

CIRJE-F-693

Realized Volatility Risk

David E. Allen
Edith Cowan University

Michael McAleer
Erasmus University Rotterdam
and Tinbergen Institute

Marcel Scharth
VU University Amsterdam
and Tinbergen Institute

December 2009; Revised in January 2010

CIRJE Discussion Papers can be downloaded without charge from:

<http://www.e.u-tokyo.ac.jp/cirje/research/03research02dp.html>

Discussion Papers are a series of manuscripts in their draft form. They are not intended for circulation or distribution except as indicated by the author. For that reason Discussion Papers may not be reproduced or distributed without the written consent of the author.

REALIZED VOLATILITY RISK

David E. Allen^a, Michael McAleer^{b,c,d} and Marcel Scharth^{c,e,*}

^a School of Accounting, Finance and Economics, Edith Cowan University

^b Econometric Institute, Erasmus University Rotterdam

^c Tinbergen Institute, The Netherlands

^d Centre for International Research on the Japanese Economy (CIRJE), Faculty of Economics, University of Tokyo

^e Department of Econometrics, VU University Amsterdam

This version: December 2009

ABSTRACT. In this paper we document that realized variation measures constructed from high-frequency returns reveal a large degree of volatility risk in stock and index returns, where we characterize volatility risk by the extent to which forecasting errors in realized volatility are substantive. Even though returns standardized by *ex post* quadratic variation measures are nearly gaussian, this unpredictability brings considerably more uncertainty to the empirically relevant *ex ante* distribution of returns. Carefully modeling this volatility risk is fundamental. We propose a dually asymmetric realized volatility (DARV) model, which incorporates the important fact that realized volatility series are systematically more volatile in high volatility periods. Returns in this framework display time varying volatility, skewness and kurtosis. We provide a detailed account of the empirical advantages of the model using data on the S&P 500 index and eight other indexes and stocks.

KEYWORDS: Realized volatility, volatility of volatility, volatility risk, value-at-risk, forecasting, conditional heteroskedasticity.

1. INTRODUCTION

The availability of high frequency stock market data and the subsequent introduction of realized volatility measures represented a substantial step forward in the accuracy with which econometric models of volatility could be evaluated (Andersen and Bollerslev, 1998) and allowed for the development of new and more precise parametric models of time varying volatility. Several researchers have looked into the properties of *ex post* volatility measures derived from high frequency data and developed time series models that invariably outperform latent variable models of the GARCH (Generalized Autoregressive Conditional Heteroskedasticity) or stochastic volatility family of models in forecasting volatility (Andersen et al., 2003), to the point that the comparison has been dropped altogether in recent papers.

Despite the successes of predictive models of realized volatility we depart from the recent literature in that turn our focus to the *unpredictable* component of realized volatility. In this work we set out to argue that a large and time varying realized volatility risk (understood as the time series volatility of realized volatility) is an essential stylized fact of index and stock

*Contact author: Marcel Scharth (mscharth@feweb.vu.nl).

Acknowledgements: the authors wish to acknowledge the insightful comments and suggestions of Marcelo Medeiros. The financial support of the Australian Research Council is gratefully acknowledged. Michael McAleer wishes to acknowledge the financial support of the National Science Council, Taiwan. Data supplied by Securities Industry Research Centre of Asia-Pacific (SIRCA).

returns that should be thoughtfully incorporated into econometric models of volatility.¹ Our main contribution is twofold. First, we provide extensive empirical and theoretical motivation showing how the stochastic structure of the innovations in volatility have fundamental implications for applications of realized volatility models. Second, we go beyond extensions of standard realized volatility models to account for conditional heteroskedasticity (e.g., Corsi et al., 2008) and bring to the forefront of our modeling approach the fact that realized volatility series exhibit a very substantial degree of (time-varying) volatility themselves. As we shall see this emphasis brings an important modeling insight for our analysis: realized volatility series are systematically more volatile in high volatility periods.

The presence of high and time-varying volatility risk brings substantially more uncertainty to the tails of the distribution of asset returns. In a standard stochastic volatility setting where return innovations (conditional on the volatility) follow a gaussian distribution the degree of volatility risk is the main determinant of the tail heaviness of the ex ante distribution of returns. Even though asset returns standardized by ex post quadratic variation measures are nearly gaussian, returns standardized by fitted or predicted values of time series volatility models are far from normal. Given the uncertainty in volatility this is expected and should not be seen as evidence against those models; explicitly modeling the higher moment is necessary. If future realized volatility is relatively unpredictable then a focus on predictive models will be insufficient for meaningfully modeling the tails of the return distribution, which in many cases (e.g., in risk management applications) is the main objective of the econometrician.

The assumptions we make about this volatility risk are accordingly consequential. Our paper is a first step in trying to fully exploit the fact that the realized volatility framework allows not only for significant advances in modeling the conditional volatility of asset returns, but also the higher moments. When realized volatility is available we do not have to rely only on rare realizations on return data to identify the extremes of the return distribution: naturally, days of very high volatility are far more frequent than days of very high volatility *and* large return shocks. Likewise, a model with time-varying return kurtosis (an implication of time varying volatility risk, for which there is conclusive evidence in our data set) which would be very hard to identify from return data only can be easily estimated in a realized volatility framework.

A related point that we will make in arguing for a stronger focus on volatility risk is that forecasting improvements brought by more sophisticated predictive realized volatility models can be swamped by the size of the volatility of realized volatility. In a representative example, we calculate that even an exponentially weighted moving average (EWMA) of realized volatility delivers predictions which are very close to more sophisticated specifications; it is crucial to account for the fact that volatility series are highly persistent, but the way this is done may have little economic relevance. We thus take the view that even though point forecasts are a very useful tool for evaluating and ranking volatility models (extensive comparisons of forecasting performance like the one performed by Hansen and Lunde (2005) are common in the literature) these statistics do not necessarily convey much information about the relative economic importance of different specifications (see Fleming et al., 2001, 2003, Chan et al., 2006, for an empirical analysis of the economic value of volatility models).

¹In this paper we privilege the term “volatility risk” used commonly in the asset pricing literature over the more awkward “volatility of volatility” denomination, which is sometimes confused with the volatility of the realized volatility estimator used for the nonparametric measurement of volatility.

In the second part of the paper we propose a new model for returns, realized volatility and volatility risk. In the first paper to consider the volatility of realized volatility, Corsi et al. (2008) extend the typical framework for modeling realized volatility by specifying a GARCH process to allow for clustering in the squared residuals of their realized volatility model. The same approach is followed by Bollerslev et al. (2009). We suggest however that this approach ignores an important empirical regularity observed from realized volatility series: the degree of volatility risk is strongly and positively related to the current level of volatility. While this finding has been suggested before in the option pricing literature (see for example Heston, 1993, Jones, 2000), this important relation has not received attention in the volatility literature. In this paper we consider a parsimonious specification where the variance of the realized volatility innovations is a linear function of the square of the volatility level, which we take to be the conditional mean of realized volatility. Another salient aspect of our model is the emphasis on extended leverage effects (following Scharth and Medeiros, 2009) which introduce realistic persistence patterns for volatility. Because of the returns/volatility, volatility/volatility risk asymmetries, we call this framework the dually asymmetric realized volatility (DARV) model.

Our empirical analysis uses high frequency data for the S&P 500 index and eight more series (between major stocks and indexes) from 1996 to 2009 to document the importance of volatility risk and analyze the performance of the dually realized volatility model when compared to other standard alternatives. We show that our volatility risk specification consistently improves forecasting performance across our series and substantially enhances the ability of the realized volatility model to account for large movements in volatility. Our results for the volatility of realized volatility are stronger than the ones obtained by Corsi et al. (2008) in that we can conclude that ignoring volatility risk has a severe adverse impact on point and density forecasting for realized volatility and the calculation of value-at-risk intervals for returns.

Other contributions to the realized volatility modeling and forecasting literature are exemplified by Andersen et al. (2003), the HAR (heterogeneous autoregressive) model of Corsi (2004), the MIDAS (mixed data sample) approach of (Ghysels et al., 2007) and the unobserved ARMA component model of Koopman et al. (2005), and Shephard and Sheppard (2009). Martens et al. (2004) develop a nonlinear (ARFIMA) model to accommodate level shifts, day-of-the-week, leverage and volatility level effects. Andersen et al. (2007) and Tauchen and Zhou (2005) argue that the inclusion of jump components significantly improves forecasting performance. McAleer and Medeiros (2008a) extend the HAR model to account for nonlinearities. Hillebrand and Medeiros (2007) also consider nonlinear models and evaluate the benefits of bootstrap aggregation (bagging) for volatility forecasting. Ghysels et al. (2007) argue that realized absolute values outperform square return-based volatility measures in predicting future increments in quadratic variation. Scharth and Medeiros (2009) introduce multiple regime models linked to asymmetric effects. Liu and Maheu (2007) derive a bayesian averaging approach for forecasting realized volatility. Bollerslev et al. (2009) propose a full system for returns, jumps and continuous time for components of price movements using realized variation measures.

This paper is structured as follows. The next section presents our empirical setup, describing our data, the realized volatility measurement method, and the stylized facts. We then proceed to discuss the properties of the ex ante distribution of the S&P500 returns when high-frequency data information is available and the relation between the stochastic structure of realized volatility shocks and the conditional distribution of returns. Section three introduces our model for realized

volatility and describes how Monte Carlo techniques can be used for translating the features of our conditional volatility, skewness and kurtosis framework into refined density forecasts for returns. In section four we consider the empirical performance of our model. Section five concludes.

2. DATA AND STYLIZED FACTS

2.1. Realized volatility and Data. Suppose that at day t the logarithmic prices of a given asset follow a continuous time diffusion:

$$dp(t + \tau) = \mu(t + \tau) + \sigma(t + \tau)dW(t + \tau), \quad 0 \leq \tau \leq 1, \quad t = 1, 2, 3, \dots$$

where $p(t + \tau)$ is the logarithmic price at time $t + \tau$, $\mu(t + \tau)$ is the drift component, $\sigma(t + \tau)$ is the instantaneous volatility (or standard deviation), and $dW(t + \tau)$ is a standard Brownian motion. Andersen et al. (2003) and Barndorff-Nielsen and Shephard (2002) showed that the daily compound returns, defined as $r_t = p(t) - p(t + 1)$, are Gaussian conditionally on $\mathcal{F}_t = \sigma(p(s), s \leq t)$, the σ -algebra (information set) generated by the sample paths of p , such that

$$r_t | \mathcal{F}_t \sim N \left(\int_0^1 \mu(t - 1 + \tau) d\tau, \int_0^1 \sigma^2(t - 1 + \tau) d\tau \right)$$

The term $IV_t = \int_0^1 \sigma^2(t - 1 + \tau) d\tau$ is known as the integrated variance, which is a measure of the day t *ex post* volatility. In this sense the integrated variance is the object of interest. In practical applications prices are observed at discrete and irregularly spaced intervals and the most widely used sampling scheme is calendar time sampling (CTS), where the intervals are equidistant in calendar time. If we set $p_{i,t}$, $i = 1, \dots, n$ to be the i th price observation during day t , realized variance is defined as $\sum_{i=1}^n r_{i,t}^2$. The realized volatility is the square-root of the realized variance and we shall denote it by RV_t . Ignoring the remaining measurement error, this *ex post* volatility measure can be modeled as an “observable” variable, in contrast to latent variable models.

In real data high frequency measures are contaminated by microstructure noise. The search for unbiased, consistent and efficient methods for measuring realized volatility has been one of the most active research topics in financial econometrics over the last years. While early references such as Andersen et al. (2001) suggest the simple selection of an arbitrary frequency to balance accuracy and dissipation of microstructure bias, a procedure known as *sparse sampling*, a number of recent articles developed estimators that dominate this procedure. In this paper we turn to the theory developed by Barndorff-Nielsen et al. (2008, 2009) and implement the consistent realized kernel estimator based on the modified Tukey-Hanning kernel. Some alternatives are the two time scales estimator of Zhang et al., 2005 and Ait-Sahalia et al., 2005, the multiscale estimator of Zhang (2006) and the preaveraging estimator of Jacod et al. (2009). See McAleer and Medeiros (2008b) and Gatheral and Oomen (2007) for a review and comparison of methods.

Our empirical analysis will focus on the realized volatility of the S&P 500 (SPX), Dow Jones (DJIA), FTSE 100, CAC 40 and Nikkei 225 indexes and the IBM, GE, Wal-Mart (WMT) and AT&T stocks. For conciseness the S&P 500 index will be at the center of our analysis, with the other series being used when appropriate to show that our results hold more generally. The raw intraday data was obtained from the Reuters Datascope Tick History database and consists of tick by tick open to close quotes filtered for possible errors. For the S&P 500 index we use the information originated in the E-Mini S&P500 futures market of the Chicago Mercantile Exchange, while

for the remaining indexes we use the actual index price series from different sources.² Following the results of Hansen and Lunde (2006), we adopt the *previous tick* method for determining prices at time marks where a quote is missing.

The period of analysis starts in January 2, 1996, and ends in June 30, 2009, providing a total of 3343 trading days in the United States. We clarify that our in-sample period used for revising the stylized facts, presenting the volatility risk findings and discussing the estimation diagnostics covers the whole sample, while the out-of-sample period used in section four runs from 2001 to the end of the sample. We need our out-of-sample period to be unusually long since the behavior of realized volatility markedly favors different kinds of models in particular years (for example, crisis periods strongly favor models with leverage effects) and a reasonable number of tail realizations are necessary to compare different alternatives for modeling volatility risk.

2.2. Stylized Facts. We begin by recalling the most important and generally accepted features of returns and realized volatility that have been identified in the literature and will inform the rest of our analysis:

- (1) the unconditional distribution of daily returns exhibits excess kurtosis;
- (2) the time variation in the conditional volatility of daily returns does not fully account for this excess kurtosis;
- (3) daily returns are not autocorrelated (except for the first order, in some cases);
- (4) daily returns standardized by realized volatility measures are almost Gaussian;
- (5) the unconditional and conditional distributions of realized variance and volatility are distinctly non-normal and extremely right-skewed;
- (6) there is strong evidence of long memory in volatility;
- (7) negative returns are associated with higher subsequent volatility (asymmetric effects).

Andersen et al. (2001) and Andersen et al. (2003) are the prime references for the stylized facts of realized volatility. We look at the S&P 500 index data for a brief empirical review of these stylized facts and the main features of the data. Figure 1 displays the time series of returns, realized volatility and log realized volatility. Table 1 presents descriptive statistics for returns, standardized returns, realized volatility and changes in realized volatility, illustrating (1), (4) and (5). Even though we do not report the descriptives for the other series we remark that one common feature for the indexes (but not the stocks) is the apparent thin tails of the scaled returns. This is not peculiar to our data set; see for example Bollerslev et al. (2009). This is probably due to jumps, which we will leave out of our core analysis but will discuss later in the paper. One striking feature of Table 1 is the extreme leptokurtosis in the realized volatility changes (ΔRV_t). In fact, only 10% of observations account for close to 80% of the variation in realized volatility across the sample. One of main objectives of this paper is to better account for these large movements.

Turning to the dynamic behavior of the volatility, the quintessential time series property observed in realized volatility has been the very high persistence in those series, which consistently exhibit empirical autocorrelation functions with hyperbolic decaying patterns; this empirical feature can be seen on Figure 2 and is referred to in a broad sense as long memory or long range dependence. Motivated by early empirical work in macroeconomics and later in the GARCH and Stochastic Volatility literature (for instance Baillie et al., 1996, Comte and Renault, 1996,

²The fully electronic E-Mini S&P500 futures contracts feature among the most liquid derivatives contracts in the world, therefore closely tracking price movements of the S&P 500 index. The index prices used for the other series are unfortunately less frequently quoted. The volatility measurements for DJIA, FTSE 100, CAC 40 and Nikkei 225 indexes are therefore of somewhat inferior quality compared to the S&P 500 index and the stocks.

and Breidt et al., 1998), the first prominent references in the realized volatility literature (e.g., Andersen et al., 2003) turned to the econometrics of fractional integration to establish a framework for modeling and forecasting realized volatility. Although no theoretical foundation has been developed to substantiate the long memory specification or elucidate the high persistence from past shocks, $I(d)$ processes emerged as a consonant description of the data generating process of volatility series, becoming a standard paradigm.

The presence of asymmetric effects has long been documented and is first illustrated by Figure 3. The first part of Figure 3 is the standard news impact curve due to Engle and Ng (1993), showing the correlation between daily returns and the next day realized volatility. In the bottom part of the figure we replace the lagged returns by the lagged monthly returns, revealing an even higher correlation for this case. This extended leverage effect has been pointed out by Scharth and Medeiros (2009) and is more starkly shown in Figure 4, which plots the time series of realized volatility and monthly returns (rescaled). The sample correlation between the two series is -0.52 . It seems that virtually all episodes of (persistently) high volatility are associated with streams of negative returns; once the index price recovers the realized volatility tends to quickly fall back to average levels.

2.3. Volatility Risk. One of the two main objectives of this paper is to extend the list of stylized facts in the econometrics of volatility with a new item: time series of realized volatility are very volatile and this volatility risk is strongly positively related to the level of volatility. Even though the noisy appearance of the time series of realized volatility (as seen for example in Figure 1) has been noted since the early days of this literature, this feature of the data has neither received much attention so far Corsi et al. (with the exception of 2008) or been considered very important. We start our analysis by asking more rigorously: how large is the volatility risk for stocks and indexes? Or equivalently, how predictable is volatility? To answer this question we start by considering the magnitude of the prediction errors that we can expect to find in a standard forecasting model for realized volatility.

The HAR (Heterogeneous Autoregressive) model proposed by Corsi (2004) is an unfolding of the Heterogeneous ARCH (HARCH) model developed by Müller et al. (1997). It is specified as a multi-component volatility model with an additive hierarchical structure, leading to an additive time series model of the realized volatility which specifies the volatility as a sum of volatility components over different horizons. Observing that we shall be working with time series models directly for the realized volatility (RV_t), consider the notation (standard for this model) $RV_{j,t} = \sum_{i=t-j+1}^t RV_{t,i}/j$. We can then write the HAR model with daily, weekly and monthly components as:

$$(1) \quad RV_t = \phi_0 + \phi_1 RV_{t-1} + \phi_2 RV_{5,t-1} + \phi_3 RV_{22,t-1} + \beta X_{t-1} + \eta_t,$$

where X_t is a vector of additional explanatory variables. We can see that the HAR specification is an $AR(22)$ model rendered parsimonious by several parameter restrictions. Simulations reported in Corsi (2004) show that the generous number of autoregressive lags renders the HAR model capable of reproducing the observed hyperbolic decay of the sample autocorrelations of realized volatility series over not long too long horizons. For its estimation simplicity the HAR model has been commonly favored in the high frequency econometrics literature (e.g., Andersen et al., 2007).

To get a first idea of the bearing of volatility risk in our data we examine the in-sample forecasting (fitting) performance of the HAR model and some simple alternatives for the S&P 500 series

(thus considering a setting where the errors of the model are minimized in the sample; stronger patterns would naturally hold in an out of sample analysis). We consider three versions of the HAR specification: one in which there are no additional explanatory variables, one in which we include the leverage effect term $I(r_{t-1} < 0)r_{t-1}$ and the equivalent variables for 5 and 22 days cumulated returns (in line with our previous discussion of asymmetric effects) and a third model in which we extend this last specification by adding the lagged realized absolute variation (RAV_{t-1}) as an explanatory variable. Ghysels et al. (2007) argues that this variable improves forecasting performance in a similar way to considering jump components in the analysis. To articulate some of the issues we also consider some ad hoc or clearly misspecified models: the lagged realized volatility as a predictor, an exponentially weighted moving average (EWMA) of realized volatility, and an ARMA(2,1) model.

Table 2 shows some statistics for one, five and twenty two days ahead forecasts. The first indication of high volatility risk comes from the fact that even for the best model (the third HAR specification, denoted HAR+AE+RAV in the table) the root mean squared error for the one day forecasts is unambiguously large (0.285) when compared for example to the median realized volatility (0.844). The results imply an average conditional coefficient of variation for realized volatility of nearly 0.3. The high forecasting R^2 s are a consequence of the high persistence in volatility and do not help us in identifying the importance of volatility risk. To keep this in perspective we adopt in this paper the practice of also reporting the forecasting R^2 for the changes in volatility (ΔRV_t). Of course, this does not alter the ranking of the models considered, but helps us to underline the relative advantages of different specifications.

This metric makes it clear that substantial fitting improvements over the ad hoc models come essentially from leverage effects, highlighting the difficulty in developing better predictive models of realized volatility: it seems that the HAR model by itself introduces little or no improvement over the EWMA and ARMA models. The benefit of the realized absolute variation predictor is also marginal in our sample. Despite the favorable results reported in Ghysels et al. (2007), Andersen et al. (2007) and other papers, we hypothesize that once more efficient non-parametric methods for measuring volatility are implemented the benefits of considering the RAV and currently available jump measures are small. More troublesome, however, is the observation that the reduction in the mean absolute error of going from a simple moving average to the best model is of only 0.016, which seems of very small economic substance compared with either the average daily volatility(0.983) or the magnitude of the forecasting errors. An analysis of longer horizon forecasts yields the same conclusions. Our exercise illustrates how forecasting improvements brought by more sophisticated realized volatility models can often be swamped by the size of the volatility of realized volatility.

The presence of high volatility risk in our data can also be seen from the percentage errors. Table 2 shows that all models are subject to large average percentage errors even for one step predictions (an average of 17.5% at the best specification). More evidence comes from Figure 5, which plots the ratio between the realized volatility observations and the forecasts over time. The plot shows that even 50% forecasting errors are not uncommon for realized volatility. This is shown more systematically by Table 3, which displays the sample statistics of this ratio for all series. The ratio is extremely skewed to the right. For the S&P 500 index, at 10% of time the actual volatility exceeds the forecast by approximately 30% and at 1% of time the actual volatility exceeded the prediction by 80%. In our view these large deviations represent a major source of risk that should be more carefully modeled in the realized volatility framework, specially for risk

management applications. In setting of out of sample uncertainty we can expect these values to be even higher.

We are now ready to state the problem at the center of our analysis, which as we will clarify in the next subsection is directly related to the presence of volatility risk (and occasionally extremely high forecasting errors in realized volatility). Even though it has been long recognized that the distributions of the stock returns scaled by realized volatility measures are approximately gaussian (e.g., Andersen et al., 2001), Table 4 and Figure 6 reveal that this is far from the case when we scale returns by the in-sample fitted values of the HAR+AE model. Table 4 shows that for all series there is substantial excess kurtosis in these scaled returns. For the indexes (but not the stocks) the distribution is pronouncedly negatively skewed due to the volatility feedback effect. In this paper, we seek to better understand and model this non-gaussianity.

Despite the sizable forecasting gains made possible by volatility models based on high frequency data, our descriptive results can be directly related to the failure of GARCH volatility models to completely account for the excess kurtosis of returns (see for example Malmsten and Teräsvirta, 2004, Carnero et al., 2004). The researcher or practitioner interested in evaluating the density of returns from the perspective of a time series model still lives in a fat tailed world and purely predictive models of volatility may have little to say about it. In this paper, we do not interpret those facts as evidence against those models, but as a consequence of two factors: volatility risk (which causes excess kurtosis in the ex ante distribution of returns) and volatility feedback (or intraday leverage) effects (causing negative skewness). In the following section, we will argue that an adequate volatility model for return density forecasting and risk management in this setting should illuminate the dynamics of the higher moments. To pursue this objective we will turn to the idea of studying the time series volatility of realized volatility following Corsi et al. (2008).

What can we say about the volatility of realized volatility? Figure 7 shows the residuals of the HAR model with leverage effects and Figure 8 display the sample autocorrelations for the squared and absolute residuals. The figures provides unambiguous evidence for the presence of conditional heteroskedasticity in realized volatility, in line with Corsi et al. (2008), Bollerslev et al. (2009) and other previous studies. The plots are also weakly suggestive of long memory, even though testing for this hypothesis more rigorously is infeasible. Figure 9 shows an interesting pattern common to all our series: when we extend the model with a GARCH(1,1) specification for the residuals, as would be natural to account for this conditional heteroskedasticity in this case, there is always a striking relation between the estimated conditional volatility of realized volatility and the fitted values of the model (the correlation between the two series in the figure is 0.42). Thus, there seems to be a close positive relation between volatility risk and the level of volatility. This is a new finding in the volatility literature, even though this relation has been explored many times before in the options pricing literature (see for example Heston, 1993, Jones, 2000).

2.4. Volatility shocks and the distribution of returns. To further motivate our model and qualify some of the assertions of the last section we now elucidate the relation between the stochastic structure of the innovations in the realized volatility and higher moments of the (ex ante) conditional distribution of asset returns. Our objective is to gain some insight into the daily conditional distribution of asset returns given the information set provided by past data. To simplify our discussion we shall focus on the skewness and kurtosis of this distribution. Different literatures such as preference theory (Dittmar, 2002), option pricing, and risk management indicate that the first four moments of the return distribution should be the relevant ones for our analysis.

We know that conditionally on the (ex post) realized volatility returns are nearly gaussian, so that a natural approach to follow is the standard stochastic volatility framework: we postulate a model for the (conditional) distribution of volatility and think in terms of a hierarchical structure where the distribution of returns conditional on volatility is gaussian. The interpretation of the model will vary depending on the quantity that is being modeled: the realized variance, the realized volatility or the log of the realized variance. We start with the simplest case. Consider the following model of returns:

$$(2) \quad \begin{aligned} r_t &= \sigma_t \varepsilon_t \\ \sigma_t^2 &= \psi_t + h_t \eta_t \end{aligned}$$

Where $\varepsilon_t \sim N(0, 1)$, $E(\eta_t) = 0$, $E(\eta_t^2) = 1$. Assume for now that ε_t and η_t are independent. ψ_t is interpreted as the conditional mean of the variance of returns, η_t is a random shock to the variance, and h_t^2 is the variance of variance. The conditional variance and kurtosis of returns are given below, showing that the excess kurtosis of returns is strictly positive in the presence of a positive variance of variance. In this case the conditional kurtosis of returns a particularly simple function of the coefficient of variation of the variance of returns.

$$(3) \quad E(r_t^2) = E(\sigma_t^2)E(\varepsilon_t^2) = \psi_t$$

$$(4) \quad \frac{E(r_t^4)}{E(r_t^2)^2} = \frac{E(\sigma_t^4)E(\varepsilon_t^4)}{\psi_t^2} = \frac{3E(\psi_t^2 + 2\psi_t h_t \eta_t + h_t^2 \eta_t^2)}{\psi_t^2} = 3 \left(1 + \frac{h_t^2}{\psi_t^2} \right)$$

In the case where we postulate a linear model for realized volatility,

$$(5) \quad \begin{aligned} r_t &= \sigma_t \varepsilon_t \\ \sigma_t &= \psi_t + h_t \eta_t, \end{aligned}$$

a little algebra shows that conditional variance and kurtosis for the return are given by:

$$(6) \quad E(r_t^2) = E(\sigma_t^2)E(\varepsilon_t^2) = E(\psi_t^2 + 2\psi_t h_t \eta_t + h_t^2 \eta_t^2) = \psi_t^2 + h_t^2$$

$$(7) \quad \frac{E(r_t^4)}{E(r_t^2)^2} = \frac{E(\sigma_t^4)E(\varepsilon_t^4)}{(\psi_t^2 + h_t^2)^2} = 3 \left(1 + \frac{4\psi_t^2 h_t^2 + 4\psi_t h_t^3 E(\eta_t^3) + h_t^4 (E(\eta_t^4) - 1)}{\psi_t^4 + 2\psi_t^2 h_t^2 + h_t^4} \right)$$

Not surprisingly, the first equation imply that ignoring time variation in the volatility of volatility will render forecasts of the conditional variance of returns biased even if the conditional mean of the realized volatility is consistently forecasted. In this case the expression for the return kurtosis is more complicated than in the last case, but contains the same essence: the kurtosis is growing on the volatility of volatility and decreasing on the level of volatility. The main difference is that now it also depends on the distribution of the standardized innovations to realized volatility, being positively related to the skewness and kurtosis of this distribution.

In Figures 10, 11 we illustrate the impact of volatility risk the distribution of returns in this setting. For low values of the volatility of volatility (or more generally, for a low conditional coefficient of variation in volatility) the distribution of returns is still very close to the gaussian

case. This is consistent with the evidence that returns standardized by the realized volatility are nearly normally distributed: since the impact of the volatility of volatility is non-linear and grows slowly with the variable, the effect of errors in the Barndorff-Nielsen et al. (2008) realized volatility estimator are not enough to generate excess kurtosis on the scaled returns. Figure 10 illustrates the consequences of the non-gaussianity on the distribution of realized volatility shocks, where we assume a positively skewed and leptokurtic distribution with parameters calibrated with the S&P 500 data. The excess kurtosis on the volatility amplify the excess kurtosis on returns for a given volatility risk level.

Finally, if we choose to model the log of the realized variance ($\log(\sigma_t^2) = \psi_t + h_t \eta_t$) we obtain:

$$(8) \quad E(r_t^2) = E(\sigma_t^2)E(\varepsilon_t^2) = \exp(\psi_t)E(\exp(h_t \eta_t))$$

$$(9) \quad \frac{E(r_t^4)}{E(r_t^2)^2} = 3 \frac{E(\exp 2h_t \eta_t)}{[E(\exp h_t \eta_t)]^2}$$

In this case both the conditional variance and the conditional kurtosis depend on the distribution of η_t . For example, if η_t is assumed normal the standard formula for the moment generating function of the normal distribution gives a conditional variance of $\exp(\psi_t + h_t^2/2)$ and a kurtosis $3 \exp(2h_t^2)$.

What about the skewness? Returning to the linear model for realized volatility, if the ε_t and η_t are independent then trivially $E(r_t^3) = E(\sigma_t^3)E(\varepsilon_t^3) = 0$, so that in this type of model the observed negative skewness on the ex ante distribution of returns must come from the negative dependence between ε_t and η_t . Writing the expression for the third moment,

$$(10) \quad E(r_t^3) = E(3\psi_t^2 h_t \eta_t \varepsilon_t^3 + 3\psi_t h_t^2 \eta_t^2 \varepsilon_t^3 + h_t^3 \eta_t^3 \varepsilon_t^3),$$

which is not particularly illuminating, but does highlight the fact that given the dependence structure between the two shocks a higher volatility risk will also increase the conditional skewness in returns.

What we take from this analysis are the ingredients for adequately modeling the empirically relevant ex ante distribution of returns in the realized volatility framework: the conditional mean of realized volatility, the volatility risk, the distribution of the shocks to volatility and the dependence structure between the shocks to returns and the shocks to volatility. The volatility risk is extremely important because is the main determinant of the excess kurtosis in the conditional distribution of returns and amplifies the negative skewness in this distribution. For risk management, option pricing and other applications, understanding and modeling these features of the conditional distribution of returns is fundamental.

The realized volatility literature so far has been mostly concerned only the first of these elements, and this is the gap the we intend to fill with the model of the next section. An exception is Bollerslev et al. (2009), which do not consider the implications of the different assumptions of the realized volatility model to the conditional distribution of returns. The main novelty of the model will be to put forward a more realistic and accurate model for the volatility risk instead of extending the standard specification of the model to allow for GARCH effects as in Corsi et al.

(2008) and Bollerslev et al. (2009) (even though our model departs from the ones considered in these papers in other fundamental aspects).

2.4.1. *Comparison with GARCH models and an important remark.* The GARCH family of models specifies the conditional variance of returns as deterministic conditional on past information, in contrast to the presence of volatility innovations in the stochastic volatility setting discussed above. In a sense our documentation of substantial volatility risk seems to reject the GARCH framework, despite the empirical successes of even the simplest GARCH models (Andersen and Bollerslev, 1998, Hansen and Lunde, 2005). Moreover, the stochastic volatility model also allows for the possibility that the shocks η_n and ε_t are correlated, giving rise to the volatility feedback effect that implies skewness in the conditional distribution of returns. However, this interpretation would not get to the heart of the volatility modeling problem.

In fact, as long as we are ultimately interested in the conditional distribution of the returns some of the features of the stochastic volatility framework can be reproduced in a GARCH setting. For example, in the stochastic volatility model the effect of increasing the volatility of volatility σ_η is to fatten the tails of the *conditional* (on past information) distribution of the returns. It has been long recognized that the gaussian GARCH model (where ε_t follows the normal distribution) fits financial data poorly, leading many researchers to assume a fat tail distribution for ε_t (typically a t-distribution). From the practical point of view the two approaches achieve the same end, though inflating the tails of ε_t in the GARCH model may look more like an extemporaneous fix to some authors. Similarly, the presence of volatility feedback effects skew the conditional distribution of the returns but assuming an asymmetrical distribution for ε_t imitates this effect in the GARCH model.

The important message is that the availability of realized volatility allows not only for significant advances in modeling the conditional volatility of returns but also the higher moments of this distribution. The reason is straightforward: since realized volatility is “observable” it is much easier to model and estimate the volatility of realized volatility than the tail heaviness parameter from return data in a GARCH model. With realized volatility we do not have to rely only on rare realizations in returns to identify the tails of the conditional return distribution. This becomes even more relevant in the presence of conditional heteroskedasticity in realized volatility documented previously, since identifying conditional kurtosis in a GARCH setting is extremely hard (Brooks et al., 2005, Creal et al., 2008).

With some long standing criticisms of volatility modeling in mind, a intuitive example further clarifies this point. Suppose we observe at a particular day a realized volatility of 10 and a return of 0. The return provides zero information about the tails of the conditional return distribution for a GARCH or other latent variable model. However if we accept that returns given volatility are normally distributed and assume that return and volatility shocks are uncorrelated then we have learned that on this particular day the “ex post 1% value at risk” was -2.326×10 , that is, a event comparable with the October 19th, 1987, crash in the Dow Jones index could have happened at the tail according to the model. Naturally, days of very high volatility are far more frequent than days of very high volatility *and* tail return shocks.

3. THE DUALY ASYMMETRIC REALIZED VOLATILITY MODEL

The dually asymmetric realized volatility (DARV) model is a first step in analyzing and incorporating the modeling qualities of a more realistic setting for the volatility risk within a standard

realized volatility model. The dual asymmetry in the model comes from leverage effects (as seen in the last section) and the positive relation between the level of volatility and the degree of volatility risk. The fundamental issue that arises in specifying the model is how to specify the relation between the volatility level and risk.

We directly model the time series of realized volatility (RV_t), as we will justify soon. To be consistent with the notation of the last section, let the conditional variance of the residuals be denoted by h_t^2 . In this paper we choose the specification $h_t^2 = \theta_0 + \theta_1 VL_t^2$, where VL_t (the volatility level) is the conditional mean of volatility ($E(RV_t|\mathcal{F}_{t-1})$, where \mathcal{F}_{t-1} is the information set at end of the previous day). Another option would be to directly allow for the asymmetry of positive or negative shocks in volatility in a GARCH model, but we have found our simpler specification to be superior. An advantage of our method is to be a straightforward way of accounting for the possibility of long memory in the volatility of realized volatility, since the volatility level will inherit this property from the RV_t series. A possibility for extending our specification would be to model the nonlinearities in the volatility level/risk relation, but more complicated specifications of this type are out of the scope of this paper.

The general specification of our model in autoregressive fractionally integrated and heterogenous autoregressive versions are:

DARV-FI Model:

$$\begin{aligned} r_t &= \mu_t + RV_t \varepsilon_t, \\ \phi(L)(1-L)^d(RV_t - \psi_t) &= \lambda_1 I(r_{t-1} < 0)r_{t-1} + \lambda_2 I(r_{5,t-1} < 0)r_{5,t-1} \\ &\quad + \lambda_3 I(r_{22,t-1} < 0)r_{22,t-1} + h_t \eta_t, \\ h_t^2 &= \theta_0 + \theta_1 VL_t^2, \end{aligned}$$

(11)

DARV-HAR Model:

$$\begin{aligned} r_t &= \mu_t + RV_t \varepsilon_t, \\ RV_t &= \phi_0 + \phi_1 RV_{t-1} + \phi_2 RV_{5,t-1} + \phi_3 RV_{22,t-1} \\ &\quad + \lambda_1 I(r_{t-1} < 0)r_{t-1} + \lambda_2 I(r_{5,t-1} < 0)r_{5,t-1} \\ &\quad + \lambda_3 I(r_{22,t-1} < 0)r_{22,t-1} + h_t \eta_t, \\ h_t^2 &= \theta_0 + \theta_1 VL_t^2, \end{aligned}$$

where r_t is the log return at day t , μ_t is the conditional mean for the returns, RV_t is the realized volatility, ε_t is i.i.d. $N(0, 1)$, ψ_t shifts the unconditional mean of realized volatility, d denotes the fractional differencing parameter, $\Phi(L)$ is a polynomial with roots outside the unit circle, L the lag operator, I is the indicator function, $r_{j,t-1}$ is a notation for the cumulated returns $\sum_{i=t-j}^{t-1} r_{t-i}$, $RV_{j,t-1} = \sum_{i=t-j}^{t-1} RV_{t-i}$, h_t is the volatility of the realized volatility, η_t is i.i.d. with $E(\eta_t) = 0$ and $E(\eta_t^2) = 1$, ε_t and η_t are allowed to be dependent and $VL_t = E(RV_t|\mathcal{F}_{t-1})$.

Following the evidence of fractional integration in realized volatility, ARFIMA models are the standard in the literature. Fractionally integrated models have been estimated for example in Andersen et al. (2003), Areal and Taylor (2002), Beltratti and Morana (2005), Deo et al. (2006), Martens et al. (2004), Thomakos and Wang (2003), among others. Nevertheless, the estimation of ARFIMA models in this context has encountered a few shortcomings. Although $I(d)$ processes

are a seemingly reasonable approximation for the data generating process of volatility series, there is no underlying theory to formally support this specification. Instead, the results of Diebold and Inoue (2001) and Granger and Hyung (2004) challenge fractional integration as the correct specification for realized volatility series by showing that long memory properties can be engendered by structural breaks or regime switching.³ Statistical tests for distinguishing between those alternatives, such as the one proposed Ohanissian et al. (2004), have been hampered by low power. Finally, Granger and Ding (1996) and Scharth and Medeiros (2009) discuss how estimates of the fractional differencing parameter are subject to excessive variation over time.

Given the lack of stronger support for a strict interpretation of fractional integration evidence and the higher computational burden in estimating and forecasting this class of models, some researchers have chosen to apply simpler time series models which are consistent with high persistence over the relevant horizons (like the HAR model of the last section), even though they do not rigorously exhibit long memory (hence being labeled 'quasi-long memory' models). Since this debate bears little relevance for our analysis we have chosen to present the dually asymmetric model in both a fractionally integrated version and a HAR version. After preliminary tests we have selected an ARFIMA(1, d ,0) model specification throughout this paper.

Bollerslev et al. (2006) and Scharth and Medeiros (2009) highlight the impact of leverage effects for the dynamics of realized volatility. The latter argues for the existence of regime switching behavior in volatility, with large falls (rises) in prices being associated with persistent regimes of high (low) variance in stock returns. The authors show that the incorporation of cumulated daily returns as an explanatory variable brings some modeling advantages by capturing this effect. While Scharth and Medeiros (2009) consider multiple regimes in a nonlinear model, we focus on a simpler linear relationship to account for the large correlation between past cumulated returns and realized volatility. We expect the use of these explanatory variables to substantially improve the ARFIMA estimation since we will be accounting for the high persistence in volatility associated to extended leverage effects (recall Figure 4).

To account for the non-gaussianity in the error terms we follow Corsi et al. (2008) and assume that the *i.i.d.* innovations η_t follow the standardized normal inverse Gaussian (which we denote by *NIG**), which is flexible enough to allow for excessive kurtosis and skewness and reproduce a number of symmetric and asymmetric distributions. A more complex approach would rely on the generalized hyperbolic distribution, which encompasses the NIG distribution and requires the estimation of an extra parameter. On the other hand typical distributions with support on the interval $(0, \infty)$, which would be a desirable feature for our case, were strongly rejected by preliminary diagnostic tests.

Some readers may be concerned that our model could in principle allow for a negative realized volatility. The immediate alternative would be to write a model for the time series of the log realized volatility. Regarding this potential problem we offer the following remarks. First, if the realized volatility is the main quantity we are concerned with our experience tells us that models for RV_t perform unambiguously better in terms of forecasting. Second, the log transformation obscures the volatility level/risk relation and makes the model for the conditional distribution of returns more complex. And most importantly, the fact that the estimated error distribution is

³Nevertheless, empirical work has found evidence of long range dependence even after accounting for possible regime changes and structural breaks in the volatility of asset returns (Lobato and Savin, 1998, Martens et al., 2004, Beltratti and Morana, 2006, Morana and Beltratti, 2004, Hyung and Franses, 2002, Scharth and Medeiros, 2009).

extremely right skewed and the conditional variance of RV shrinks with the level of the variable eliminates this concern for all practical purposes in our data.

Finally, to model the asymmetry in the conditional return distribution we let η_t and ε_t be dependent and model this dependence via a bivariate Clayton copula/. The copula approach is a straightforward way to account for non-linearities in this dependence relation and has the important advantage of not requiring the joint estimation of the return and volatility equations. Let $U = \Phi(\varepsilon_t)$ and $V = 1 - \Upsilon(\nu_t)$, where $\Phi(\cdot)$ and $\Upsilon(\cdot)$ are the corresponding normal and NIG^* cdfs for ε_t and η_t respectively. The joint CDF or copula of U and V is given by:

$$(12) \quad C_\kappa \equiv P(U \leq u, V \leq v) = (u^{-\kappa} + v^{-\kappa} - 1)^{-1/\kappa}$$

In this simple copula specification, returns and volatility are negatively correlated and display lower tail dependence (days of very low returns and very high volatility are linked, where the strength of this association is given by the parameter κ).

Other Comments. To reduce bias on our estimators and avoid distortions of the error distribution, we control the mean of the dependent variable for day of the week and holiday effects using dummies. Martens et al. (2004) and Scharth and Medeiros (2009) show that volatility sometimes tend to be lower on Mondays and Fridays, while substantially less volatility is observed around certain holidays.

In contrast with Bollerslev et al. (2009), which also considers a full system for returns, realized volatility and the volatility of volatility, our model does not consider jump components in the realized volatility.⁴ As we have mentioned in Section 2 the use of jumps does not seem to bring important forecasting advantages when the analysis is based on the efficient realized kernel estimator for volatility. On the other hand, the inclusion of a jump equation would substantially increase the complexity of the model, requiring us to model and estimate the joint distribution of return, volatility and jump shocks. For predicting and simulating the model multiple periods ahead this is a substantial burden. Since the ultimate interest lies in the conditional distribution of returns a parsimonious alternative is to ignore the distinction between continuous and jump components in realized volatility and carefully model the distribution of returns given realized volatility (considering the possible impact of jumps on it). For the objectives of this paper the assumption of normality of the return shocks seems reasonable (specially for the stocks).

3.1. Estimation and Density Forecasting. We estimate the two versions of the dually asymmetric realized volatility model by maximum likelihood. The fact that the conditional volatility of volatility h_t depends on the conditional mean of the realized volatility brings no issues for the estimation. However, a full maximum likelihood procedure for the ARFIMA model Sowell (e.g., 1992) is unavailable under the assumptions of conditional heteroskedasticity and the NIG distribution for the errors η_t . We then follow the standard approach in the literature and turn to a consistent approximate maximum likelihood procedure where the fractional differencing operator $(1 - L)^d$ is replaced by a truncation of the respective binomial expansion.⁵ The use of this approximate estimator does not impact in any way the main arguments of this paper. For reference the log-likelihood function is given by:

⁴The literature on the non-parametric measurement of the jump components includes Andersen et al. (2007), Tauchen and Zhou (2005) and Barndorff-Nielsen and Shephard (2006), among others.

⁵ $(1 - L)^d = 1 - dL + \frac{d(d-1)L^2}{2!} - \frac{d(d-1)(d-2)L^3}{3!} + \dots$

$$\begin{aligned}
 \ell(\hat{d}, \hat{\phi}, \hat{\psi}, \hat{\lambda}, \hat{\theta}, \hat{\alpha}, \hat{\beta}; RV_{1...T}, X_{1...T}) &= T \log(\hat{\alpha}) - T \log(\pi) + \sum_{t=1}^T \log \left[K_1(\hat{\alpha} \hat{\delta} (1 + \hat{y}_t^2)^{1/2}) \right] \\
 (13) \quad &- 0.5 \sum_{t=1}^T \log(1 + \hat{y}_t^2) + T \hat{\delta} (\hat{\alpha}^2 - \hat{\beta}^2)^{1/2} + \hat{\delta} \hat{\beta} \sum_{t=1}^T \hat{y}_t \\
 &- 0.5 \sum_{t=1}^T \log(\hat{h}_t)
 \end{aligned}$$

where X collects the additional explanatory variables, α and β are the tail heaviness and asymmetry parameters of the standardized NIG distribution. $\gamma = (\hat{\alpha}^2 - \hat{\beta}^2)^{1/2}$ and $\hat{y}_t = \frac{\hat{\eta}_t / \hat{h}_t - \hat{\omega}}{\hat{\delta}}$, where ω and δ are the location and scale parameters associated with the standardized NIG distributed with parameters α and β .

The copula specification for the joint distribution of return and volatility innovations allows us to estimate the copula by maximum likelihood in a separate stage once we have obtained estimates for η_t and ε_t from the marginal models. For simplicity we estimate the mean of returns μ_t by the sample mean (since the daily expected return is very small μ_t is immaterial for our analysis), so that $\hat{\varepsilon}_t = (r_t - \hat{\mu}) / RV_t$.

Unfortunately, an analytical solution for the return density implied by our flexible normal variance-mean mixture hypothesis (realized volatility is distributed normal inverse gaussian and returns given volatility are normally distributed) is not available. Except for a few cases such as one day ahead point forecasts for realized volatility, many quantities of interest based on our model have to be obtained by simulation. We consider the following Monte Carlo method which can be easily implemented and made accurate with realistic computational power. Conditional on information up to day t we implement the following general procedure for simulating joint paths for returns and volatility (where \sim is used to denote a simulated quantity):

- (1) In the first step the functional form of the model is used for the evaluation of forecasts \widehat{RV}_{t+1} and \hat{h}_{t+1} conditional on past realized volatility observations, returns, and other variables.
- (2) Using the estimated copula we randomly generate S pairs of return ($\tilde{\varepsilon}_{t+1,j}$, $j = 1, \dots, S$) and volatility ($\tilde{\eta}_{t+1,j}$, $j = 1, \dots, S$) shocks with the according marginal distributions. Antithetic variables are used to balance the return innovations for location and scale.
- (3) We obtain S simulated volatilities through $\widetilde{RV}_{t+1,j} = \widehat{RV}_{t+1} + h_t \tilde{\eta}_{t+1,j}$, $j = 1, \dots, S$. Each of these volatilities generate a returns $\tilde{r}_{t+1,j} = \hat{\mu}_t + \widetilde{RV}_{t+1,j} \tilde{\varepsilon}_{t+1,j}$
- (4) This procedure can be iterated in the natural way to generate multiple paths for returns and realized volatility.

4. EMPIRICAL ANALYSIS

The rest of this paper will be dedicated to studying the modeling qualities of the dually asymmetric realized volatility model in our data. We consider five alternative specifications chosen to illuminate the improvements introduced by different elements of the model: the homoskedastic ARFIMA(1,d,0) model with and without (extended) leverage effects, the ARFIMA(1,d,0)-GARCH model with and without leverage effects and the HAR-GARCH model with leverage effects. We leave the simpler HAR specifications out of the analysis as they are essentially redundant to the

fractionally integrated counterparts. We consider the dependence between the return and volatility innovations on all specifications. Recall from Section 2 that our out of sample period covers eight and half years from January 2001 to June 2009 (2106 trading days for the US indexes and stocks). All out of sample implementations re-estimate the models quarterly using the full past data to calculate the desired statistics.

Important questions that we consider in this section are: what is the substance of time-varying volatility risk for standard realized volatility models? Is the standard realized volatility model significantly improved by our volatility risk specification when compared to the standard GARCH approach? What is the importance of (extended) leverage effects? Does using an ARFIMA or HAR model make a salient difference in this setting?

4.1. Full Sample Parameter Estimates and Diagnostics. Tables 5 and 6 show the parameter estimates for all our specifications for the S&P 500 series. While most of the estimates are unexceptional and in line with the previous literature, we draw attention to two noteworthy results. First, in the ARFIMA setting considering either conditional heteroskedasticity or extended leverage effects substantially change our estimates for the fractional differencing parameter and the unconditional mean of realized volatility. Second, the leverage effect coefficients are significantly larger in the dually asymmetric estimations in comparison to other models. This interaction is likely to be consequential for our forecasting results.

Table 7 displays a variety of estimation diagnostics. Not surprisingly, the inclusion of leverage effects and time varying volatility risk considerably improve the fit of the specifications according to the Schwarz criterion and other standard statistics. The first piece of evidence in favor of the DARV model also comes from this analysis: our specification for the volatility risk unambiguously improve the fit of the model compared to the specifications with GARCH effects, though both alternatives seem to appropriately account for the autocorrelation in the squared residuals. On the other hand, an adverse result affecting all specifications come from the (small) sample autocorrelation in the residuals. Reversing this result would require ad hoc modifications in our setting, leaving some role for more complex models or structural breaks to capture these dynamics.

The ability to correctly model the conditional distribution of realized volatility is fundamental for the cardinal issues of this paper. To investigate this problem we implement a Kolmogorov-Smirnov test for the hypothesis that the standardized residuals are well described by the estimated NIG distribution. The two versions of the DARV model are easily consistent with this hypothesis, while the alternative models are either strongly rejected or susceptible to the choice of significance level.

4.2. Point Forecasts. We now turn to out of sample forecasts. As we have argued in Section 2, the set of realistic assumptions for the behavior of realized volatility imply that if our main objective is to model for the conditional distribution of returns then an excessive focus on the point forecasting abilities of different volatility models may be inappropriate: the conditional mean of volatility is far from enough to describe the tails of the return distribution. Without a model for the realized volatility risk we do not have an expressive model for the returns. Moreover, the time series volatility of realized volatility is so high that it is extremely hard to obtain economically substantive improvements in predicting realized volatility.

However, this should not be confused with the argument that forecasting does not matter as the conditional mean of volatility is approximately the conditional volatility of returns itself. Out of sample predictions have been the main basis of comparison in the volatility literature and are

the subject of extensive analysis (e.g., Hansen and Lunde, 2005). Forecasting is a very useful tool for studying and ranking volatility models, even though it may not be very informative about the relative modeling qualities of various alternatives: because volatility is so persistent even a simple moving average will have a similar performance to more theoretically sound models.

The evaluation of forecasts is based on the mean absolute error (MAE), the root mean squared error (RMSE) and the estimation of the Mincer-Zarnowitz regression:

$$RV_t = \alpha + \beta \widetilde{RV}_{t|t-1,i} + \varepsilon_{t,i}$$

where RV_t is the observed realized volatility on day t and $\widetilde{RV}_{t|t-1,i}$ is the one-step-ahead forecast of model i for the volatility on day t . If the model i is correctly specified then $\alpha = 0$ and $\beta = 1$. We report the R^2 of the regression as a measure of the ability of the model to track variance over time and test of superior predictive ability (SPA) test developed by Hansen (2005). The null hypothesis is that a given model is not inferior to any other competing models in terms of a given loss function.

The point forecasting statistics for the S&P 500 series are displayed in Table 8, where we consider one, five and twenty two days ahead predictions. The results for the other series are arranged in Table 9 and are limited to one period ahead predictions for conciseness. In 9 we report the R^2 for changes in volatility (we have discussed the motivation for this in Section 2.3) and the respective SPA test for the RMSE (in parenthesis). The foremost message that the results bring is again that asymmetric effects are essential for improving forecasting performance in realized volatility: forecasts are significantly improved for all series when leverage effects are included. The results for the S&P 500 series suggest however that this advantage is decreasing in the forecasting horizon.

In line with the full sample results, the dually asymmetric model outperforms the standard ARFIMA-GARCH and HAR-GARCH models in one day ahead forecasting for all series, even though the difference is only significant at the 5% level in the SPA test for the FTSE and AT&T series (the difference is also significant at the 5% level in a meta test across all series, which we do not report in the table). This improvement in forecasting is also supported by the longer horizon results for the S&P 500 series, which reveal statistically significant differences. Finally, the results indicate no expressive divergence between the HAR and ARFIMA specifications.

4.3. Volatility Risk. While we emphasize the positive evidence from the forecasting exercise for our volatility risk model we stress again that the difference in forecasting performance by itself is unlikely to be economically substantial, though in line with improvements generally reported in the volatility literature. Our main question is whether the dually symmetric specification introduce a better model for volatility risk, and consequently to the tails of the conditional distribution of returns. We consider this problem in this section.

To answer this question we need to define an informative metric for how well different models are able to describe the relevant dimension of the conditional distribution of realized volatility. Since the time series volatility of realized volatility is latent and dependent on the specification for the conditional mean of the series, a meaningful direct analysis of volatility risk forecasting is infeasible. For this reason we investigate conditional forecasts of realized volatility. Our approach consists in calculating ex post empirical quantiles for the daily realized volatility changes (ΔRV_t) in the 2001-2009 period and calculating (out of sample) forecasts for the change in volatility given that it exceeds the relevant quantile.

We provide two motivations for this method. First, since the true conditional realized volatility quantiles are unobservable the use of the ex post quantiles is a straightforward way of obtaining an uniform conditioning case for comparing different models by their performance in the upper tail of the distribution. And most importantly, analyzing whether the dually asymmetric model is better capable of accounting for the largest movements in volatility observed in our data goes at the heart of our problem of better describing volatility risk and the tails of the conditional return distribution. Again, because the volatility innovations are unobservable the use of ΔRV_t is adequate for comparing the models. To complement this analysis we also consider conditional forecasts based on ex post realized volatility quantiles themselves. We interpret the results of this section as being the main empirical evidence for the DARV model, since they directly address the issue of volatility risk.

The results for the S&P 500 index are organized in Tables 10 and 11, while the finding for the remaining series are summarized in Table 12. For the S&P 500 we consider forecasts conditional on the change in volatility and the volatility exceeding the 80th, 90th, 95th and 99th percentiles, while for the other series we only consider the 90th percentile. As expected the models with constant volatility risk perform extremely poorly compared to the heteroskedastic models, again highlighting the importance of time varying realized volatility risk. More importantly, the results strongly support the dually asymmetric model. With the only exception of the CAC index the DARV model improves the conditional forecasts for the changes in volatility, in most cases substantially. The same pattern holds for the forecasts conditional on the realized volatility quantile.

From the S&P 500 we can see that these results are even more striking at the 99th percentile, where in contrast with the DARV models the GARCH specifications have almost no forecasting power: the improvement in RMSE of going from the ARFIMA-GARCH model to the DARV model is about the same as going from constant volatility risk to the ARFIMA-GARCH model.

4.4. Value-at-Risk. To conclude our empirical analysis we implement a brief value-at-risk analysis for the S&P 500 index. Even though this exercise is not particularly informative about the modeling qualities of the different specifications studied in this paper, we consider it important to check whether our models yield plausible results for this standard risk management metric. In addition, we also wish to use this section to further illustrate the possible pitfalls of excessively relying on point forecasts and ignoring volatility risk. To do so we introduce as a reference a more standard way of calculating value-at-risk measures, namely considering $r_t \sim N(0, \widetilde{RV}_t)$ (where \widetilde{RV}_t is the forecasted realized volatility). We label this approach (incorrect for our models) the point forecast method, in contrast with the appropriate Monte Carlo method of section 3.1.

The evaluation of value-at-risk forecasts is based on the likelihood ratio tests for unconditional coverage and independence of Christoffersen (1998), where conditional skewness is allowed for in all models. Our analysis is similar to Beltratti and Morana (2005), who study the benefits of value-at-risk with long memory. Let $\widehat{q}_{t|t-1}^i(\alpha)$ be the $(1 - \alpha)$ interval forecast of model i for day t conditional on information on day $t - 1$. In our application, we consider 1%, 2.5% and 5% value-at-risk measures, i.e., $\alpha = 0.01, 0.025$ and 0.05 , respectively. We construct the sequence of coverage failures for the lower α tail as:

$$F_{t|t-1} = \begin{cases} 1 & \text{if } r_{t+1} < \widehat{q}_{t+1|t}^i(\alpha) \\ 0 & \text{if } r_{t+1} > \widehat{q}_{t+1|t}^i(\alpha) \end{cases}$$

where r_t is the return observed on day t . The unconditional coverage (UC) is a test of the null $\mathbb{E}(F_{t+1|t}) = \alpha$ against $\mathbb{E}(F_{t+1|t}) \neq \alpha$. The test of independence is constructed against a first-order Markov alternative. Finally, let z be the predicted cumulative density function evaluated at the observed returns that are below the value-at-risk. If the model is well specified then we should expect that the sample average of z is close to $\alpha/2$, so that we use this as a proxy for checking whether the models generate adequate expected shortfall values.

The value-at-risk performance of the models are organized and presented in Table 13. The results show that as expected from our analysis the method of calculating VaRs based only on the point forecast of volatility is severely biased towards underestimating the value-at-risk, failing to provide adequate coverage at all intervals. The Monte Carlo method in turn significantly reduces or eliminates the problem of excess violations for all models (even though the exercise has no power for ranking them). However, most models are rejected for the 5% Value-at-Risk. For reference, Table 14 reports the predicted cumulative density function from all the models calculated at the lowest returns observed in our sample. Despite the fact that the 2007-2009 financial crisis brought realized volatility to unprecedented levels in the data we do not observe catastrophic failures of our value-at-risk intervals (even for the misspecified models), supporting the robustness arguments of Section 2.4.1.

5. CONCLUSION

In this paper we have documented that realized variation measures constructed from high-frequency returns reveal a large degree of volatility risk in stock and index returns, where we characterize volatility risk by the extent to which forecasting errors in realized volatility are substantive. Even though returns standardized by *ex post* quadratic variation measures are nearly gaussian, this unpredictability brings considerably more uncertainty to the empirically relevant *ex ante* distribution of returns. We have demonstrated how the study of volatility risk (or equivalently, the volatility of realized volatility) is essential for developing better models of the conditional distribution of returns, as this concept is inexorably related to the higher moments of the return distribution under the standard stochastic volatility setting. We have argued that the availability of realized volatility allows not only for significant advances in modeling the conditional volatility of returns but also the higher moments.

Far from exhausting the analysis of the empirical properties of this volatility risk, we have documented the close positive relation between the volatility of realized volatility and the level of volatility. To account for this fact we propose the dually asymmetric realized volatility (DARV) model and present extensive empirical evidence that by recognizing that realized volatility series are systematically more volatile in high volatility periods we are able to improve the out of sample performance of realized volatility models. Particularly in predicting the possibility of large movements and extremes in daily volatility (using conditional forecasts) we have found the difference to be substantial in many cases, with direct consequences to risk management and pricing applications.

To keep our discussion concise we have left out some important issues that can be explored in future work. We select two examples. First, in practice advances in realized volatility modeling may not be translated so neatly into improvements in modeling the conditional distribution of returns. Two aspects of the link between realized volatility and returns should be studied more carefully. The assumption that returns standardized by realized volatility are approximately

normal and independent seems to be inadequate for some series. Is there a role for jumps in adjusting the distribution? Do the problems in measuring realized volatility make this relation less straightforward? We have also only considered a simple model for the dependence between return and volatility innovations. Second, we have mostly analyzed the performance of different models in one day ahead applications. Because financial quantities are so persistent many incongruent models are misleadingly competitive at very short horizons. More emphasis should be placed in investigating whether different models are consistent with a realistic longer horizon dynamics. Our analysis suggest that to do so we may need a more solid understanding of asymmetric effects.

REFERENCES

- ANDERSEN, T. AND T. BOLLERSLEV (1998): "Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts," *International Economic Review*, 39, 885–906.
- ANDERSEN, T., T. BOLLERSLEV, AND F. DIEBOLD (2007): "Roughing it up: Including Jump Components in the Measurement, Modeling and Forecasting of Return Volatility," *Review of Economics and Statistics*, 89, 701–720.
- ANDERSEN, T., T. BOLLERSLEV, F. DIEBOLD, AND H. EBENS (2001): "The Distribution of Realized Stock Return Volatility," *Journal of Financial Economics*, 61, 43–76.
- ANDERSEN, T., T. BOLLERSLEV, F. DIEBOLD, AND P. LABYS (2003): "Modeling and Forecasting Realized Volatility," *Econometrica*, 71, 579–625.
- AREAL, N. AND S. R. TAYLOR (2002): "The Realized Volatility of FTSE-100 Futures Prices," *Journal of Futures Markets*, 22, 627–648.
- AÏT-SAHALIA, Y., P. MYKLAND, AND L. ZHANG (2005): "Ultra High Frequency Volatility Estimation with Dependent Microstructure Noise," NBER Working Papers 11380, National Bureau of Economic Research.
- BAILLIE, R., T. BOLLERSLEV, AND H. O. MIKKELSEN (1996): "Fractionally integrated generalized autoregressive conditional heteroskedasticity," *Journal of Econometrics*, 74, 3–30.
- BARNDORFF-NIELSEN, O., P. HANSEN, A. LUNDE, AND N. SHEPHARD (2008): "Designing realised kernels to measure the ex-post variation of equity prices in the presence of noise," *Econometrica*, 76, 1481–1536.
- (2009): "Subsampling Realized Kernels," *Journal of Econometrics*, forthcoming.
- BARNDORFF-NIELSEN, O. AND N. SHEPHARD (2002): "Econometric analysis of realized volatility and its use in estimating stochastic volatility models," *Journal of the Royal Statistical Society B*, 64, 253–280.
- (2006): "Econometrics of Testing for Jumps in Financial Economics using Bipower Variation," *Journal of Financial Econometrics*, 4, 1–30.
- BELTRATTI, A. AND C. MORANA (2005): "Statistical Benefits of Value-At-Risk with Long Memory," *Journal of Risk*, 7.
- (2006): "Breaks and Persistency: Macroeconomic Causes of Stock Market Volatility," *Journal of Econometrics*, 131, 151–177.
- BOLLERSLEV, T., U. KRETSCHMER, C. PIGORSCH, AND G. TAUCHEN (2009): "A Discrete-Time Model for Daily S&P500 Returns and Realized Variations: Jumps and Leverage Effects," *Journal of Econometrics*, forthcoming.
- BOLLERSLEV, T., J. LITVINOVA, AND G. TAUCHEN (2006): "Leverage and Volatility Feedback Effects in High-Frequency Data," *Journal of Financial Econometrics*, 4, 353–384.

- BREIDT, F., N. CRATO, AND P. DE LIMA (1998): "The detection and estimation of long memory in stochastic volatility," *Journal of Econometrics*, 83, 325–348.
- BROOKS, C., S. BURKE, S. HERAVI, AND G. PERSAUD (2005): "Autoregressive Conditional Kurtosis," *Journal of Financial Econometrics*, 3, 399–421.
- CARNERO, M. A., D. PEÑA, AND E. RUIZ (2004): "Persistence and kurtosis in GARCH and stochastic volatility models," *Journal of Financial Econometrics*, 2, 319–342.
- CHAN, W., R. JHA, AND M. KALIMPALLI (2006): "The economic value of trading with realized volatility in the S&P 500 index options market," Working paper.
- CHRISTOFFERSEN, P. (1998): "Evaluating Interval Forecasts," *International Economic Review*, 39, 841–62.
- COMTE, F. AND E. RENAULT (1996): "Long memory continuous time models," *Journal of Econometrics*, 73, 101–149.
- CORSI, F. (2004): "A Simple Long Memory Model of Realized Volatility," Manuscript, University of Southern Switzerland.
- CORSI, F., U. KRETSCHMER, S. MITTNIK, AND C. PIGORSCH (2008): "The Volatility of Realized Volatility," *Econometric Reviews*, 27, 46–78.
- CREAL, D., S. KOOPMAN, AND A. LUCAS (2008): "A General Framework for Observation Driven Time-Varying Parameter Models," Tinbergen Institute Discussion Papers 08-108/4, Tinbergen Institute.
- DEO, R., C. HURVICH, AND Y. LU (2006): "Forecasting Realized Volatility Using a Long-Memory Stochastic Volatility," *Journal of Econometrics*, 131, 29–58.
- DIEBOLD, F. AND A. INOUE (2001): "Long Memory and Regime Switching," *Journal of Econometrics*, 105, 131–159.
- DITTMAR, R. (2002): "Nonlinear Pricing Kernels, Kurtosis Preference, and Evidence from the Cross Section of Equity Returns," *Journal of Finance*, 57, 369–403.
- ENGLE, R. F. AND V. K. NG (1993): "Measuring and Testing the Impact of News on Volatility," *Journal of Finance*, 48, 1749–1778.
- FLEMING, J., C. KIRBY, AND B. OSTDIEK (2001): "The Economic Value of Volatility Timing," *Journal of Finance*, 56, 329–352.
- (2003): "The Economic Value of Volatility Timing Using Realized Volatility," *Journal of Financial Economics*, 67, 473–509.
- GATHERAL, J. AND R. OOMEN (2007): "Zero Intelligence Variance Estimation," Working paper, Warwick Business School.
- GHYSELS, E., A. SINKO, AND R. VALKANOV (2007): "MIDAS Regressions: Further Results and New Directions," *Econometric Reviews*, 26, 53–90.
- GRANGER, C. AND Z. DING (1996): "Varieties of Long Memory Models," *Journal of Econometrics*, 73, 61–77.
- GRANGER, C. AND N. HYUNG (2004): "Occasional Structural Breaks and Long Memory with an Application to the S&P 500 Absolute Stock Returns," *Journal of Empirical Finance*, 11, 399–421.
- HANSEN, P. (2005): "A Test for Superior Predictive Ability," *Journal of Business and Economic Statistics*, 23, 365–380.
- HANSEN, P. AND A. LUNDE (2005): "A Forecast Comparison of Volatility Models: Does Anything beat a GARCH(1,1) Model?" *Journal of Applied Econometrics*, 20, 873–889.

- (2006): “Realized variance and market microstructure noise (with discussion),” *Journal of Business and Economic Statistics*, 24, 127–218.
- HESTON, S. (1993): “A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options,” *Review of Financial Studies*, 6, 327–343.
- HILLEBRAND, E. AND M. MEDEIROS (2007): “Forecasting Realized Volatility Models: The Benefits of Bagging and Nonlinear Specifications,” Discussion paper, Pontifical Catholic University of Rio de Janeiro.
- HYUNG, N. AND P. FRANCES (2002): “Inflation rates: long-memory, level shifts, or both?” Report 2002-08, Econometric Institute, Erasmus University.
- JACOD, J., Y. LI, P. MYKLAND, M. PODOLSKIJ, AND M. VETTER (2009): “Microstructure noise in the continuous case: The pre-averaging approach,” *Stochastic Processes and their Applications*, 119, 2249–2276.
- JONES, C. (2000): “The Dynamics of Stochastic Volatility: Evidence from Underlying and Options Markets,” Working paper, Simon School of Business, University of Rochester.
- KOOPMAN, S., B. JUNGBACKER, AND E. HOL (2005): “Forecasting Daily Variability of the S&P 100 Stock Index Using Historical, Realised and Implied Volatility Measurements,” *Journal of Empirical Finance*, 23, 445–475.
- LIU, C. AND J. MAHEU (2007): “Forecasting Realized Volatility: A Bayesian Model Averaging Approach,” Working paper.
- LOBATO, I. N. AND N. E. SAVIN (1998): “Real and Spurious Long-Memory Properties of Stock-Market Data,” *Journal of Business and Economic Statistics*, 16, 261–268.
- MALMSTEN, H. AND T. TERÄSVIRTA (2004): “Stylized Facts of Financial Time Series and Three Popular Models of Volatility,” Working Paper Series in Economics and Finance 563, Stockholm School of Economics.
- MARTENS, M., D. VAN DIJK, AND M. DE POOTER (2004): “Modeling and Forecasting S&P 500 Volatility: Long Memory, Structural Breaks and Nonlinearity,” Discussion Paper 04-067/4, Tinbergen Institute.
- MCALÉER, M. AND M. MEDEIROS (2008a): “A Multiple Regime Smooth Transition Heterogeneous Autoregressive Model for Long Memory and Asymmetries,” *Journal of Econometrics*, 104–119.
- (2008b): “Realized Volatility: A Review,” *Econometric Reviews*, 27, 10–45.
- MÜLLER, U., M. DACOROGNA, R. DAV, R. OLSEN, O. PICTET, AND J. VON WEIZSACKER (1997): “Volatilities of Different Time Resolutions – Analysing the Dynamics of Market Components,” *Journal of Empirical Finance*, 4, 213–239.
- MORANA, C. AND A. BELTRATTI (2004): “Structural Change and Long Range Dependence in Volatility of Exchange Rates: Either, Neither or Both?” *Journal of Empirical Finance*, 11, 629–658.
- OHANISSIAN, A., J. RUSSELL, AND R. TSAY (2004): “True or Spurious Long Memory in Volatility: Does it Matter for Pricing Options?” Working paper, Graduate School of Business, University of Chicago.
- SCHARTH, M. AND M. MEDEIROS (2009): “Asymmetric Effects and Long Memory in the Volatility of Dow Jones Stocks,” *International Journal of Forecasting*, 25, 304–327.
- SHEPHARD, N. AND K. SHEPPARD (2009): “Realising the future: forecasting with high frequency based volatility (HEAVY) models,” *Journal of Applied Econometrics*, forthcoming.

- SOWELL, F. (1992): “Maximum Likelihood Estimation of Stationary Univariate Fractionally Integrated Time Series Models,” *Journal of Econometrics*, 165–188.
- TAUCHEN, G. AND H. ZHOU (2005): “Identifying Realized Jumps on Financial Markets,” Working paper, Duke University, Department of Economics.
- THOMAKOS, D. AND T. WANG (2003): “Realized Volatility in the Futures Market,” *Journal of Empirical Finance*, 10, 321–353.
- ZHANG, L. (2006): “Efficient Estimation of Stochastic Volatility Using Noisy Observations: A Multi-Scale Approach,” *Bernoulli*, 2, 1019–1043.
- ZHANG, L., P. MYKLAND, AND Y. AÏT-SAHALIA (2005): “A Tale of Two Time Scales: Determining Integrated Volatility with Noisy High-Frequency Data,” *Journal of the American Statistical Association*, 100, 1394–1411.

TABLE 1. DESCRIPTIVE STATISTICS: S&P 500

Statistic	r_t	RV_t	r_t/RV_t	ΔRV_t
Mean	0.012	0.984	0.042	0.000
Std. Deviation	1.343	0.605	1.013	0.359
Skewness	-0.186	3.417	0.040	-0.913
Kurtosis	10.570	26.342	2.743	65.184
Min	-9.470	0.212	-3.296	-6.840
$Q_{0.1}$	-1.443	0.478	-1.282	-0.308
$Q_{0.25}$	-0.610	0.606	-0.658	-0.135
$Q_{0.75}$	0.654	1.153	0.711	0.127
$Q_{0.9}$	1.368	1.584	1.382	0.301
Max	10.957	9.673	3.230	5.280

TABLE 2. IN-SAMPLE FORECASTING FOR SIMPLE MODELS: S&P 500.

The table reports the in sample forecasting results for the S&P 500 daily realized volatility for the period between 1996 and 2008. RMSE is the root mean squared error, MAE the mean absolute error, MAPE the mean absolute percentage error, R^2 is the R-squared of a linear regression of the actual realized volatility on the forecasts.

Model	R^2	1 Day			$R^2(\Delta)$
		RMSE	MAE	MAPE	
RV_{t-1}	0.680	0.359	0.208	0.210	-
EWMA	0.728	0.319	0.190	0.192	0.217
ARMA(2,1)	0.730	0.315	0.188	0.194	0.229
HAR	0.730	0.315	0.187	0.193	0.230
HAR + AE	0.776	0.287	0.176	0.184	0.360
HAR+AE+RAV	0.779	0.285	0.175	0.183	0.370

Model	R^2	5 Days (cumulated)			$R^2(\Delta)$
		RMSE	MAE	MAPE	
$RV_{5,t-1}$	0.758	1.406	0.869	0.165	-
EWMA	0.771	1.362	0.840	0.160	0.060
ARMA(2,1)	0.771	1.323	0.825	0.164	0.128
HAR	0.781	1.293	0.805	0.158	0.154
HAR + AE	0.805	1.219	0.763	0.151	0.254

Model	R^2	22 Days (cumulated)			$R^2(\Delta)$
		RMSE	MAE	MAPE	
$RV_{22,t-1}$	0.611	7.019	4.443	0.201	-
EWMA	0.679	6.383	4.098	0.187	0.190
ARMA(2,1)	0.671	6.457	4.142	0.190	0.284
HAR	0.691	6.260	4.023	0.182	0.292
HAR + AE	0.695	6.214	4.052	0.187	0.295

TABLE 3. DESCRIPTIVE STATISTICS FOR RV_t/\widehat{RV}_t RATIOS

	Mean	Std. Dev.	Skewness	Kurtosis	$Q_{0.75}$	$Q_{0.9}$	$Q_{0.95}$	$Q_{0.99}$
S&P 500	1.00	0.25	1.86	12.17	1.11	1.28	1.43	1.84
DJIA	1.01	0.30	1.94	15.52	1.15	1.35	1.53	1.95
FTSE	1.00	0.33	2.97	27.12	1.13	1.37	1.55	2.14
CAC	1.00	0.25	2.14	21.84	1.12	1.29	1.41	1.76
Nikkei	1.00	0.27	1.19	6.72	1.14	1.33	1.47	1.86
IBM	1.00	0.22	1.13	6.96	1.11	1.25	1.38	1.72
GE	1.00	0.20	0.85	5.30	1.11	1.25	1.37	1.63
WMT	1.00	0.24	1.25	7.82	1.12	1.29	1.41	1.74
AT&T	1.00	0.27	1.67	10.25	1.12	1.32	1.48	1.87

TABLE 4. DESCRIPTIVE STATISTICS FOR RETURNS STANDARDIZED BY IN-SAMPLE REALIZED VOLATILITY FITTED VALUES

	Mean	Std. Dev.	Skewness	Kurtosis	$Q_{0.01}$
S&P 500	-0.019	1.061	-0.392	4.305	-2.780
DJIA	0.019	1.090	-0.340	3.944	-2.867
FTSE	-0.009	1.147	-0.177	3.694	-2.905
CAC	-0.011	1.053	-0.182	3.346	-2.654
Nikkei	-0.054	1.083	-0.152	3.715	-2.832
IBM	0.035	1.033	-0.076	4.505	-2.539
GE	-0.025	1.012	0.026	3.897	-2.479
WMT	-0.039	1.001	0.066	3.909	-2.432
AT&T	-0.037	1.028	0.024	4.397	-2.523

TABLE 5. ESTIMATED PARAMETERS (S&P 500): ARFIMA MODELS

The table shows parameter estimates for different restrictions of the model: $(1 - \phi_1 L)(1 - L)^d(RV_t - \psi_t) = \lambda_1 I(r_{t-1} < 0)r_{t-1} + \lambda_2 I(r_{5,t-1} < 0)r_{5,t-1} + \lambda_3 I(r_{22,t-1} < 0)r_{22,t-1} + h_t \eta_t$, $h_t^2 = \theta_0 + \theta_1 V L_t^2 + \theta_2 h_{t-1}^2 + \theta_3 \nu_{t-1}^2$, $\eta_t \sim NIG^*(\alpha, \beta), (\Phi(\varepsilon_t), 1 - \Upsilon(\eta_t)) \sim C_{clayton}^\kappa$

Parameter	ARFIMA		ARFIMA+AE		ARFIMA-GARCH		ARFIMA+AE-GARCH		DARV-FI	
ψ	0.989	(0.029)	0.653	(0.028)	0.970	(0.053)	0.546	(0.028)	0.400	(0.025)
d	0.340	(0.011)	0.261	(0.008)	0.464	(0.012)	0.352	(0.013)	0.367	(0.015)
ϕ_1	0.064	(0.017)	0.033	(0.019)	-0.075	(0.017)	-0.047	(0.020)	-0.074	(0.019)
λ_1	-	-	-0.055	(0.006)	-	-	-0.049	(0.005)	-0.072	(0.006)
λ_2	-	-	-0.018	(0.003)	-	-	-0.012	(0.003)	-0.023	(0.004)
λ_3	-	-	-0.014	(0.002)	-	-	-0.011	(0.002)	-0.013	(0.002)
θ_0	0.105	(0.005)	0.089	(0.006)	0.002	(0.000)	0.001	(0.000)	0.013	(0.001)
θ_1	-	-	-	-	-	-	-	-	0.101	(0.007)
θ_2	-	-	-	-	0.845	(0.018)	0.852	(0.015)	-	-
θ_3	-	-	-	-	0.127	(0.017)	0.121	(0.015)	-	-
α	0.906	(0.046)	0.855	(0.052)	1.841	(0.125)	1.663	(0.117)	1.800	(0.168)
β	0.548	(0.046)	0.479	(0.049)	1.075	(0.106)	0.910	(0.095)	1.037	(0.140)
κ	0.169	(0.024)	0.157	(0.023)	0.207	(0.024)	0.228	(0.024)	0.254	(0.025)

TABLE 6. ESTIMATED PARAMETERS (S&P 500): HAR MODELS

The table shows parameter estimates for different restrictions of the model: $RV_t = \phi_0 + \phi_1 RV_{t-1} + \phi_2 RV_{5,t-1} + \phi_3 RV_{22,t-1} + \lambda_1 I(r_{t-1} < 0)r_{t-1} + \lambda_2 I(r_{5,t-1} < 0)r_{5,t-1} + \lambda_3 I(r_{22,t-1} < 0)r_{22,t-1} + h_t \eta_t$, $h_t^2 = \theta_0 + \theta_1 VL_t^2 + \theta_2 h_{t-1}^2 + \theta_3 \nu_{t-1}^2$, $\eta_t \sim NIG^*(\alpha, \beta)$, $(\Phi(\varepsilon_t), 1 - \Upsilon(\eta_t)) \sim C_{clayton}^\kappa$

Parameter	HAR/AE-GARCH		DARV (HAR)	
ϕ_0	0.090	(0.008)	0.087	(0.009)
ϕ_1	0.231	(0.016)	0.250	(0.016)
ϕ_2	0.357	(0.017)	0.362	(0.022)
ϕ_3	0.273	(0.016)	0.244	(0.018)
λ_1	-0.052	(0.005)	-0.072	(0.006)
λ_2	-0.016	(0.003)	-0.024	(0.004)
λ_3	-0.007	(0.002)	-0.009	(0.002)
θ_0	0.001	(0.000)	0.001	(0.001)
θ_1	-	-	0.054	(0.003)
θ_2	0.853	(0.014)	-	-
θ_3	0.122	(0.014)	-	-
α	1.752	(0.136)	1.669	(0.149)
β	1.014	(0.116)	0.931	(0.121)
κ	0.242	(0.024)	0.272	(0.025)

TABLE 7. ESTIMATION DIAGNOSTICS (S&P 500)

The table show a variety of diagnostics for the estimations presented in Tables 5 and 6. In the table $\hat{\nu}_t$ denote the estimated residuals and $\hat{\eta} = \hat{\nu}_t / \hat{h}_t$ = the standardized residuals. For the statistical tests in the table the p-values are reported. K-S denotes the Kolmogorov-Smirnov Test. z is the probability integral transform of the standardized residuals using the estimated NIG distribution

	ARFIMA	ARFIMA +AE	ARFIMA GARCH	ARFIMA+AE GARCH	DARV (FI)	HAR+AE GARCH	DARV (HAR)
Log-Likelihood	292.34	455.29	757.00	882.31	955.65	870.12	915.38
R^2	0.688	0.723	0.739	0.778	0.789	0.778	0.787
BIC	-455.59	-749.20	-1360.69	-1579.05	-1717.64	-1546.58	-1629.04
Std. Dev. ($\hat{\nu}_t$)	0.323	0.295	0.312	0.288	0.281	0.288	0.283
Ljung-Box (1) ($\hat{\nu}_t$)	0.000	0.000	0.000	0.021	0.009	0.325	0.000
Ljung-Box (5) ($\hat{\nu}_t$)	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Ljung-Box (10) ($\hat{\nu}_t$)	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Skewness ($\hat{\eta}_t$)	4.58	3.86	1.63	1.68	2.02	1.64	1.95
Kurtosis ($\hat{\eta}_t$)	65.02	52.03	9.82	10.47	14.24	9.88	13.74
ARCH (1) ($\hat{\eta}_t$)	0.000	0.000	0.306	0.397	0.693	0.422	0.692
ARCH (5) ($\hat{\eta}_t$)	0.000	0.000	0.910	0.876	0.315	0.843	0.340
ARCH (10) ($\hat{\eta}_t$)	0.000	0.000	0.293	0.510	0.158	0.636	0.194
K-S Test ($\hat{\eta}_t$)	0.000	0.000	0.051	0.076	0.716	0.015	0.913
Ljung-Box (1) ($z - \bar{z}$)	0.073	0.362	0.000	0.039	0.927	0.251	0.181
Ljung-Box (1) ($z - \bar{z}$) ²	0.000	0.000	0.199	0.250	0.994	0.283	0.716
Ljung-Box (1) ($z - \bar{z}$) ³	0.003	0.011	0.002	0.115	0.998	0.619	0.357
Ljung-Box (1) ($z - \bar{z}$) ⁴	0.000	0.000	0.274	0.208	0.965	0.479	0.857

TABLE 8. FORECASTING RESULTS: S&P 500.

The table reports the out-of-sample forecasting results for the S&P 500 daily realized volatility for the period between Jan/2001 and Jun/2009, where each model is re-estimated quarterly and used for one day ahead predictions. The specification for the conditional mean and conditional heteroskedasticity are separated by dashes. AE means that the model is estimated with asymmetric effects. DARV denotes the dually asymmetric realized volatility model. RMSE is the root mean squared error, MAE the mean absolute error. R^2 is the R-squared of a linear regression of the actual realized volatility on the forecasts. $R^2(\Delta)$ is the R-squared of a linear regression of the observed realized volatility change ($RV_t - RV_{t-1}$) on the forecasts. SPA is the p -value of the Superior Predictive Ability test developed by Hansen (2005). The null hypothesis is that a given model is not inferior to any other competing models in terms of a given loss function.

1 Day							
Model	R^2	RMSE	MAE	MAPE	$R^2(\Delta)$	SPA in MSE	SPA in R^2
ARFIMA	0.767	0.342	0.189	0.205	0.178	0.001	0.002
ARFIMA+AE	0.817	0.301	0.174	0.191	0.317	0.001	0.000
ARFIMA-GARCH	0.790	0.316	0.167	0.166	0.244	0.034	0.033
ARFIMA+AE-GARCH	0.829	0.283	0.157	0.160	0.380	0.573	0.585
DARV (FI)	0.834	0.277	0.156	0.161	0.407	0.719	0.727
HAR+AE-GARCH	0.827	0.285	0.159	0.166	0.374	0.405	0.416
DARV (HAR)	0.831	0.280	0.158	0.165	0.395	0.400	0.382
5 Days (cumulated)							
Model	R^2	RMSE	MAE	MAPE	$R^2(\Delta)$	SPA in MSE	SPA in R^2
ARFIMA	0.790	1.593	0.936	0.195	0.089	0.008	0.003
ARFIMA+AE	0.819	1.494	0.875	0.183	0.132	0.009	0.001
ARFIMA-GARCH	0.820	1.348	0.789	0.148	0.123	0.005	0.007
ARFIMA+AE-GARCH	0.834	1.296	0.749	0.141	0.165	0.009	0.012
DARV (FI)	0.841	1.269	0.733	0.139	0.201	0.538	0.536
HAR+AE-GARCH	0.842	1.265	0.740	0.143	0.173	0.014	0.015
DARV (HAR)	0.847	1.244	0.728	0.141	0.213	0.787	0.761
22 Days (cumulated)							
Model	R^2	RMSE	MAE	MAPE	$R^2(\Delta)$	SPA in MSE	SPA in R^2
ARFIMA	0.651	8.246	5.062	0.235	0.137	0.008	0.002
ARFIMA+AE	0.683	7.862	5.528	0.283	0.202	0.000	0.000
ARFIMA-GARCH	0.701	7.183	4.238	0.178	0.220	0.043	0.043
ARFIMA+AE-GARCH	0.714	7.031	4.066	0.170	0.272	0.010	0.011
DARV (FI)	0.720	6.953	3.994	0.167	0.283	0.252	0.290
HAR+AE-GARCH	0.732	6.806	4.059	0.178	0.309	0.010	0.007
DARV (HAR)	0.736	6.752	4.020	0.176	0.332	0.576	0.573

TABLE 9. FORECASTING RESULTS: OTHER SERIES

The table reports the out-of-sample forecasting results the realized volatility of the others series in the period between Jan/2001 and Jun/2009, where each model is re-estimated quarterly and used for one day ahead predictions. The specification for the conditional mean and conditional heteroskedasticity are separated by dashes. AE means that the model is estimated with asymmetric effects. DARV denotes the dually asymmetric realized volatility model. The table reports the R-squared of a linear regression of the actual realized volatility change on the forecasts ($R^2(\Delta)$). The parenthesis gives the p -value of the Superior Predictive Ability test developed by Hansen (2005) for null hypothesis that a given model is not inferior to any other competing alternatives in MSE.

	ARFIMA		ARFIMA+AE		HAR+AE		
	ARFIMA	+AE	GARCH	GARCH	DARV(FI)	GARCH	
DJIA	0.207 (0.000)	0.269 (0.004)	0.262 (0.018)	0.345 (0.142)	0.381 (0.802)	0.346 (0.128)	0.377 (0.659)
FTSE	0.236 (0.001)	0.314 (0.001)	0.259 (0.001)	0.347 (0.023)	0.368 (0.819)	0.334 (0.002)	0.346 (0.004)
CAC	0.199 (0.001)	0.265 (0.005)	0.232 (0.008)	0.278 (0.167)	0.301 (0.862)	0.269 (0.013)	0.283 (0.026)
Nikkei	0.213 (0.037)	0.256 (0.096)	0.223 (0.040)	0.266 (0.510)	0.270 (0.834)	0.258 (0.120)	0.259 (0.091)
IBM	0.193 (0.000)	0.232 (0.000)	0.254 (0.000)	0.290 (0.423)	0.296 (0.604)	0.294 (0.354)	0.301 (0.809)
GE	0.161 (0.004)	0.199 (0.007)	0.206 (0.007)	0.259 (0.111)	0.281 (0.851)	0.251 (0.051)	0.275 (0.576)
WMT	0.269 (0.002)	0.287 (0.001)	0.296 (0.078)	0.321 (0.051)	0.334 (0.789)	0.316 (0.130)	0.325 (0.340)
AT&T	0.211 (0.000)	0.221 (0.000)	0.237 (0.002)	0.252 (0.019)	0.259 (0.896)	0.251 (0.005)	0.259 (0.720)

TABLE 10. CONDITIONAL FORECASTS (S&P 500): LARGE REALIZED VOLATILITY CHANGES.

The table reports out of sample conditional forecasting results for the S&P 500 daily realized volatility for the period between Jan/2001 and Jun/2009, where each model is re-estimated quarterly and used for one day ahead predictions. The forecasts are conditional on the change in volatility ΔRV_t exceeding the defined ex post empirical percentile (calculated within the out of sample years). The specification for the conditional mean and conditional heteroskedasticity are separated by dashes. AE means that the model is estimated with asymmetric effects. DARV denotes the dually asymmetric realized volatility model. RMSE is the root mean squared error. R^2 is the R-squared of a linear regression of the realized volatility change on the forecasts.

ΔRV_t Model	> 80 th Percentile		> 90 th Percentile		> 95 th Percentile		> 99 th Percentile	
	R^2	MSE	R^2	RMSE	R^2	RMSE	R^2	RMSE
ARFIMA	0.092	0.399	0.080	0.504	0.070	0.637	0.000	1.068
ARFIMA+AE	0.035	0.402	0.013	0.513	0.000	0.658	0.034	1.114
ARFIMA-GARCH	0.254	0.353	0.255	0.446	0.231	0.569	0.059	1.002
ARFIMA+AE-GARCH	0.246	0.354	0.248	0.448	0.214	0.577	0.039	1.015
DARV (FI)	0.300	0.343	0.317	0.428	0.285	0.549	0.198	0.900
HAR+AE-GARCH	0.242	0.355	0.244	0.449	0.209	0.578	0.037	1.016
DARV (HAR)	0.309	0.343	0.324	0.426	0.286	0.549	0.166	0.909

TABLE 11. CONDITIONAL FORECASTS (S&P 500): HIGH REALIZED VOLATILITY.

The table reports out of sample conditional forecasting results for the S&P 500 daily realized volatility for the period between Jan/2001 and Jun/2009, where each model is re-estimated quarterly and used for one day ahead predictions. The forecasts are conditional on the realized volatility exceeding the defined ex post empirical percentile (calculated within the out of sample years). The specification for the conditional mean and conditional heteroskedasticity are separated by dashes. AE means that the model is estimated with asymmetric effects. DARV denotes the dually asymmetric realized volatility model. RMSE is the root mean squared error. R^2 is the R-squared of a linear regression of the actual realized volatility on the forecasts.

RV_t	> 80 th Percentile	> 90 th Percentile	> 95 th Percentile	> 99 th Percentile
Model	R^2 MSE	R^2 RMSE	R^2 RMSE	R^2 RMSE
ARFIMA	0.542 0.633	0.397 0.818	0.198 1.044	0.000 1.452
ARFIMA+AE	0.663 0.544	0.563 0.700	0.473 0.874	0.134 1.403
ARFIMA-GARCH	0.584 0.582	0.455 0.744	0.299 0.919	0.004 1.305
ARFIMA+AE-GARCH	0.672 0.515	0.574 0.658	0.493 0.789	0.069 1.227
DARV (FI)	0.685 0.499	0.592 0.639	0.520 0.754	0.360 1.055
HAR+AE-GARCH	0.669 0.515	0.568 0.658	0.486 0.784	0.067 1.223
DARV (HAR)	0.682 0.502	0.589 0.646	0.517 0.758	0.347 1.049

TABLE 12. CONDITIONAL FORECASTING RESULTS: OTHER SERIES

The table reports out of sample conditional forecasting results for the other realized volatility series for the period between Jan/2001 and Jun/2009, where each model is re-estimated quarterly and used for one day ahead predictions. The forecasts are conditional on the change in realized volatility and the realized volatility exceeding the defined ex post empirical percentile (calculated within the out of sample years). The specification for the conditional mean and conditional heteroskedasticity are separated by dashes. AE means that the model is estimated with asymmetric effects. DARV denotes the dually asymmetric realized volatility model. The Table reports the R^2 of a linear regression of the actual values on the forecasts.

	ARFIMA	ARFIMA+AE	ARFIMA-GARCH	ARFIMA+AE-GARCH	DARV(FI)	HAR+AE-GARCH	DARV(HAR)
R^2 of $E(\Delta RV_t \Delta RV_t > Q_{0.9})$:							
DJIA	0.100	0.074	0.147	0.136	0.224	0.141	0.210
FTSE	0.084	0.106	0.138	0.135	0.211	0.124	0.191
CAC	0.036	0.073	0.245	0.219	0.169	0.206	0.192
Nikkei	0.038	0.029	0.123	0.108	0.153	0.105	0.158
IBM	0.008	0.014	0.219	0.226	0.287	0.229	0.291
GE	0.170	0.120	0.270	0.264	0.285	0.253	0.294
WMT	0.059	0.082	0.105	0.109	0.170	0.116	0.165
AT&T	0.016	0.020	0.021	0.019	0.026	0.023	0.033
R^2 of $E(RV_t RV_t > Q_{0.9})$:							
DJIA	0.333	0.463	0.389	0.506	0.542	0.494	0.540
FTSE	0.114	0.234	0.210	0.278	0.311	0.266	0.300
CAC	0.078	0.162	0.213	0.202	0.220	0.188	0.218
Nikkei	0.466	0.490	0.506	0.526	0.526	0.503	0.512
IBM	0.431	0.466	0.462	0.485	0.498	0.486	0.497
GE	0.435	0.478	0.463	0.499	0.520	0.480	0.511
WMT	0.135	0.186	0.186	0.216	0.241	0.206	0.228
AT&T	0.288	0.307	0.310	0.317	0.337	0.339	0.357

TABLE 13. VALUE AT RISK ANALYSIS.

The table reports the out of sample value-at-risk results for the S&P 500 daily realized volatility for the period between Jan/2001 and Jun/2009, where each model is re-estimated quarterly and used for calculating 1%, 2.5% and 5% value-at-risk thresholds by the Monte Carlo Method described in section 3.1 and the ad hoc point forecasting method where $r_t \sim N(0, \widehat{RV})$. The specification for the conditional mean and conditional heteroskedasticity are separated by dashes. AE means that the model is estimated with asymmetric effects. DARV denotes the dually asymmetric realized volatility model. The column failures indicates the proportion of days when returns over the next day in the α lower tail of the predicted distribution. UC and IND are the p -values of the likelihood ratio tests for unconditional coverage and independence (against a first order Markov alternative) developed by Christoffersen (1998) (the joint test is omitted to save space). ES is the average of the empirical cumulative density function of returns at the VaR failures.

	1% VaR						
	Monte Carlo				Forecast VaR		
	Failures	UC	IND	ES	Failures	UC	IND
ARFIMA	0.007	0.162	0.643	0.006	0.021	0.000	0.951
ARFIMA+AE	0.008	0.358	0.599	0.006	0.023	0.000	0.943
ARFIMA-GARCH	0.009	0.815	0.536	0.005	0.025	0.000	0.806
ARFIMA+AE-GARCH	0.009	0.646	0.556	0.005	0.025	0.000	0.773
DARV (FI)	0.009	0.815	0.536	0.005	0.025	0.000	0.806
HAR+AE-GARCH	0.010	0.990	0.515	0.006	0.025	0.000	0.773
DARV (HAR)	0.009	0.646	0.556	0.004	0.026	0.000	0.740

	2.5% VaR						
	Monte Carlo				Forecast VaR		
	Failures	UC	IND	ES	Failures	UC	IND
ARFIMA	0.023	0.510	0.943	0.015	0.044	0.000	0.238
ARFIMA+AE	0.024	0.817	0.839	0.014	0.042	0.000	0.310
ARFIMA-GARCH	0.030	0.161	0.480	0.014	0.047	0.000	0.143
ARFIMA+AE-GARCH	0.030	0.125	0.950	0.014	0.046	0.000	0.193
DARV (FI)	0.030	0.161	0.932	0.013	0.046	0.000	0.193
HAR+AE-GARCH	0.029	0.204	0.506	0.013	0.047	0.000	0.155
DARV (HAR)	0.030	0.125	0.455	0.013	0.045	0.000	0.207

	5% VaR						
	Monte Carlo				Forecast VaR		
	Failures	UC	IND	ES	Failures	UC	IND
ARFIMA	0.054	0.447	0.654	0.027	0.071	0.000	0.837
ARFIMA+AE	0.056	0.250	0.032	0.027	0.072	0.000	0.011
ARFIMA-GARCH	0.060	0.035	0.130	0.026	0.076	0.000	0.941
ARFIMA+AE-GARCH	0.062	0.013	0.007	0.026	0.073	0.000	0.027
DARV (FI)	0.061	0.022	0.009	0.025	0.072	0.000	0.034
HAR+AE-GARCH	0.060	0.044	0.013	0.025	0.076	0.000	0.004
DARV (HAR)	0.059	0.055	0.014	0.025	0.075	0.000	0.001

TABLE 14. Robustness: Forecasted Return CDF at the Lowest Observed Returns (S&P 500)

The table reports the forecasted cumulative density functions (using the Monte Carlo Method) evaluated at the ten lowest observed returns in the period between Jan/2001 and Jun/2009. The specification for the conditional mean and conditional heteroskedasticity are separated by dashes. AE means that the model is estimated with asymmetric effects. DARV denotes the dually asymmetric realized volatility model.

Date	Return	RV_t	ARFIMA	ARFIMA	ARFIMA	ARFIMA+AE	DARV(FI)	HAR+AE	DARV(HAR)
				+AE	GARCH	GARCH		GARCH	
29-Sep-08	-9.219	4.845	0.002	0.001	0.003	0.002	0.002	0.003	0.002
07-Oct-08	-5.911	4.017	0.007	0.009	0.019	0.024	0.026	0.023	0.027
09-Oct-08	-7.922	4.393	0.008	0.009	0.021	0.022	0.025	0.022	0.027
15-Oct-08	-9.470	3.665	0.018	0.013	0.044	0.039	0.031	0.047	0.043
22-Oct-08	-6.295	3.678	0.113	0.145	0.129	0.149	0.164	0.161	0.174
05-Nov-08	-5.412	2.499	0.057	0.053	0.066	0.067	0.065	0.086	0.084
19-Nov-08	-6.311	3.532	0.033	0.027	0.043	0.038	0.042	0.042	0.048
20-Nov-08	-6.948	5.858	0.035	0.068	0.045	0.077	0.100	0.083	0.106
01-Dec-08	-9.354	2.562	0.005	0.004	0.014	0.011	0.010	0.015	0.013
20-Jan-09	-5.426	2.505	0.019	0.016	0.023	0.019	0.025	0.014	0.019

REALIZED VOLATILITY RISK

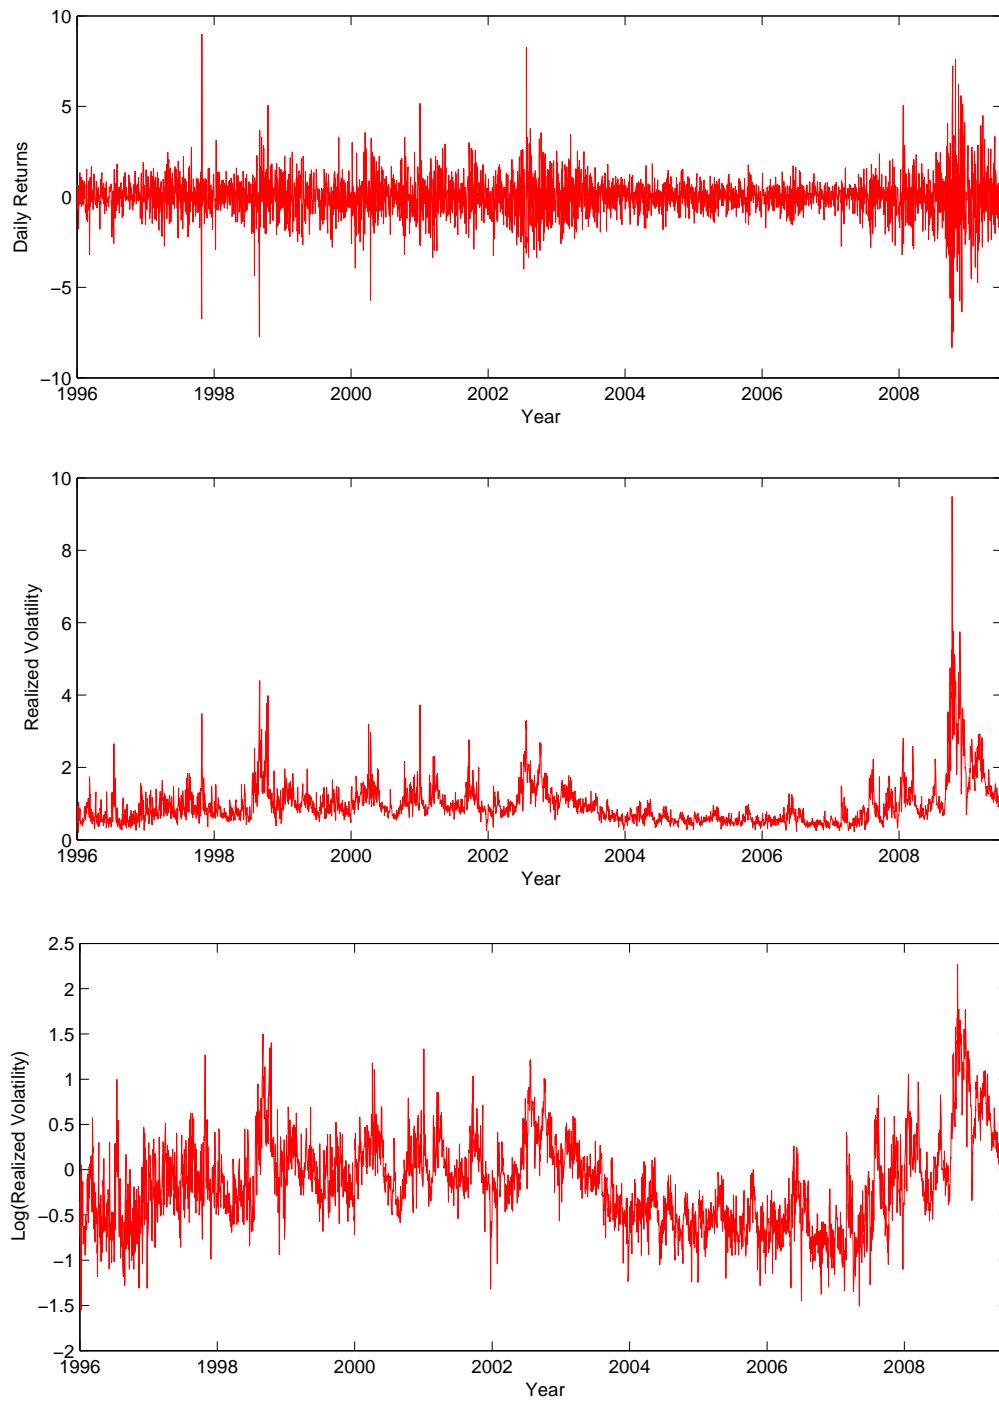


FIGURE 1. Time Series of returns (top), realized volatility (middle) and log realized volatility (bottom) for the S&P 500 index.

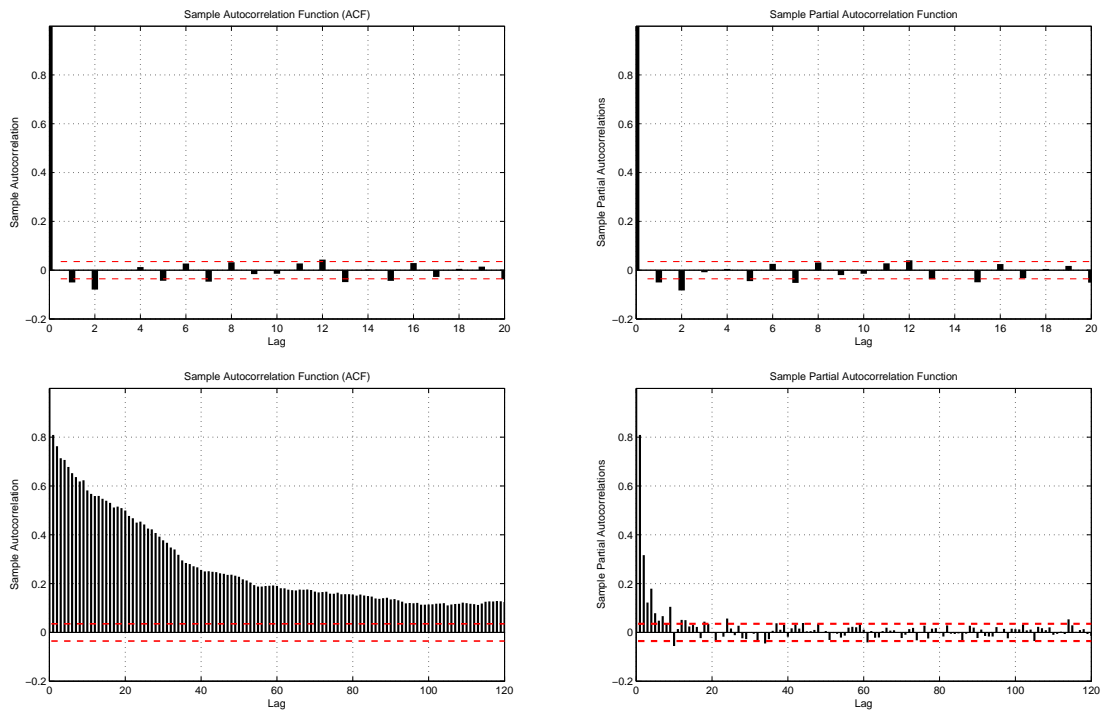


FIGURE 2. Sample autocorrelations and partial autocorrelations of returns (top) and realized volatility (bottom). The dashed lines give the standard Bartlett 95% confidence interval.

REALIZED VOLATILITY RISK




FIGURE 3. Scatter plots of lagged returns (top) and lagged monthly returns against realized volatility. The solid lines are regression lines, similar to news impact curves




FIGURE 4. Realized Volatility (top) and Monthly Returns (bottom) for the S&P 500 index.




FIGURE 5. In-Sample Percentage Errors for the HAR model with leverage effects.




FIGURE 6. Empirical Density for Returns Standardized by Realized Volatility Predictions.




FIGURE 7. Residuals series of the HAR model with leverage effects.




FIGURE 8. Sample autocorrelations for the squared (left) and absolute (right) residuals of the HAR model with leverage effects.




FIGURE 9. GARCH standard deviation series (top) and realized volatility fitted values (bottom) for the HAR model with leverage effects.

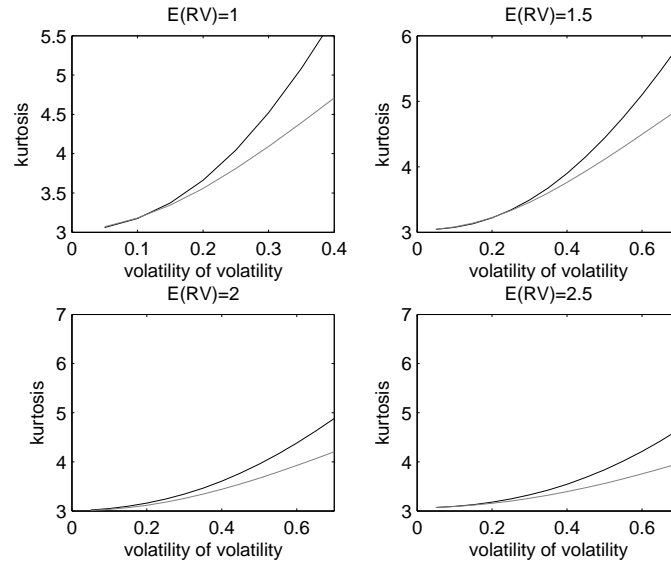


FIGURE 10. Kurtosis of the simulated distribution under the assumption that shocks to realized have the NIG (upper line) and normal distributions.




FIGURE 11. Densities for the simulated distributions (no volatility feedback).

REALIZED VOLATILITY RISK




FIGURE 12. Densities for the simulated distributions (with volatility feedback).




FIGURE 13. S&P500 Estimated Volatility of Realized Volatility.