

NoVaS Transformations: Flexible Inference for Volatility Forecasting*

Dimitris N. Politis[†] Dimitrios D. Thomakos[‡]

January 28, 2012

Abstract

In this paper we present several new findings on the NoVaS transformation approach for volatility forecasting introduced by Politis (2003a,b, 2007). In particular: (a) we present a new method for accurate volatility forecasting using NoVaS ; (b) we introduce a “time-varying” version of NoVaS and show that the NoVaS methodology is applicable in situations where (global) stationarity for returns fails such as the cases of local stationarity and/or structural breaks and/or model uncertainty; (c) we conduct an extensive simulation study on the forecasting ability of the NoVaS approach under a variety of realistic data generating processes (DGP); and (d) we illustrate the forecasting ability of NoVaS on a number of real datasets and compare it to realized and range-based volatility measures. Our empirical results show that the NoVaS -based forecasts lead to a much ‘tighter’ distribution of the forecasting performance measure. Perhaps our most remarkable finding is the *robustness* of the NoVaS forecasts in the context of structural breaks and/or other non-stationarities of the underlying data. Also striking is that forecasts based on NoVaS invariably outperform those based on the benchmark $GARCH(1,1)$ even when the true DGP *is* $GARCH(1,1)$ when the sample size is moderately large, e.g. 350 daily observations.

Keywords: ARCH, forecasting, GARCH, local stationarity, robustness, structural breaks, volatility.

*Earlier results from this research were presented in seminars at the Departments of Economics of the University of California at San Diego, University of Cyprus, and the University of Crete, as well as several conferences. We would like to thank Elena Andreou, conference and seminar participants for useful comments and suggestions. Many thanks are also due to an anonymous referee for a most constructive report, and to the Editors, Xiaohong Chen and Norman Swanson, for all their hard work in putting this volume together.

[†]Department of Mathematics and Department of Economics, University of California, San Diego, USA. Email: politis@math.ucsd.edu

[‡]Department of Economics, University of Peloponnese, Greece, and Rimini Center for Economic Analysis, Italy. Email: thomakos@uop.gr

1 Introduction

Accurate forecasts of the volatility of financial returns is an important part of empirical financial research. In this paper we present a number of new findings on the NoVaS transformation approach to volatility prediction. The NoVaS methodology was introduced by Politis (2003a,b, 2007) and further expanded in Politis and Thomakos (2008). The name of the method is an acronym for ‘Normalizing and Variance Stabilizing’ transformation. NoVaS is based on exploratory data analysis ideas, it is model-free, data-adaptive and—as the paper at hand hopes to demonstrate—especially relevant when making forecasts in the context of underlying data generating processes (DGPs) that exhibit non-stationarities (e.g. locally stationary time series, series with parameter breaks or regime switching etc.). In general, NoVaS allows for a flexible approach to inference, and is also well suited for application to short time series.

The original development of the NoVaS approach was made in Politis (2003a,b, 2007) having as its ‘spring board’ the popular ARCH model with normal innovations. In these papers, the main application was forecasting squared returns (as a proxy for forecasting volatility), and the evaluation of forecasting performance was addressed via the L_1 -norm (instead of the usual MSE) since the case was made that financial returns might not have finite 4th moment.

In the paper at hand we further investigate the performance of NoVaS in a pure forecasting context.¹ First, we present a method for *bona fide* volatility forecasting, extending the original NoVaS notion of forecasting squared returns. Second, we conduct a very comprehensive simulation study about the relative forecasting performance of NoVaS: we consider a wide variety of volatility models as data generating processes (DGPs), and we compare the forecasting performance of NoVaS with that of a benchmark *GARCH*(1,1) model. We introduce the notion of a “time-varying” NoVaS approach and show that is especially relevant in these cases where the assumption of global stationarity fails. The results of our simulations show that NoVaS forecasts lead to a much ‘tighter’ distribution of the forecasting performance measure (mean absolute deviation of the forecast errors), when compared to the benchmark model, for all DGPs we consider. This finding is especially relevant in the context of volatility forecasting for risk management. We further illustrate the use of NoVaS for a number of real datasets and compare the forecasting performance of NoVaS-based volatility forecasts with realized and range-based volatility measures, which are frequently used in assessing the performance of volatility forecasts.

The literature on volatility modeling, forecasting and the evaluation of volatility forecasts is very large and varied in topics covered. Possibly related to the paper at hand is the work by Hansen (2006) in which the problem of forming predictive intervals is addressed using a semiparametric, transformation-based approach. Hansen works with a set of (standardized) residuals from a parametric model, and then uses the empirical distribution function of these residuals to compute conditional quantiles that can be used in forming prediction intervals. The main similarity between Hansen’s work and NoVaS is that both approaches use a transformation of the original data and the empirical distribution to make forecasts. The main difference,

¹See also Politis and Thomakos (2008).

however, is that Hansen works in the context of a (possibly misspecified) model whereas NoVaS is totally model-free.

We can only selectively mention here some recent literature related to the forecasting problems we address: Mikosch and Starica (2004) for change in structure in volatility time series and GARCH modeling; Meddahi (2001) for an eigenfunction volatility modeling approach; Peng and Yao (2003) for robust LAD estimation of GARCH models; Poon and Granger (2003) for assessing the forecasting performance of various volatility models; Hansen, Lunde and Nason (2003) on selecting volatility models; Andersen, Bollerslev and Meddahi (2004, 2005) on analytic evaluation of volatility forecasts and the use of realized volatilities in evaluating volatility forecasts; Ghysels and Forsberg (2007) on the use and predictive power of absolute returns; Francq and Zakoïan (2005), Lux and Morales-Arias (2010) and Choi, Yu and Zivot (2010) on switching regime GARCH models, structural breaks and long memory in volatility; Hillebrand (2005) on GARCH models with structural breaks; Hansen and Lunde (2005, 2006) for comparing forecasts of volatility models against the standard *GARCH*(1,1) model and for consistent ranking of volatility models and the use of an appropriate series as the ‘true’ volatility; Ghysels, Santa Clara and Valkanov (2006) for predicting volatility by mixing data at different frequencies and Ghysels and Sohn (2009) for the type of power variation that predicts well volatility in the context of mixed data frequencies. Andersen, Bollerslev and Diebold (2007) for modeling realized volatility when jump components are included; Chen, Gerlach and Lin (2008) examine volatility forecasting in the context of threshold models coupled with volatility measurement based on intra-day range. The whole line of work of Andersen, Bollerslev, Diebold and their various co-authors on realized volatility and volatility forecasting is nicely summarized in their review article “Volatility and Correlation Forecasting”, in the *Handbook of Economic Forecasting*, see Andersen *et al.* (2006). Bandi and Russell (2008) discuss the selection of optimal sampling frequency in realized volatility estimation and forecasting; Patton and Sheppard (2008) discuss the evaluation of volatility forecasts while Patton and Sheppard (2009) present results on optimal combinations of realized volatility estimators in the context of volatility forecasting. Fryzlewicz, Sapatinas and Subba-Rao (2006, 2007) and Dahlhaus and Subba-Rao (2006, 2007) all work in the context of local stationarity and a new class of ARCH processes with slowly varying parameters. Of course this list is by no means complete.

The rest of the paper is organized as follows: in Section 2 we briefly review the general development of the NoVaS approach; in Section 3 we present the design of our simulation study and discuss the simulation results on forecasting performance; in Section 4 we present empirical applications of NoVaS using real-world data; finally, in Section 5 we offer some concluding remarks.

2 Review of the NoVaS Methodology

In this section we present a brief overview of the NoVaS transformation, the implied NoVaS distribution, the methods for distributional matching and NoVaS forecasting. For a more com-

prehensive review of the NoVaS methodology see Politis and Thomakos (2008).

2.1 NoVaS transformation and implied distribution

Let us consider a zero mean, strictly stationary time series $\{X_t\}_{t \in \mathbb{Z}}$ corresponding to the returns of a financial asset. We assume that the basic properties of X_t correspond to the ‘stylized facts’² of financial returns:

1. X_t has a non-Gaussian, approximately symmetric distribution that exhibits excess kurtosis.
2. X_t has time-varying conditional variance (volatility), denoted by $h_t^2 \stackrel{\text{def}}{=} \mathbb{E}[X_t^2 | \mathcal{F}_{t-1}]$ that exhibits strong dependence, where $\mathcal{F}_{t-1} \stackrel{\text{def}}{=} \sigma(X_{t-1}, X_{t-2}, \dots)$.
3. X_t is dependent although it possibly exhibits low or no autocorrelation which suggests possible nonlinearity.

These well-established properties affect the way one models and forecasts financial returns and their volatility and form the starting point of the NoVaS methodology.

The first step in the NoVaS transformation is variance stabilization to address the time-varying conditional variance property of the returns. We construct an empirical measure of the *time-localized* variance of X_t based on the information set $\mathcal{F}_{t|t-p} \stackrel{\text{def}}{=} \{X_t, X_{t-1}, \dots, X_{t-p}\}$

$$\gamma_t \stackrel{\text{def}}{=} G(\mathcal{F}_{t|t-p}; \alpha, \mathbf{a}), \gamma_t > 0 \forall t \quad (1)$$

where α is a scalar control parameter, $\mathbf{a} \stackrel{\text{def}}{=} (a_0, a_1, \dots, a_p)^\top$ is a $(p+1) \times 1$ vector of control parameters and $G(\cdot; \alpha, \mathbf{a})$ is to be specified.³ The function $G(\cdot; \alpha, \mathbf{a})$ can be expressed in a variety of ways, using a parametric or a semiparametric specification. To keep things simple we assume that $G(\cdot; \alpha, \mathbf{a})$ is additive and takes the following form:

$$G(\mathcal{F}_{t|t-p}; \alpha, \mathbf{a}) \stackrel{\text{def}}{=} \alpha s_{t-1} + \sum_{j=0}^p a_j g(X_{t-j}) \quad (2)$$

$$s_{t-1} = (t-1)^{-1} \sum_{j=1}^{t-1} g(X_j)$$

with the implied restrictions (to maintain positivity for γ_t) that $\alpha \geq 0$, $a_i \geq 0$, $g(\cdot) > 0$ and $a_p \neq 0$ for identifiability. Although other choices are possible, the natural choices for $g(z)$ are $g(z) = z^2$ or $g(z) = |z|$. With these designations, our empirical measure of the time-localized variance becomes a combination of an unweighted, recursive estimator s_{t-1} of the unconditional variance of the returns $\sigma^2 = \mathbb{E}[X_1^2]$, or of the mean absolute deviation of the returns $\delta = \mathbb{E}|X_1|$, and a weighted average of the current⁴ and the past p values of the squared or absolute returns.

Using $g(z) = z^2$ results in a measure that is reminiscent of an $ARCH(p)$ model which was employed in Politis (2003a,b, 2007). The use of absolute returns, i.e. $g(z) = |z|$ has also been

²Departures from the assumption of these ‘stylized facts’ have been discussed in Politis and Thomakos (2008); in this paper, we are mostly concerned about departures/breaks in stationarity—see Section 2.4 in what follows.

³See the discussion about the calibration of α and \mathbf{a} in the next section.

⁴The necessity and advantages of including the current value is elaborated upon by Politis (2003a,b,2004,2007).

advocated for volatility modeling; see e.g. Ghysels and Forsberg (2007) and the references therein. Robustness in the presence of outliers is an obvious advantage of absolute vs. squared returns. In addition, note that the mean absolute deviation is *proportional* to the standard deviation for the symmetric distributions that will be of current interest.

The second step in the NoVaS transformation is to use γ_t in constructing a studentized version of the returns, akin to the standardized innovations in the context of a parametric (e.g. GARCH-type) model. Consider the series W_t defined as:

$$W_t \equiv W_t(\alpha, \mathbf{a}) \stackrel{\text{def}}{=} \frac{X_t}{\phi(\gamma_t)} \quad (3)$$

where $\phi(z)$ is the time-localized standard deviation that is defined relative to our choice of $g(z)$, for example $\phi(z) = \sqrt{z}$ if $g(z) = z^2$ or $\phi(z) = z$ if $g(z) = |z|$. The aim now is to choose the NoVaS parameters in such a way as to make W_t follow as closely as possible a chosen target distribution that is easier to work with. The natural choice for such a distribution is the normal—hence the ‘normalization’ in the NoVaS acronym; other choices (such as the uniform) are also possible in applications, although perhaps not as intuitive. Note that by solving for X_t in equation (3), and using the fact that γ_t depends on X_t , it follows that we have the *implied* model representation:

$$X_t = U_t A_{t-1} \quad (4)$$

where U_t is the series obtained from the transformed series W_t in (3) and is required for forecasting—see Politis and Thomakos (2008). The component A_{t-1} depends only on past square or absolute returns, similar to the ARCH component of a GARCH model.

Remark 1. Politis (2003b, 2004, 2007) makes the case that financial returns seem to have finite second moment but infinite 4th moments. In that case, the normal target does not seem to be compatible with the choice of absolute returns—and the same is true of the uniform target—as it seems that the case $g(z) = |z|$ might be better suited for data that do not have a finite second moment. Nevertheless, there is always the possibility of encountering such extremely heavy-tailed data, e.g. in emerging markets, for which the absolute returns might be helpful.⁵ The set-up of potentially infinite 4th moments has been considered by Hall and Yao (2003) and Berkes and Horvath (2004) among others, and has important implications on an issue crucial in forecasting, namely the choice of loss function for evaluating forecast performance. The most popular criterion for measuring forecasting performance is the mean-squared error (MSE) which, however, is inapplicable in forecasting squared returns (and volatility) when the 4th moment is infinite. In contrast, the mean absolute deviation (MAD) is as intuitive as the MSE but does not suffer from this deficiency, and can thus be used in evaluating the forecasts of either squared or absolute returns and volatility; this L_1 loss criterion will be our preferred choice in this paper.⁶

⁵This might well be the case of the EFG dataset of Section 4 in what follows.

⁶See also the recent paper by Hansen and Lunde (2006) about the relevance of MSE in evaluating volatility forecasts.

2.2 NoVaS distributional matching

We next turn to the issue of optimal selection of the NoVaS parameters. The free parameters are p (the NoVaS order), and (α, \mathbf{a}) . The parameters α and \mathbf{a} are constrained to be non-negative to ensure the same for the variance. In addition, motivated by unbiasedness considerations, Politis (2003a,b, 2007) suggested the convexity condition $\alpha + \sum_{j=0}^p a_j = 1$. Finally, thinking of the coefficients a_i as local smoothing weights, it is intuitive to assume $a_i \geq a_j$ for $i > j$. We now discuss in detail the case when $\alpha = 0$; see Remark 2 for the case of nonzero α . A suitable scheme that satisfies the above conditions is given by exponential weights in Politis (2003a,b, 2007):

$$a_j = \begin{cases} 1/\sum_{j=0}^p \exp(-bj) & \text{for } j = 0 \\ a_0 \exp(-bj) & \text{for } j = 1, 2, \dots, p \end{cases} \quad (5)$$

where b is the exponential rate. We require the calibration of two parameters: a_0 and b . In this connection, let $\boldsymbol{\theta} \stackrel{\text{def}}{=} (p, b) \mapsto (\alpha, \mathbf{a})$, and denote the studentized series as $W_t \equiv W_t(\boldsymbol{\theta})$ rather than $W_t \equiv W_t(\alpha, \mathbf{a})$. For any given value of the parameter vector $\boldsymbol{\theta}$ we need to evaluate the ‘closeness’ of the marginal distribution of W_t with the target distribution.

Many different objective functions could be used for this. Let us denote such an objective function by $D_n(\boldsymbol{\theta})$, that obeys $D_n(\boldsymbol{\theta}) \geq 0$ and consider the following algorithm given in Politis (2003a, 2007):

- Let p take a very high starting value, e.g., let $p_{\max} \approx n/4$.
- Let $\alpha = 0$ and consider a discrete grid of b values, say $B \stackrel{\text{def}}{=} (b_{(1)}, b_{(2)}, \dots, b_{(M)})$, $M > 0$. Find the optimal value of b , say b^* , that minimizes $D_n(\boldsymbol{\theta})$ over $b \in B$, and compute the optimal parameter vector \mathbf{a}^* using equation (5).
- Trim the value of p by removing (i.e., setting to zero) the a_j parameters that do not exceed a pre-specified threshold, and re-normalize the remaining parameters so that their sum equals one.

The solution then takes the general form:

$$\boldsymbol{\theta}_n^* \stackrel{\text{def}}{=} \underset{\boldsymbol{\theta}}{\operatorname{argmin}} D_n(\boldsymbol{\theta}) \quad (6)$$

Such an optimization procedure will always have a solution in view of the intermediate value theorem and is discussed in the previous work on NoVaS.⁷ In empirical applications with financial returns it is usually sufficient to consider kurtosis-matching and thus to have $D_n(\boldsymbol{\theta})$ to take the form:

$$D_n(\boldsymbol{\theta}) \stackrel{\text{def}}{=} \left| \frac{\sum_{t=1}^n (W_t - \bar{W}_n)^4}{n s_n^4} - \kappa^* \right| \quad (7)$$

⁷This part of the NoVaS application appears similar at the outset to the Minimum Distance Method (MDM) of Wolfowitz (1957). Nevertheless, their objectives are quite different since the latter is typically employed for parameter estimation and testing whereas in NoVaS there is little interest in parameters—the focus lying on effective forecasting.

where $\bar{W}_n \stackrel{\text{def}}{=} (1/n) \sum_{t=1}^n W_t$ denotes the the sample mean, $s_n^2 \stackrel{\text{def}}{=} (1/n) \sum_{t=1}^n (W_t - \bar{W}_n)^2$ denotes the sample variance of the $W_t(\boldsymbol{\theta})$ series, and κ^* denotes the theoretical kurtosis coefficient of the target distribution. For the normal distribution $\kappa^* = 3$.

Remark 2. The discussion so far was under the assumption that the parameter α , that controls the weight given to the recursive estimator of the unconditional variance, is zero. If desired one can select a non-zero value by doing a direct search over a discrete grid of possible values while obeying the summability condition $\alpha + \sum_{j=0}^p a_j = 1$. For example, one can choose the value of α that optimizes out-of-sample predictive performance; see Politis (2003a,b, 2007) for more details.

2.3 NoVaS Forecasting

Once the NoVaS parameters are calibrated one can compute volatility forecasts. In fact, as Politis (2003a,b, 2007) has shown, one can compute forecasts for different functions of the returns, including higher powers (with absolute value or not). The choice of an appropriate forecasting loss function, both for producing and for evaluating the forecasts, is crucial for maximizing forecasting performance. Per our Remark 1, we focus on the L_1 loss function for producing the forecasts and the mean absolute deviation (MAD) of the forecast errors for assessing forecasting performance. After optimization of the NoVaS parameters we now have both the optimal transformed series $W_t^* = W_t(\boldsymbol{\theta}_n^*)$ but also the series U_t^* , the optimized version of the component of the *implied* model of equation (4). For a complete discussion of how one obtains NoVaS forecasts see Politis and Thomakos (2008) . In this section we present new results on NoVaS volatility forecasting.

Consider first the case where forecasting is performed based on squared returns. In Politis and Thomakos (2008) it is explained in detail that we require two components to forecast squared returns: one component is the conditional median of U_{n+1}^{2*} series and the other is the (known at time n) component A_n^{2*} . The rest of the procedure depends on the dependence properties of the studentized series W_n^* and the target distribution. From our experience, what has invariably been observed with financial returns is that their corresponding W_n^* series appears—for all practical purposes—to be uncorrelated.⁸ If the target distribution is the normal then—by the approximate normality of its *joint* distributions—the W_n^* series would be independent as well. The series U_n^* would inherit the W_n^* s independence by equations (3) and (4), and therefore the best estimate of the conditional median of U_{n+1}^{2*} is the unconditional sample median. Based on the above discussion we are now able to obtain volatility forecasts \hat{h}_{n+1}^2 in a variety of ways: (a) we can use the forecasts of squared (or absolute) returns; (b) we can use only the component of the conditional variance A_n^{2*} for $\phi(z) = \sqrt{z}$ or A_n for $\phi(z) = z$, akin to a GARCH approach; (c) we can combine (a) and (b) and use the forecast of the empirical measure $\hat{\gamma}_{n+1}$.

⁸This is an empirical finding; if, however, the W_n^* series is not independent then a slightly different procedure involving a (hopefully) linear predictor would be required—see Politis (2003a, 2007) and Politis and Thomakos (2008) for details.

The volatility forecast based on (a) above would be:

$$\widehat{h}_{n+1,1}^2 \equiv \widehat{X}_{n+1}^2 \stackrel{\text{def}}{=} \widehat{\text{Med}} [U_n^{2*}] A_n^{2*}. \quad (8)$$

When using (b) the corresponding forecast would just be the power of the A_n^* component, something very similar to an $ARCH(\infty)$ forecast:

$$\widehat{h}_{n+1,2}^2 \stackrel{\text{def}}{=} A_n^{2*}. \quad (9)$$

However, the most relevant and appropriate volatility forecast in the NoVaS context should be based on (c), i.e. on a forecast of the estimate of the time-localized variance measure $\widehat{\gamma}_{n+1}$, which was originally used to initiate the NoVaS procedure in equation (1). What is important to note is that forecasting based on $\widehat{\gamma}_{n+1}$ is neither forecasting of squared returns nor forecasting based on past information alone. It is, in fact, a linear combination of the two, thus incorporating elements from essentially two approaches. Combining equations (1), (2), (3), (4), (8) and (9) it is straightforward to show that $\widehat{\gamma}_{n+1}$ can be expressed as:

$$\begin{aligned} \widehat{\gamma}_{n+1} &\equiv \widehat{h}_{n+1,3}^2 \stackrel{\text{def}}{=} \left\{ a_0^* \widehat{\text{Med}} [U_n^{2*}] + 1 \right\} A_n^{2*} \\ &= a_0^* \widehat{h}_{n+1,1}^2 + \widehat{h}_{n+1,2}^2. \end{aligned} \quad (10)$$

Equation (10) is our new proposal for volatility forecasting using NoVaS. In his original work, Politis (2003a) used equation (8), and in effect conducted forecasting of the one-step-ahead squared returns via NoVaS. By contrast, equation (10) is a *bona fide* predictor of the one-step-ahead *volatility*, i.e., the conditional variance. For this reason, equation (10) will be the formula used in what follows, our simulations and real data examples.

Forecasts using absolute returns are constructed in a similar fashion, the only difference being that we will be forecasting *directly* standard deviations \widehat{h}_{n+1} and not variances. It is straightforward to show that the forecast based on (c) would be given by:

$$\begin{aligned} \widehat{\gamma}_{n+1} &\equiv \widehat{h}_{n+1,3} \stackrel{\text{def}}{=} \left\{ a_0^* \widehat{\text{Med}} [|U_n^*|] + 1 \right\} |A_n^*| \\ &= a_0^* \widehat{h}_{n+1,1} + \widehat{h}_{n+1,2} \end{aligned} \quad (11)$$

with $\widehat{h}_{n+1,1}$ and $\widehat{h}_{n+1,2}$ being identical expressions to equations (8) and (9) which use the absolute value transformation.

2.4 Departures from the assumption of stationarity: local stationarity and structural breaks

Consider the case of a very long time series $\{X_1, \dots, X_n\}$, e.g., a daily series of stock returns spanning a decade. It may be unrealistic to assume that the stochastic structure of the series has stayed invariant over such a long stretch of time. A more realistic model might assume a slowly-changing stochastic structure, i.e., a locally stationary model as given by Dahlhaus (1997).

Recent research has tried to address this issue by fitting time-varying GARCH models to the data but those techniques have not found global acceptance yet, in part due to their extreme computational cost. Fryzlewicz, Sapatinas and Subba-Rao (2006, 2007) and Dahlhaus

and Subba-Rao (2006, 2007b) all work in the context of local stationarity for a new class of ARCH processes with slowly varying parameters.

Surprisingly, NoVaS is flexible enough to accommodate such smooth/slow changes in the stochastic structure. All that is required is a *time-varying* NoVaS fitting, i.e., selecting/calibrating the NoVaS parameters on the basis of a *rolling* window of data as opposed to using the entire available past. Interestingly, as will be apparent in our simulations, the time-varying NoVaS method works well even in the presence of *structural breaks* that would typically cause a breakdown of traditional methods unless explicitly taken into account. The reason for this robustness is the simplicity in the NoVaS estimate of local variance: it is just a linear combination of (present and) past squared returns. Even if the coefficients of the linear combination are not optimally selected (which may happen in the neighborhood of a break), the linear combination remains a reasonable estimate of local variance.

By contrast, the presence of structural breaks can throw off the (typically nonlinear) fitting of GARCH parameters. Therefore, a GARCH practitioner must always be on the look-out for structural breaks, essentially conducting a hypothesis test before each application. While there are several change point tests available in the literature, the risk of non-detection of a change point can be a concern. Fortunately, the NoVaS practitioner does not have to worry about structural breaks because of the aforementioned *robustness* of the NoVaS approach.

3 NoVaS Forecasting Performance: A Simulation Analysis

It is of obvious interest to compare the forecasting performance of NoVaS-based volatility forecasts with the standard benchmark model, the $GARCH(1, 1)$, under a variety of different underlying DGPs. Although there are numerous models for producing volatility forecasts, including direct modeling of realized volatility series, it is not clear which of these models should be used in any particular situation, and whether they can always offer substantial improvements over the GARCH benchmark. In the context of a simulation, we will be able to better see the relative performance of NoVaS -based volatility forecasts versus GARCH-based forecasts and, in addition, we will have available the true volatility measure for forecast evaluation. This latter point, the availability of an appropriate series of true volatility, is important since in practice we do not have such a series of true volatility. The proxies range from realized volatility—generally agreed to be one of the best (if not the best) such measure—, to range-based measures, and to squared returns. We use such proxies in the empirical examples of the next section.

3.1 Simulation Design

We consider a variety of models as possible DGPs.⁹ Each model $j = 1, 2, \dots, M (= 7)$ is simulated over the index $i = 1, 2, \dots, N (= 500)$ with time indices $t = 1, 2, \dots, T (= 1250)$. The sample size T amounts to about 5 years of daily data. The parameter values for the models are chosen so as to reflect annualized volatilities between about 8% to 25%, depending on the model being used. For each model we simulate a volatility series and the corresponding returns series based on the standard representation:

$$\begin{aligned} X_{t,ij} &\stackrel{\text{def}}{=} \mu_j + h_{t,ij} Z_{t,ij} \\ h_{t,ij}^2 &\stackrel{\text{def}}{=} h_j(h_{t-1,ij}^2, X_{t-1,ij}^2, \boldsymbol{\theta}_{tj}) \end{aligned} \quad (12)$$

where $h_j(\cdot)$ changes depending on the model being simulated.

The seven models simulated are: a standard GARCH, a GARCH with discrete breaks (B-GARCH), a GARCH with slowly varying parameters (TV-GARCH), a Markov switching GARCH (MS-GARCH), a smooth transition GARCH (ST-GARCH), a GARCH with an added deterministic function (D-GARCH) and a stochastic volatility model (SV-GARCH). Note that the parameter vector $\boldsymbol{\theta}_t$ will be time-varying for the Markov switching model, the smooth transition model, the time-varying parameters model and the discrete breaks model. For the simulation we set $Z_t \sim t_{(3)}$, standardized to have unit variance.¹⁰

We next present the volatility equations of the above models. For ease of notation we drop the i and j subscripts when presenting the models. The first model we simulate is a standard *GARCH*(1, 1) with volatility equation given by:

$$h_t^2 = \omega + \alpha h_{t-1}^2 + \beta (X_{t-1} - \mu)^2 \quad (13)$$

The parameter values were set to $\alpha = 0.9$, $\beta = 0.07$ and $\omega = 1.2e - 5$, corresponding to an annualized volatility of 10%. The mean return was set to $\mu = 2e - 4$ (same for all models, except the MS-GARCH) and the volatility series was initialized with the unconditional variance.

The second model we simulate is a *GARCH*(1, 1) with discrete changes (breaks) in the variance parameters. These breaks depend on changes in the annualized unconditional variance, ranging from about 8% to about 22% and we assume two equidistant changes per year for a total of $B = 10$ breaks. The model form is identical to the *GARCH*(1, 1) above:

$$h_t^2 = \omega_b + \alpha_b h_{t-1}^2 + \beta_b (X_{t-1} - \mu)^2, \quad b = 1, 2, \dots, B \quad (14)$$

The α_b parameters were drawn from a uniform distribution in the interval $[0.8, 0.99]$ and the β_b parameters were computed as $\beta_b = 1 - \alpha_b - c$, for c either 0.015 or 0.02. The ω_b parameters were computed as $\omega_b = \sigma_b^2(1 - \alpha_b - \beta_b)/250$, where σ_b^2 is the annualized variance.

⁹In our design we do not just go for a limited number of DGPs but for a wide variety and we also generate a large number of observations, totalling over 4 million, across models and replications. Note that the main computational burden is the numerical (re)optimization of the *GARCH* model over 300K times across all simulations - and that involves (re)optimization only every 20 observations!

¹⁰We fix the degrees of freedom to their true value of 3 during estimation and forecasting, thus giving GARCH a relative advantage in estimation.

The third model we simulate is a $GARCH(1, 1)$ with slowly varying variance parameters, of a nature very similar to the time-varying ARCH models recently considered by Dahlhaus and Subba-Rao (2006, 2007). The model is given by:

$$h_t^2 = \omega(t) + \alpha(t)h_{t-1}^2 + \beta(t)(X_{t-1} - \mu)^2 \quad (15)$$

where the parameters satisfy the finite unconditional variance assumption $\alpha(t) + \beta(t) < 1$ for all t . The parameters functions $\alpha(t)$ and $\beta(t)$ are sums of sinusoidal functions of different frequencies ν_k of the form $c(t) = \sum_{k=1}^K \sin(2\pi\nu_k t)$, for $c(t) = \alpha(t)$ or $\beta(t)$. For $\alpha(t)$ we set $K = 4$ and $\nu_k = \{1/700, 1/500, 1/250, 1/125\}$ and for $\beta(t)$ we set $K = 2$ and $\nu_k = \{1/500, 1/250\}$. That is, we set the persistence parameter function $\alpha(t)$ to exhibit more variation than the parameter function $\beta(t)$ that controls the effect of squared returns.

The fourth model we simulate is a two-state Markov Switching $GARCH(1, 1)$ model, after Francq and Zakoian (2005). The form of the model is given by:

$$h_t^2 = \sum_{s=1}^2 \mathbf{1}\{P(S_t = s)\} [\omega_s + \alpha_s h_{t-1}^2 + \beta_s (X_{t-1} - \mu_s)^2] \quad (16)$$

In the first regime (high persistence and high volatility state) we set $\alpha_1 = 0.9$, $\beta_1 = 0.07$ and $\omega_1 = 2.4e - 5$, corresponding to an annualized volatility of 20%, and $\mu_1 = 2e - 4$. In the second regime (low persistence and low volatility state) we set $\alpha_2 = 0.7$, $\beta_2 = 0.22$ and $\omega_2 = 1.2e - 4$ corresponding to an annualized volatility of 10%, and $\mu_2 = 0$. The transition probabilities for the first regime are $p_{11} = 0.9$ and $p_{12} = 0.1$ while for the second regime we try two alternative specifications $p_{21} = \{0.3, 0.1\}$ and $p_{22} = \{0.7, 0.9\}$.

The fifth model we simulate is a (logistic) smooth transition $GARCH(1, 1)$; see Taylor (2004) and references therein for a discussion on the use of such models. The form the model takes is given by:

$$h_t^2 = \sum_{s=1}^2 Q_s(X_{t-1}) [\omega_s + \alpha_s h_{t-1}^2 + \beta_s (X_{t-1} - \mu_s)^2] \quad (17)$$

where $Q_1(\cdot) + Q_2(\cdot) = 1$ and $Q_s = [1 + \exp(-\gamma_1 X_{t-1}^{\gamma_2})]^{-1}$ is the logistic transition function. The parameters $\alpha_s, \beta_s, \omega_s$ and μ_s are set to the same values as in the previous MS-GARCH model. The parameters of the transition function are set to $\gamma_1 = 12.3$ and $\gamma_2 = 1$.

The sixth model we simulate is a $GARCH(1, 1)$ model with an added smooth deterministic function yielding a *locally stationary* model as a result. For the convenient case of a linear function we have that the volatility equation is the same as in the standard $GARCH(1, 1)$ model in equation (13) while the return equation takes the following form:

$$X_t = \mu + [a - b(t/T)] h_t Z_t \quad (18)$$

To ensure positivity of the resulting variance we require that $(a/b) > (t/T)$. Since $(t/T) \in (0, 1]$ we set $a = \alpha + \beta = 0.97$ and $b = (\beta/\alpha) \approx 0.078$ so that the positivity condition is satisfied for all t .

Finally, the last model we simulate is a stochastic volatility model with the volatility equation expressed in logarithmic terms and taking the form of an autoregression with normal innovations. The model now takes the form:

$$\log h_t^2 = \omega + \alpha \log h_{t-1}^2 + w_t, \quad w_t \sim \mathcal{N}(0, \sigma_w^2) \quad (19)$$

and we set the parameter values to $\alpha = 0.95$, $\omega \approx -0.4$ and $\sigma_w = 0.2$.

For each simulation run i and for each model j we split the sample into two parts $T = T_0 + T_1$, where T_0 is the estimation sample and T_1 is the forecast sample. We consider two values for T_0 , namely 250 or 900, which correspond respectively to about a year and three and a half years of daily data. We roll the estimation sample T_1 times and thus generate T_1 out-of-sample forecasts. In estimation the parameters are re-estimated (for GARCH) or updated (for NoVaS) every 20 observations (about one month for daily data). We always forecast the volatility of the corresponding return series we simulate (eqs. (10) and (11)) and evaluate it with the known, one-step ahead simulated volatility. NoVaS forecasts are produced for using a normal target distribution and both squared and absolute returns. The nomenclature used in the tables is as follows:

1. SQNT, NoVaS forecasts made using squared returns and normal target.
2. ABNT, NoVaS forecasts made using absolute returns and normal target.
3. GARCH, L_2 -based GARCH forecasts.
4. M-GARCH, L_1 -based GARCH forecasts.

The naïve forecast benchmark is the sample variance of the rolling estimation sample. Therefore, for each model j being simulated we produce a total of $F = 4$ forecasts; the forecasts are numbered $f = 0, 1, 2, \dots, F$ with $f = 0$ denoting the naïve forecast. We then have to analyze T_1 forecast errors $e_{t,ijf} \stackrel{\text{def}}{=} h_{t+1,ij}^2 - \hat{h}_{t+1,ijf}^2$. Using these forecast errors we compute the mean absolute deviation for each model, each forecast method and each simulation run as:

$$m_{ijf} = MAD_{ijf} \stackrel{\text{def}}{=} \frac{1}{T_1} \sum_{t=T_0+1}^T |e_{t,ijf}| \quad (20)$$

The values $\{m_{ijf}\}_{i=1, \dots, N; j=1, \dots, M; f=0, \dots, F}$ now become our data for meta-analysis. We compute various descriptive statistics about their distribution (across i , the independent simulation runs and for each f the different forecasting methods) like mean (\bar{x}_f in the tables), std. deviation ($\hat{\sigma}_f$ in the tables), min, the 10%, 25%, 50%, 75%, 90% quantiles and max (Q_p in the tables, $p = 0, 0.1, 0.25, 0.5, 0.75, 0.9, 1$). For example, we have that:

$$\bar{x}_{jf} \stackrel{\text{def}}{=} \frac{1}{N} \sum_{i=1}^N m_{ijf} \quad (21)$$

We also compute the percentage of times that the relative (to the benchmark) MAD 's of the NoVaS forecasts are better than the GARCH forecasts. Define $m_{ij,N} \stackrel{\text{def}}{=} m_{ijf}/m_{ij0}$, $f = 1, 2$

to be the ratio of the MAD of any of the NoVaS forecasts relative to the benchmark and $m_{ij,G} \stackrel{\text{def}}{=} m_{ijf}/m_{ij0}$, $f = 3, 4$ to be the ratio of the MAD of the two GARCH forecasts relative to the benchmark. That is, for each model j and forecasting method f we compute (dropping the j model subscript):

$$\widehat{P}_f \stackrel{\text{def}}{=} \frac{1}{N} \sum_{i=1}^N \mathbf{1}(m_{ij,N} \leq m_{ij,G}). \quad (22)$$

Then, we consider the total number of times that any NoVaS forecasting method had a smaller relative MAD compared to the relative MAD of the GARCH forecasts and compute also $\widehat{P} \stackrel{\text{def}}{=} \cup_f \widehat{P}_f$ as the union across. So \widehat{P}_f , for $f = 1, 2$ corresponds to the aforementioned methods NoVaS methods SQNT and ABNT respectively and \widehat{P} corresponds to their union.

3.2 Discussion of Simulation Results

The simulation helps compare the NoVaS forecasts to the usual GARCH forecasts, i.e., L_2 -based GARCH forecasts, and also to the M-GARCH forecasts, i.e., L_1 -based GARCH forecasts, the latter being recommended by Politis (2003a, 2004, 2007). All simulations results, that is the statistics of the MAD 's of equation (20) and the probabilities of equation (22), are compacted in three tables, Table 1 through Table 3. In Tables 1 and 2 we have the statistics for the MAD 's (Table 1 has the case of 1000 forecasts (smaller estimation sample) while Table 2 has the case of 350 forecasts (larger estimation sample). Table 3 has the statistics on the probabilities.

The main result that emerges from looking at these Tables is the very good and competitive performance of NoVaS forecasts, even when the true DGP is GARCH (DGP1 in the tables).¹¹ While it would seem intuitive that GARCH forecasts would have an advantage in this case we find that *any* of the NoVaS methods (SQNT, ABNT) is seen to outperform both GARCH and M-GARCH in *all* measured areas: mean of the MAD distribution (\bar{x}_f , mean error), tightness of MAD distribution ($\hat{\sigma}_f$ and the related quantiles), and finally the % of times NoVaS MAD was better. Actually, in this setting, the GARCH forecasts are vastly underperforming as compared to the Naive benchmark. The best NoVaS method here is the SQNT that achieves a mean error \bar{x}_f almost half of that of the benchmark, and with a much tighter MAD distribution. Comparing Tables 1 and 2 sheds more light in this situation: it appears that a training sample of size 250 is just too small for GARCH to work well; with a training sample of size 900 the performance of GARCH is greatly improved, and GARCH manages to beat the benchmark in terms of mean error (but not variance). SQNT NoVaS however is *still* the best method in terms of mean error

¹¹The phenomenon of poor performance of GARCH forecasting when the DGP is actually GARCH may seem puzzling and certainly deserves further study. Our experience based on the simulations suggests that the culprit is the occasional instability of the numerical MLE used to fit the GARCH model (computations performed in R using an explicit log-likelihood function with R optimization routines). Although in most trials the GARCH fitted parameters were accurate, every so often the numerical MLE gave grossly inaccurate answers which, of course, affect the statistics of forecasting performance. This instability was less pronounced when the fitting was done based on a large sample (case of 900). Surprisingly, a training sample as large as 250 (e.g. a year of daily data) was not enough to ward off the negative effects of this instability in fitting (and forecasting) based on the GARCH model.

and variance; it beats M-GARCH in terms of the \hat{P}_1 percentage, and narrowly underperforms as compared to GARCH in this criterion. All in all, SQNT NoVaS volatility forecasting appears to beat GARCH forecasts when the DGP *is* GARCH—a remarkable finding. Furthermore, GARCH apparently requires a very large training sample in order to work well; but with a sample spanning 3-4 years questions of non-stationarity may arise that will be addressed in what follows.

When the DGP is a GARCH with discrete breaks (B-GARCH, DGP2 in the tables) it is apparent here that ignoring possible structural breaks when fitting a GARCH model can be disastrous. The GARCH forecasts vastly underperform compared to the Naive benchmark with either small (Table 1) or big training sample (Table 2). Interestingly, *both* NoVaS methods are better than the benchmark with SQNT seemingly the best again. The SQNT method is better than either GARCH method at least 86% of the time. It should be stressed here that NoVaS does *not* attempt to estimate any breaks; it applies totally automatically, and is seemingly unperturbed by structural breaks. When we have a DGP of a GARCH with slowly varying parameters (TV-GARCH) the results are similar to the previous case except that the performance of GARCH is a little better as compared to the benchmark—but only when given a big training sample (compare Tables 1 and 2 for DGP3). However, still *both* NoVaS methods are better than either GARCH method. The best is again SQNT. Either of those beats either GARCH method at least 88% of the time (Table 3). For the Markov switching GARCH (MS-GARCH)(DGPs 4a and 4b in the tables) the results are essentially the same as with DGP2: GARCH forecasts vastly underperform the Naive benchmark with either small or big training sample. Again *all* NoVaS methods are better than the benchmark with SQNT being the best.

For the fifth DGP, the smooth transition GARCH (ST-GARCH)(DGP5 in the tables) the situation is more like the first one (where the DGP is plain GARCH); with a large enough training sample, GARCH forecasts are able to beat the benchmark, and be competitive with NoVaS. Still, however, SQNT NoVaS is best, not only because of smallest mean error but also in terms of tightness of *MAD* distribution. The results are also similar to the next DGP, GARCH with deterministic function (D-GARCH)(DGP6 in the tables), where given a large training sample, GARCH forecasts are able to beat the benchmark, and be competitive with NoVaS. Again, SQNT NoVaS is best, not only because of smallest mean error but also in terms of tightness of *MAD* distribution. Finally, for the last DGP, stochastic volatility model (SV-GARCH) (DGP7 in the tables) a similar behavior to the above two cases is found, but although (with a big training sample) GARCH does well in terms of mean error, note the large spread of the *MAD* distribution.

The results from the simulations can be summarized as follows:

- GARCH forecasts are extremely off-the-mark when the training sample is not large (of the order of 2-3 years of daily data). Note that large training sample sizes are prone to be problematic if the stochastic structure of the returns changes over time.
- Even given a large training sample, NoVaS forecasts are best; this holds *even when the true DGP is actually GARCH!*

- Ignoring possible breaks (B-GARCH), slowly varying parameters (TV-GARCH), or a Markov switching feature (MS-GARCH) when fitting a GARCH model can be disastrous in terms of forecasts. In contrast, NoVaS forecasts seem unperturbed by such gross non-stationarities.
- Ignoring the presence of a smooth transition GARCH (ST-GARCH), a GARCH with an added deterministic function (D-GARCH), or a stochastic volatility model (SV-GARCH) does not seem as crucial at least when the the implied nonstationarity features are small and/or slowly varying.
- Overall, it seems that SQNT NoVaS is the volatility forecasting method of choice since it is the best in all examples except TV-GARCH (in which case it is a close second to ABNT NoVaS).

4 Empirical Application

In this section we provide an empirical illustration of the application and potential of the NoVaS approach using four real datasets. In judging the forecasting performance for NoVaS we consider different measures of ‘true’ volatility, including realized and range-based volatility.

4.1 Data and Summary Statistics

Our first dataset consists of monthly returns and associated realized volatility for the S&P500 index, with the sample extending from February 1970 to May 2007 for a total of $n = 448$ observations. The second dataset consists of monthly returns and associated realized, range-based volatility for the stock of Microsoft (MSFT). The sample period is from April 1986 to August 2007 for a total of $n = 257$ observations. For both these datasets the associated realized volatility was constructed by summing daily squared returns (for the S&P500 data) or daily range-based volatility (for the MSFT data). Specifically, if we denote by $r_{t,i}$ the i^{th} daily return for month t then the monthly realized volatility is defined as $\sigma_t^2 \stackrel{\text{def}}{=} \sum_{i=1}^m r_{t,i}^2$, where m is the number of days. For the calculation of the realized range-based volatility denote by $H_{t,i}$ and $L_{t,i}$ the daily high and low prices for the i^{th} day of month t . The daily range-based volatility is defined as in Parkinson (1980) as $\sigma_{t,i}^2 \stackrel{\text{def}}{=} [\ln(H_{t,i}) - \ln(L_{t,i})]^2 / [4 \ln(2)]$; then, the corresponding monthly realized measure would be defined as $\sigma_t^2 \stackrel{\text{def}}{=} \sum_{i=1}^m \sigma_{t,i}^2$. Our third dataset consists of daily returns and realized volatility for the US dollar/Japanese Yen exchange rate for a sample period between 1997 and 2005 for a total of $n = 2236$ observations. The realized volatility measure was constructed as above using intraday returns. The final dataset we examine is the stock of a major private bank in the Athens Stock Exchange, EFG Eurobank. The sample period is from 1999 to 2004 for a total of $n = 1403$ observations. For lack of intraday returns we use the daily range-based volatility estimator as defined before.

Descriptive statistics of the returns for all four of our datasets are given in Table 4. We are

mainly interested in the kurtosis of the returns, as we will be using kurtosis-based matching in performing NoVaS . All series have unconditional means that are not statistically different from zero and no significant serial correlation, with the exception of the last series (EFG) that has a significant first order serial correlation estimate. Also, all four series have negative skewness which is, however, statistically insignificant except for the monthly S&P500 and MSFT series where it is significant at the 5% level. Finally, all series are characterized by heavy tails with kurtosis coefficients ranging from 5.04 (monthly S&P500) to 24.32 (EFG). The hypothesis of normality is strongly rejected for all series.

In Figures 1 to 8 we present graphs for the return series, the corresponding volatility and log volatility, the quantile-quantile (QQ) plot for the returns and four recursive moments. The computation of the recursive moments is useful for illustrating the potential unstable nature that may be characterizing the series. Figures 1 and 2 are for the monthly S&P500 returns, Figures 3 and 4 are for monthly MSFT returns, Figures 5 and 6 are for the daily USD/Yen returns and Figures 7 and 8 are for the daily EFG returns. Of interest are the figures that plot the estimated recursive moments. In Figure 2 we see that the mean and standard deviation of the monthly S&P500 returns are fairly stable while the skewness and kurtosis exhibit breaks. In fact, the kurtosis exhibits the tendency to rise in jolts/shocks and does not retreat to previous levels thereby indicating that there might not be a finite fourth moment for this series. Similar observations can be made for the other four series as far as recursive kurtosis goes. This is especially relevant about our argument that NoVaS can handle such possible global non-stationarities.

4.2 NoVaS Optimization and Forecasting Specifications

Our NoVaS in-sample analysis is performed for two possible combinations of target distribution and variance measures, i.e. squared and absolute returns using a normal target, as in the simulation analysis. We use the exponential NoVaS algorithm as discussed in section 2, with $\alpha = 0.0$, a trimming threshold of 0.01 and $p_{\max} = n/4$. The objective function for optimization is kurtosis-matching, i.e. $D_n(\boldsymbol{\theta}) = |\mathcal{K}_n(\boldsymbol{\theta})|$, as in equation (7) — robustness to deviations from these baseline specification is also discussed below. The results of our in-sample analysis are given in Table 5. In the table we present the optimal values of the exponential constant b^* , the first coefficient a_0^* , the implied optimal lag length p^* , the value of the objective function $D_n(\boldsymbol{\theta}^*)$ and two measures of distributional fit. The first is the QQ correlation coefficient for the original series, QQ_X , and the second is the QQ correlation coefficient for the transformed series $W_t(\boldsymbol{\theta}^*)$ series, QQ_W . These last two measures are used to gauge the ‘quality’ of the attempted distributional matching before and after the application of the NoVaS transformation.

Our NoVaS out-of-sample analysis is reported in Tables 6, 7, 8 and 9. All forecasts are based on a rolling sample whose length n_0 differs according to the series examined: for the monthly S&P500 series we use $n_0 = 300$ observations; for the monthly MSFT series we use $n_0 = 157$ observations; for EFG series we use $n_0 = 900$ observations; for the daily USD/Yen series we use $n_0 = 1250$ observations. The corresponding evaluation samples are $n_1 = \{148, 100, 986, 503\}$ for the four series respectively. Note that our examples cover a variety of different lengths, ranging

from 157 observations for the MSFT series to 1250 observations for the USD/Yen series. All forecasts we make are ‘honest’ out-of-sample forecasts: they use only observations prior to the time period to be forecasted. The NoVaS parameters are re-optimized as the window rolls over the entire evaluation sample (every month for the monthly series and every 20 observations for the daily series). We forecast volatility both by using absolute or squared returns (depending on the specification), as described in the section on NoVaS forecasting, and by using the empirical variance measure $\hat{\gamma}_{n+1}$ - see eqs. (10) and (11).¹² To compare the performance of the NoVaS approach we estimate and forecast using a standard *GARCH*(1,1) model for each series, assuming a $t_{(\nu)}$ distribution with degrees of freedom estimated from the data. The parameters of the model are re-estimated as the window rolls over, as described above. As noted in Politis (2003a,b, 2007), the performance of GARCH forecasts is found to be improved under an L_1 rather than L_2 loss. We therefore report standard mean forecasts as well as median forecasts from the GARCH models. We always evaluate our forecasts using the ‘true’ volatility measures given in the previous section and report several measures of forecasting performance. This is important as a single evaluation measure may not always provide an accurate description of the performance of competing models.

We first calculate the mean absolute deviation (MAD) and root mean-squared (RMSE) of the forecast errors $e_t \stackrel{\text{def}}{=} \sigma_t^2 - \hat{\sigma}_t^2$, given by:

$$MAD(e) \stackrel{\text{def}}{=} \frac{1}{n_1} \sum_{t=n_0+1}^n |e_t|, \quad RMSE(e) \stackrel{\text{def}}{=} \sqrt{\frac{1}{n_1} \sum_{t=n_0+1}^n (e_t - \bar{e})^2} \quad (23)$$

where $\hat{\sigma}_t^2$ denotes the forecast for any of the methods/models we use. As a Naive benchmark we use the (rolling) sample variance. We then calculate the Diebold and Mariano (1995) test for comparing forecasting models. We use the absolute value function in computing the relevant statistic and so we can formally compare the MAD rankings of the various models.

Finally, we calculate and report certain statistics based on the forecasting unbiasedness regression (also known as ‘Mincer-Zarnowitz regression’). This regression can be expressed in several ways and we use the following representation:

$$e_t = a + b\hat{\sigma}_t^2 + \zeta_t \quad (24)$$

where ζ_t is the regression error. Under the hypothesis of forecast unbiasedness we expect to have $\mathbb{E}[e_t | \mathcal{F}_{t-1}] = 0$ and therefore we expect $a = b = 0$ (and $\mathbb{E}[\zeta_t | \mathcal{F}_{t-1}] = 0$ as well.) Furthermore, the R^2 from the above regression is an indication as to how much of the forecast error variability can still be explained by the forecast. For any two competing forecasting models A and B we say that model A is superior to model B if $R_A^2 < R_B^2$, i.e. if we can make no further improvements in our forecast.

Our forecasting results are summarized in Tables 6 and 7 for the MAD and RMSE rankings and in Tables 8 and 9 for the Diebold-Mariano test and forecasting unbiasedness regressions. Similar results were obtained when using a recursive sample and are available on request.

¹²All NoVaS forecasts were made without applying an explicit predictor as all $W_t(\theta^*)$ series were found to be uncorrelated.

4.3 Discussion of Results

We begin our discussion with the in-sample results and, in particular, the degree of normalization achieved by NoVaS . Looking at the value of the objective function in Table 5 we see that it is zero to three decimals for practically all cases. Therefore, NoVaS is very successful in reducing the excess kurtosis in the original return series. In addition, the quantile-quantile correlation coefficient is very high (in excess of 0.99 in all cases examined, frequently being practically one). One should compare the two QQ measures of *before and after* the NoVaS transformation to see the difference that the transformation has on the data. The case of the EFG series is particularly worth mentioning as that series has the highest kurtosis: we can see from the table that we get a QQ correlation coefficient in excess of 0.998; this is a very clear indication that the desired distributional matching has been achieved for all practical purposes. A visual confirmation of the differences in the distribution of returns before and after NoVaS transformations is given in Figures 9 to 12. In these figures we have QQ plots for all the series and four combinations of return distributions, including the uniform for visual comparison. It is apparent from these figures that normalization has been achieved in all cases examined. Finally, a second noticeable in-sample result is the optimal lag length chosen by the different NoVaS specifications. In particular, we see from Table 16 that the optimal lag length is greater when using squared returns than when using absolute returns. As expected, longer lag lengths are associated with a smaller a_0^* coefficient.

We now turn to the out-of-sample results on the forecasting performance of NoVaS , which are summarized in Tables 6, 7, 8 and 9. The results are slightly different across the series we examine but the overall impression is that the NoVaS-based forecasts are superior to the GARCH forecasts, based on the combined performance of all evaluation measures. We discuss these in turn.

If we look at the MAD results in Table 6 the NoVaS forecasts outperform both the Naive benchmark and the GARCH-based forecasts. Note that the use of squared returns gives better results in the two series with the smallest sample kurtosis (S&P500 and USD/Yen series) while the use of absolute returns gives better results in the two series with the highest kurtosis (MSFT and EFG series). Its also worthwhile to note that the most drastic performance improvement, vis-a-vis the benchmark, can be seen for the MSFT series (smallest sample size) and the EFG series (highest kurtosis).¹³ This is important since we expected NoVaS to perform well in both these cases: the small sample size makes inference difficult while high kurtosis can be the result of non-stationarities in the series. Finally, the results are similar if we consider the RMSE ranking in Table 7. Based on these two descriptive evaluation measures the NoVaS forecasts outperform the benchmark and GARCH models.

To examine whether there are statistically significant differences between the NoVaS and GARCH forecasts and the benchmark, we next consider the results from the application of the

¹³Note also the performance improvement from the use of the median GARCH vs. the mean GARCH forecasts for the MSFT series. Recall that our simulation results showed that the performance of a GARCH model could be way off the mark if the training sample was small; here we use only 157 observations for training the MSFT series and the GARCH forecasts cannot outperform even the Naive benchmark.

Diebold-Mariano (1995) test for comparing forecasting performance. Looking at Table 7 we can see that there are statistically significant differences between the NoVaS forecasts and the Naive benchmark for the S&P500 series and the MSFT series, with the NoVaS forecasts being significantly better.¹⁴ For the other two series the test does not indicate a (statistically) superior performance of any of the other models compared to the benchmark.

Our empirical results so far clearly indicate that the NoVaS forecasts offer improvements in forecasting performance, both over the Naive benchmark and the GARCH models. We next discuss the results from the forecasting unbiasedness regressions of equation (24), where we try to see whether the forecasts are correct ‘on average’ and whether they make any systematic mistakes. We start by noting that the estimates from a regression like equation (24) suffer from bias since the regressor used, $\hat{\sigma}_t^2$, is estimated and not measured directly. Therefore we should be interpreting our results with some caution and connect them with our previous discussion. Looking at Table 9 we can see that in many cases the constant term a is estimated to be (numerically close to) zero, although it is statistically significant. The slope parameter b estimates show that there is still bias in the direction of the forecasts, either positive or negative, but the NoVaS estimates of b are in general much lower than those of the benchmark and the GARCH models, with the exception of the MSFT series. Furthermore, for the S&P500 and the EFG series the slope parameter is not statistically significant, at the 10% level, indicating a possibly unbiased NoVaS forecast. The R^2 values from these regressions are also supportive of the NoVaS forecasts (remember that low values are preferred over high values): the corresponding R^2 values from the NoVaS forecasts are lower than both the benchmark and the GARCH values by at least 30%. Note that for the S&P500 series where the value of the R^2 of the benchmark is lower than the corresponding NoVaS value, we also have a (numerically) large value for the slope parameter b for the benchmark compared to NoVaS. The only real problem with the R^2 from these regressions is for the MSFT series which we discuss below in Remark 4. All in all the results from Table 9 support the superior performance of NoVaS against its competitors and show that is a much less biased forecasting procedure.

Remark 3. Can we obtain further improvements using the NoVaS methodology? In particular, how do changes in the value of the α parameter affect the forecasting performance? This is an empirically interesting question since our results can be affected both by the small sample size and the degree of kurtosis in the data. The MSFT series exhibits both these problems and it is thus worthwhile to see whether we can improve our results by allowing the unconditional estimator of the variance to enter the calculations.¹⁵ We repeated our analysis for the MSFT series using $\alpha = 0.5$ and our results improved dramatically. The MAD and RMSE values from the ABNT NoVaS method dropped from 0.551 to 0.360 and from 0.951 to 0.524 respectively, with the Diebold-Mariano test still indicating a statistically significant performance over the Naive benchmark. In addition, the results from the forecasting unbiasedness regression are now better than the benchmark for the ABNT NoVaS method: the estimate of the slope parameter

¹⁴For the MSFT series the benchmark forecasts are also significantly better than the GARCH forecasts.

¹⁵Changing the value of α did not result in improvements in the other three series.

b is -0.145 and not statistically significant while the R^2 value is 0.010 compared to 0.012 for the benchmark.

In summary, our results are especially encouraging because they reflect on the very idea of the NoVaS transformation: a model-free approach that can account for different types of potential DGPs, that include breaks, switching regimes and lack of higher moments. NoVaS is successful in overcoming the parametrization and estimation problems that one would encounter in models that have variability and uncertainty not only in their parameters but also in their functional form. Of course our results are specific to the datasets examined and, it is true, we made no attempt to consider other types of parametric volatility models. But this is one of the problems that NoVaS attempts to solve: we have no *a priori* guidance as to which parametric volatility model to choose, be it simple GARCH, exponential GARCH, asymmetric GARCH and so on. With NoVaS we face no such problem as the very concept of a model does not enter into consideration.

5 Concluding Remarks

In this paper we have presented several findings on the NoVaS transformation approach for volatility forecasting introduced by Politis (2003a,b, 2007) and extended in Politis and Thomakos (2007). It was shown that NoVaS can be a flexible method for forecasting volatility of financial returns that is simple to implement, and robust against non-stationarities.

In particular, we focused on a new method for volatility forecasting using NoVaS and conducted an extensive simulation to study its forecasting performance under different DGPs. It was shown that the NoVaS methodology remains successful in situations where (global) stationarity fails such as the cases of local stationarity and/or structural breaks, and invariably outperforms the GARCH benchmark for all non-GARCH DGPs. Remarkably, the NoVaS methodology was found to outperform the GARCH forecasts *even* when the underlying DGP *is* itself a (stationary) GARCH as long as the sample size is only moderately large. It was also found that NoVaS forecasts lead to a much ‘tighter’ distribution of the forecasting performance measure used (the *MAD*) for all DGPs considered. Our empirical illustrations using four real datasets are also very supportive of the excellent forecasting performance of NoVaS compared to the standard GARCH forecasts.

Extensions of the current work include, among others, the use of the NoVaS approach on empirical calculations of value at risk (VaR), the generalization to more than one assets and the calculation of NoVaS correlations, and further extensive testing on the out-of-sample forecasting performance of the proposed method. Some of the above are pursued by the authors.

References

- [1] Andersen, T.G., Bollerslev, T., Christoffersen, P.F., and F. X. Diebold, 2006. “Volatility and Correlation Forecasting” in G. Elliott, C.W.J. Granger, and Allan Timmermann (eds.), *Handbook of Economic Forecasting*, Amsterdam: North-Holland, pp. 778-878.
- [2] Andersen, T.G., Bollerslev, T. and Meddahi, N., 2004. “Analytic evaluation of volatility forecasts”, *International Economic Review*, vol. 45, pp. 1079-1110.
- [3] Andersen, T.G., Bollerslev, T. and Meddahi, N., 2005. “Correcting the Errors: Volatility Forecast Evaluation Using High-Frequency Data and Realized Volatilities”, *Econometrica*, vol. 73, pp. 279-296.
- [4] Bandi, F.M. and J.R. Russell, 2008. “Microstructure noise, realized variance, and optimal sampling”, *Review of Economic Studies*, vol. 75, pp. 339-369.
- [5] Berkes, I. and L. Horvath, 2004. “The efficiency of the estimators of the parameters in GARCH processes”, *Annals of Statistics*, 32, pp. 633-655.
- [6] Chen, K., Gerlach, R. and Lin, E.W.M., 2008. “Volatility forecasting using threshold heteroscedastic models of the intra-day range”, *Computational Statistics & Data Analysis*, vol. 52, pp. 2990-3010.
- [7] Choi, K., Yu, W.-C. and E. Zivot, 2010. “Long memory versus structural breaks in modeling and forecasting realized volatility”, *Journal of International Money and Finance*, vol. 29, pp. 857-875.
- [8] Dahlhaus, R. (1997), “Fitting time series models to nonstationary processes”, *Annals of Statistics*, 25 pp. 1-37.
- [9] Dahlhaus, R. and S. Subba-Rao, 2006. “Statistical Inference for Time-Varying ARCH Processes”, *Annals of Statistics*, vol. 34, pp. 1075-1114.
- [10] Dahlhaus, R. and S. Subba-Rao, 2007. “A Recursive Online Algorithm for the Estimation of Time Varying ARCH Parameters”, *Bernoulli*, vol 13, pp. 389-422.
- [11] Diebold, F. X. and R. S. Mariano, 1995. “Comparing Predictive Accuracy”, *Journal of Business and Economic Statistics*, vol. 13, pp. 253-263.
- [12] Francq, C. and J-M. Zakoian, 2005. “L2 Structures of Standard and Switching-Regime GARCH Models”, *Stochastic Processes and Their Applications*, 115, pp. 1557-1582.
- [13] Fryzlewicz, P., Sapatinas, T. and S. Subba-Rao, 2006. “A Haar-Fisz Technique for Locally Stationary Volatility Estimation”, *Biometrika*, vol. 93, pp. 687-704.
- [14] Fryzlewicz, P., Sapatinas, T. and S. Subba-Rao, 2008. “Normalized Least Squares Estimation in Time-Varying ARCH Models”, *Annals of Statistics*, vol. 36, pp. 742-786.

- [15] Ghysels, E. and L. Forsberg, 2007. “Why Do Absolute Returns Predict Volatility So Well?”, *Journal of Financial Econometrics*, vol. 5, pp. 31-67.
- [16] Ghysels, E. and B. Sohn, 2009. “Which power variation predicts volatility well?”, *Journal of Empirical Finance*, vol. 16, pp. 686-700.
- [17] Ghysels, E., P. Santa-Clara, and R. Valkanov, 2006. “Predicting Volatility: How to Get Most Out of Returns Data Sampled at Different Frequencies”, *Journal of Econometrics*, vol. 131, pp. 59-95.
- [18] Hall, P. and Q. Yao, 2003. “Inference in ARCH and GARCH Models with heavy-tailed errors”, *Econometrica*, 71, pp. 285-317.
- [19] Hansen, B., 2006. “Interval Forecasts and Parameter Uncertainty”, *Journal of Econometrics*, vol. 127, pp. 377-398.
- [20] Hansen, P. R. and A. Lunde, 2005. “A forecast comparison of volatility models: does anything beat a $GARCH(1, 1)$?”, *Journal of Applied Econometrics*, 20(7), pp. 873-889.
- [21] Hansen, P. R. and A. Lunde, 2006. “Consistent ranking of volatility models”, *Journal of Econometrics*, 131, pp. 97-121.
- [22] Hansen, P.R., Lunde, A. and Nason, J.M., 2003. “Choosing the best volatility models: the model confidence set approach”, *Oxford Bulletin of Economics and Statistics*, vol. 65, pp. 839-861.
- [23] Hillebrand, E. 2005. “Neglecting Parameter Changes in GARCH Models”, *Journal of Econometrics*, 129, pp. 121-138.
- [24] Lux, T. and L. Morales-Arias, 2010. “Forecasting volatility under fractality, regime switching, long-memory and t-innovations”, *Computational Statistics & Data Analysis*, vol. 54, pp. 2676-2692.
- [25] Meddahi, N., 2001. “An eigenfunction approach for volatility modeling”, Technical report, CIRANO Working paper 2001s-70, University of Montreal.
- [26] Mikosch, T. and C. Starica, 2004. “Change of Structure in Financial Time Series, Long Range Dependence and the GARCH model”, CAF Working Paper Series, No. 58.
- [27] Parkinson, M., 1980. “The Extreme Value Method for Estimating the Variance of the Rate of Return”, *Journal of Business*, 53, pp.6168.
- [28] Patton, A., 2011. “Volatility forecast evaluation and comparison using imperfect volatility proxies”, *Journal of Econometrics*, vol. 160, pp. 246-256.
- [29] Patton, A. and K. Sheppard, 2008. “Evaluating volatility and correlation forecasts”, in T. G. Andersen *et al.*, (Eds.), *Handbook of Financial Time Series*, Springer Verlag.

- [30] Patton, A. and K. Sheppard, 2009. "Optimal combinations of realized volatility estimators", *International Journal of Forecasting*, vol. 25, pp. 218-238.
- [31] Peng, L. and Q. Yao, 2003. "Least absolute deviations estimation for ARCH and GARCH models", *Biometrika*, 90, pp. 967-975.
- [32] Politis, D.N., 2003a. "Model-Free Volatility Prediction", UCSD Dept. of Economics Discussion Paper 2003-16.
- [33] Politis, D.N., 2003b. "A Normalizing and Variance-Stabilizing Transformation for Financial Time Series, in *Recent Advances and Trends in Nonparametric Statistics*, M.G. Akritas and D.N. Politis, (Eds.), Elsevier: North Holland, pp. 335-347.
- [34] Politis, D.N., 2004. "A heavy-tailed distribution for ARCH residuals with application to volatility prediction", *Annals of Economics and Finance*, vol. 5, pp. 283-298.
- [35] Politis, D.N., 2007. "Model-free vs. model-based volatility prediction", *J. Financial Econometrics*, vol. 5, pp. 358-389.
- [36] Politis, D. and D. Thomakos, 2008. "Financial Time Series and Volatility Prediction using NoVaS Transformations", in *Forecasting in the Presence of Parameter Uncertainty and Structural Breaks*, D. E. Rapach and M. E. Wohar (Eds.), Emerald Group Publishing Ltd.
- [37] Poon, S. and C. Granger, 2003. "Forecasting Volatility in Financial Markets: A Review", *Journal of Economic Literature*, 41, pp. 478539.
- [38] Taylor, J., 2004. "Volatility Forecasting using Smooth Transition Exponential Smoothing", *International Journal of Forecasting*, vol. 20, pp. 273-286.
- [39] Wolfowitz, A., 1957. "The Minimum Distance Method", *Annals of Mathematical Statistics*, 28, pp. 75-88.

Table 1. Summary of simulation results across DGP and models, $T_1 = 1,000$

\bar{x}_f	DGP1	DGP2	DGP3	DGP4a	DGP4b	DGP5	DGP6	DGP7
Naive	0.24	0.43	0.31	0.36	0.48	0.32	0.16	0.26
SQNT	0.14	0.17	0.14	0.20	0.18	0.15	0.12	0.21
ABNT	0.21	0.28	0.15	0.30	0.26	0.24	0.18	0.23
GARCH	2.64	29.10	1.70	1.33	3.21	2.05	1.62	1.50
M-GARCH	1.56	16.15	1.02	0.88	1.91	1.25	0.98	0.95
$\hat{\sigma}_f$	DGP1	DGP2	DGP3	DGP4a	DGP4b	DGP5	DGP6	DGP7
Naive	0.33	0.96	0.53	0.42	2.34	0.34	0.17	0.16
SQNT	0.08	0.47	0.23	0.12	0.15	0.07	0.04	0.13
ABNT	0.09	0.47	0.16	0.14	0.15	0.10	0.05	0.11
GARCH	13.43	385.48	14.11	3.04	23.07	10.15	9.01	8.74
M-GARCH	7.39	212.13	7.78	1.68	12.71	5.60	4.96	4.81
$Q_{0.10}$	DGP1	DGP2	DGP3	DGP4a	DGP4b	DGP5	DGP6	DGP7
Naive	0.09	0.13	0.12	0.15	0.13	0.12	0.08	0.17
SQNT	0.09	0.10	0.06	0.14	0.12	0.11	0.10	0.15
ABNT	0.16	0.17	0.09	0.23	0.19	0.19	0.15	0.18
GARCH	0.10	0.15	0.10	0.17	0.13	0.12	0.09	0.18
M-GARCH	0.16	0.18	0.11	0.24	0.19	0.18	0.14	0.22
$Q_{0.50}$	DGP1	DGP2	DGP3	DGP4a	DGP4b	DGP5	DGP6	DGP7
Naive	0.15	0.22	0.19	0.24	0.23	0.21	0.10	0.23
SQNT	0.11	0.12	0.09	0.17	0.15	0.13	0.10	0.19
ABNT	0.19	0.20	0.11	0.27	0.23	0.22	0.16	0.22
GARCH	0.34	0.50	0.20	0.41	0.31	0.26	0.21	0.33
M-GARCH	0.29	0.40	0.17	0.37	0.30	0.26	0.20	0.32
$Q_{0.90}$	DGP1	DGP2	DGP3	DGP4a	DGP4b	DGP5	DGP6	DGP7
Naive	0.45	0.71	0.51	0.61	0.62	0.62	0.28	0.32
SQNT	0.19	0.21	0.19	0.26	0.24	0.20	0.15	0.26
ABNT	0.28	0.36	0.20	0.37	0.33	0.32	0.22	0.28
GARCH	3.53	4.19	1.51	2.88	2.83	2.53	1.78	2.71
M-GARCH	2.04	2.51	0.91	1.79	1.69	1.53	1.13	1.62

Notes:

1. DGP i denotes the i^{th} data generating process as follows: 1 for GARCH, 2 for B-GARCH, 3 for TV-GARCH, 4a and 4b for MS-GARCH, 5 for ST-GARCH, 6 for D-GARCH and 7 for SV-GARCH.
2. Table entries give statistics of the MAD of the forecast errors over 500 replications and $T_1 = 1,000$ denotes the number of forecasts generated for computing MAD in each replication.
3. \bar{x}_f denotes the sample mean, $\hat{\sigma}_f$ denotes the sample std. deviation and Q_p denotes the p^{th} sample quantile of the MAD distribution over 500 replications.
4. Naïve denotes forecasts based on the rolling sample variance, SQNT (ABNT) denotes NoVaS forecasts based on a normal target distribution and squared (absolute) returns, GARCH and M-GARCH denote L_2 and L_1 based forecasts from a standard GARCH model.

Table 2. Summary of simulation results across DGP and models, $T_1 = 350$

\bar{x}_f	DGP1	DGP2	DGP3	DGP4a	DGP4b	DGP5	DGP6	DGP7
Naive	0.26	0.39	0.31	0.37	0.47	0.31	0.13	0.26
SQNT	0.14	0.10	0.13	0.20	0.20	0.15	0.11	0.22
ABNT	0.21	0.22	0.15	0.32	0.27	0.25	0.17	0.24
GARCH	0.22	0.65	0.20	2.70	5.56	0.19	0.12	0.24
M-GARCH	0.24	0.47	0.20	1.65	3.21	0.24	0.15	0.27
$\hat{\sigma}_f$	DGP1	DGP2	DGP3	DGP4a	DGP4b	DGP5	DGP6	DGP7
Naive	0.39	0.87	0.58	0.70	1.95	0.42	0.19	0.33
SQNT	0.13	0.09	0.30	0.16	0.30	0.12	0.05	0.36
ABNT	0.13	0.32	0.19	0.33	0.26	0.17	0.06	0.28
GARCH	0.75	4.99	0.37	42.77	84.17	0.31	0.22	0.98
M-GARCH	0.49	2.75	0.38	23.68	46.39	0.27	0.14	0.58
$Q_{0.10}$	DGP1	DGP2	DGP3	DGP4a	DGP4b	DGP5	DGP6	DGP7
Naive	0.07	0.12	0.13	0.11	0.11	0.10	0.04	0.16
SQNT	0.09	0.07	0.06	0.13	0.11	0.10	0.10	0.13
ABNT	0.15	0.12	0.09	0.21	0.18	0.17	0.14	0.16
GARCH	0.04	0.07	0.08	0.08	0.07	0.06	0.04	0.13
M-GARCH	0.09	0.09	0.10	0.14	0.12	0.12	0.08	0.16
$Q_{0.50}$	DGP1	DGP2	DGP3	DGP4a	DGP4b	DGP5	DGP6	DGP7
Naive	0.14	0.21	0.19	0.22	0.20	0.20	0.08	0.22
SQNT	0.11	0.08	0.08	0.16	0.14	0.12	0.10	0.19
ABNT	0.18	0.15	0.11	0.25	0.21	0.21	0.15	0.21
GARCH	0.10	0.13	0.12	0.15	0.13	0.12	0.07	0.18
M-GARCH	0.17	0.15	0.13	0.23	0.19	0.19	0.13	0.23
$Q_{0.90}$	DGP1	DGP2	DGP3	DGP4a	DGP4b	DGP5	DGP6	DGP7
Naive	0.48	0.56	0.49	0.64	0.67	0.56	0.24	0.34
SQNT	0.20	0.13	0.19	0.27	0.27	0.21	0.13	0.28
ABNT	0.29	0.28	0.20	0.40	0.37	0.30	0.20	0.30
GARCH	0.35	0.37	0.28	0.45	0.42	0.34	0.18	0.26
M-GARCH	0.33	0.34	0.29	0.47	0.46	0.34	0.20	0.34

Notes:

1. DGP i denotes the i^{th} data generating process as follows: 1 for GARCH, 2 for B-GARCH, 3 for TV-GARCH, 4a and 4b for MS-GARCH, 5 for ST-GARCH, 6 for D-GARCH and 7 for SV-GARCH.
2. Table entries give statistics of the MAD of the forecast errors over 500 replications and $T_1 = 1,000$ denotes the number of forecasts generated for computing MAD in each replication.
3. \bar{x}_f denotes the sample mean, $\hat{\sigma}_f$ denotes the sample std. deviation and Q_p denotes the p^{th} sample quantile of the MAD distribution over 500 replications.
4. Naïve denotes forecasts based on the rolling sample variance, SQNT (ABNT) denotes NoVaS forecasts based on a normal target distribution and squared (absolute) returns, GARCH and M-GARCH denote L_2 and L_1 based forecasts from a standard GARCH model.

Table 3. Summary of simulation results across DGP and models
percentage of times that NoVaS forecasts are better than the benchmarks

DGP	Benchmark	\hat{P}_1	\hat{P}_2	\hat{P}	\hat{P}_1	\hat{P}_2	\hat{P}
DGP1	GARCH	0.93	0.66	0.93	0.43	0.13	0.43
	M-GARCH	1.00	0.74	1.00	0.86	0.35	0.86
DGP2	GARCH	0.98	0.76	0.98	0.86	0.35	0.86
	M-GARCH	0.99	0.87	0.99	0.96	0.42	0.96
DGP3	GARCH	0.98	0.85	1.00	0.89	0.52	0.98
	M-GARCH	0.99	0.98	1.00	0.96	0.91	0.99
DGP4a	GARCH	0.94	0.62	0.94	0.42	0.14	0.42
	M-GARCH	1.00	0.73	1.00	0.85	0.30	0.86
DGP4b	GARCH	0.90	0.60	0.90	0.45	0.18	0.46
	M-GARCH	1.00	0.75	1.00	0.87	0.36	0.89
DGP5	GARCH	0.91	0.55	0.91	0.47	0.14	0.47
	M-GARCH	1.00	0.67	1.00	0.91	0.31	0.92
DGP6	GARCH	0.76	0.55	0.76	0.24	0.09	0.24
	M-GARCH	1.00	0.61	1.00	0.77	0.19	0.77
DGP7	GARCH	0.90	0.70	0.91	0.36	0.17	0.40
	M-GARCH	0.97	0.99	1.00	0.84	0.73	0.91

Notes:

1. DGP i denotes the i^{th} data generating process as follows: 1 for GARCH, 2 for B-GARCH, 3 for TV-GARCH, 4a and 4b for MS-GARCH, 5 for ST-GARCH, 6 for D-GARCH and 7 for SV-GARCH.
2. Table entries give the proportion of times that the NoVaS MAD relative to the naïve benchmark was smaller than the GARCH MAD relative to the same benchmark, see equation (22) in the main text.

Table 4. Descriptive Statistics for Empirical Series

Series	n	\bar{x}	$\hat{\sigma}$	\mathcal{S}	\mathcal{K}	\mathcal{N}	$\hat{r}(1)$
S&P500, monthly	448	1.01%	4.35%	-0.37	5.04	0.00	0.00
MSFT, monthly	257	0.00%	1.53%	-1.75	9.00	0.00	-0.10
USD/Yen, daily	2236	-0.00%	0.72%	-0.70	8.52	0.00	0.00
EFG, daily	1403	-0.07%	2.11%	-1.24	24.32	0.00	0.14

Notes:

1. n denotes the number of observations, \bar{x} denotes the sample mean, $\hat{\sigma}$ denotes the sample standard deviation, \mathcal{S} denotes the sample skewness, \mathcal{K} denotes the sample kurtosis.
2. \mathcal{N} is the p-value of the Cramer-Von Misses test for normality of the underlying series.
3. $\hat{r}(1)$ denotes the estimate of the first order serial correlation coefficient.

Table 5. Full-sample NoVaS Summary Measures

Type	b^*	$D_n(\theta^*)$	a_0^*	p^*	QQ_X	QQ_W
S&P500 monthly						
SQNT	0.039	0.000	0.052	34	0.989	0.996
ABNT	0.070	0.000	0.078	27	0.989	0.996
MSFT monthly						
SQNT	0.175	0.000	0.171	15	0.916	0.988
ABNT	0.251	0.000	0.231	12	0.916	0.986
USD/Yen daily						
SQNT	0.062	0.000	0.071	29	0.978	0.999
ABNT	0.121	0.000	0.124	20	0.978	0.999
EFG daily						
SQNT	0.089	0.007	0.096	24	0.943	0.999
ABNT	0.171	0.000	0.166	16	0.943	0.999

Notes:

1. SQNT, ABNT denote NoVaS made forecasts based on square and absolute returns and a normal target distribution.
2. b^* , a_0^* and p^* denote the optimal exponential constant, first coefficient and implied lag length.
3. $D_n(\theta^*)$ is the value of the objective function based on kurtosis matching.
4. QQ_X and QQ_W denote the QQ correlation coefficient of the original series and the transformed series respectively.

Table 6. Mean Absolute Deviation (MAD) of Forecast Errors

Series	Naïve	SQNT	ABNT	Mean	Median
				GARCH	GARCH
S&P500, monthly	0.152	0.118	0.134	0.139	0.157
MSFT, monthly	1.883	1.030	0.551	43.28	23.67
USD/Yen, daily	0.026	0.016	0.018	0.022	0.016
EFG, daily	0.251	0.143	0.120	0.225	0.141

Table 7. Root Mean-Squared (RMSE) of Forecast Errors

Series	Naïve	SQNT	ABNT	Mean	Median
				GARCH	GARCH
S&P500, monthly	0.243	0.206	0.206	0.224	0.232
MSFT, monthly	0.530	1.552	0.951	162.0	89.17
USD/Yen, daily	0.031	0.028	0.028	0.030	0.029
EFG, daily	0.227	0.208	0.194	0.211	0.212

Notes:

1. All forecasts computed using a rolling evaluation sample.
2. The evaluation sample used for computing the entries of the tables is as follows: 148 observations for the monthly S&P500 series, 100 observations for the monthly MSFT series, 986 observations for the daily USD/Yen series and 503 observations for the daily EFG series.
3. Table entries are the values of the evaluation measure (MAD for Table 18 and RMSE for Table 19) multiplied by 100 (S&P500 and MSFT monthly series) and by 1000 (USD/Yen and EFG daily series) respectively.
4. SQNT, ABNT denote NoVaS made forecasts based on square and absolute returns and normal target distribution.
5. Mean and median GARCH forecasts denote forecasts made with a GARCH model and an underlying t error distribution with degrees of freedom estimated from the data.
6. The Naive forecast is based on the rolling sample variance.

**Table 8. Diebold-Mariano Test for Difference in Forecasting Performance
NoVaS and GARCH against the Naive benchmark**

Series	SQNT	ABNT	Mean GARCH	Median GARCH
<u>S&P500, monthly</u>				
Test value	3.369	1.762	1.282	-0.414
p-value	0.000	0.078	0.200	0.679
<u>MSFT, monthly</u>				
Test value	2.931	7.022	-2.671	-2.559
p-value	0.003	0.000	0.007	0.010
<u>USD/Yen, daily</u>				
Test value	0.101	0.083	0.037	0.096
p-value	0.919	0.933	0.971	0.924
<u>EFG, daily</u>				
Test value	1.077	1.301	0.259	1.095
p-value	0.281	0.190	0.795	0.274

Notes:

1. See Tables 17 and 18 for column nomenclature.
2. The entries of Table 19 are the test and p-values for the Diebold-Mariano (1995) test for comparing forecasting accuracy. The tests use the absolute value function for the calculation of the statistic and are expressed relative to the Naive benchmark.
3. Positive values indicate that the competing model is superior, negative values that the Naive benchmark is superior.

Table 9. Forecast Unbiasedness Regressions

Series	Naïve	SQNT	ABNT	Mean GARCH	Median GARCH
<u>S&P500, monthly</u>					
Estimates	(-0.003,1.824)	(0.000,0.317)	(0.000,0.879)	(-0.002,1.685)	(-0.002,3.879)
p-values	(0.597,0.540)	(0.527,0.055)	(0.344,0.000)	(0.000,0.000)	(0.000,0.000)
R^2	0.003	0.025	0.111	0.118	0.177
<u>MSFT, monthly</u>					
Estimates	(-0.025,0.242)	(0.004,-0.859)	(0.004,-0.729)	(0.007,-1.000)	(0.007,-1.000)
p-values	(0.000,0.276)	(0.000,0.000)	(0.000,0.000)	(0.000,0.000)	(0.000,0.000)
R^2	0.012	0.871	0.689	1.000	1.000
<u>USD/Yen, daily</u>					
Estimates	(0.000,-1.099)	(0.000,-0.476)	(0.000,0.355)	(0.000,-0.803)	(0.000,0.642)
p-values	(0.000,0.000)	(0.000,0.000)	(0.000,0.000)	(0.000,0.000)	(0.000,0.000)
R^2	0.188	0.055	0.017	0.136	0.029
<u>EFG, daily</u>					
Estimates	(0.000,-0.767)	(0.000,-0.378)	(0.000,0.058)	(0.000,0.138)	(0.000,0.567)
p-values	(0.017,0.000)	(0.000,0.000)	(0.000,0.518)	(0.038,0.318)	(0.038,0.025)
R^2	0.072	0.062	0.001	0.002	0.002

Notes:

1. See Tables 17 and 18 for column nomenclature.
2. The entries of Table 20 are the coefficient estimates (\hat{a}, \hat{b}) (first line), corresponding p-values (second line) and R^2 (third line) from the forecast unbiasedness regression $e_t = a + b\hat{\sigma}_t^2 + \zeta_t$.
3. Under the hypothesis of forecast unbiasedness we must have $a = b = 0$ and $R^2 \rightarrow 0$. For any two competing models A and B for which we have that $R_A^2 < R_B^2$ we say that model A is superior to model B .

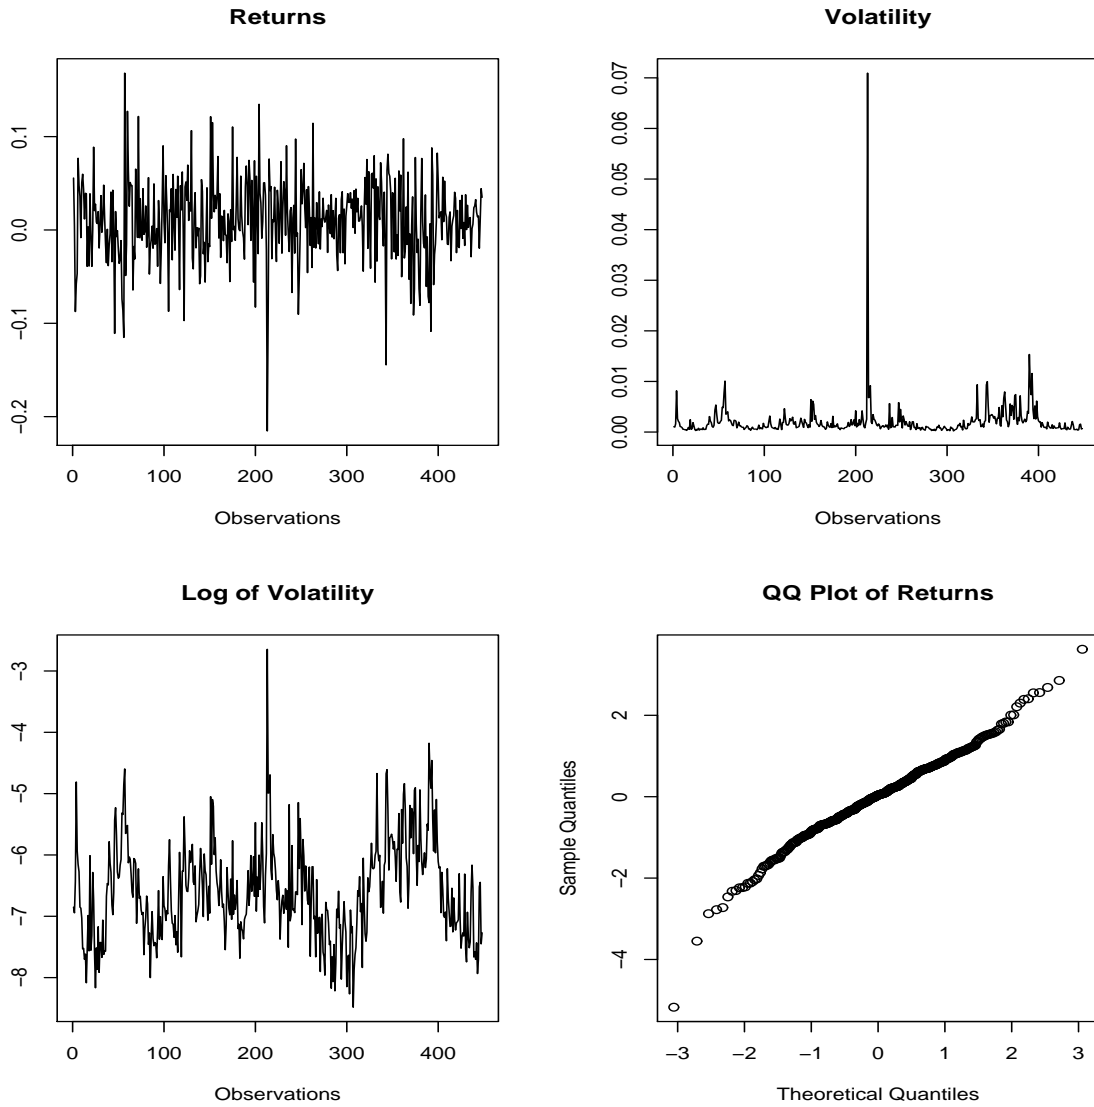


Figure 1: Return, volatility and QQ plots for the monthly S&P500 series

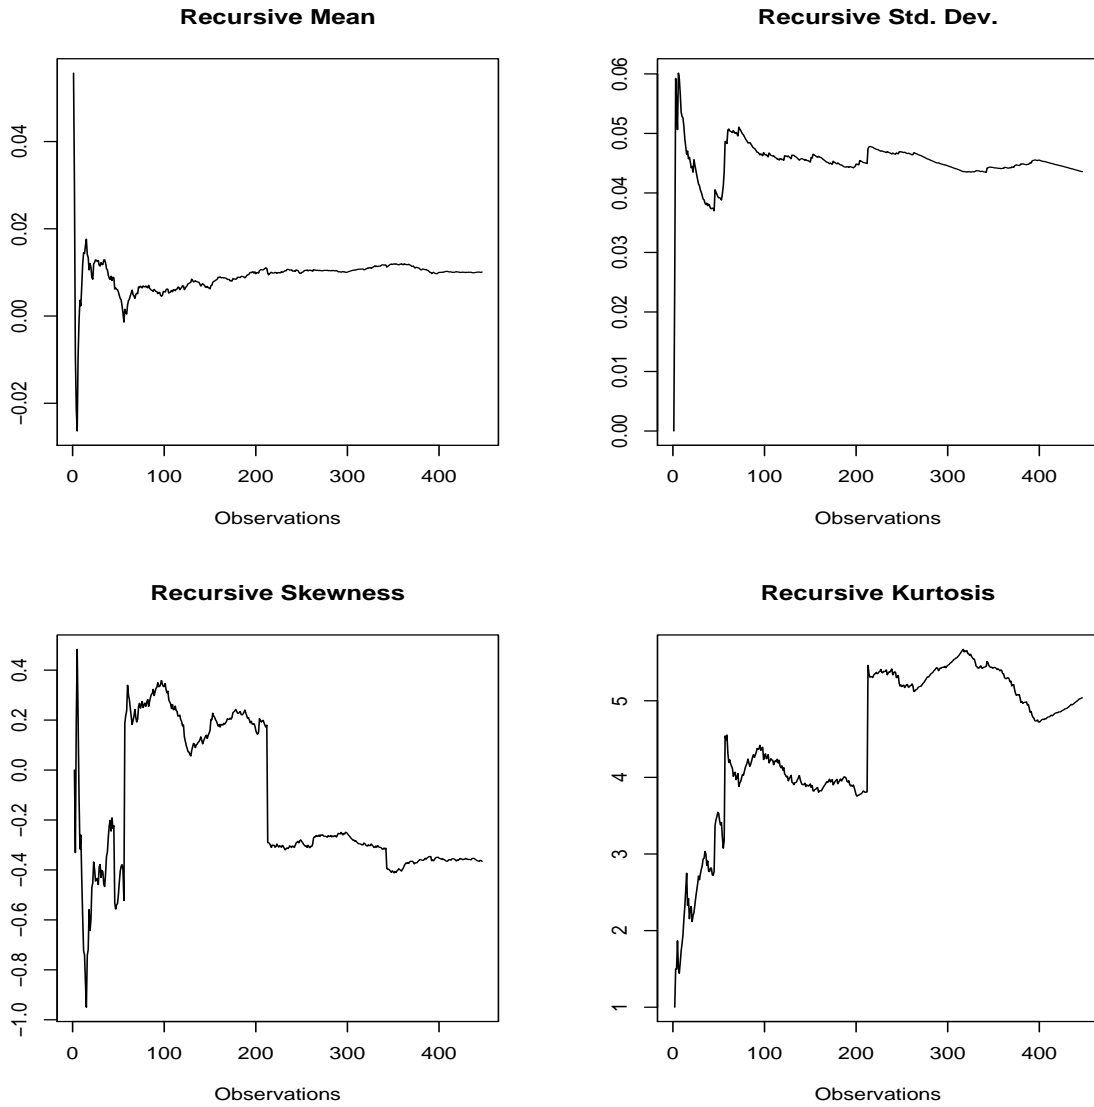


Figure 2: Recursive moments for the monthly S&P500 series

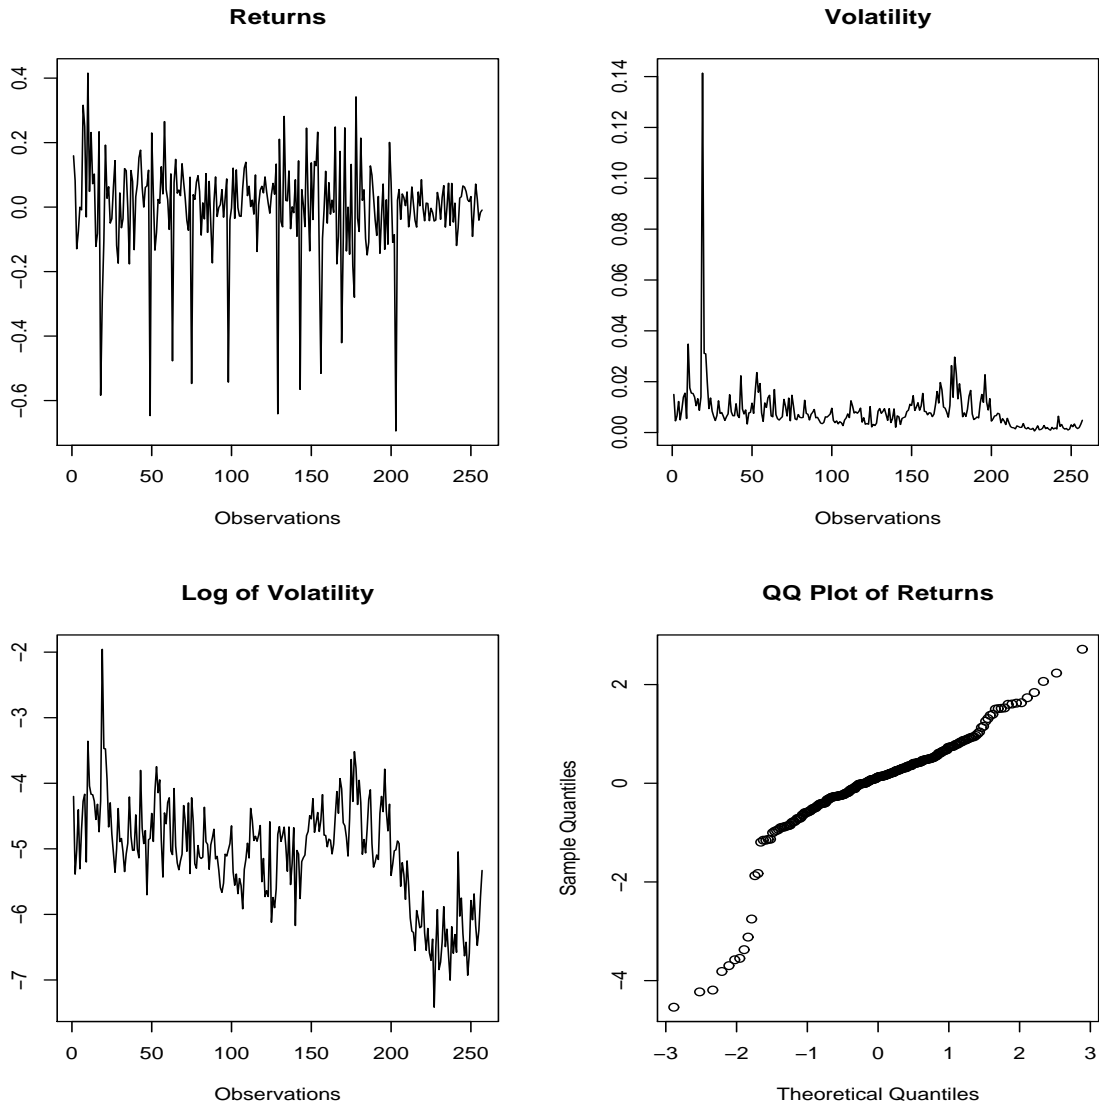


Figure 3: Return, volatility and QQ plots for the monthly MSFT series

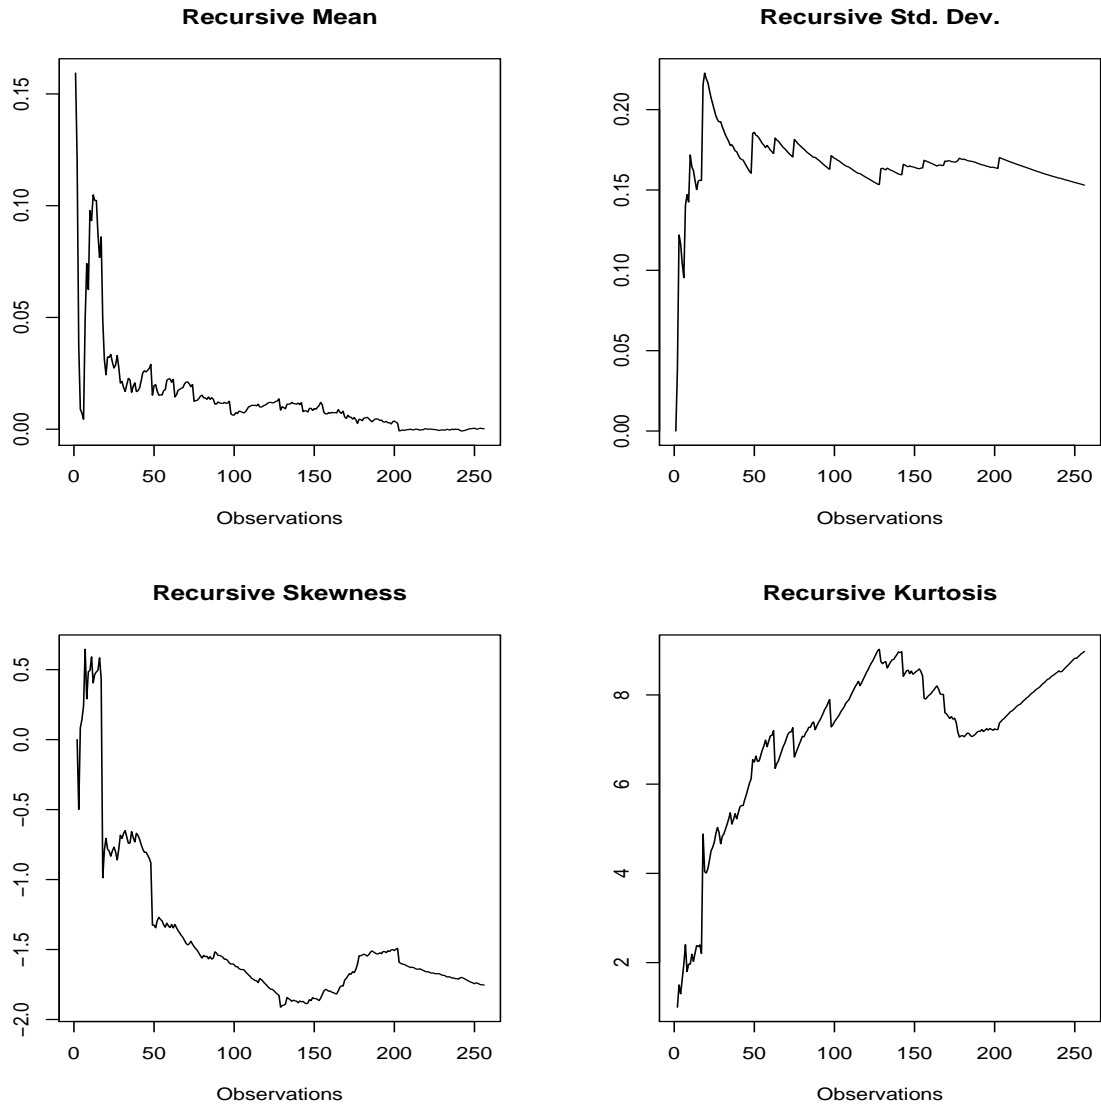


Figure 4: Recursive moments for the monthly MSFT series

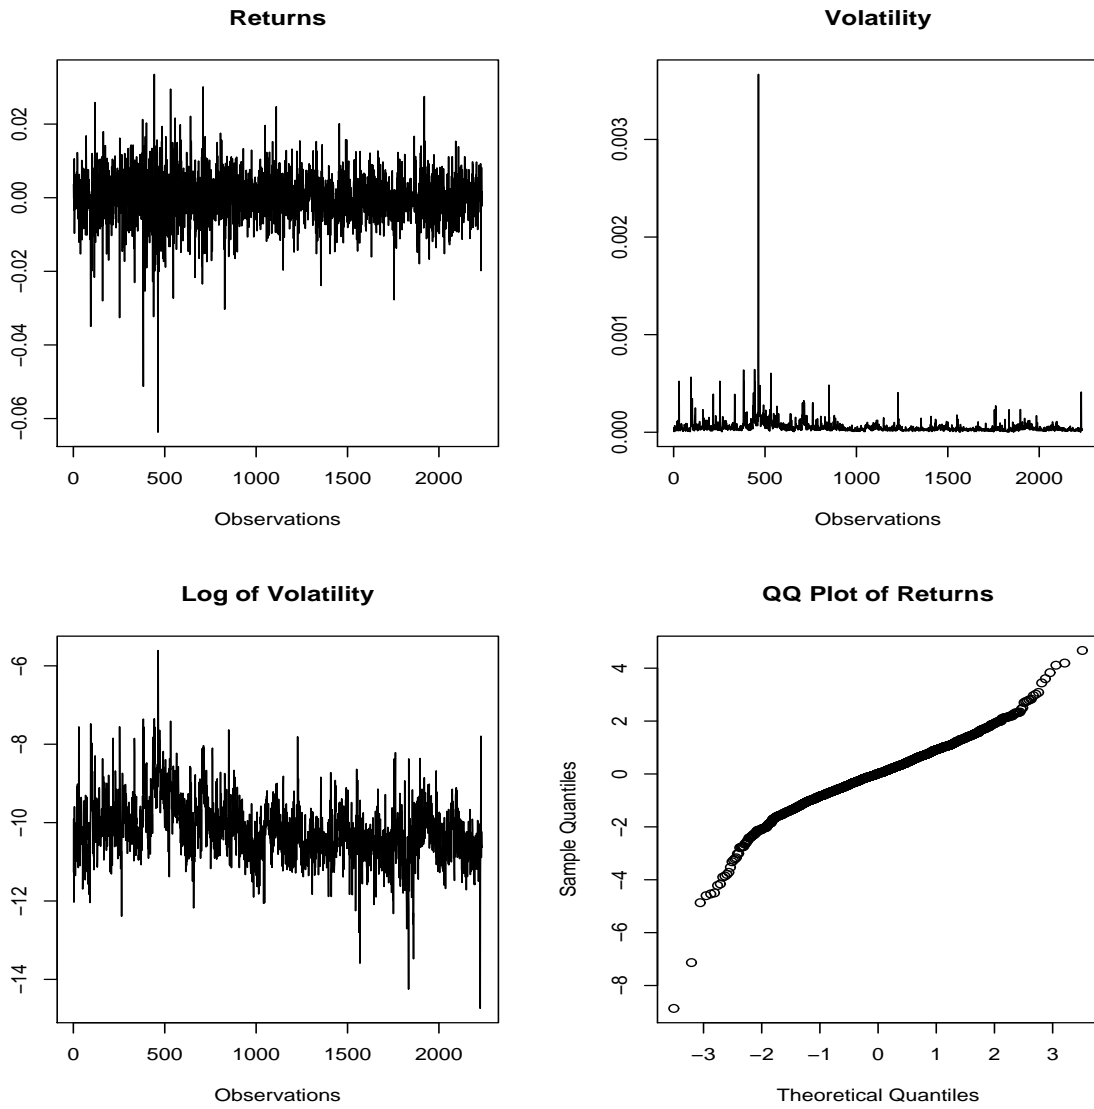


Figure 5: Return, volatility and QQ plots for the daily USD/Yen series

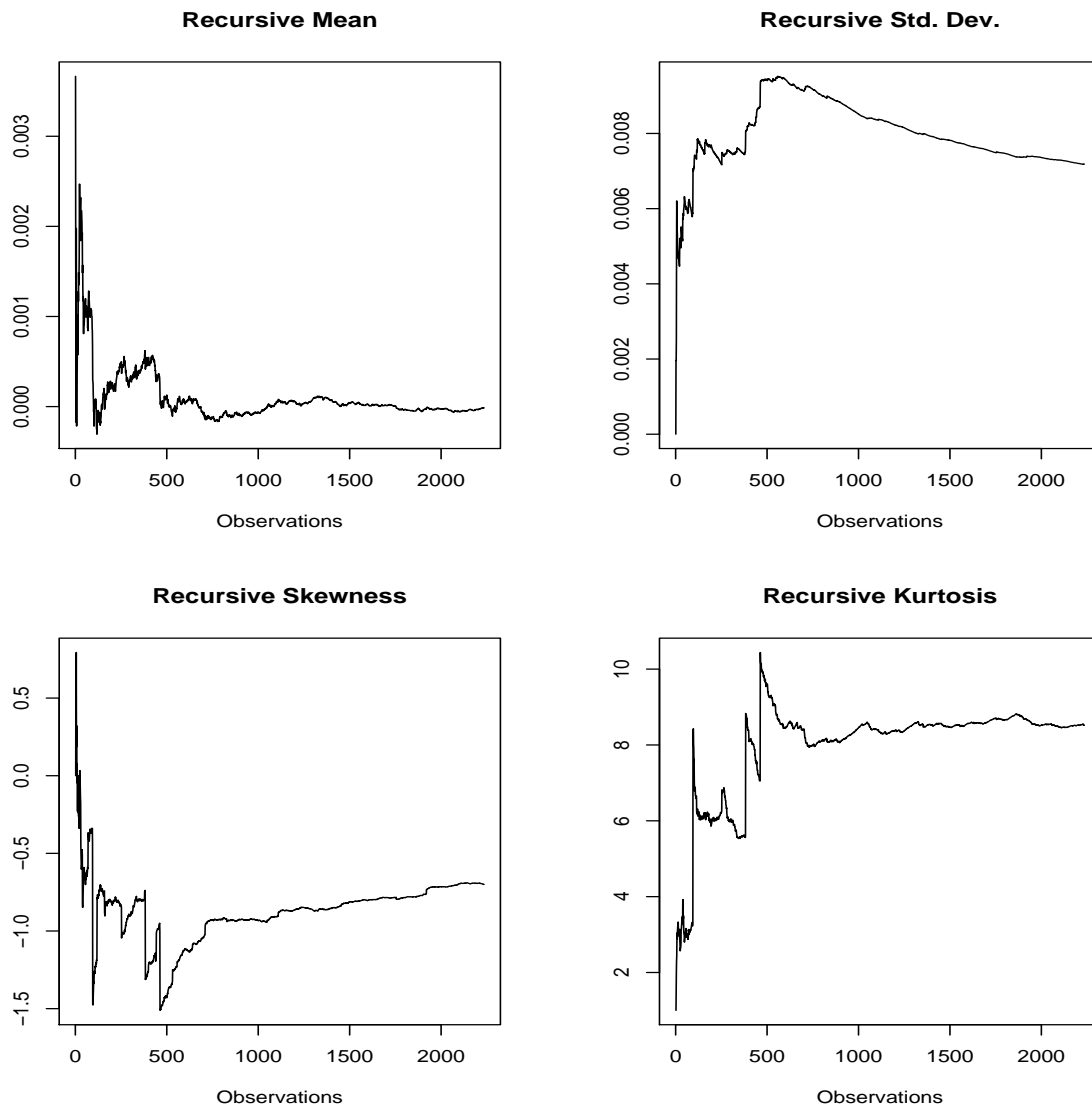


Figure 6: Recursive moments for the daily USD/Yen series

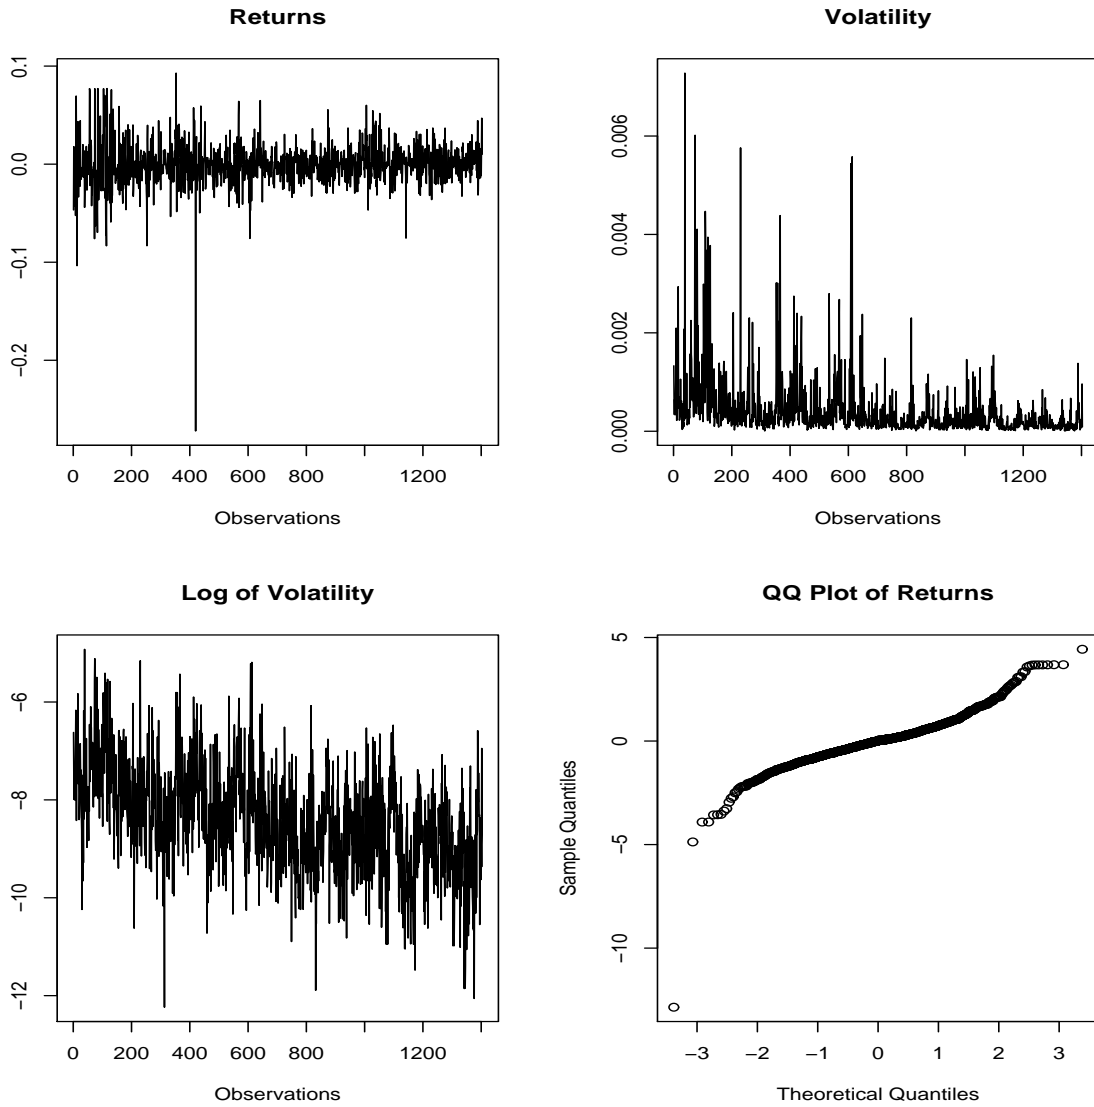


Figure 7: Return, volatility and QQ plots for the daily EFG series

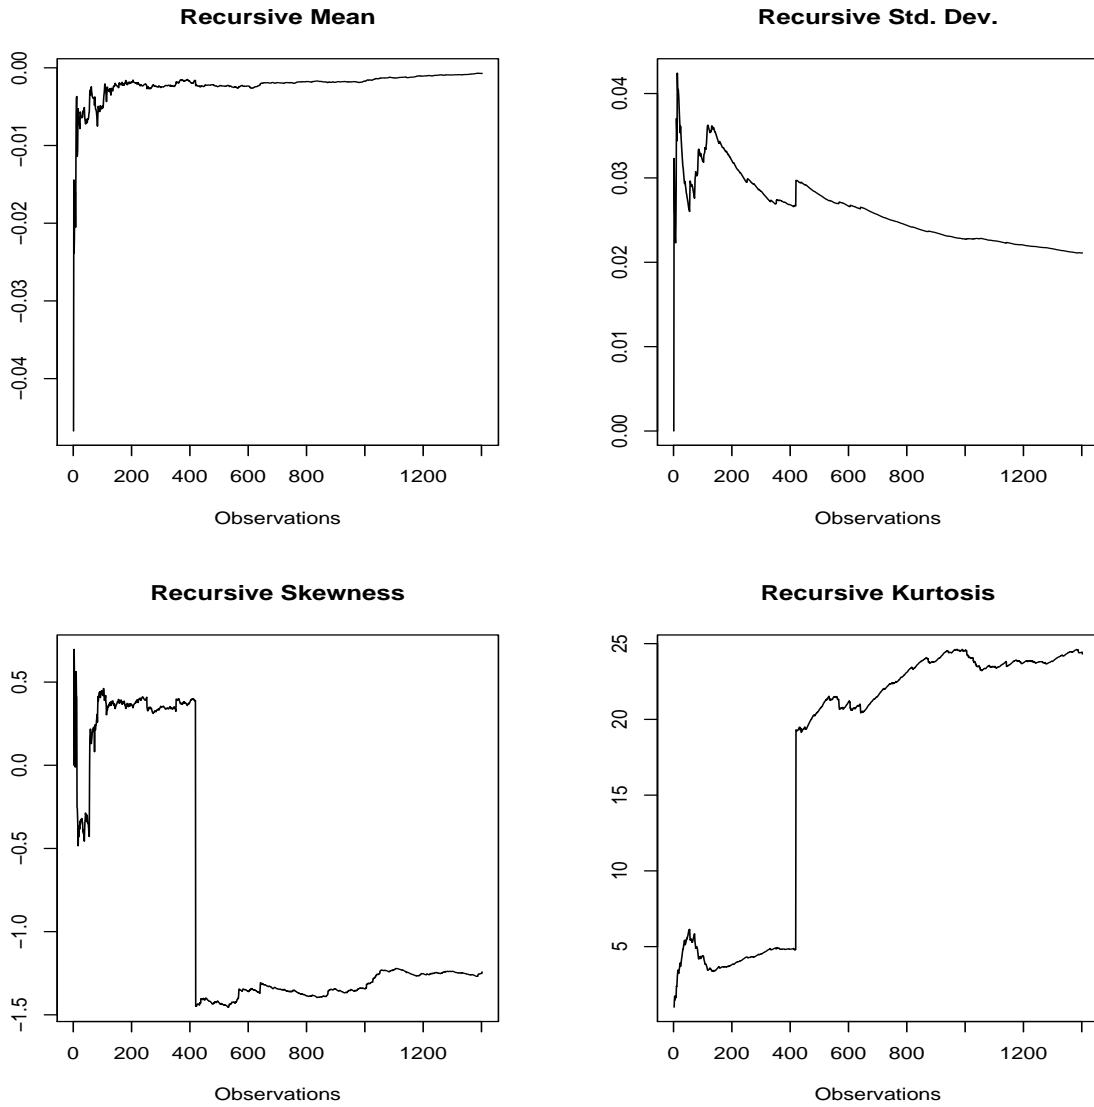


Figure 8: Recursive moments for the daily EFG series

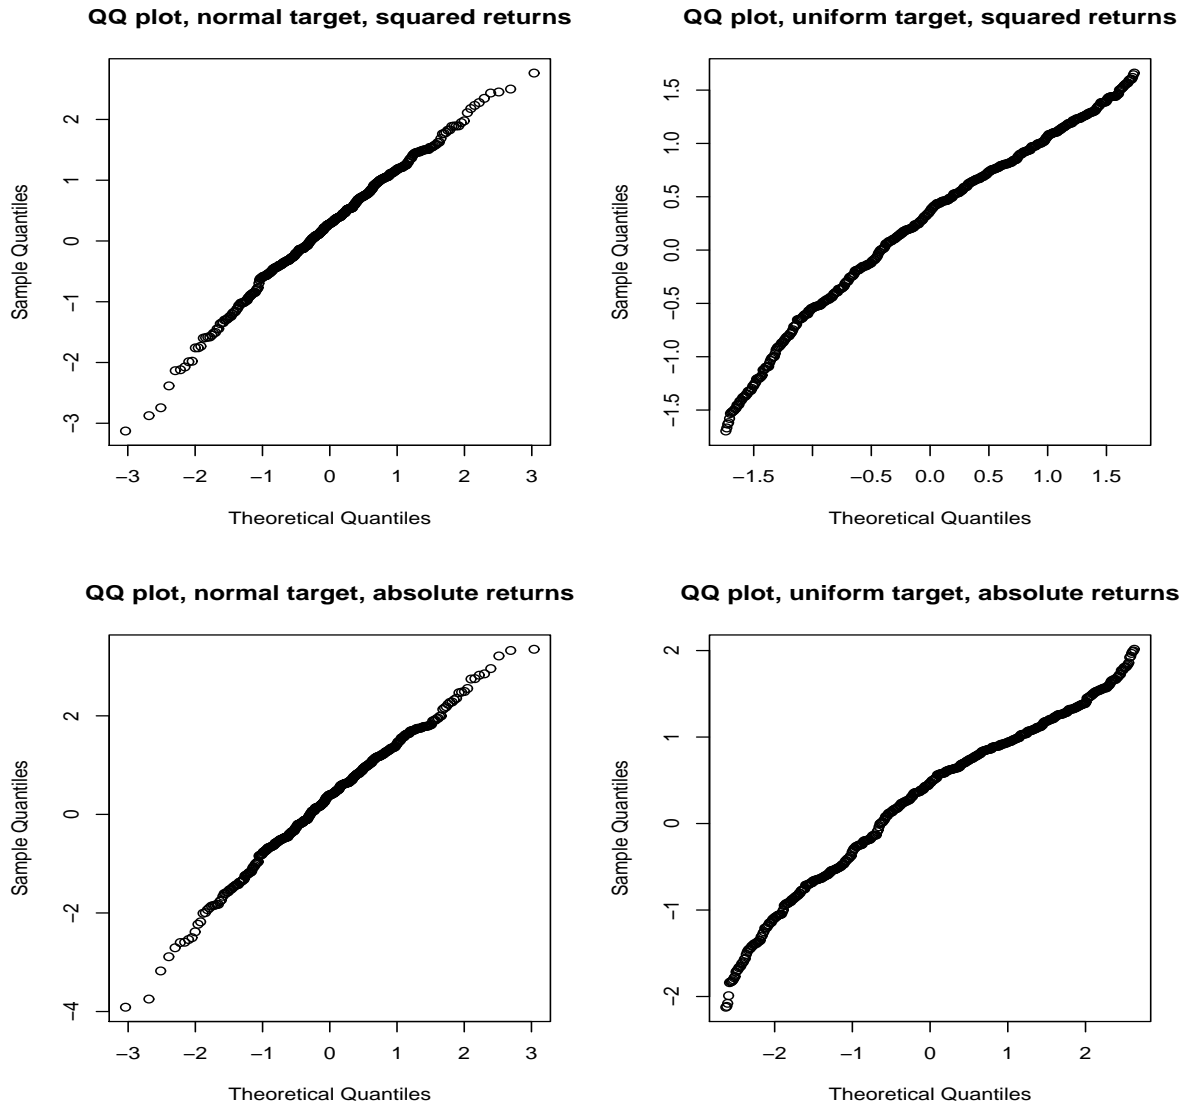


Figure 9: QQ plots of the NoVaS -transformed W series for the monthly S&P500 series

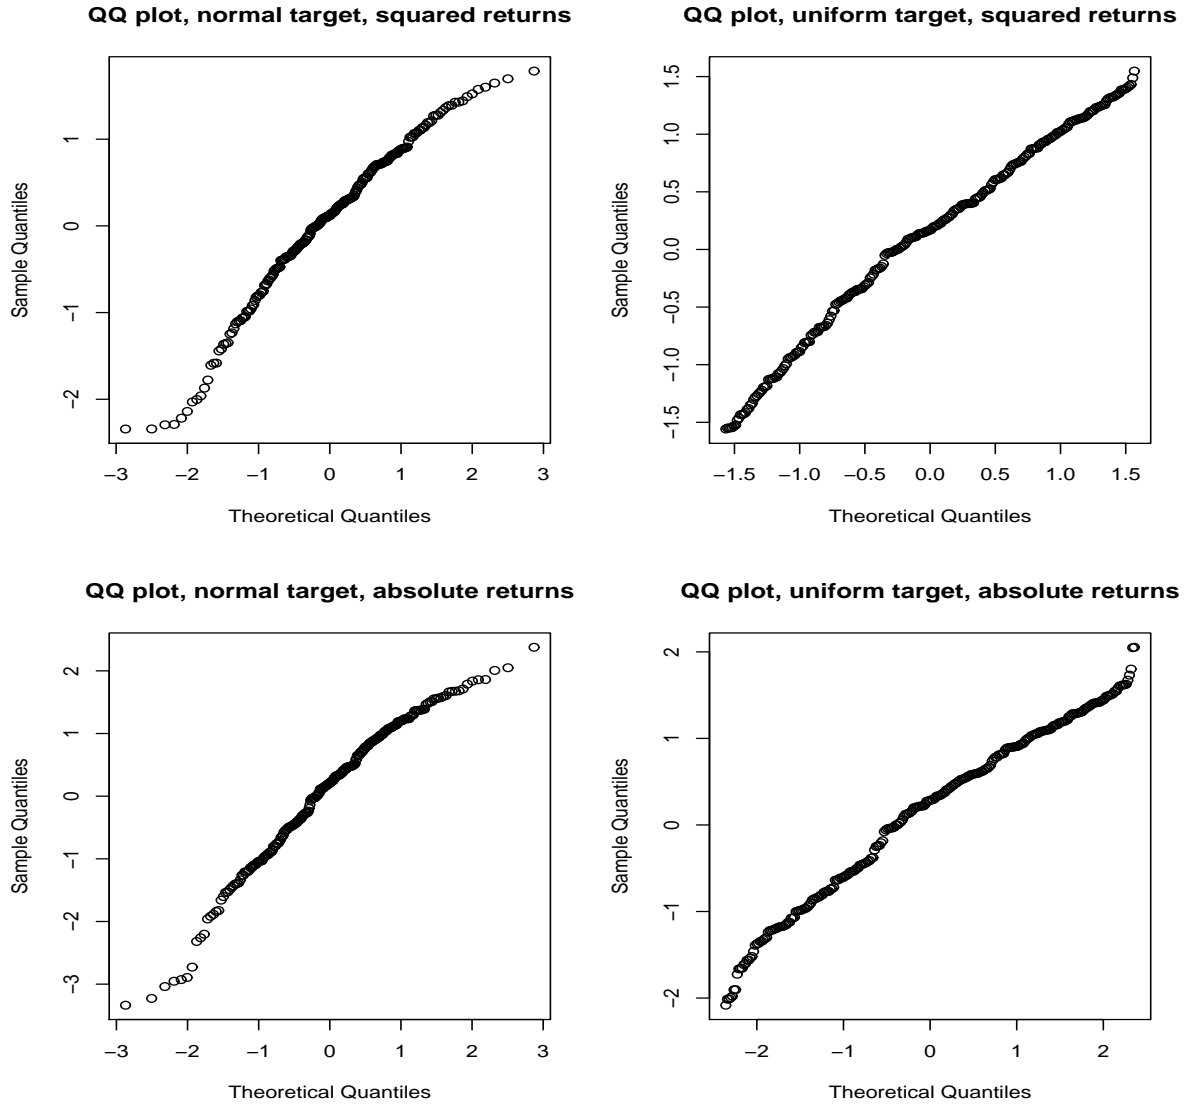


Figure 10: QQ plots of the NoVaS -transformed W series for the monthly MSFT series

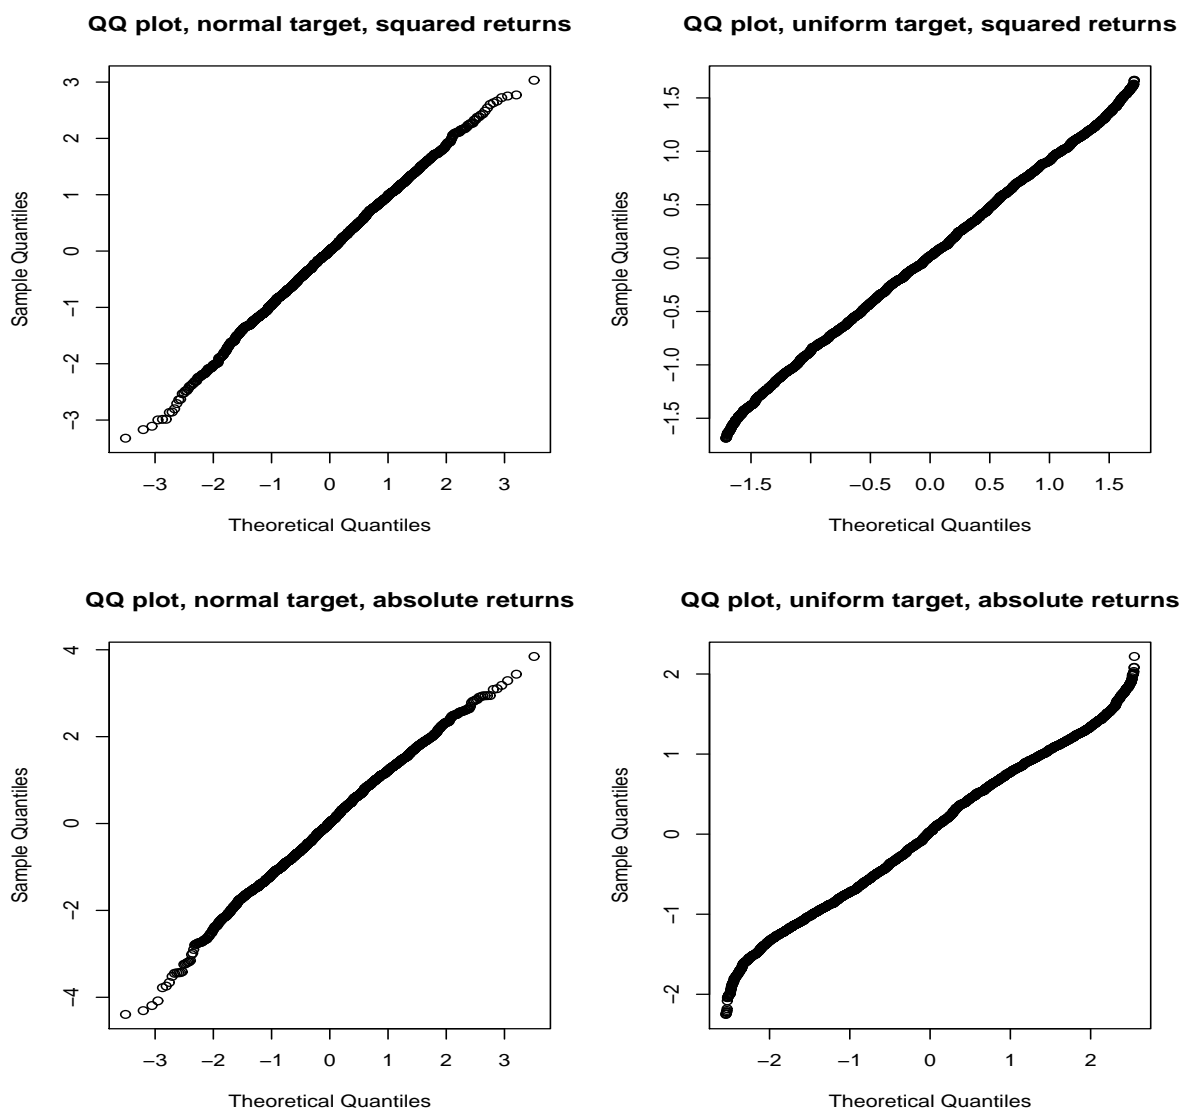


Figure 11: QQ plots of the NoVaS -transformed W series for the daily USD/Yen series

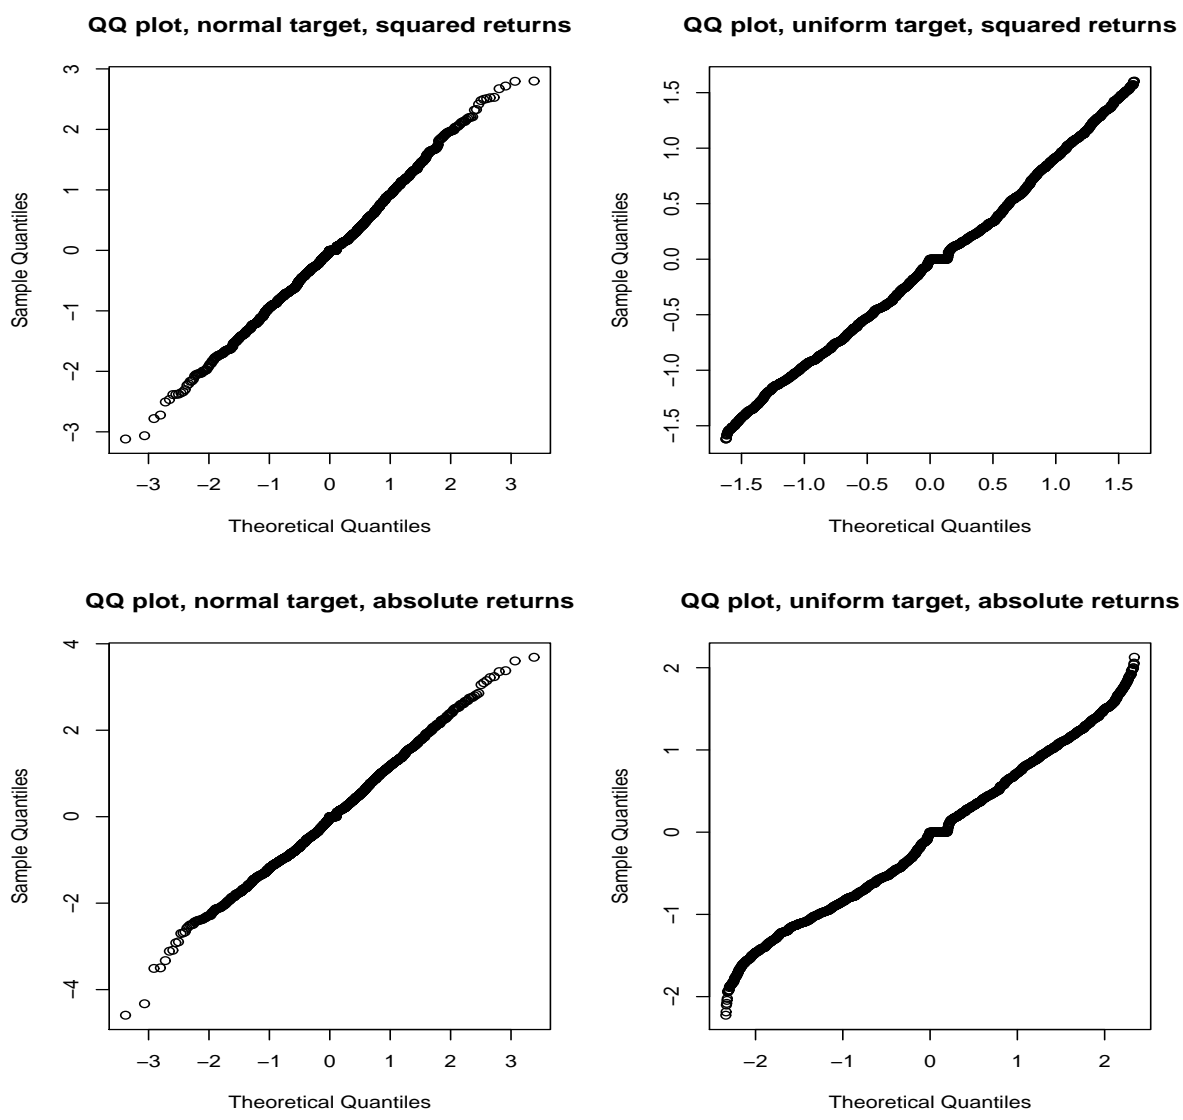


Figure 12: QQ plots of the NoVaS -transformed W series for the daily EFG series