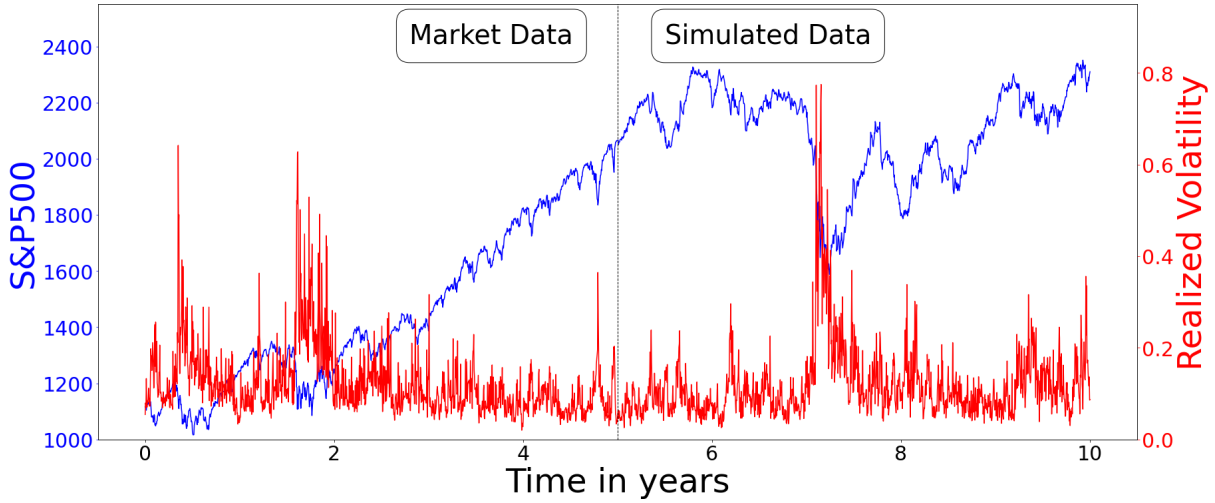

Figure 6: Example of joint dynamics between the S&P500 and its volatility. The first 5 years consist of real data, while the subsequent 5 years are generated from 3-parameter ARPDV model with $\alpha = 1.315, b_1 = -3.42, s = 0.316$.

Appendix C Exponential-OU process

We know than if Y follows

$$\mathrm{d}Y_t = \eta(\bar{Y} - Y_t)\mathrm{d}t + \upsilon \mathrm{d}Z_t,$$

the conditional distribution of Y_t given $\mathcal{F}_0$ is (Mejía Vega 2018):

$$Y_t \sim \mathcal{N}\left(\bar{Y} + (Y_0 - \bar{Y})e^{-\eta t},\ \frac{\upsilon^2}{2\eta}(1 - e^{-2\eta t})\right).$$

Therefore, if $\bar{Y} = -\dfrac{\upsilon^2}{4\eta}$, by setting $s^2 = \dfrac{\upsilon^2}{2\eta}$ we have:

$$\lim_{t\eta \to +\infty} \exp(Y_t) \sim \mathcal{LN}\big(-0.5s^2, s^2\big).$$

Appendix D The model used for numerical experiments

Using the approximation method outlined in [13], the model 4 can be approximated by the following Markovian model:

$$\begin{cases} \dfrac{\mathrm{d}P_t}{P_t} & = \sigma_t(\lambda \mathrm{d}t + \mathrm{d}B_t), \\ \mathrm{d}\mathrm{M}_t & = \mathbf{1}_n \cdot \left(\dfrac{\mathrm{d}P_t}{P_t} + (\bar{\sigma} - \sigma_t)\mathrm{d}t \right) - \Lambda \odot \mathrm{M}_t \mathrm{d}t, \\ \sigma_t & = \bar{\beta} \cdot \langle \mathrm{W}, \mathrm{M}_t \rangle, \end{cases}$$

where W is the vector of weights $(w_i)_{1 \leq i \leq n}$ and Λ_j the vector of discount coefficients $(\gamma_i)_{1 \leq i \leq n}$, with $(w_i, \gamma_i)_{1 \leq i \leq n}$ given by

$$\mathrm{W} = \begin{bmatrix} w_1 \\ ... \\ w_n \end{bmatrix}, \quad \Lambda = \begin{bmatrix} \gamma_1 \\ ... \\ \gamma_n \end{bmatrix},$$

such as

$$w_i = \frac{c\big(1 - \mathrm{x}^{-\alpha}\big)\mathrm{x}^{\alpha(i - \frac{n}{2})}}{\alpha \Gamma(\alpha)}, \qquad \gamma_i = \frac{\alpha(\mathrm{x}^{1+\alpha} - 1)\mathrm{x}^{i-1-\frac{n}{2}}}{(1+\alpha)(\mathrm{x}^{\alpha} - 1)}, \tag{7}$$

where $\Gamma(.)$ is the gamma function.

For the numerical experiments conducted in section 3, we utilize a 20-factor approximation ($n = 20$). In addition, to perform simulations from this model, we use the following explicit Euler discretization scheme:

$$\begin{cases} P_{t+\Delta t} & = P_t\Big(1 + \sigma_t(\lambda \Delta t + \sqrt{\Delta t} X_{t+\Delta t})\Big), \\ \mathrm{M}_{t+\Delta t} & = \mathrm{M}_t \odot \big(\mathbf{1}_n - \Lambda \cdot \Delta t\big) + \mathbf{1}_n \cdot \left(\sigma_t \sqrt{\Delta t} \left(\rho X_{t+\Delta t} + \sqrt{1 - \rho^2} Z_{t+\Delta t} \right) + (\bar{\sigma} - \sigma_t)\Delta t \right), \\ \sigma_{t+\Delta t} & = \beta \cdot \langle \mathrm{W}, \mathrm{M}_{t+\Delta t} \rangle, \end{cases}$$

where X_t and Z_t are i.i.d realizations from a standard normal distribution.

It is worth to note that in order to ensure the stability of the scheme, all coordinates of Λ must lower than $\frac{1}{\Delta t}$. This condition is verified for the specification considered in section 3.