

Forecasting Volatility with Time-Varying Leverage and Volatility of Volatility Effects

Leopoldo Catania^a, Tommaso Proietti^b

^a*Aarhus University and CREATES*

^b*University of Rome Tor Vergata and CREATES*

Abstract

The prediction of volatility is of primary importance for business applications in risk management, asset allocation and pricing of derivative instruments. This paper proposes a novel measurement model which takes into consideration the possibly time-varying interaction of realized volatility and asset returns, according to a bivariate model aiming at capturing the main stylised facts: (i) the long memory of the volatility process, (ii) the heavy-tailedness of the returns distribution, and (iii) the negative dependence of volatility and daily market returns. We assess the relevance of “volatility in volatility” and time-varying “leverage” effects in the out-of-sample forecasting performance of the model, and evaluate the density forecasts of the future level of market volatility. The empirical results illustrate that our specification can outperform the benchmark HAR–RV, both in terms of point and density forecasts.

Keywords: realized volatility, forecasting, leverage effect, volatility in volatility.

1. Introduction

Market volatility plays a predominant role in applied finance. Starting from the seminal work of Andersen et al. (2001), during the last decade the possibility of estimating market volatility using intradaily returns opened the way for a myriad of new applications in risk management, asset allocation and pricing of derivative instruments. Meanwhile, the properties of realized volatility time series have been deeply investigated in a series of studies that have documented the long memory feature of the series. In a highly influential paper, Corsi (2009) proposed the heterogeneous autoregressive model (HAR–RV), which accounts for the long memory via a long autoregression, parsimoniously parameterized in terms of three parameters.

Besides the long memory of volatility, the aspect of “volatility of volatility” is also of interest. Specifically, there is empirical evidence of time-varying second conditional moment of realized volatility measures. For instance, Bollerslev et al. (2009) investigate this aspect with reference to the temporal variation of expected returns. Their results suggest that accounting for volatility of volatility is of primary importance for forecasting. Other recent applications deal with the concept of “quarticity”, also related to “volatility of volatility” and its implications in prediction; see for example Bollerslev et al. (2016).

Interestingly, a number of new contributions that jointly model realized volatility and market returns have been also proposed. The idea behind these models is to properly account for another important characteristic of realized volatility that concerns its negative dependence with daily market returns (i.e. the so called “leverage effect”). The most notable contributions are Shephard and Sheppard (2010) and Takahashi et al. (2009), respectively. While the multivariate extensions of this framework are nontrivial, contributions in this direction are due to Hansen et al. (2014), Noureldin et al. (2012) and Janus et al. (2016).¹

In this paper, we propose a new bivariate model of realized volatility and market returns accounting for the three above well established features. In particular, (i) the conditional mean of the realized volatility series can show long range dependence, (ii) the volatility of volatility is allowed to vary over time according to an observation driven process, (iii) the conditional correlation of volatility and asset returns is allowed to be nonzero and time-varying, and finally (iv) both returns and the realized volatility measure follow an heavy-tailed distribution. Our specification nests the model with no leverage and a constant volatility of volatility. Moreover, we consider the alternative specification such that the conditional volatility is short memory.

We build our new model in a fully observation-driven framework and depart from the Realized Volatility–Stochastic Volatility (RV–SV) type of models introduced by Takahashi et al. (2009). Indeed, our proposal is rooted in the recent advances in score-driven (SD) models,² see Creal et al. (2013) and Harvey (2013). SD models have been found to offer a good trade-off between flexibility and complexity of time-series models and to be able to address a series of relevant empirical problems in a variety of fields, while remaining in a fully observation-driven framework.³ For further details on the SD approach, see the paper of Koopman et al. (2016) and the papers listed on the associated website www.gasmodels.com.

The forecasting ability of the proposed model is assessed through an extensive empirical application using publicly available data from the Oxford-Man Institute of Quantitative Finance Realized Library. Namely, we investigate the usefulness of accounting for the aforementioned features for point and density prediction of future market volatility. Our results show that the model can outperform the benchmark HAR–RV model, both in terms of point and density forecasts.

The paper is organized in the following manner. Section 2 introduces the proposed model and details the estimation procedure. Section 3 details filtering procedure for time-varying parameters. Section 4 reports an analysis on the shape of the response of future volatility with respect to current shocks. Section 5 reports the forecasting empirical application. Finally, Section 6 concludes and suggests further possible developments.

¹Of remarkable interest are also the working papers of Asai et al. (2016), Hansen et al. (2016) and Lucas and Opschoor (2016).

²Also known as Generalized Autoregressive Score (GAS) and Dynamic Conditional Score (DCS) models.

³This means that model estimation and inference can benefit from the availability of the likelihood function in closed form.

2. A new joint model for market returns and realized volatility

Let $y_t \in \mathbb{R}$ and $x_t \in \mathbb{R}$ be the financial return and its associated realized log-volatility measure (logarithm of the realized variance) at time $t \in \mathbb{N}$. Conditional on $\mathcal{F}_{t-1} = \sigma(y_{t-s}, x_{t-s}, s > 0)$, the bivariate vector $\mathbf{z}_t = (y_t, x_t)'$ is modelled as follows:

$$\mathbf{z}_t | \mathcal{F}_{t-1} \sim \mathcal{T}(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t, \nu), \quad \boldsymbol{\mu}_t = \begin{bmatrix} 0 \\ \mu_t \end{bmatrix}, \quad \boldsymbol{\Sigma}_t = \begin{bmatrix} \exp(2\mu_t) & c_t \\ c_t & q_t \end{bmatrix} \quad (1)$$

where $\mathcal{T}$ denotes the bivariate Student's t distribution with ν degrees of freedom, with probability density functions (*pdf*)

$$p(\mathbf{z}_t | \mathcal{F}_{t-1}) = \frac{\Gamma((\nu+2)/2)}{\Gamma(\nu/2)(\nu-2)\pi|\boldsymbol{\Sigma}_t|^{1/2}} \left(1 + \frac{1}{\nu-2}(\mathbf{z}_t - \boldsymbol{\mu}_t)' \boldsymbol{\Sigma}_t^{-1}(\mathbf{z}_t - \boldsymbol{\mu}_t)\right)^{-(\nu+2)/2}, \quad (2)$$

and $\boldsymbol{\mu}_t = E(\mathbf{z}_t | \mathcal{F}_{t-1})$, $\boldsymbol{\Sigma}_t = \text{Var}(\mathbf{z}_t | \mathcal{F}_{t-1})$.

A distinctive feature of our specification is that $\mu_t \in \mathbb{R}$, the conditional mean of the realized log-volatility, enters the expression of the conditional variance of the returns series, so that $\text{Var}(y_t | \mathcal{F}_{t-1}) = \exp(2\mu_t)$. We set the conditional mean of y_t equal to zero, i.e., we assume that financial returns are a martingale difference sequence with respect to $\mathcal{F}_{t-1}$. The dynamic model for the evolution of the parameters μ_t , ρ_t , and q_t , that will be specified in the next section, implies that they are measurable with respect to $\mathcal{F}_{t-1}$. For the degrees of freedom we assume $\nu > 2$. Our specification also implies that both y_t and x_t , conditionally on $\mathcal{F}_{t-1}$, are univariate Student's t distributed, see Kotz and Nadarajah (2004, pp. 15–16).

The conditional covariance between y_t and x_t models the leverage effect. It is parameterized as $c_t = \exp(\mu_t)\sqrt{q_t}\rho_t$, where $\rho_t \in (-1, 1)$ is the time-varying conditional correlation. In the financial econometrics literature there is no consensus on the definition of the leverage effect. The traditional one dates back at least to Crane (1959) and relates equity returns to the changes in the level of their volatility, i.e., when the stock price falls, the future volatility increases and vice-versa. The theoretical background behind this reasoning has been principally developed by financial economists starting from the seminal work of Black and Scholes (1973) and grounding on the well-known Modigliani–Miller's framework, see Black (1976) and Christie (1982).

Statistical models have incorporated this notion in different ways. In the GARCH framework (Engle, 1982; Bollerslev, 1986) the leverage effect is incorporated as an asymmetric response of the conditional volatility process to past returns, i.e., when $y_t < 0$, $\partial \text{Var}(y_t | \mathcal{F}_{t-1}) / \partial y_{t-1}$ is larger than when $y_t \geq 0$; see, for example, Nelson (1991), Glosten et al. (1993) and Zakoian (1994). In the stochastic volatility (SV) framework (Taylor, 1986), the negative covariance between the disturbances of the returns and log-volatility equations has different implications according as to whether the dynamic model for the log-volatility process is formulated in contemporaneous form, as in Harvey and Shephard (1996), or in future form, as in Jacquier et al. (2004). As shown by Yu (2005), only when the covariance is contemporaneous y_t is a martingale

difference sequence and the specification is consistent with the efficient market hypothesis of Fama (1970). Our definition of the leverage effect is consistent with the arguments of Yu (2005).

3. The dynamic model for the time-varying parameters

Our specification of the updating mechanism of the parameters μ_t , ρ_t and q_t , is based on the recent Score Driven (SD) framework introduced by Creal et al. (2013) and Harvey (2013). SD models build a dynamic updating equation exploiting the information contained in the score of conditional distribution of the data. This way of introducing time-variation has been proven to be very flexible and has been used in a number of different applied settings, see for example the paper by Koopman et al. (2016).

It is convenient to formulate a SD updating model for a smooth transformation of the original parameters, which makes the range of their values unconstrained. Denote $\boldsymbol{\theta}_t = (\mu_t, \rho_t, q_t)'$, $\boldsymbol{\theta}_t \in \mathbb{R} \times (-1, 1) \times \mathbb{R}^+$, and let $\tilde{\boldsymbol{\theta}}_t = (\tilde{\mu}_t, \tilde{\rho}_t, \tilde{q}_t)'$ be a transformation of $\boldsymbol{\theta}_t$, $\tilde{\boldsymbol{\theta}}_t \in \mathbb{R}^3$, such that $\Lambda(\tilde{\boldsymbol{\theta}}_t) = \boldsymbol{\theta}_t$ is the continuous $\mathcal{F}_{t-1}$ -measurable mapping function $\Lambda(\cdot)$

$$\Lambda : \begin{cases} \mu_t &= \tilde{\mu}_t \\ \rho_t &= (1 - \exp(-\tilde{\rho}_t)) / (1 + \exp(-\tilde{\rho}_t)) \cdot \\ q_t &= \exp(\tilde{q}_t) \end{cases} \quad (3)$$

The Jacobian matrix, $\mathcal{J}(\tilde{\boldsymbol{\theta}}_t) = \partial \Lambda(\tilde{\boldsymbol{\theta}}_t) / \partial \tilde{\boldsymbol{\theta}}_t$, associated to this transformation is:

$$\mathcal{J}(\tilde{\boldsymbol{\theta}}_t) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & v(\tilde{\rho}_t) & 0 \\ 0 & 0 & \exp(\tilde{q}_t) \end{pmatrix}, \quad (4)$$

where $v(\tilde{\rho}_t) = 2 \exp(-\tilde{\rho}_t) / (1 + \exp(-\tilde{\rho}_t))^2$.

The updating equation for $\boldsymbol{\theta}_t$ is then:

$$\boldsymbol{\theta}_{t+1} = \Lambda(\tilde{\boldsymbol{\theta}}_{t+1}) \quad (5)$$

$$\tilde{\boldsymbol{\theta}}_{t+1} = (\mathbf{I}_3 - \mathbf{B})\boldsymbol{\kappa} + \mathbf{B}\tilde{\boldsymbol{\theta}}_t + \mathbf{A}(L)\tilde{\mathbf{s}}_t, \quad (6)$$

where $\mathbf{I}_3$ is the 3×3 identity matrix and $\tilde{\mathbf{s}}_t = \partial \log p(\mathbf{z}_t | \mathcal{F}_{t-1}) / \partial \tilde{\boldsymbol{\theta}}_t$ is the score of the conditional distribution of $\mathbf{z}_t$ with respect to the transformed parameters $\tilde{\boldsymbol{\theta}}_t$. This quantity is recovered via the chain rule as:

$$\tilde{\mathbf{s}}_t = \mathcal{J}(\tilde{\boldsymbol{\theta}}_t)' \mathbf{s}_t, \quad (7)$$

where $\mathbf{s}_t = \partial \log p(\mathbf{z}_t | \mathcal{F}_{t-1}) / \partial \boldsymbol{\theta}_t$, the score of $p(\mathbf{z}_t | \mathcal{F}_{t-1})$ with respect to $\boldsymbol{\theta}_t$, $\mathbf{s}_t = (s_t^\mu, s_t^\rho, s_t^q)'$, with

$$s_t^\mu = \eta_t \left[z_t^y + \frac{(x_t - \mu_t)}{q_t} - \rho_t \frac{y_t(1 + x_t - \mu_t)}{q_t \exp(\mu_t)} \right] - 1, \quad (8)$$

$$s_t^\rho = \frac{\eta_t}{q_t} \left[\frac{(x_t - \mu_t)^2}{q_t} - \rho_t \frac{y_t(x_t - \mu_t)}{\exp(\mu_t)} \right] - \frac{1}{q_t}, \quad (9)$$

$$s_t^q = \eta_t \left[-\frac{\rho_t \ell_t}{1 - \rho_t^2} + \frac{y_t(x_t - \mu_t)}{\exp(\mu_t) q_t} \right] + \frac{\rho_t}{1 - \rho_t^2}, \quad (10)$$

and, setting $z_t^y = y_t^2 / \exp(2\mu_t)$ and $z_t^x = (x_t - \mu_t)^2 / q_t$,

$$\eta_t = 2\xi_t / (1 - \rho_t^2), \quad \xi_t = \frac{\nu + 2}{2(\nu - 2)\omega_t}, \quad \omega_t = 1 + \frac{\ell_t}{(\nu - 2)(1 - \rho_t^2)}, \quad \ell_t = z_t^y + z_t^x - \frac{2\rho_t y_t(x_t - \mu_t)}{\exp(\mu_t q_t)}. \quad (11)$$

Note that under correct model specification, $\mathbf{s}_t$ evolves as a martingale difference sequence implying that $E(\tilde{\boldsymbol{\theta}}_t) = \boldsymbol{\kappa}$.

In Equation (5), the matrices $\mathbf{A}(L)$ and $\mathbf{B}$ and the vector $\boldsymbol{\kappa} = (\kappa^\mu, \kappa^\rho, \kappa^q)'$ determine the evolution of $\tilde{\boldsymbol{\theta}}_t$ and its long run level, respectively. Depending on the parametrization of $\mathbf{A}$ and $\mathbf{B}$ we can specify different updating schemes for $\tilde{\boldsymbol{\theta}}_t$. Specifically, $\mathbf{B}$ determines the response of $\tilde{\boldsymbol{\theta}}_{t+1}$ to its previous value and influences the persistence of the process. In this paper, we specify a diagonal structure for $\mathbf{B}$, i.e. $\mathbf{B} = \text{diag}(b^\mu, b^\rho, b^q)$, which implies that possible autoregressive interactions between parameters are ruled out. Furthermore, in order to avoid an explosive behaviour of $\tilde{\boldsymbol{\theta}}_t$, we impose $|b^j| < 1$ for $j = \mu, \rho, q$.

As for the lag polynomial matrix $\mathbf{A}(L)$, we consider the specification

$$\mathbf{A}(L) = \text{diag}(a^\mu(1 - L)^{-d}, a^\rho, a^q). \quad (12)$$

which implies that μ_t has an ARFIMA(1, d ,0) representation with memory parameter d . To enforce weak stationarity of $\tilde{\boldsymbol{\theta}}_t$ one assume $d < 1/2$. A key stylized fact of log-realized volatility is the long memory property of the process. This is well known from the seminal contribution of Corsi (2009) and further developments, see for instance Proietti (2016). Our way of introducing long memory in the evolution of μ_t follows the approach of Janus et al. (2016). The transformed correlation and volatility of volatility parameters follows a first order autoregressive evolution.

Given a bivariate time-series of financial returns and realized log-volatilities $\{(y_t, x_t)', t = 1, \dots, T\}$ model parameters $\boldsymbol{\xi} = (\kappa^\mu, \kappa^\rho, \kappa^q, a^\mu, a^\rho, a^q, d, b^\mu, b^\rho, b^q, \nu)'$ are estimated via numerical maximization of the likelihood function:

$$\hat{\boldsymbol{\xi}} = \underset{\boldsymbol{\xi} \in \Xi}{\text{argmax}} \sum_{t=1}^T \log p(\mathbf{z}_t | \mathcal{F}_{t-1}; \boldsymbol{\xi}), \quad (13)$$

where we have highlighted the dependence of $p(\mathbf{z}_t | \mathcal{F}_{t-1})$ from $\boldsymbol{\xi}$. In the implementation of (13), we set the presample values of μ_t equal to its unconditional mean and truncate the binomial expansion of the fractional polynomial $(1 - L)^{-d}$ at lag 1,000.

4. The News Impact Curve

Understanding the reaction of the conditional volatility process to past observations is of primary importance for policy makers and market participants. In the financial econometrics literature, the News Impact Curve (NIC), introduced by Engle and Ng (1993) for GARCH models, is the key tool used to investigate this important aspect. The NIC plots the volatility at time $t + 1$ against a shock which hits the system at time t . The leverage effect implies an asymmetric reaction of the volatility to past shocks, so that the NIC for models that incorporate this feature is asymmetric. As stated before, GARCH models incorporate the leverage effect in a rather simplistic way, that is, estimating an additional parameter which increases the volatility proportionally to some function of the past negative shock, as for example the approaches of Nelson (1991) and Glosten et al. (1993). Differently, as we reported in the previous section, SV models incorporate leverage effect through a correlation parameter between the disturbances that affect the measurement and state equations of the state space formulation. Unfortunately, for SV models the NIC is hard to derive since $p(y_t|\mathcal{F}_{t-1})$, and its moments, are not available in closed form. Solutions exploiting computational intensive simulation techniques are available (Takahashi et al., 2013), but rarely implemented.

Our framework provides an easy way to derive the NIC and shed lights about the three different sources of variation of the conditional volatility process: i) the level of leverage, via ρ_t ; ii) the fatness of the tail of the conditional distribution, measured by tail index ν ; iii) the volatility of volatility component, q_t . The three drivers influence the NIC in different ways and offer intuitive insights about the volatility structure, namely:

- i) The correlation between y_t and x_t determines the asymmetric behaviour of the volatility response. This is consistent with the usual interpretation of the leverage effect. However, as it will be clear later, the direction of the asymmetry also depends on the realization of the realized log-volatility at time t , x_t .
- ii) The degree of freedom parameter helps in robustifying the update, that is, if ν is low, large shocks are tapered, so that they exert little influence on tomorrow's volatility. This mechanism is well-known in the Score Driven literature and several authors have pointed out its benefit from an applied perspective, see *e.g.*, Creal et al. (2013), Harvey (2013), and Harvey and Luati (2014). The reasoning behind this behaviour is that, if there is statistical evidence of fat-tails, extreme observations are treated as realizations from the tails of the distribution, and hence, they do not imply a strong signal in terms of increased volatility.
- iii) The volatility of volatility influences the amplitude of the update. The signal is increased if the volatility of volatility is high and reduced if it is low.

To illustrate the behaviour of the NIC we simplify our model assuming that $\tilde{\mu}_t$ is short memory ($d = 0$), so that $\mathbf{A}(L) = \mathbf{A} = \text{diag}(a^\mu, a^\rho, a^q)$; as in Engle and Ng (1993), we start by setting the unrestricted parameters $\tilde{\theta}_t$ to their unconditional level, i.e. $\tilde{\rho} = \kappa^\rho/(1 - b^\rho)$, $\tilde{q} = \kappa^q/(1 - b^q)$, and $\tilde{\mu} = \kappa^\mu/(1 - b^\mu)$, and denote by ρ, q, μ , their transformation via $\Lambda(\cdot)$. Differently from the NIC defined in Engle and Ng (1993), in our case we observe two “shocks” hitting the system and letting the conditional volatility level to be updated from $\sigma_t = \exp(\mu_t)$ to $\sigma_{t+1} = \exp(\mu_{t+1})$. Formally, given the vector of new observations at time t , $\mathbf{z}_t = (y_t, x_t)'$, the NIC is defined as:

$$NIC(y_t, x_t; \mu, \rho, q, \nu) = \sigma \exp(a^\mu s_t^\mu), \quad (14)$$

where $\sigma = \exp(\mu)$ and s_t^μ , defined in (8), is evaluated at (μ, ρ, q) .

We now graphically investigate the different shapes of the NIC with respect to different values of model parameters and new observations. Figure B.1 displays $NIC(y_t, x_t; 0, \rho, 1, \nu)$ as a function of y_t for the following selected values of x_t : i) $x_t = -5.5$, implying that realized volatility decreases, ii) $x_t = 0$ (no changes in volatility) iii) $x_t = 5.5$ (increase in volatility). The results are reported for different values of ρ : panels (a), (b) and (c) refer to the case of a negative correlation between y_t and x_t ($\rho = -0.7$). Panels (d), (e) and (f) refer to the case of no correlation between y_t and x_t ($\rho = 0$). Panels (g), (h) and (i) are relative to the case of positive correlation between y_t and x_t ($\rho = 0.7$). We consider the sensitivity to the degree of freedom parameter ν in Figure B.1, which considers the following cases: very fat tails $\nu = 3$ (black, continuous), moderate fat tails $\nu = 6$ (red, dashed), and relatively thin tails $\nu = 9$ (green, dotted).

[Insert Figure B.1 about here.]

The effects of ρ and ν on the NIC are very evident from Figure B.1. As expected, ρ determines the levels of asymmetry of the NIC, while ν its behaviour when extreme observations are considered. Looking at panels (d), (e) and (f), we observe that NIC is symmetric in the case $\rho = 0$. The value of x_t determines the level of the curve: for values of y_t near zero, the model increases (decreases) the volatility if $x_t > 0$ ($x_t < 0$) compared to the case $x_t = \mu$. The cases $\rho_t = 0.7$ and $\rho_t = -0.7$ provide interesting insights about the relation between the volatility level at time $t + 1$ and the observations $(y_t, x_t)'$. Since in these cases returns and volatility are correlated, the concordance between their sign also influences the update.⁴ Intuitively, if we observe that $\text{sign}(y_t x_t) = +1$ and the correlation is negative, this is a signal of increased uncertainty since such event is less likely than $\text{sign}(y_t x_t) = -1$ if $\rho < 0$. The model translates this unlikely event in an increase of the volatility as it is evident from panels (a) and (c). Differently, as shown by panels (g) and (i), if $\rho > 0$ the opposite holds. Panels (b) and (h) are more difficult to interpret and there is no a clear reasoning explaining the asymmetric behaviour of NIC in that circumstances.

⁴Actually it is the sign of the difference between y_t and x_t from their averages which we assume to be zeros in both cases.

Figure B.2 displays $NIC(\bar{y}, x_t; 0, \rho, 1, \nu)$ as a function of x_t , for selected values of the remaining arguments. For returns we consider the following cases: i) a large negative return, $y_t = -7$; ii) zero return, $y_t = 0$; and iii) a large positive return, $y_t = 7$. As before, results are reported for different values of ρ and ν . Again, ρ determines the asymmetry of the curve, while ν its behaviour with respect to extreme values of y_t and x_t . The symmetric case is still recovered when $\rho = 0$. The robustness of the filter is evident from all the plots. Specifically, we note that panels (b), (e) and (h) are consistent with Figure 1 of Harvey and Luati (2014), where the authors study the robustness of Score Driven filters for location models. In our framework their results are extended to the case when scale and location are simultaneously updated.

The main conclusions of this graphical analysis is that our specification generates a flexible and interpretable NIC and that asymmetric response is not only due to the correlation parameter ρ_t , but also to the state of the system via the levels of x_t , y_t and the sign of their product.

[Insert Figure B.2 about here.]

5. Empirical application

Our empirical application deals with the ability of our model to predict realized log-volatility. For our analysis we use the very well-known Oxford-Man Institute's realized library, see Gerd et al. (2009), which is freely available at <http://realized.oxford-man.ox.ac.uk/data>, from which we collect the daily returns and the realized log-volatility series for 21 equity indices listed in the data set. In order to be robust to possible microstructure noise present in the data, x_t is the logarithmic transformation of the realized kernel volatility measure, proposed by Barndorff-Nielsen et al. (2008). The length of the time series ranges from $T = 3605$ to $T = 4272$ and covers the period from 3 January 2000 to 11 October 2016. Financial returns have been demeaned to be consistent with the specification given in (1).

Table A.1 presents some descriptive statistics. Columns two to five report the sample correlation between financial returns and their associated realized log-volatility for sub-periods of 4 years. The empirical correlations are almost always negative suggesting the presence of leverage effect for the considered time series. Furthermore, the correlation is heterogeneous across financial indices.

[Insert Table A.1 about here.]

The last four columns report the empirical standard deviation of the realized log-volatility series across different years. Still we observe heterogeneity across financial returns and moderate evidence of time-variation. For all the series we observe higher values of the empirical standard deviation during the turbulent period 2008/2011.

5.1. Full sample analysis

We first report a full sample analysis of the different nested specifications of our model, arising, for instance, by constraining $d = 0$, or imposing the ρ_t and/or q_t are time invariant (setting $a^\rho = b^\rho = 0$ and/or $a^q = b^q = 0$, respectively).

We label our encompassing specification Score Driven Fractional Integrated Realized Volatility (SDFIx) model. The specification enforcing $d = 0$ will be labelled Score Driven Realized Volatility (SDx). In both cases, nested versions recovered by imposing constraints on the coefficients of the system matrices **A** and **B** are indicated with the $-\bar{j}$ label, $j = \rho, q$. For example, the SDx model featuring a time invariant ρ is denoted by SDx- $\bar{\rho}$.

Table A.2 lists the 6 different model specifications we define. All the specifications are estimated by maximum likelihood using the full available sample. Model selection is carried out using the Bayesian Information Criterion (BIC).

[Insert Table A.2 about here.]

Table A.3 reports the BIC for the specifications listed in Table A.2. The specification which is chosen most often according to BIC is SDx- $\bar{\rho}$, i.e. the model without fractional integration and time-invariant correlation. Exceptions are the IXC2 series (the NASDAQ Composite index) and the RUT2 series (the RUSSELL 2000 index) for which the specification with lower BIC is SDx, i.e. the most general model without fractional integration. Selection according to AIC (not reported) suggest a different story and provides mixed evidence across different specifications.

[Insert Table A.3 about here.]

[Insert Figure B.3 about here.]

Figure B.3 illustrates the filtered parameters for the SDx model estimated on the IXC2 series. The period starts on 3 January 2000 and ends 11 October 2016. Blue and red vertical bands indicate periods of US and European recessions (usually associated with periods of financial turmoil). Panel (a) depicts the filtered log-volatility level μ_t and compares it with the observed realized log-volatility. The robustness of the SD filter is evident from the plot. Indeed, the filtration does not seem to be affected by the extreme observations which are present in the data. Volatility has increased considerably during the years 2008 – 2009, which are characterised by the Global Financial Crisis (GFC). For observing levels of volatility comparable to those of the GFC, we need to get back to the early 2000 Dot-com bubble. Panel (b) reports the filtered volatility of volatility parameter, q_t . We find that the volatility of volatility oscillates between a long run average of about 0.24 and presents higher values when also the volatility level is high. Finally, panel (c) reports the correlation coefficient associated with the leverage effect, $\rho_t = \text{cor}(y_t, x_t)$. Interestingly, the estimated

correlation coefficient is almost zero during the GFC and negative during the recovery phase. This finding may suggest that the leverage effect is more pronounced during tranquil periods. However, for most series there is no substantive evidence for a time varying correlation.

5.2. Out-of-sample predictive analysis

Our aim is to predict the one step ahead log-volatility measure and investigate if including different features such as: i) fractional integration, ii) time-varying leverage effect, and iii) time-varying volatility of volatility is of any help. To this end, we perform a classical forecasting analysis using the recursive method of forecasting (Marcellino et al., 2006). Specifically, for each series we consider the first $S = 2000$ observations as the in sample period, and then we compute one step ahead rolling forecasts for the last $H = T - S$ observations. The estimation window is of fixed length. The model parameters are estimated using only the data belonging to the in sample period and updated each time a new observation becomes available until the end of the sample. Since the lengths of the series are different, also the lengths of the out of sample period, H , varies across the series, ranging between $H = 1605$ and $H = 2272$.

The predictive performances of the specifications listed in Table A.2 are compared with that of the HAR-RV model by Corsi (2009), which represents a valid benchmark for realized volatility series. The HAR-RV estimates a constrained autoregressive model of order 22, parsimoniously parameterized in terms of 3 coefficients. Specifically, the model is defined as:

$$x_t = \beta_0 + \beta^{(d)}x_{t-1} + \beta^{(w)}x_{t-1}^{(w)} + \beta^{(m)}x_{t-1}^{(m)} + \eta\varepsilon_t, \quad (15)$$

where $\varepsilon_t \stackrel{iid}{\sim} \mathcal{N}(0, 1)$ and:

$$x_{t-1}^{(w)} = \frac{1}{5} \sum_{s=1}^5 x_{t-s} \quad (16)$$

$$x_{t-1}^{(m)} = \frac{1}{22} \sum_{s=1}^{22} x_{t-s}, \quad (17)$$

are the weekly and monthly simple moving averages of the log-volatility.

The SDFIx model and nested specifications defined in Table A.2 provide the one step ahead conditional distribution of the bivariate vector $(y_{t+1}, x_{t+1})'$, which is defined in (1). Starting from the joint distribution of $(y_{t+1}, x_{t+1})'$ we straightforwardly compute the marginal distribution of x_{t+1} conditionally on $\mathcal{F}_t$. Also HAR-RV provides the one step ahead distribution of x_{t+1} , which is Gaussian with mean $\hat{x}_{t+1} = \hat{\beta}_0 + \hat{\beta}^{(d)}x_{t-1} + \hat{\beta}^{(w)}x_{t-1}^{(w)} + \hat{\beta}^{(m)}x_{t-1}^{(m)}$ and variance $\hat{\eta}^2$, where $\hat{\beta}_0, \hat{\beta}^{(d)}, \hat{\beta}^{(w)}$, and $\hat{\eta}^2$ are the ordinary least squares estimates of $\beta_0, \beta^{(d)}, \beta^{(w)}$, and η^2 , respectively.

Since all the models provide the full distribution of x_{t+1} , we compare the predictive ability of different models both in terms of point and density forecasts. The point forecasts are evaluated via the mean square error (MSE), $H^{-1} \sum_{h=1}^H (\widehat{x_{T+h}} - x_{T+h})^2$. The calibration and accuracy of the one-step-ahead predictive

densities are assessed via the Continuous Ranked Probability Score (CRPS), introduced by Matheson and Winkler (1976). The CRPS is a *proper* scoring rule,⁵ defined as:

$$CRPS_{t+1} = \int_{\mathbb{R}} [F(z | \mathcal{F}_t) - \mathbf{I}\{x_{t+1} < z\}]^2 dz, \quad (18)$$

where $F(z | \mathcal{F}_t)$ is the one step ahead cumulative density function (*cdf*) evaluated in z . Evidently, the CRPS measures the distance between the predicted *cdf*, $F(x_{t+1} | \mathcal{F}_t)$, and the empirical *cdf* represented as a step function in x_{t+1} . Averaging the CRPS over the out-of-sample period provides the quantity at the base of our comparative analysis. Models with lower averaged CRPS are preferred. Analytical solutions of (18) are generally not available. However, if the predictive distribution is Gaussian, as in HAR-RV, with mean μ and variance σ^2 , for an ex-post realization x_{t+1} CRPS is given by:

$$CRPS_{t+1} = \sigma \left[2\phi\left(\frac{x_{t+1} - \mu}{\sigma}\right) + \frac{x_{t+1} - \mu}{\sigma} \left(2\Phi\left(\frac{x_{t+1} - \mu}{\sigma}\right) - 1 \right) - \frac{1}{\sqrt{\pi}} \right], \quad (19)$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ are the *pdf* and *cdf* of a standard Gaussian distribution, respectively. For the SDFIx model and nested specifications CRPS can be approximated via:

$$CRPS_{t+1} = \frac{1}{2} \mathbb{E}^{p_x} [|X - X'| | \mathcal{F}_t] - \mathbb{E}^{p_x} [|X - x_t| | \mathcal{F}_t], \quad (20)$$

where X and X' are independent copies of a random variable with *pdf* $p_x(\cdot)$ conditional on $\mathcal{F}_t$, that is, the one step ahead predicted distribution delivered by the model.

5.2.1. Point forecast results

Table A.4 lists the MSE evaluated over the out of sample period for models listed in Table A.2. Results are reported relative to the HAR-RV model.

[Insert Table A.4 about here.]

Results clearly favour the new SDFIx model against the HAR-RV. Specifically, the SDx- $\bar{\rho}$ is preferred 13 times over 21 and is the best performing model providing gains over the benchmark in the order of around 15%. After SDx- $\bar{\rho}$, the SDFIx- $\bar{\rho}$ specification is preferred 6 times, suggesting that for some models the definition of fractional integrated dynamics helps in forecasting the one step ahead realized log-volatility. Finally, for HSI2 and NSEI, the SDx model reports the lowest MSE. Interestingly, for the Asian indices the inclusion of time-varying correlation helps in the forecast. We can conclude that a time-varying leverage does not result in improved volatility point forecasts for equity indices. Fractional integration is important for some series, but results are mixed. Overall, the improvements over the benchmark are very evident and

⁵Given a random variable $y \in \mathbb{R}$ with continuous probability density function f , the scoring rule $S(f, y)$ is said to be *proper* if and only if $\mathbb{E}_f[S(f, y)] = \int_{\mathbb{R}} f(y)S(f, y)dy \leq \int_{\mathbb{R}} f(y)S(g, y)dy = \mathbb{E}_f[S(g, y)]$ for all density functions f and g .

of high magnitude. Results are statistically significant according to the test of Diebold and Mariano (1995) (not reported) computed at the 1% confidence interval.

In order to track the relative performance of the predictors over the out-of sample test period, as in Paye (2012), we compute the forecast error quadratic loss differential:

$$\Delta SE_{i,t+1} = (x_{t+1} - \hat{x}_{HAR,t+1})^2 - (x_{t+1} - \hat{x}_{i,t+1})^2, \quad (21)$$

where $\hat{x}_{HAR,t+1}$ is the forecast of HAR–RV and $\hat{x}_{i,t}$ is the forecast of model i , with i indexing the models listed in Table A.2. Figure B.4 displays the series of cumulative $\Delta SE_{i,t+1}$ for $i = \{\text{SDFIx}, \text{SDFIx}-\bar{\rho}\}$. Periods where the plot line slopes upward represent periods in which model i outperform the HAR–RV benchmark. Conversely, downward-sloping segments indicate periods of outperformance of the benchmark. We can see that SDFIx and SDFIx– $\bar{\rho}$ outperform the benchmark over the entire period. For some series like AORD2, GSPTSE, HSI2 and RUT2 we observe periods where the plot displays downward-sloping segments. Nevertheless, these are not related to the phase of the business cycle. Finally, the comparison between SDFIx and SDFIx– $\bar{\rho}$ also suggests that the leverage effect does not play a relevant role in improving the forecasting performance.

[Insert Figure B.4 about here.]

5.2.2. Density forecast results

Table A.5 reports the average CRPS over the test period for the models considered in Table A.2; the results are relative to the CRPS score of the HAR–RV model.

[Insert Table A.4 about here.]

Consistently with the MSE results, SDx models and nested specifications provide better results than the HAR–RV benchmark for all the series considered. The accuracy gains are generally around 7%. The SDx– $\bar{\rho}$ is still preferred most of the times (10 out of 21), however, for density predictions the results are more heterogeneous across the rival specifications. Indeed, we observe that for IBEX2, the most general specification, SDFIx, reports the lowest average CRPS.

Similar to the point forecast analysis we compute the loss differential in terms of CRPS with respect to the benchmark. To this end, we define the $\Delta CRPS$ as follow:

$$\Delta CRPS_{t+1} = CRPS_{HAR,t+1} - CRPS_{i,t+1}, \quad (22)$$

where $CRPS_{HAR,t+1}$ is the CRPS associated to the HAR–RV model and $CRPS_{i,t+1}$ to model i . Figure B.5 displays the series of cumulative $\Delta CRPS_{t+1}$ for $i = \{\text{SDFIx}, \text{SDFIx}-\bar{\rho}\}$. The interpretation of upward- and downward-sloping segments is unchanged. The plot is very similar to the point forecast analysis, that is, SDFIx and SDFIx– $\bar{\rho}$ outperform the benchmark over the full out of sample period and there is no particular gain in accommodating a time-varying correlation.

6. Conclusive remarks and possible extension

We have proposed a new model for a bivariate time series consisting of financial returns and realized log-volatility measures. The model is able to replicate several new features in a coherent framework such as: i) long memory or persistence of the volatility process, ii) the presence of time-varying leverage effect, and iii) the time-varying volatility of volatility. We have studied in details the shape of the News Impact Curve, and have shown that the new model is able to represent the leverage effect as usually intended in the literature. However, in our framework the definition of the leverage effect also depends on the concordance between the signs of financial return and its realized log-volatility measure as well as on the level of correlation.

Finally, the empirical evidence has provided strong support in favour of our proposal. An extensive forecast analysis based on the Oxford-Man Institute’s realized library data set has shown that from both point and density forecasting perspectives, our specifications could outperform a valid and solid univariate predictor of realized volatility. The exercise has also revealed that the conditional correlation between returns and log-realized variance is a stable feature.

As a direction for future research we foresee extensions that can encompass additional stylized facts, such as skewness. An interesting proposition is to consider a skewed- t distribution for our bivariate system.

References

- Andersen, T. G., Bollerslev, T., Diebold, F. X., and Labys, P. (2001). The distribution of realized exchange rate volatility. *Journal of the American statistical association*, 96(453):42–55.
- Asai, M., Chang, C.-L., and McAleer, M. (2016). Realized matrix-exponential stochastic volatility with asymmetry, long memory and spillovers.
- Barndorff-Nielsen, O. E., Hansen, P. R., Lunde, A., and Shephard, N. (2008). Designing realized kernels to measure the ex post variation of equity prices in the presence of noise. *Econometrica*, 76(6):1481–1536.
- Black, F. (1976). Studies of stock price volatility changes. In *Proceedings of the 1976 Meeting of the Business and Economic Statistics Section*, pages 177–181. American Statistical Association, Washington, D.C.
- Black, F. and Scholes, M. (1973). The pricing of options and corporate liabilities. *The journal of political economy*, pages 637–654.
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307 – 327.
- Bollerslev, T., Patton, A. J., and Quaedvlieg, R. (2016). Exploiting the errors: A simple approach for improved volatility forecasting. *Journal of Econometrics*, 192(1):1–18.
- Bollerslev, T., Tauchen, G., and Zhou, H. (2009). Expected stock returns and variance risk premia. *Review of Financial studies*, 22(11):4463–4492.
- Christie, A. A. (1982). The stochastic behavior of common stock variances: Value, leverage and interest rate effects. *Journal of financial Economics*, 10(4):407–432.
- Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics*, pages 174–196.
- Crane, B. (1959). *The Sophisticated Investor: A Guide to Stock-market Profits*. Simon and Schuster.
- Creal, D., Koopman, S. J., and Lucas, A. (2013). Generalized autoregressive score models with applications. *Journal of Applied Econometrics*, 28(5):777–795.

- Diebold, F. X. and Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & economic statistics*, 20(1):134–144.
- Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica*, 50(4):987–1007.
- Engle, R. F. and Ng, V. K. (1993). Measuring and testing the impact of news on volatility. *The journal of finance*, 48(5):1749–1778.
- Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The journal of Finance*, 25(2):383–417.
- Gerd, H., Lunde, A., Shephard, N., and Sheppard, K. (2009). Oxford-man institute’s realized library, version 0.2, oxford-man institute, university of oxford.
- Glosten, L. R., Jagannathan, R., and Runkle, D. E. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance*, 48(5):1779–1801.
- Hansen, P. R., Janus, P., and Koopman, S. J. (2016). Realized wishart-garch: A score-driven multi-asset volatility model.
- Hansen, P. R., Lunde, A., and Voev, V. (2014). Realized beta garch: a multivariate garch model with realized measures of volatility. *Journal of Applied Econometrics*, 29(5):774–799.
- Harvey, A. and Luati, A. (2014). Filtering with heavy tails. *Journal of the American Statistical Association*, 109(507):1112–1122.
- Harvey, A. C. (2013). *Dynamic models for volatility and heavy tails: with applications to financial and economic time series*, volume 52. Cambridge University Press.
- Harvey, A. C. and Shephard, N. (1996). Estimation of an asymmetric stochastic volatility model for asset returns. *Journal of Business & Economic Statistics*, 14(4):429–434.
- Jacquier, E., Polson, N. G., and Rossi, P. E. (2004). Bayesian analysis of stochastic volatility models with fat-tails and correlated errors. *Journal of Econometrics*, 122(1):185–212.
- Janus, P., Lucas, A., and Opschoor, A. (2016). New heavy models for fat-tailed returns and realized covariance kernels. *Journal of Business & Economic Statistics (forthcoming)*.
- Koopman, S. J., Lucas, A., and Scharth, M. (2016). Predicting time-varying parameters with parameter-driven and observation-driven models. *Review of Economics and Statistics*, 98(1):97–110.
- Kotz, S. and Nadarajah, S. (2004). *Multivariate t-distributions and their applications*. Cambridge University Press.
- Lucas, A. and Opschoor, A. (2016). Fractional integration and fat tails for realized covariance kernels and returns.
- Marcellino, M., Stock, J. H., and Watson, M. W. (2006). A comparison of direct and iterated multistep ar methods for forecasting macroeconomic time series. *Journal of econometrics*, 135(1):499–526.
- Matheson, J. E. and Winkler, R. L. (1976). Scoring rules for continuous probability distributions. *Management Science*, 22(10):1087–1096.
- Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica: Journal of the Econometric Society*, pages 347–370.
- Noureldin, D., Shephard, N., and Sheppard, K. (2012). Multivariate high-frequency-based volatility (heavy) models. *Journal of Applied Econometrics*, 27(6):907–933.
- Paye, B. S. (2012). dj vol: Predictive regressions for aggregate stock market volatility using macroeconomic variables. *Journal of Financial Economics*, 106(3):527 – 546.
- Proietti, T. (2016). Component-wise representations of long-memory models and volatility prediction. *Journal of Financial Econometrics*, pages 668–692.
- Shephard, N. and Sheppard, K. (2010). Realising the future: forecasting with high-frequency-based volatility (heavy) models. *Journal of Applied Econometrics*, 25(2):197–231.
- Takahashi, M., Omori, Y., and Watanabe, T. (2009). Estimating stochastic volatility models using daily returns and realized

- volatility simultaneously. *Computational Statistics & Data Analysis*, 53(6):2404–2426.
- Takahashi, M., Omori, Y., and Watanabe, T. (2013). News impact curve for stochastic volatility models. *Economics Letters*, 120(1):130–134.
- Taylor, S. J. (1986). *Modelling Financial Times Series*. Wiley.
- Yu, J. (2005). On leverage in a stochastic volatility model. *Journal of Econometrics*, 127(2):165–178.
- Zakoian, J.-M. (1994). Threshold heteroskedastic models. *Journal of Economic Dynamics and control*, 18(5):931–955.

Appendix A. Tables

Series	Correlation				Volatility of Volatility			
	2000/2003	2004/2007	2008/2011	2012/2016	2000/2003	2004/2007	2008/2011	2012/2016
AEX	-0.09	-0.17	-0.14	-0.18	0.46	0.35	0.49	0.40
AORD2	-0.05	-0.16	-0.13	-0.13	0.39	0.43	0.49	0.37
BVSP	-0.08	-0.21	-0.08	0.01	0.31	0.35	0.50	0.32
DJI2	-0.07	-0.13	-0.09	-0.15	0.37	0.32	0.56	0.40
FCHI2	-0.07	-0.14	-0.13	-0.14	0.38	0.35	0.45	0.36
FTSE2	-0.09	-0.12	-0.10	-0.13	0.41	0.38	0.48	0.33
FTSEMIB	-0.11	-0.21	-0.15	-0.18	0.43	0.35	0.47	0.35
FTSTI	-0.01	-0.15	0.02	-0.03	0.29	0.27	0.46	0.22
GDAXI2	-0.14	-0.22	-0.09	-0.18	0.41	0.34	0.49	0.37
GSPTSE	-0.01	-0.26	-0.16	-0.15	0.36	0.35	0.52	0.39
HSI2	-0.14	-0.16	-0.03	-0.09	0.34	0.38	0.44	0.32
IBEX2	-0.11	-0.14	-0.09	-0.19	0.36	0.35	0.40	0.35
IXIC2	-0.06	-0.20	-0.09	-0.21	0.42	0.31	0.52	0.36
KS11	-0.05	-0.21	-0.13	-0.10	0.35	0.36	0.50	0.31
MXX	-0.05	-0.11	-0.06	-0.10	0.47	0.39	0.45	0.29
N2252	-0.01	-0.10	-0.09	-0.14	0.29	0.37	0.49	0.43
NSEI	-0.09	-0.29	-0.15	-0.16	0.40	0.42	0.49	0.33
RUT2	-0.08	-0.16	-0.09	-0.17	0.39	0.35	0.51	0.39
SPX2	-0.07	-0.12	-0.10	-0.18	0.38	0.35	0.56	0.41
SSMI	-0.03	-0.18	-0.07	-0.22	0.45	0.29	0.47	0.33
STOXX50E	-0.10	-0.21	-0.13	-0.16	0.44	0.36	0.46	0.43

Table A.1: Columns two to five report the empirical correlation coefficient between daily financial returns and their associated realized log-volatility measure across different years. The last four columns report the standard deviation of the log-volatility across different years. Data is obtained from the Oxford-Man Institute's realized library Gerd et al. (2009) freely available at: <http://realized.oxford-man.ox.ac.uk/data>. The first column reports the ticker associated to the equity index as found in the data set.

<i>With Fractional Integrated Dynamics</i>	
SDFIx	The most general specification defined in Section 2.
SDFIx- $\bar{q}$	SDFIx with constant volatility of volatility, $q_t = q$.
SDFIx- $\bar{\rho}$	SDFIx with constant correlation, $\rho_t = \rho$.
<i>Without Fractional Integrated Dynamics</i>	
SDx	SDFIx with $d^\mu = 0$ and $d^q = 0$.
SDx- $\bar{q}$	SDx with constant volatility of volatility, $q_t = q$.
SDx- $\bar{\rho}$	SDx with constant correlation, $\rho_t = \rho$.

Table A.2: List of model specifications. Starting from the most general Score Driven Fractional Integrated Realized Volatility (SDFIx) model, nested specifications are recovered by applying constraints to the system matrices of the dynamic equation (5).

	<i>With Fractional Integrated Dynamics</i>			<i>Without Fractional Integrated Dynamics</i>		
Index	SDFIx	SDFIx- $\bar{q}$	SDFIx- $\bar{\rho}$	SDx	SDx- $\bar{q}$	SDx- $\bar{\rho}$
AEX	-27759.92	-27738.81	-27770.96	-27770.36	-27747.17	-27787.08
AORD2	-27577.62	-27576.06	-27590.10	-27591.34	-27584.40	-27606.31
BVSP	-21773.42	-21781.35	-21783.07	-21789.97	-21789.68	-21799.65
DJI2	-26760.76	-26704.14	-26775.51	-26777.44	-26712.48	-26792.19
FCHI2	-26932.65	-26904.10	-26949.40	-26948.30	-26912.46	-26965.02
FTSE2	-29589.02	-29587.17	-29605.39	-29605.72	-29595.51	-29622.09
FTSEMIB	-26055.62	-26043.79	-26071.28	-26072.32	-26052.14	-26087.98
FTSTI	-28774.05	-28697.24	-28786.90	-28790.58	-28705.50	-28803.42
GDAXI2	-26558.99	-26528.92	-26579.47	-26575.05	-26537.27	-26589.69
GSPTSE	-24446.28	-24457.88	-24460.39	-24462.66	-24466.07	-24476.77
HSI2	-25852.87	-25846.55	-25865.49	-25865.30	-25854.81	-25879.17
IBEX2	-26686.75	-26653.52	-26701.11	-26700.88	-26661.87	-26717.53
IXIC2	-26071.18	-26053.94	-26032.47	-26087.87	-26062.29	-26049.15
KS11	-26876.77	-26876.19	-26883.11	-26888.12	-26884.51	-26899.05
MXX	-22884.82	-22889.40	-22900.63	-22901.13	-22897.74	-22914.89
N2252	-25420.06	-25396.65	-25437.27	-25434.13	-25404.96	-25449.47
NSEI	-23100.72	-23082.87	-23116.25	-23117.10	-23091.06	-23132.55
RUT2	-22334.58	-22262.54	-22320.96	-22351.26	-22270.88	-22337.65
SPX2	-26296.62	-26268.81	-26295.53	-26308.23	-26277.15	-26310.96
SSMI	-30423.29	-30414.71	-30436.91	-30439.95	-30423.05	-30453.59
STOXX50E	-25406.86	-25357.14	-25430.25	-25423.10	-25365.49	-25439.50

Table A.3: BIC evaluated over the full sample for the model specifications defined in Table A.2. Sample sizes vary across the series and range from $T = 3605$ to $T = 4272$ covering the period from 3 January 2000 to 11 October 2016. Gray cells indicate models with lower BIC values, and hence preferred from a likelihood perspective.

Index	SDFI _x	SDFI _{x-$\bar{q}$}	SDFI _{x-$\bar{\rho}$}	SD _x	SD _{x-$\bar{q}$}	SD _{x-$\bar{\rho}$}	HAR-RV
AEX	0.854	0.854	0.853	0.853	0.853	0.853	1.000
AORD2	0.950	0.950	0.950	0.950	0.950	0.950	1.000
BVSP	0.896	0.895	0.895	0.896	0.896	0.896	1.000
DJI2	0.903	0.903	0.898	0.903	0.903	0.898	1.000
FCHI2	0.858	0.858	0.858	0.857	0.857	0.857	1.000
FTSE2	0.845	0.845	0.844	0.846	0.846	0.846	1.000
FTSEMIB	0.850	0.850	0.848	0.850	0.850	0.848	1.000
FTSTI	0.879	0.879	0.876	0.878	0.878	0.876	1.000
GDAXI2	0.874	0.874	0.874	0.873	0.873	0.872	1.000
GSPTSE	0.912	0.912	0.912	0.912	0.912	0.912	1.000
HSI2	0.949	0.950	0.949	0.947	0.947	0.948	1.000
IBEX2	0.811	0.811	0.811	0.810	0.810	0.810	1.000
IXIC2	0.827	0.827	0.811	0.827	0.827	0.811	1.000
KS11	0.846	0.845	0.846	0.843	0.843	0.843	1.000
MXX	0.930	0.929	0.926	0.931	0.931	0.928	1.000
N2252	0.873	0.873	0.872	0.873	0.873	0.872	1.000
NSEI	0.830	0.830	0.831	0.829	0.829	0.830	1.000
RUT2	0.909	0.909	0.903	0.908	0.908	0.902	1.000
SPX2	0.882	0.882	0.871	0.883	0.883	0.873	1.000
SSMI	0.845	0.845	0.840	0.845	0.845	0.840	1.000
STOXX50E	0.861	0.860	0.861	0.860	0.860	0.860	1.000

Table A.4: Mean squared error evaluated over the out of sample period for models listed in Table A.2. Results are reported relative to the HAR-RV model of Corsi (2009). A value lower than 1 indicates outperformance with respect to the benchmark. Gray cells indicate models with lower value across the different specifications.

Index	SDFI _x	SDFI _{x-$\bar{q}$}	SDFI _{x-$\bar{\rho}$}	SD _x	SD _{x-$\bar{q}$}	SD _{x-$\bar{\rho}$}	HAR-RV
AEX	0.922	0.921	0.921	0.921	0.921	0.921	1.000
AORD2	0.978	0.978	0.978	0.978	0.977	0.978	1.000
BVSP	0.943	0.943	0.943	0.944	0.944	0.943	1.000
DJI2	0.950	0.950	0.948	0.950	0.950	0.948	1.000
FCHI2	0.920	0.920	0.920	0.919	0.919	0.919	1.000
FTSE2	0.915	0.916	0.915	0.916	0.916	0.916	1.000
FTSEMIB	0.919	0.919	0.918	0.918	0.918	0.917	1.000
FTSTI	0.936	0.936	0.936	0.937	0.936	0.935	1.000
GDAXI2	0.928	0.928	0.928	0.928	0.928	0.928	1.000
GSPTSE	0.955	0.955	0.955	0.955	0.954	0.955	1.000
HSI2	0.964	0.964	0.964	0.963	0.963	0.962	1.000
IBEX2	0.890	0.891	0.891	0.891	0.891	0.891	1.000
IXIC2	0.905	0.906	0.897	0.906	0.906	0.897	1.000
KS11	0.915	0.915	0.915	0.915	0.915	0.915	1.000
MXX	0.961	0.961	0.958	0.961	0.962	0.960	1.000
N2252	0.930	0.930	0.930	0.930	0.930	0.930	1.000
NSEI	0.909	0.909	0.909	0.908	0.908	0.909	1.000
RUT2	0.951	0.951	0.948	0.950	0.950	0.947	1.000
SPX2	0.941	0.941	0.935	0.941	0.941	0.935	1.000
SSMI	0.915	0.915	0.912	0.915	0.915	0.912	1.000
STOXX50E	0.920	0.920	0.920	0.920	0.920	0.920	1.000

Table A.5: Average Continuous Ranked Probability Score evaluated over the out of sample period for models listed in Table A.2. Results are reported relative to the HAR-RV model of Corsi (2009). Gray cells indicate the minimum values across the rival specifications.

Appendix B. Figures




Figure B.1: News Impact Curve with respect to y_t (horizontal axis) and selected values of x_t (RV_t): i) decrease of log-volatility $x_t = -5.5$, ii) no changes in volatility $x_t = \mu$, iii) increase in volatility, $x_t = 5.5$. Models' parameters are fixed at their unconditional level similar to what is usually found in empirical applications with financial returns in percentage points: $\mu_t = \mu = 0$, $q_t = q = 1$. The system matrix that premultiplies the score is fixed to $\mathbf{A} = 0.01 \times \mathbf{I}_3$. Panels (a), (b) and (c) report the case of negative correlation between y_t and x_t ($\rho = -0.7$). Panels (d), (e) and (f) report the case of absence of correlation between y_t and x_t ($\rho = 0$). Panels (g), (h) and (i) report the case of positive correlation between y_t and x_t ($\rho = 0.7$). NICs are displayed for different values of the degree of freedom parameter, $\nu = 3$ (black, continuous), $\nu = 6$ (red, dashed), $\nu = 9$ (green, dotted). The horizontal axis indicates the financial return at time t , y_t , while the vertical axis indicates the corresponding volatility level at time $t+1$, $\sigma_{t+1} = \exp(\mu_{t+1})$. The orange horizontal line is reported as a benchmark and indicates the volatility level observed in the case of $(y_t, x_t)' = (0, \mu)'$ and $\rho = 0$.




Figure B.2: News Impact Curve with respect to x_t (horizontal axis) and selected values of y_t : i) large negative return, $y_t = -7$, ii) zero return, $y_t = 0$, iii) large positive return, $y_t = 7$. Models' parameters are fixed at their unconditional level similar to what is usually found in empirical applications with financial returns in percentage points: $\mu_t = \mu = 0$, $q_t = q = 1$. The system matrix that premultiplies the score is fixed to $\mathbf{A} = 0.01 \times \mathbf{I}_3$. Panels (a), (b) and (c) report the case of negative correlation between y_t and x_t ($\rho = -0.7$). Panels (d), (e) and (f) report the case of absence of correlation between y_t and x_t ($\rho = 0$). Panels (g), (h) and (i) report the case of positive correlation between y_t and x_t ($\rho = 0.7$). NICs are displayed for different values of the degree of freedom parameter, $\nu = 3$ (black, continuous), $\nu = 6$ (red, dashed), $\nu = 9$ (green, dotted). The horizontal axis indicates the financial return at time t , y_t , while the vertical axis indicates the corresponding volatility level at time $t + 1$, $\sigma_{t+1} = \exp(\mu_{t+1})$. The orange horizontal line is reported as a benchmark and indicates the volatility level observed in the case of $(y_t, x_t)' = (0, \mu)'$ and $\rho = 0$.




Figure B.3: Filtered parameters for the SDx model estimated on the IXIC2 (the NASDAQ Composite index) series. The period starts on the 3 January 2000 and ends the 11 October 2016. Panel (a) depicts the filtered log-volatility level μ_t (solid line) and the observed realized log-volatility (gray points). Panel (b) reports the filtered volatility of volatility, q_t . Panel (c) reports the correlation coefficient $\rho_t = \text{cor}(y_t, x_t)$. Blue and red vertical bands indicate periods of US and European recessions according to the OECD and FRED Recession Indicators, respectively.




Figure B.4: Time series plot of the cumulative squared error difference between the SDFIx and HAR-RV models (*black solid line*) and SDFIx- $\bar{p}$ and HAR-RV models (*red dashed line*). Blue and red vertical bands indicate periods of US and European recessions according to the OECD and FRED Recession Indicators, respectively. The series STOX50E is not reported for space reasons, but is qualitatively similar to other plots.




Figure B.5: Time series plot of the cumulative CRPS difference (ΔCRPS) between the SDFIx and HAR-RV models (*black solid line*) and SDFIx- $\bar{p}$ and HAR-RV models (*red dashed line*). Blue and red vertical bands indicate periods of US and European recessions according to the OECD and FRED Recession Indicators, respectively. The series STOXX50E is not reported for space reasons, but is qualitatively similar to other plots.