

REALISED VARIANCE FORECASTING UNDER BOX-COX TRANSFORMATIONS

Nick Taylor

Accounting and Finance
Discussion Paper 16 / 4

10 June 2016



School of Economics, Finance and Management
University of Bristol
8 Woodland Road
Bristol BS8 1TN
United Kingdom

Realised variance forecasting under Box-Cox transformations

Nick Taylor^a

^a*School of Economics, Finance and Management, University of Bristol, Bristol, BS8 1TN, UK.*

Abstract

The benefits associated with modeling Box-Cox transformed realised variance data are assessed. In particular, the quality of realised variance forecasts with and without this transformation applied are examined in an out-of-sample forecasting competition. Using various realised variance measures, data transformations, volatility models and assessment methods, and controlling for data mining issues, the results indicate that data transformations can be economically and statistically significant. Moreover, the quartic transformation appears to be the most effective in this regard. The conditions under which the effectiveness of using transformed data varies are identified.

Keywords: Realised variance, forecasting competitions, Box-Cox transformation, loss function, reality check.

JEL: C22, C53, C58, G17.

Email address: `nick.taylor@bristol.ac.uk` (Nick Taylor)

1. Introduction

Power transformations, and more generally Box-Cox (BC) transformations, have long been recognised as an effective way of achieving well specified models with symmetric errors and stable error variance; see Tukey (1957) and Box and Cox (1964). More recently, attention has focused on assessing the out-of-sample performance of time-series models applied to BC transformed data. For instance, Bårdsena and Lütkepohl (2011), Lütkepohl and Xu (2012), Proietti and Lütkepohl (2013) and Mayr and Ulbricht (2015) demonstrate that out-of-sample forecasts based on models applied to BC transformed macroeconomic series can be more accurate than those based on using the raw series (cf. Nelson and Granger, 1979). Inspired by these results we consider whether BC transformations are useful within the context of forecasting future realised variances.

The use of transformations in the context of realised variance is not new. Indeed, applications of models to log transformed realised variance is common practise; see, e.g., Andersen et al. (2003), Andersen et al. (2007), Corsi (2009), Hansen et al. (2012) and Koopman and Scharth (2013).¹ More recently, BC transformations have been considered in this context; see, e.g., Gonçalves and Meddahi (2011), Weigand (2014), and Nugroho and Morimoto (2016).² We add to this literature by examining the out-of-sample performance of a variety of contemporary models applied to BC transformed realised variance in comparison to those applied to raw realised variance. In contrast to previous studies we focus is on whether the BC transformation itself is beneficial to volatility forecasters.

The question of whether to transform the raw realised variance series will depend on the loss function used to assess forecasting performance. We follow the extant literature and consider the mean square (MS) and quasi-likelihood (QLIK) error loss functions applied to the raw realised variance series. These belong to the Bregman loss function family, and are thus consistent in the sense that the conditional mean is the optimal forecast (Banerjee et al., 2005, Gneiting, 2011, and Patton, 2015).³ Under these loss functions the optimal forecast is obtained by minimising any

¹The use of log transformed realised variance is based on previous findings that show that these data are Gaussian distributed; see, e.g., Andersen et al. (2001a, 2001b).

²Transformations are not always applied. For instance, Bollerslev et al. (2016) augment the popular long-memory heterogenous autoregressive (HAR) model of Corsi (2009), but decide not to apply the log transformation as in the original HAR model.

³The Bregman loss function family also possess the quality that for correctly specified models with nested information sets, the ‘*ranking of these forecasts by MSE is sufficient for their ranking by any Bregman loss function*’

Bregman loss function. If we consider a model with parameters estimated by minimising the MS loss function, then it follows that this forecast is optimal. However, the models themselves may not be ‘suited’ to the raw (possibly highly non-Gaussian) data. It is possible that models applied to transformed data may be superior because they more closely match the true data generating process. Consequently, although the parameters are not optimised with respect to the loss function, their forecasts may be superior because of the suitability of the model to the data. It is this trade-off (parameter optimisation versus model suitability) that we are examining in the current paper by considering whether or not to transform realised variance data.

Using a comprehensive set of ten S&P 500 index realised variance measures, we examine whether the use of BC transformations has value to forecasters. The results indicate that such transformations can improve forecasts of future realised variance across a range of models and under both the MS and QLIK loss functions. Moreover, quality differences between forecasts based on modeling raw and transformed data can be significant even when one controls for data mining by using the reality check statistical tests proposed by White (2000). Of the BC transformations that we consider it is the quartic root (and not the log) transformation that delivers the best results. Finally, we demonstrate that the benefits of BC transformation are not evenly spread over the realised variance measures. Indeed, for some measures no benefits are found – a result that we demonstrate is driven by the degree of skewness in the raw realised variance series.

The rest of the paper is organised as follows. The next section contains a description of the methodologies employed and is followed by the empirical results. The final section concludes.

2. Methodologies

This section contains the models and methods used to construct forecasts, and the means by which the relative quality of the forecasts is assessed.

2.1. Forecast construction: The problem

Let x_t be the raw data that we wish to forecast, in our case realised variance, $x_t > 0$ and $t = 1, 2, \dots, T$. As x_t is likely to be highly non-Gaussian we model the BC transformed data given

(Patton, 2015).

by

$$y_t = f(x_t; \lambda) = \begin{cases} \frac{x_t^\lambda - 1}{\lambda}, & \lambda \neq 0, \\ \ln x_t, & \lambda = 0. \end{cases} \quad (1)$$

It follows that $x_t = g(y_t; \lambda) = f^{-1}(y_t; \lambda)$.⁴ Suppose the forecaster models y_t and obtains h -step ahead forecasts given by the conditional mean of y_{t+h} , that is, $E[y_{t+h}|\mathcal{F}_t]$, where \mathcal{F}_t is the forecaster's information set. Moreover, suppose that we require the conditional mean of x_{t+h} , that is, $E[x_{t+h}|\mathcal{F}_t]$, or equivalently $E[g(y_{t+h}; \lambda)|\mathcal{F}_t]$.

2.2. Forecast construction: The solution(s)

One obvious solution would be to take $g(E[y_{t+h}|\mathcal{F}_t]; \lambda)$, henceforth referred to as the *naive adjustment* forecast. However, Jensen's inequality tells us that for convex functions (like g) $g(E[y_{t+h}|\mathcal{F}_t]; \lambda) \leq E[g(y_{t+h}; \lambda)|\mathcal{F}_t]$. Therefore an expression for $E[g(y_{t+h}; \lambda)|\mathcal{F}_t]$ is required. To this end, taking a Taylor series expansion about the conditional mean of y_{t+h} (denoted $\mu_{t+h|t}$) and taking conditional expectations we have

$$E[g(y_{t+h}; \lambda)|\mathcal{F}_t] = g(\mu_{t+h|t}; \lambda) \left(1 + \sum_{k=1}^{\infty} g_k(\mu_{t+h|t}; \lambda) \mu_{k,t+h|t} \right), \quad (2)$$

where $\mu_{k,t+h|t}$ is the k th conditional moment of y_{t+h} about its conditional mean, and

$$g_k(\mu_{t+h|t}; \lambda) = \frac{1 - \lambda(k-1)}{k(1 + \lambda\mu_{t+h|t})} g_{k-1}(\mu_{t+h|t}; \lambda),$$

with $g_0 = 1$. This is henceforth referred to as the *full adjustment* forecast. Simplifications of the expression in (2) are possible. One could assume that y_{t+h} has a Gaussian distribution, or one could ignore all moments except the mean and variance. Adjustments based on these assumptions are provided in Table 1, and deliver forecasts henceforth referred to as the *Gaussian adjustment* and *second-order adjustment* forecasts.⁵

Insert Table 1 here

⁴Note that y_t represents raw realised variance when $\lambda = 1$.

⁵Subsets of the results in Table 1 have been derived previously. Granger and Newbold (1976) derive the solution associated with the log transformation under the Gaussian distribution assumption via Hermite polynomial expansions, Pankratz and Dudley (1987) obtain the solution when $\lambda = 1/N$, where N is a positive integer, and Proietti and Riani (2009) derive a more general result associated with all λ values under the Gaussian distribution assumption.

The implementation of the above formulae require further augmentation in order for them to become practical to implement. First, in the subsequent empirical application the full adjustment method employs (2) with a truncation such that only the first ten conditional moments are considered. Extending beyond this point has no effect on the accuracy of forecasts. Second, higher conditional moments (that is, the second conditional moment and higher) are estimated using their unconditional sample counterparts.

2.3. Models

A number of realised variance (and transformations thereof) models are available. In addition to a simple first-order autoregressive (AR) model, we consider the following models.

2.3.1. The HAR model

The first model considered is the popular (and successful) heterogeneous AR (HAR) model proposed by Corsi (2009). This model provides a parsimonious representation of realised variance that attempts to capture the long memory observed in previous studies; see, e.g., Anderson et al. (2001a, 2001b, 2003). Conditional expectations of demeaned transformed (daily frequency) realised variance based on the HAR model take the following form:

$$\mathbb{E}[y_t | \mathcal{F}_{t-1}] = \gamma_1 y_{t-1} + \gamma_2 \sum_{i=1}^5 y_{t-i} + \gamma_3 \sum_{j=1}^{22} y_{t-j}. \quad (3)$$

This can be written as a restricted infinite-order autoregressive (AR(∞)) model such that

$$\mathbb{E}[y_t | \mathcal{F}_{t-1}] = \sum_{i=1}^{\infty} \Pi_i^* y_{t-i}, \quad (4)$$

where

$$\Pi_i^* = \begin{cases} \gamma_1 + \gamma_2 + \gamma_3, & \text{for } i = 1, \\ \gamma_2 + \gamma_3, & \text{for } i = 2, \dots, 5, \\ \gamma_3 & \text{for } i = 6, \dots, 22, \end{cases}$$

and zero otherwise. It follows from the chain-rule of forecasting that h -step forecasts are given by

$$\mathbb{E}[y_{t+h} | \mathcal{F}_t] = \sum_{i=1}^{\infty} \Pi_i^* \mathbb{E}[y_{t+h-i} | \mathcal{F}_t], \quad (5)$$

where $E[y_{t+j}|\mathcal{F}_t] = y_{t+j}$ for $j \leq 0$.⁶

2.3.2. The HEAVY model

A number of competitors to the HAR model exist. To incorporate improved measures of realised variance into a conditional GARCH-type specification, a first-order version of the high frequency based volatility (HEAVY) model proposed by Shephard and Sheppard (2010) takes the form

$$E[y_t|\mathcal{F}_{t-1}] = \alpha_1 y_{t-1} + \beta_1 E[y_{t-1}|\mathcal{F}_{t-2}], \quad (6a)$$

$$E[z_t|\mathcal{F}_{t-1}] = \alpha_2 y_{t-1} + \beta_2 E[z_{t-1}|\mathcal{F}_{t-2}]. \quad (6b)$$

Here z_t represents an alternative BC transformed measure of realised variance (given by the BC transformed squared daily return). The specification can be written more compactly as

$$E[\mathbf{y}_t|\mathcal{F}_{t-1}] = \mathbf{A}\mathbf{y}_{t-1} + \mathbf{B}E[\mathbf{y}_{t-1}|\mathcal{F}_{t-2}], \quad (7)$$

where

$$\mathbf{y}_t = \begin{bmatrix} y_t \\ z_t \end{bmatrix}, \quad \mathbf{A} = \begin{bmatrix} \alpha_1 & 0 \\ \alpha_2 & 0 \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} \beta_1 & 0 \\ 0 & \beta_2 \end{bmatrix}.$$

This equation can also be written in VARMA(1,1) form,

$$E[\mathbf{y}_t|\mathcal{F}_{t-1}] = (\mathbf{A} + \mathbf{B})\mathbf{y}_{t-1} - \mathbf{B}\boldsymbol{\epsilon}_{t-1}, \quad (8)$$

where $\boldsymbol{\epsilon}_t = \mathbf{y}_t - E[\mathbf{y}_t|\mathcal{F}_{t-1}]$, or as a restricted VAR(∞) process such that

$$E[\mathbf{y}_t|\mathcal{F}_{t-1}] = \sum_{i=1}^{\infty} \boldsymbol{\Pi}_i^* \mathbf{y}_{t-i}, \quad (9)$$

where $\boldsymbol{\Pi}_i^* = \mathbf{A}^i \mathbf{B}^{i-1}$. It also follows that h -step forecasts are given by

$$E[\mathbf{y}_{t+h}|\mathcal{F}_t] = \sum_{i=1}^{\infty} \boldsymbol{\Pi}_i^* E[\mathbf{y}_{t+h-i}|\mathcal{F}_t], \quad (10)$$

⁶Audrino and Knaus (2016) demonstrate that the restrictive specification of the HAR model is not inferior to an unrestricted AR model in which the specification is determined via the *lasso* selection criterion.

where $E[\mathbf{y}_{t+j}|\mathcal{F}_t] = \mathbf{y}_{t+j}$ for $j \leq 0$. Our interest is in the first element of this vector-valued conditional expectation (that is, $E[y_{t+h}|\mathcal{F}_t]$). An integrated version of this model is possible by imposing the restriction that $\beta_1 = 1 - \alpha_1$. This version is henceforth referred to as the IHEAVY model.

2.3.3. The RealGARCH model

As an alternative to the above model, a simplified version of the first-order RealGARCH model (ignoring leverage effects) proposed by Hansen et al. (2012) can be written as

$$E[y_t|\mathcal{F}_{t-1}] = \delta E[z_t|\mathcal{F}_{t-1}], \quad (11a)$$

$$E[z_t|\mathcal{F}_{t-1}] = \alpha_2 y_{t-1} + \beta_2 E[z_{t-1}|\mathcal{F}_{t-2}]. \quad (11b)$$

Substituting (11b) into (11a) and rearranging we obtain

$$E[y_t|\mathcal{F}_{t-1}] = \delta \alpha_2 y_{t-1} + \beta_2 E[y_{t-1}|\mathcal{F}_{t-2}], \quad (12a)$$

$$E[z_t|\mathcal{F}_{t-1}] = \alpha_2 y_{t-1} + \beta_2 E[z_{t-1}|\mathcal{F}_{t-2}]. \quad (12b)$$

This can also be written in the matrix form given by (7), but now

$$\mathbf{y}_t = \begin{bmatrix} y_t \\ z_t \end{bmatrix}, \quad \mathbf{A} = \begin{bmatrix} \delta \alpha_2 & 0 \\ \alpha_2 & 0 \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} \beta_2 & 0 \\ 0 & \beta_2 \end{bmatrix}.$$

It necessarily follows that this can be expressed as the VARMA(1,1) process in (8) or the restricted VAR(∞) process as in (9). Comparing these versions of the HEAVY and RealGARCH models we see that the latter is actually a restricted version of the former (with the restrictions involving β_1 and β_2). An integrated version of this model is possible by imposing the restriction that $\beta_2 = 1 - \delta \alpha_2$. This version is henceforth referred to as the IRealGARCH model.

2.4. Estimation details

The parameters associated with the above models are estimated such that the sum of squared errors is minimised.⁷ An increasing window (IW) and a rolling window (RW) of past data are used to generate h -step ahead out-of-sample forecasts (that is, conditional means denoted $\hat{x}_{t+h|t}$). The realised variance forecasts based on models of raw data are compared to the realised variance forecasts based on transformed data (with the formulae in Table 1 used to convert the transformed forecasts to realised variance forecasts).

2.5. Performance assessment: The loss function

The performance of the various forecasting methods is assessed using the following homogenous Bregman loss function family proposed in the context of realised variance forecasting by Patton (2011):

$$L(x_{t+h}, \hat{x}_{t+h|t}; b) = \begin{cases} \frac{1}{(b+1)(b+2)}(x_{t+h}^{b+2} - \hat{x}_{t+h|t}^{b+2}) - \frac{1}{b+1}\hat{x}_{t+h|t}^{b+1}(x_{t+h} - \hat{x}_{t+h|t}), & \text{for } b \notin \{-1, -2\}, \\ \hat{x}_{t+h|t} - x_{t+h} + x_{t+h} \ln\left(\frac{x_{t+h}}{\hat{x}_{t+h|t}}\right), & \text{for } b = -1, \\ \frac{x_{t+h}}{\hat{x}_{t+h|t}} - \ln\left(\frac{x_{t+h}}{\hat{x}_{t+h|t}}\right) - 1, & \text{for } b = -2. \end{cases} \quad (13)$$

Here $b = 0$ corresponds to MS loss, and $b = -2$ corresponds to QLIK loss. Under this Bregman loss function family, Patton (2015) demonstrates that the performance rank of a forecasting method can vary over b in the presence of misspecified models, parameter estimation error, or non-nested information sets. Thus we consider performance under both the MS and QLIK loss functions.

2.6. Performance assessment: Hypothesis tests

The null hypothesis is that models based on transformed data have no superior predictive ability over those based on non-transformed data. The alternative hypothesis is that the former do have superior ability. We use differences in the means of the above loss function values to test this null.

⁷For the HAR model ordinary least squares (OLS) is used. For the HEAVY, IHEAVY, RealGARCH and IRealGARCH models the parameters are estimated using the Constrained Optimisation (CO) package in GAUSS v.11. CO solutions are obtained using the Newton-Raphson and cubic/quadratic step length methods.

Formally, we use the reality check approach to the null given by

$$H_0 : \max_{k=1,\dots,K} \mathbb{E}[L_{0,t+h} - L_{k,t+h}] \leq 0, \quad (14a)$$

against the alternative

$$H_1 : \max_{k=1,\dots,K} \mathbb{E}[L_{0,t+h} - L_{k,t+h}] > 0, \quad (14b)$$

where $L_{0,t+h}$ is the forecast loss associated with the benchmark model, and $L_{k,t+h}$ is the forecast loss associated with the k th competitor model.

In our case we examine the predictive performance of the set of models based on BC transformed data against a benchmark model based on raw data. The null hypothesis is that no model based on transformed data outperforms the benchmark model based on raw data (for a given loss function), against the alternative that at least one model based on transformed data outperforms the benchmark model. To avoid data snooping bias we use the block bootstrap procedure proposed by White (2000) to test the above hypothesis.

3. Results

This section contains the empirical results associated with the relative performance of models based on raw and transformed realised variance data.

3.1. Data

We consider ten realised variance measures associated with the S&P 500 index. The first four measures (denoted RV1, RV2, RV3 and RV4) are based on the following estimators: two standard realised variance estimators based on 5 and 10-minute frequency returns (Andersen et al., 2001a, and Barndorff-Nielsen and Shephard, 2002); the jump-robust bipower variation estimator based on 5-minute frequency returns (Barndorff-Nielsen and Shephard, 2004); and the downside risk semivariance estimator based on 5-minute frequency returns (Barndorff-Nielsen et al., 2010). In addition four multiscale versions of these four estimators are considered (denoted RV5, RV6, RV7 and RV8) in which a 1-minute subsample is used (Zhang et al., 2005). Finally, we consider the microstructure noise-robust realised kernel estimator (Barndorff-Nielsen et al. 2008), and the jump-robust median truncated realised variance estimator (Andersen et al., 2012). These are denoted

RV9 and RV10. All series were collected from the Oxford-Man Institute of Quantitative Finance Realized Library (<http://realized.oxford-man.ox.ac.uk/data>). The data span the period from January 1, 2000, to December 31, 2015.

3.2. Preliminary analysis

The estimated coefficients and standard errors associated with the HAR model are provided in Table 2. All results in this table are based on the full sample of data. Panel A provides these results when realised variance is used. The other panels contain results when transformed realised variance is used. In particular, we consider square root transformed data (panel B), quartic root transformed data (panel C), and log transformed data (panel D). In addition, the R^2 statistics are provided. Two sets of these statistics are provided. The first correspond to those observed in the regression. The second set are calculated by first transforming the fitted values into raw data form and then calculating the R^2 values based on these raw fitted values. In doing this we are able to compare measures of fit over the different data used.

Insert Table 2 here

The results indicate that the three coefficients in the HAR model applied to realised variance are significant – a result that highlights the persistent nature of these data. Moreover, the model provides a good fit to the data. For instance, using the realised variance measure based on 5-minute frequency returns (RV1), the R^2 statistic equals 53.816%. When BC transformed data are used the fit increases dramatically indicating improved suitability to these data. For instance, the corresponding R^2 statistic equals 69.030% when log transformed data are used. Transforming the fitted log transformed data back to fitted realised variance values using the naive adjustment method delivers a R^2 statistic of 51.565%. This value rises to 54.904% when the full adjustment method is used. Importantly this value exceeds that observed when realised variance is modeled. Thus use of the HAR model with BC transformed data delivers a superior representation of the data (both in raw and transformed form). Similar results hold for the other realised variance measures and BC transformations.

3.3. Out-of-sample performance

The previous analysis demonstrates that BC transformations have virtue in an in-sample estimation setting. However, the true test of the approach is within an out-of-sample context. To this end, the AR, HAR, HEAVY, IHEAVY, RealGARCH and IRealGARCH models are estimated using non-transformed and BC transformed realised variance data. The initial estimation window is from January 1, 2000 to December 31, 2005. The estimation window is increased by one observation until the December 31, 2015 observation is used. Increasing and rolling windows of data are used (that is, IW and RW estimation), with out-of-sample 1 to 5-step ahead forecasts generated at each point.⁸ The MS loss associated with each model relative to (divided by) that associated with the HAR model (using IW estimation) applied to non-transformed data are provided in Table 3. In addition to daily horizon forecasts, we also consider weekly horizon forecasts based on integrated 1 to 5-step ahead forecasts. Extant results are provided in Table 4.

Insert Tables 3 and 4 here

The results in Table 3 indicate that there is variation in the performance of models applied to the raw data, with the AR model delivering the least accurate forecasts (entries above unity) and the IHEAVY and IRealGARCH models delivering the most accurate forecasts (entries below unity). For instance, applying these models to the realised variance measure based on 5-minute frequency returns (RV1) and using IW estimation delivers relative MS losses of 1.197, 0.962 and 0.952, respectively. Comparing the results associated with the IW and RW estimation methods we see that the former is universally superior. This relative ranking of the models does not vary considerably over the realised variance measures and loss functions; however, absolute performance does appear to vary over this space.

The underlying question is whether it is better to model raw or transformed data in order to deliver improved forecasts of the raw data. Applying the models to the BC transformed data delivers more accurate forecasts than those based on raw data, with the quartic and log transformations appearing best.⁹ For instance, the log transformation version of the HAR, HEAVY and RealGARCH models (using IW estimation) delivers relative MS losses associated with RV1

⁸We follow Corsi (2009) and adopt a 1000-day rolling window.

⁹Given the superiority of the full adjustment method we now focus exclusively on this method.

of 0.882, 0.877 and 0.890. It is also noticeable that the relative performance of the models is now more uniform. By contrast, there is still considerable variation in the degree of benefit over the realised variance measures.

Similar results are observed in Table 4: models based on BC transformed data perform better than those based on raw data. However, relative performance (with respect to the HAR model using IW estimation applied to raw data) is even more acute. For instance, the HAR, HEAVY and RealGARCH models based on log transformed data (using IW estimation) now deliver relative MS losses associated with RV1 of 0.781, 0.788 and 0.794. Thus modeling log transformed data delivers a meaningful improvement in forecasting performance. While the extent of this improvement does depend of the realised variance measure used, it is never detrimental to model BC transformed data.

3.4. Hypothesis testing

The effects noted in Tables 3 and 4 are large in magnitude and hence appear economically significant. However, we can go further and subject these to a statistical test. The results in Tables 5 and 6 contain the p -values associated with the bootstrap reality check tests described in subsection 2.3. The results are based on a comparison of all models applied to BC transformed data (for instance, all models based on log transformed data with IW and RW estimated parameters) with the best model applied to raw data. Daily and weekly horizon results under the MS and QLIK loss function assumptions are provided in Tables 5 and 6, respectively.

Insert Tables 5 and 6 here

A number of findings are apparent in Table 5 (daily horizon). The null that at least one model based on BC transformed data has superior predictive ability to the best model based on raw data can be rejected at the 10% level in a large number of instances. However, there is variation in the rejection rates over the realised variance measures, over the loss functions, and over the BC transformation parameters. It is also noteworthy that the quartic root transformation ($\lambda = 1/4$) performs extremely well.¹⁰ Indeed, it delivers the best forecasts under the MS loss function, and is

¹⁰This result is not at odds with other studies. For instance, Nugroho and Morimoto (2016) find that a λ value around 0.1 is optimal in their in-sample analysis of a stochastic volatility model applied to BC transformed realised variance.

no worse than the commonly-used log transformation under the QLIK loss function. For instance, under the former loss function the average p -value is 0.083, compared to 0.156 achieved by the log transformation, while under QLIK loss the respective values are 0.000 and 0.001.¹¹

The results in Table 6 (weekly horizon) show that the quartic root transformation is still useful, with more rejections of the null (at the 10% level) than those associated with the log transformation under both loss functions. Moreover, the log transformation delivers forecasts that are least attractive. In particular, the average p -values under this transformation equal 0.176 (MS loss) and 0.394 (QLIK loss). This compares to respective average p -values of 0.103 and 0.045 for the quartic root transformation. Thus over daily and weekly horizons the quartic transformation performs relatively well.

3.5. Performance determinants

The above results show that the BC transformation does not work well for all realised variance measures. The motivation for use of the BC transformation was that the underlying data are likely to be non-Gaussian (with high positive skewness). It follows that the benefits from such a transformation are likely to increase (decrease) as skewness increases (decreases). To investigate this prediction we consider the relationship between the mean losses associated with each model applied to each transformed realised variance measure (relative to the mean losses associated with each model applied to each raw realised variance measure), *and* the unconditional sample skewness associated with each raw realised variance measure. The scatter plots in Figures 1 depict this relationship when the squared, quartic and log transformed versions of the HAR model (using IW estimation) are used, under the MS and QLIK loss functions and for daily and weekly horizons. In addition, the OLS fitted values associated with these scatter plots are also presented.

Insert Figure 1 here

The plots depict a clear negative relationship - that is, more (less) skewness is associated with lower (higher) losses when using transformed data. Moreover, the unreported p -values associated with the OLS slope coefficients from regressions of relative mean loss on skewness are uniformly close

¹¹The higher rejection rates noted under QLIK loss may reflect the higher test power observed under this loss function (Patton and Sheppard, 2009).

to zero. This negative relationship holds true under the MS and QLIK loss function assumptions, though the slope is less steep under the latter assumption. Moreover, it persists when weekly horizons are considered.

It is also interesting to consider whether the benefits of using BC transformed data are constant over time. Figure 2 provides plots of relative mean losses against time when the HAR model (using IW estimation) is applied to quartic root transformed realised variance under the MS and QLIK loss functions and for daily and weekly horizons. Time-variation in the mean forecast losses is achieved by smoothing forecast losses using a Gaussian kernel smoothing estimator with window size equal to 66 observations.

Insert Figure 2 here

Under the MS loss function assumption the plots indicate that major benefits are available around the high volatility episode observed during the financial crisis. By contrast, under the QLIK loss function assumption the benefits appear more evenly distributed over time with no noticeable benefit observed around the financial crisis. The QLIK loss function penalises under-prediction more than over-prediction, while the MS loss function is symmetric. It follows that in comparison to the forecasts associated with models applied to raw data, the forecasts associated with models applied to transformed data are generally more accurate but tend to under-predict future realised variance.

4. Conclusions

The costs/benefits of using forecasts based on models applied to BC transformed realised variance are examined. The findings are summarized as follows. First, forecasts based on various models applied to BC transformations of realised variance tend to be more accurate than those based on various models applied to raw realised variance. Second, the benefits can be significant in an economic and statistical sense. Third, the commonly-used log transformation does not appear to deliver the best results in terms of statistical significance. Rather it is the quartic transformation that exhibits the best quality in this regard. Finally, relative forecast accuracy varies over the realised variance measures, with data skewness driving this cross-sectional variation. Moreover,

performance varies over time and appears to be a function of market conditions (primarily volatility levels).

The results have implications for researchers and market practitioners. Care is required when examining the performance of proposed models in that BC transformations can have a large impact on results. We advise that forecasts of future realised variance require some form of BC transformation, with the quartic transformation our recommendation. Moreover, the nature of the data should be considered as not all realised variance measures are suited to data transformation. In particular, only measures with high skewness appear ripe for transformation; see Lütkepohl and Xu (2012) for similar conditional findings in the context of macroeconomic forecasting.

References

- Andersen, T., Bollerslev, T. & Diebold, F. (2007). Roughing it up: Including jump components in the measurement, modeling and forecasting of return volatility. *Review of Economics and Statistics*, 89, 707–720.
- Andersen, T., Bollerslev, T., Diebold, F. & Ebens, H. (2001a). The distribution of realized stock return volatility. *Journal of Financial Economics*, 61, 43–76.
- Andersen, T., Bollerslev, T., Diebold, F. & Labys, P. (2001b). The distribution of exchange rate volatility. *Journal of the American Statistical Association*, 96, 42–55.
- Andersen, T., Bollerslev, T., Diebold, F. & Labys, P. (2003). Modeling and forecasting realized volatility. *Econometrica*, 71, 579–625.
- Andersen, T., Dobrev, D. & Schaumburg, E. (2012). Jump-robust volatility estimation using nearest neighbor truncation. *Journal of Econometrics*, 169, 75–93.
- Audrino, F. & Knaus, S. (2016). Lassoing the HAR model: A model selection perspective on realized volatility dynamics. *Econometric Reviews*, forthcoming.
- Banerjee, A., Guo, X. & Wang, H. (2005). On the optimality of conditional expectation as a Bregman predictor. *IEEE Transactions on Information Theory*, 51, 2664–2669.
- Bårdsena, G. & Lütkepohl, H. (2011). Forecasting levels of log variables in vector autoregressions. *International Journal of Forecasting*, 27, 1108–1115.
- Barndorff-Nielsen, O., Hansen, P., Lunde, A. & Shephard, N. (2008). Designing realised kernels to measure the ex-post variation of equity prices in the presence of noise. *Econometrica*, 76, 1481–1536.
- Barndorff-Nielsen, O., Kinnebrouck, S. & Shephard, N. (2010). Measuring downside risk: Realised semivariance. In T. Bollerslev, J. Russell, and M. Watson (Eds.), *Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle*. Oxford University Press.

- Barndorff-Nielsen, O. & Shephard, N. (2002). Econometric analysis of realised volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society, Series B*, 64, 253–280.
- Barndorff-Nielsen, O. & Shephard, N. (2004). Power and Bipower Variation with Stochastic Volatility and Jumps, *Journal of Financial Econometrics*, 2, 1–37.
- Bollerslev, T., Patton, A. & Quaedvlieg, R. (2016). Exploiting the errors: A simple approach for improved volatility forecasting. *Journal of Econometrics*, forthcoming.
- Box G. & Cox, D. (1964). An analysis of transformations (with discussion). *Journal of the Royal Statistical Society, Series B*, 26, 211–246.
- Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics* 7, 174–196.
- Gneiting, T. (2011). Making and evaluating point forecasts, *Journal of the American Statistical Association*, 106, 746–762.
- Gonçalves, S. & Meddahi, N. (2011). Box-Cox transforms for realized volatility. *Journal of Econometrics*, 160, 129–144.
- Granger, C. & Newbold, P. (1976). Forecasting transformed time series. *Journal of the Royal Statistical Society, Series B*, 38, 189–203.
- Hansen, P., Huang, Z. & Shek, H. (2012). Realized GARCH: A joint model for returns and realized measures of volatility. *Journal of Applied Econometrics*, 27, 877–906.
- Koopman, S. & Scharth, M. (2013). The analysis of stochastic volatility in the presence of daily realized measures. *Journal of Financial Econometrics*, 11, 76–115.
- Lütkepohl, H. & Xu, F. (2012). The role of the log transformation in forecasting economic variables. *Empirical Economics*, 42, 619–638.
- Mayr, J. & Ulbricht, D. (2015). Log versus level in VAR forecasting: 42 million empirical answers - Expect the unexpected. *Economics Letters*, 126, 40–42.

- Nelson, H. & Granger, C. (1979). Experience with using the Box-Cox transformation when forecasting economic time series. *Journal of Econometrics*, 10, 57–69.
- Nugroho, D. & Morimoto, T. (2016). Box-Cox realized asymmetric stochastic volatility models with generalized Students t-error distributions. *Journal of Applied Statistics*, forthcoming.
- Pankratz, A. & Dudley, U. (1987). Forecasts of power-transformed series. *Journal of Forecasting*, 6, 239–248.
- Patton, A. (2011). Volatility forecast comparison using imperfect volatility proxies. *Journal of Econometrics*, 160, 246–256.
- Patton, A. (2015). Evaluating and comparing possibly misspecified forecasts. Duke University Working Paper.
- Patton, A. & Sheppard, K. (2009). Optimal combinations of realised volatility estimators. *International Journal of Forecasting*, 25, 218–238.
- Proietti, T. & Lütkepohl, H. (2013). Does the Box-Cox transformation help in forecasting macroeconomic time series? *International Journal of Forecasting*, 29, 88–99.
- Proietti, T. & Riani, M. (2009). Transformations and seasonal adjustment, *Journal of Time Series Analysis*, 30, 47–69.
- Shephard, N. & Sheppard, K. (2010). Realising the future: Forecasting with high frequency based volatility (HEAVY) models, *Journal of Applied Econometrics*, 25, 197–231.
- Tukey, J. (1957). On the comparative anatomy of transformations, *Annals of Mathematical Statistics*, 28, 602–632.
- Weigand, R. (2014). Matrix Box-Cox models for multivariate realized volatility. University of Regensburg Working Papers in Business, Economics and Management Information Systems.
- White, H. (2000). A reality check for data snooping. *Econometrica*, 68, 1097–1126.
- Zhang, L., Mykland, P. & Aït-Sahalia, Y. (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 – Conditional expectations under BC transformations

Adjustment	Conditional expectation
Panel A. General form	
Naive	$g(\mu_{t+h t}; \lambda)$
Second-order approximation	$g(\mu_{t+h t}; \lambda) \left(1 + \frac{\sigma_{t+h t}^2(1-\lambda)}{2(1+\lambda\mu_{t+h t})^2} \right)$
Normal approximation	$g(\mu_{t+h t}; \lambda) \left(1 + \sum_{k=1}^{\infty} g_{2k}(\mu_{t+h t}; \lambda)(2k-1)!!\sigma_{t+h t}^{2k} \right)$
Full	$g(\mu_{t+h t}; \lambda) \left(1 + \sum_{k=1}^{\infty} g_k(\mu_{t+h t}; \lambda)\mu_{k,t+h t} \right)$
Panel B. Specific form: N th root ($\lambda = 1/N$), where N is a positive integer	
Naive	$(1 + \mu_{t+h t}/N)^N$
Second-order approximation	$(1 + \mu_{t+h t}/N)^N \left(1 + \frac{(N-1)\sigma_{t+h t}^2}{2N(1+\mu_{t+h t}/N)^2} \right)$
Normal approximation	$(1 + \mu_{t+h t}/N)^N \left(1 + \sum_{k=1}^{\lfloor N/2 \rfloor} \frac{(N-1)!(2k-1)!!\sigma_{t+h t}^{2k}}{(2k)!(N-2k)!N^{2k-1}(1+\mu_{t+h t}/N)^{2k}} \right)$
Full	$(1 + \mu_{t+h t}/N)^N \left(1 + \sum_{k=1}^N \frac{(N-1)!\mu_{k,t+h t}}{k!(N-k)!N^{k-1}(1+\mu_{t+h t}/N)^k} \right)$
Panel C. Specific form: Log ($\lambda = 0$)	
Naive	$\exp(\mu_{t+h t})$
Second-order approximation	$\exp(\mu_{t+h t}) \left(1 + \frac{\sigma_{t+h t}^2}{2} \right)$
Normal approximation	$\exp(\mu_{t+h t}) \exp\left(\frac{\sigma_{t+h t}^2}{2}\right)$
Full	$\exp(\mu_{t+h t}) \left(1 + \sum_{k=1}^{\infty} \frac{1}{k!} \mu_{k,t+h t} \right)$

Notation:

$$g(\mu_{t+h|t}; \lambda) = (1 + \lambda\mu_{t+h|t})^{1/\lambda}$$

$$g_k(\mu_{t+h|t}; \lambda) = \frac{1-\lambda(k-1)}{k(1+\lambda\mu_{t+h|t})} g_{k-1}(\mu_{t+h|t}; \lambda) \text{ and } g_0 = 1$$

$\mu_{t+h|t}$ is the h -step ahead conditional mean

$\mu_{k,t+h|t}$ is the h -step ahead k th conditional moment about the conditional mean

$\sigma_{t+h|t}$ ($= \sqrt{\mu_{2,t+h|t}}$) is the h -step ahead conditional standard deviation

The notation $\lfloor \cdot \rfloor$ and $!!$ represent the floor and double factorial functions, respectively

Notes: This table contains expressions for the h -step ahead conditional expectations of the raw data (x_t) as a function of the h -step ahead conditional moments of the BC transformed data (y_t).

Table 2 – In-sample HAR model parameter estimates and fit

Determinants	Realised Variance Measure									
	1	2	3	4	5	6	7	8	9	10
Panel A: Raw data										
Daily RV	0.270 (0.102)	0.289 (0.069)	0.256 (0.138)	0.156 (0.074)	0.197 (0.149)	0.219 (0.140)	0.184 (0.159)	0.164 (0.125)	0.219 (0.133)	0.170 (0.169)
Weekly RV	0.350 (0.117)	0.329 (0.085)	0.426 (0.166)	0.343 (0.116)	0.407 (0.136)	0.419 (0.147)	0.428 (0.155)	0.423 (0.131)	0.385 (0.125)	0.438 (0.166)
Monthly RV	0.165 (0.042)	0.167 (0.043)	0.120 (0.053)	0.207 (0.055)	0.150 (0.043)	0.135 (0.046)	0.136 (0.046)	0.153 (0.043)	0.155 (0.044)	0.146 (0.051)
R^2	53.816	53.524	57.410	42.987	49.933	52.891	49.447	48.288	50.340	50.576
Panel B: BC transformed data ($\lambda = 1/2$)										
Daily RV	0.361 (0.043)	0.300 (0.038)	0.453 (0.056)	0.220 (0.038)	0.446 (0.072)	0.414 (0.059)	0.446 (0.080)	0.318 (0.061)	0.405 (0.057)	0.480 (0.096)
Weekly RV	0.364 (0.062)	0.406 (0.057)	0.332 (0.071)	0.369 (0.056)	0.326 (0.081)	0.349 (0.072)	0.332 (0.088)	0.384 (0.071)	0.348 (0.071)	0.305 (0.100)
Monthly RV	0.149 (0.039)	0.151 (0.040)	0.115 (0.038)	0.215 (0.043)	0.127 (0.037)	0.128 (0.038)	0.121 (0.039)	0.158 (0.040)	0.136 (0.038)	0.123 (0.039)
R^2	69.897	66.926	75.218	57.646	74.917	73.558	74.766	67.491	72.872	77.062
R^2 (naive adj.)	54.291	53.595	57.393	43.171	50.100	53.167	49.416	48.784	51.009	49.604
R^2 (full adj.)	54.480	53.837	57.512	43.582	50.209	53.300	49.524	48.989	51.144	49.686
Panel C: BC transformed data ($\lambda = 1/4$)										
Daily RV	0.337 (0.032)	0.273 (0.030)	0.459 (0.031)	0.215 (0.025)	0.465 (0.032)	0.425 (0.032)	0.477 (0.035)	0.315 (0.030)	0.397 (0.031)	0.523 (0.037)
Weekly RV	0.397 (0.040)	0.436 (0.039)	0.332 (0.040)	0.377 (0.036)	0.324 (0.041)	0.350 (0.041)	0.319 (0.044)	0.379 (0.039)	0.369 (0.040)	0.289 (0.043)
Monthly RV	0.144 (0.028)	0.150 (0.029)	0.116 (0.026)	0.218 (0.031)	0.122 (0.026)	0.127 (0.026)	0.116 (0.026)	0.169 (0.028)	0.131 (0.027)	0.110 (0.025)
R^2	71.090	67.938	76.908	58.925	77.523	75.825	77.837	68.059	74.806	80.159
R^2 (naive adj.)	53.689	52.488	57.366	41.708	50.591	53.244	49.805	48.217	51.248	49.899
R^2 (second-order adj.)	54.467	53.478	57.867	43.269	51.012	53.771	50.213	49.129	51.800	50.205
R^2 (Gaussian adj.)	54.469	53.481	57.869	43.277	51.013	53.772	50.214	49.132	51.801	50.205
R^2 (full adj.)	54.503	53.521	57.888	43.359	51.033	53.797	50.233	49.160	51.826	50.220
Panel D: BC transformed data ($\lambda = 0$)										
Daily RV	0.320 (0.027)	0.266 (0.026)	0.440 (0.026)	0.209 (0.022)	0.448 (0.026)	0.413 (0.026)	0.466 (0.026)	0.292 (0.023)	0.376 (0.027)	0.502 (0.025)
Weekly RV	0.398 (0.033)	0.423 (0.034)	0.332 (0.033)	0.362 (0.031)	0.327 (0.033)	0.346 (0.034)	0.318 (0.033)	0.367 (0.031)	0.376 (0.033)	0.304 (0.031)
Monthly RV	0.150 (0.024)	0.160 (0.024)	0.126 (0.022)	0.225 (0.025)	0.128 (0.021)	0.135 (0.022)	0.122 (0.021)	0.185 (0.023)	0.136 (0.022)	0.109 (0.020)
R^2	69.030	65.836	74.859	56.375	75.794	73.982	76.377	64.386	72.871	78.437
R^2 (naive adj.)	51.565	49.912	56.136	38.383	50.208	52.171	49.577	45.667	50.075	50.103
R^2 (second-order adj.)	54.641	53.665	58.396	43.699	52.003	54.395	51.288	49.590	52.409	51.493
R^2 (Gaussian adj.)	54.779	53.854	58.482	44.048	52.066	54.481	51.346	49.813	52.506	51.537
R^2 (full adj.)	54.904	53.963	58.570	44.184	52.150	54.586	51.425	49.848	52.611	51.594

Notes: This table contains the parameter estimates and standard errors associated with the HAR model applied to raw and BC transformed realised variance. Two sets of R^2 statistics are provided. The first correspond to those observed when the models are estimated. The second set are calculated by first transforming the fitted values into raw data form and then calculating the R^2 values based on these fitted values. Four versions of the second set are provided, each one corresponding to a different way of converting (adjusting) the transformed fitted values to their raw data form equivalents.

Table 3 – Out-of-sample forecast losses (daily horizon)

		Realised Variance Measure									
Model	Estimation	1	2	3	4	5	6	7	8	9	10
Panel A: Forecast losses using raw data											
AR	IW	1.197	1.177	1.185	1.205	1.226	1.215	1.216	1.222	1.225	1.180
HAR		1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
HEAVY		0.979	0.990	0.968	0.988	0.944	0.975	0.933	0.963	0.955	0.893
IHEAVY		0.962	0.987	0.925	0.977	0.865	0.931	0.848	0.905	0.897	0.740
RealGARCH		0.965	0.960	0.995	0.994	0.931	0.952	0.920	0.878	0.934	0.945
IRealGARCH		0.952	0.972	0.936	0.978	0.853	0.909	0.834	0.796	0.884	0.790
AR	RW	1.310	1.327	1.344	1.262	1.499	1.464	1.496	1.458	1.474	1.287
HAR		1.083	1.111	1.149	1.065	1.187	1.200	1.218	1.203	1.160	1.087
HEAVY		1.059	1.086	1.092	1.038	1.096	1.127	1.086	1.105	1.096	0.957
IHEAVY		0.986	1.020	0.966	0.999	0.895	0.965	0.878	0.959	0.930	0.749
RealGARCH		1.161	1.127	1.195	1.092	1.168	1.232	1.144	1.102	1.176	1.080
IRealGARCH		1.041	1.059	1.042	1.053	0.958	1.042	0.935	0.909	1.004	0.853
Panel B: Forecast losses using BC transformed data ($\lambda = 1/2$)											
AR	IW	1.065	1.090	1.041	1.121	0.940	1.024	0.938	0.954	0.961	0.846
HAR		0.895	0.943	0.863	0.924	0.760	0.842	0.757	0.773	0.784	0.689
HEAVY		0.886	0.941	0.850	0.921	0.744	0.838	0.738	0.772	0.772	0.664
IHEAVY		0.884	0.938	0.845	0.913	0.740	0.833	0.732	0.765	0.769	0.652
RealGARCH		0.883	0.946	0.846	0.919	0.732	0.831	0.729	0.758	0.762	0.661
IRealGARCH		0.882	0.923	0.839	0.907	0.726	0.824	0.721	0.746	0.759	0.647
AR	RW	1.100	1.131	1.103	1.131	1.040	1.115	1.045	1.044	1.047	0.916
HAR		0.931	0.967	0.914	0.947	0.833	0.906	0.834	0.829	0.846	0.745
HEAVY		0.913	0.957	0.894	0.942	0.808	0.892	0.804	0.822	0.825	0.713
IHEAVY		0.902	0.946	0.872	0.928	0.781	0.867	0.772	0.800	0.804	0.679
RealGARCH		0.908	0.948	0.885	0.938	0.792	0.880	0.785	0.787	0.811	0.726
IRealGARCH		0.900	0.944	0.866	0.921	0.767	0.857	0.756	0.763	0.794	0.692
Panel C: Forecast losses using BC transformed data ($\lambda = 1/4$)											
AR	IW	1.026	1.081	0.991	1.099	0.863	0.966	0.863	0.896	0.891	0.760
HAR		0.884	0.945	0.839	0.922	0.719	0.816	0.716	0.753	0.752	0.633
HEAVY		0.878	0.944	0.829	0.918	0.708	0.813	0.702	0.752	0.745	0.616
IHEAVY		0.873	0.937	0.824	0.906	0.705	0.809	0.698	0.742	0.741	0.612
RealGARCH		0.885	0.957	0.831	0.916	0.705	0.816	0.697	0.749	0.746	0.609
IRealGARCH		0.878	0.948	0.824	0.900	0.700	0.808	0.692	0.736	0.741	0.604
AR	RW	1.031	1.085	1.010	1.092	0.900	0.997	0.905	0.921	0.919	0.797
HAR		0.892	0.940	0.859	0.924	0.749	0.838	0.751	0.767	0.771	0.662
HEAVY		0.881	0.937	0.844	0.921	0.732	0.831	0.730	0.764	0.759	0.640
IHEAVY		0.875	0.930	0.836	0.908	0.725	0.822	0.721	0.752	0.752	0.633
RealGARCH		0.883	0.948	0.838	0.919	0.718	0.824	0.711	0.755	0.753	0.632
IRealGARCH		0.878	0.940	0.831	0.904	0.710	0.814	0.701	0.741	0.745	0.624
Panel D: Forecast losses using BC transformed data ($\lambda = 0$)											
AR	IW	1.016	1.100	0.957	1.101	0.807	0.925	0.805	0.878	0.850	0.698
HAR		0.882	0.946	0.826	0.915	0.700	0.802	0.695	0.745	0.740	0.609
HEAVY		0.877	0.944	0.819	0.911	0.693	0.802	0.685	0.745	0.736	0.597
IHEAVY		0.891	0.964	0.835	0.922	0.700	0.813	0.692	0.748	0.743	0.601
RealGARCH		0.890	0.970	0.826	0.905	0.696	0.810	0.683	0.740	0.744	0.595
IRealGARCH		0.898	0.979	0.835	0.920	0.698	0.815	0.688	0.744	0.747	0.596
AR	RW	1.033	1.139	0.966	1.123	0.816	0.944	0.815	0.893	0.867	0.707
HAR		0.883	0.941	0.834	0.918	0.709	0.808	0.709	0.746	0.744	0.618
HEAVY		0.876	0.941	0.825	0.914	0.702	0.808	0.697	0.748	0.740	0.604
IHEAVY		0.899	0.968	0.852	0.935	0.715	0.827	0.713	0.760	0.752	0.619
RealGARCH		0.888	0.965	0.834	0.916	0.704	0.816	0.695	0.750	0.746	0.608
IRealGARCH		0.901	0.975	0.845	0.941	0.712	0.824	0.704	0.761	0.751	0.623

Notes: This table contains the mean forecast losses for each model relative to the HAR model (using IW estimation applied to raw data). The MS loss function is assumed.

Table 4 – Out-of-sample forecast losses (weekly horizon)

		Realised Variance Measure									
Model	Estimation	1	2	3	4	5	6	7	8	9	10
Panel A: Forecast losses using raw data											
AR	IW	1.332	1.352	1.264	1.473	1.390	1.360	1.321	1.297	1.391	1.298
HAR		1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
HEAVY		1.021	1.002	0.983	0.987	1.021	1.008	0.980	1.001	1.031	0.940
IHEAVY		0.888	0.933	0.790	0.912	0.609	0.789	0.568	0.670	0.690	0.323
RealGARCH		1.013	0.953	1.032	1.058	1.044	0.997	1.003	0.867	1.025	1.136
IRRealGARCH		0.965	0.989	0.853	1.043	0.646	0.822	0.596	0.625	0.734	0.380
AR	RW	1.805	1.894	2.055	1.608	3.743	2.849	3.831	2.840	3.362	2.047
HAR		1.346	1.412	1.665	1.245	1.976	1.918	2.095	1.994	1.800	1.344
HEAVY		1.497	1.536	1.734	1.239	2.294	2.095	2.254	2.066	2.148	1.389
IHEAVY		0.916	0.984	0.843	0.943	0.645	0.837	0.601	0.742	0.728	0.331
RealGARCH		1.758	1.669	2.034	1.395	2.954	2.616	2.858	2.220	2.719	1.922
IRRealGARCH		1.078	1.133	0.960	1.175	0.753	0.989	0.696	0.751	0.868	0.412
Panel B: Forecast losses using BC transformed data ($\lambda = 1/2$)											
AR	IW	1.104	1.250	0.954	1.423	0.662	0.970	0.625	0.825	0.755	0.377
HAR		0.776	0.844	0.694	0.835	0.481	0.677	0.458	0.531	0.545	0.279
HEAVY		0.794	0.854	0.710	0.830	0.500	0.693	0.472	0.537	0.563	0.293
IHEAVY		0.798	0.855	0.720	0.808	0.503	0.698	0.473	0.530	0.567	0.281
RealGARCH		0.787	0.851	0.704	0.833	0.489	0.684	0.465	0.530	0.552	0.292
IRRealGARCH		0.802	0.839	0.716	0.819	0.493	0.688	0.468	0.524	0.561	0.283
AR	RW	1.099	1.183	0.988	1.344	0.802	1.047	0.775	0.868	0.854	0.459
HAR		0.857	0.934	0.796	0.891	0.592	0.797	0.582	0.625	0.644	0.342
HEAVY		0.887	0.956	0.823	0.899	0.644	0.840	0.621	0.658	0.693	0.368
IHEAVY		0.836	0.908	0.757	0.838	0.548	0.749	0.514	0.575	0.614	0.298
RealGARCH		0.860	0.931	0.796	0.887	0.616	0.811	0.591	0.606	0.664	0.373
IRRealGARCH		0.851	0.917	0.763	0.852	0.541	0.741	0.505	0.550	0.613	0.310
Panel C: Forecast losses using BC transformed data ($\lambda = 1/4$)											
AR	IW	1.162	1.328	0.980	1.498	0.632	0.974	0.592	0.859	0.754	0.338
HAR		0.759	0.827	0.671	0.834	0.449	0.654	0.424	0.516	0.518	0.248
HEAVY		0.767	0.830	0.682	0.822	0.459	0.664	0.433	0.515	0.525	0.256
IHEAVY		0.775	0.835	0.698	0.794	0.471	0.674	0.445	0.507	0.534	0.263
RealGARCH		0.763	0.827	0.679	0.827	0.453	0.659	0.429	0.515	0.521	0.253
IRRealGARCH		0.768	0.827	0.689	0.797	0.460	0.664	0.437	0.506	0.526	0.257
AR	RW	1.130	1.264	0.962	1.431	0.642	0.970	0.603	0.829	0.758	0.342
HAR		0.788	0.863	0.698	0.845	0.476	0.684	0.452	0.531	0.544	0.262
HEAVY		0.802	0.872	0.717	0.851	0.498	0.709	0.472	0.546	0.562	0.276
IHEAVY		0.802	0.869	0.727	0.812	0.502	0.709	0.475	0.531	0.562	0.279
RealGARCH		0.794	0.869	0.707	0.843	0.483	0.693	0.457	0.534	0.549	0.272
IRRealGARCH		0.793	0.854	0.707	0.814	0.479	0.684	0.451	0.518	0.545	0.271
Panel D: Forecast losses using BC transformed data ($\lambda = 0$)											
AR	IW	1.290	1.451	1.086	1.591	0.679	1.059	0.628	0.940	0.827	0.369
HAR		0.758	0.820	0.670	0.830	0.444	0.652	0.418	0.519	0.514	0.246
HEAVY		0.758	0.818	0.677	0.812	0.449	0.657	0.424	0.514	0.516	0.251
IHEAVY		0.822	0.895	0.740	0.856	0.482	0.710	0.457	0.536	0.552	0.266
RealGARCH		0.757	0.823	0.673	0.827	0.445	0.654	0.420	0.519	0.514	0.247
IRRealGARCH		0.808	0.872	0.726	0.884	0.472	0.696	0.451	0.547	0.543	0.261
AR	RW	1.333	1.504	1.135	1.611	0.709	1.123	0.657	0.956	0.869	0.360
HAR		0.781	0.851	0.687	0.849	0.456	0.671	0.429	0.528	0.530	0.248
HEAVY		0.788	0.855	0.712	0.851	0.475	0.697	0.450	0.543	0.542	0.257
IHEAVY		0.870	0.947	0.791	0.915	0.518	0.761	0.495	0.577	0.587	0.287
RealGARCH		0.794	0.863	0.703	0.874	0.471	0.692	0.442	0.547	0.541	0.259
IRRealGARCH		0.846	0.906	0.752	0.943	0.502	0.730	0.474	0.581	0.569	0.285

Notes: This table contains the mean forecast losses for each model relative to the HAR model (using IW estimation applied to raw data). The MS loss function is assumed.

Table 5 – Reality check results (daily horizon)

Bregman loss function	Realised Variance Measure										Av. p -value
	1	2	3	4	5	6	7	8	9	10	
Panel A: Test p -values using BC transformed data ($\lambda = 1/2$)											
MS loss	0.145	0.336	0.099	0.029	0.065	0.168	0.093	0.393	0.039	0.134	0.150
QLIK loss	0.015	0.909	0.086	0.855	0.191	0.320	0.843	0.326	0.319	1.000	0.486
Panel B: Test p -values using BC transformed data ($\lambda = 1/4$)											
MS loss	0.056	0.421	0.018	0.004	0.067	0.052	0.048	0.070	0.077	0.021	0.083
QLIK loss	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Panel C: Test p -values using BC transformed data ($\lambda = 0$)											
MS loss	0.109	0.872	0.026	0.043	0.077	0.069	0.070	0.148	0.112	0.038	0.156
QLIK loss	0.000	0.004	0.000	0.002	0.000	0.000	0.000	0.005	0.003	0.000	0.001

Notes: This table contains the p -values associated with reality check tests of the null hypothesis that no model based on BC transformed data outperforms the best model based on raw data (the benchmark model), against the alternative that at least one model based on BC transformed data outperforms the benchmark model. The final column contains the average (Av.) p -values across all realised variance measures.

Table 6 – Reality check results (weekly horizon)

Bregman loss function	Realised Variance Measure										Av. p -value
	1	2	3	4	5	6	7	8	9	10	
Panel A: Test p -values using BC transformed data ($\lambda = 1/2$)											
MS loss	0.095	0.244	0.143	0.320	0.039	0.098	0.047	0.081	0.034	0.178	0.128
QLIK loss	0.006	0.003	0.000	0.142	0.000	0.001	0.000	0.025	0.000	0.000	0.017
Panel B: Test p -values using BC transformed data ($\lambda = 1/4$)											
MS loss	0.087	0.249	0.074	0.326	0.048	0.064	0.043	0.037	0.061	0.035	0.103
QLIK loss	0.036	0.017	0.002	0.272	0.001	0.015	0.002	0.109	0.000	0.000	0.045
Panel C: Test p -values using BC transformed data ($\lambda = 0$)											
MS loss	0.190	0.332	0.199	0.470	0.069	0.135	0.055	0.149	0.069	0.095	0.176
QLIK loss	0.916	0.274	0.438	1.000	0.066	0.136	0.060	0.967	0.058	0.025	0.394

Notes: This table contains the p -values associated with reality check tests of the null hypothesis that no model based on BC transformed data outperforms the best model based on raw data (the benchmark model), against the alternative that at least one model based on BC transformed data outperforms the benchmark model. The final column contains the average (Av.) p -values across all realised variance measures.

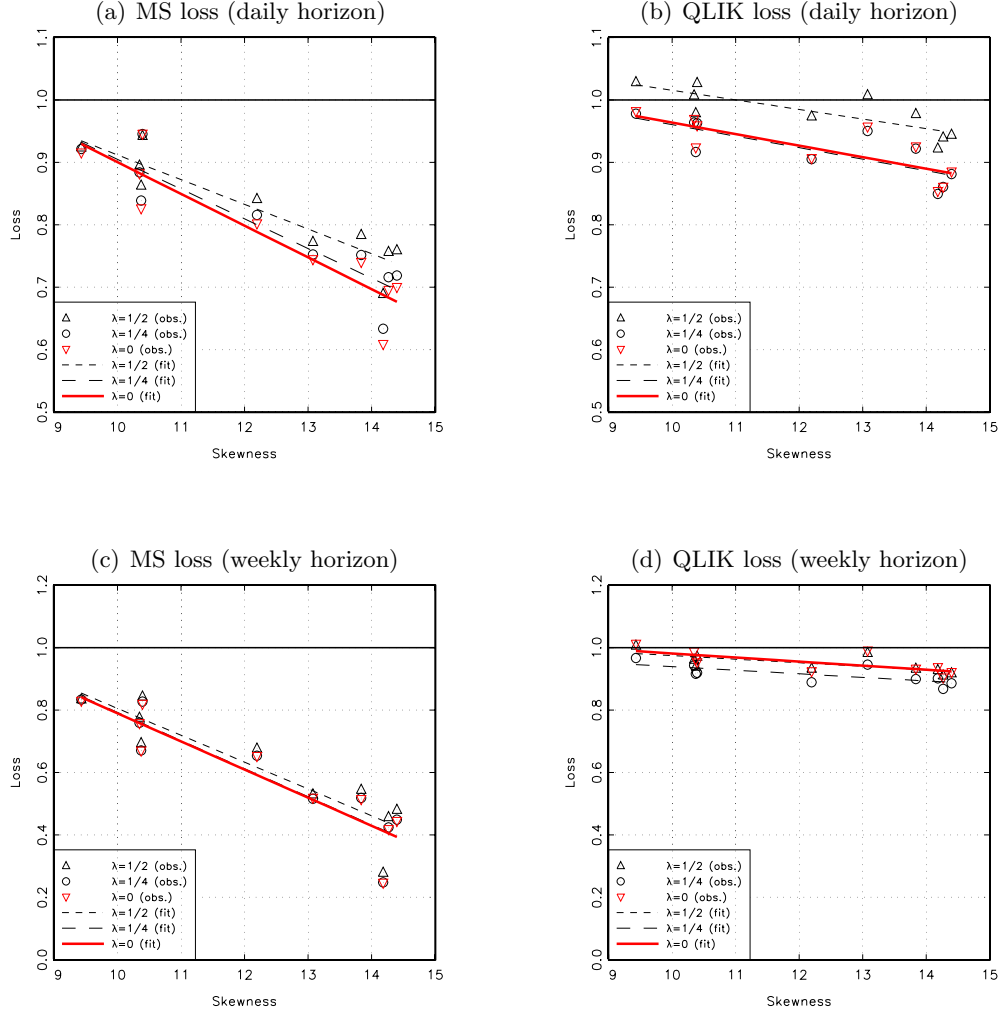


Figure 1 – Forecast losses and skewness

This figure contains plots of the mean forecast losses for BC transformed HAR models relative to the HAR model (both using IW estimation) against unconditional sample skewness for each realised variance measure. The fit is given by the OLS cross-sectional regression of mean forecast loss on skewness.

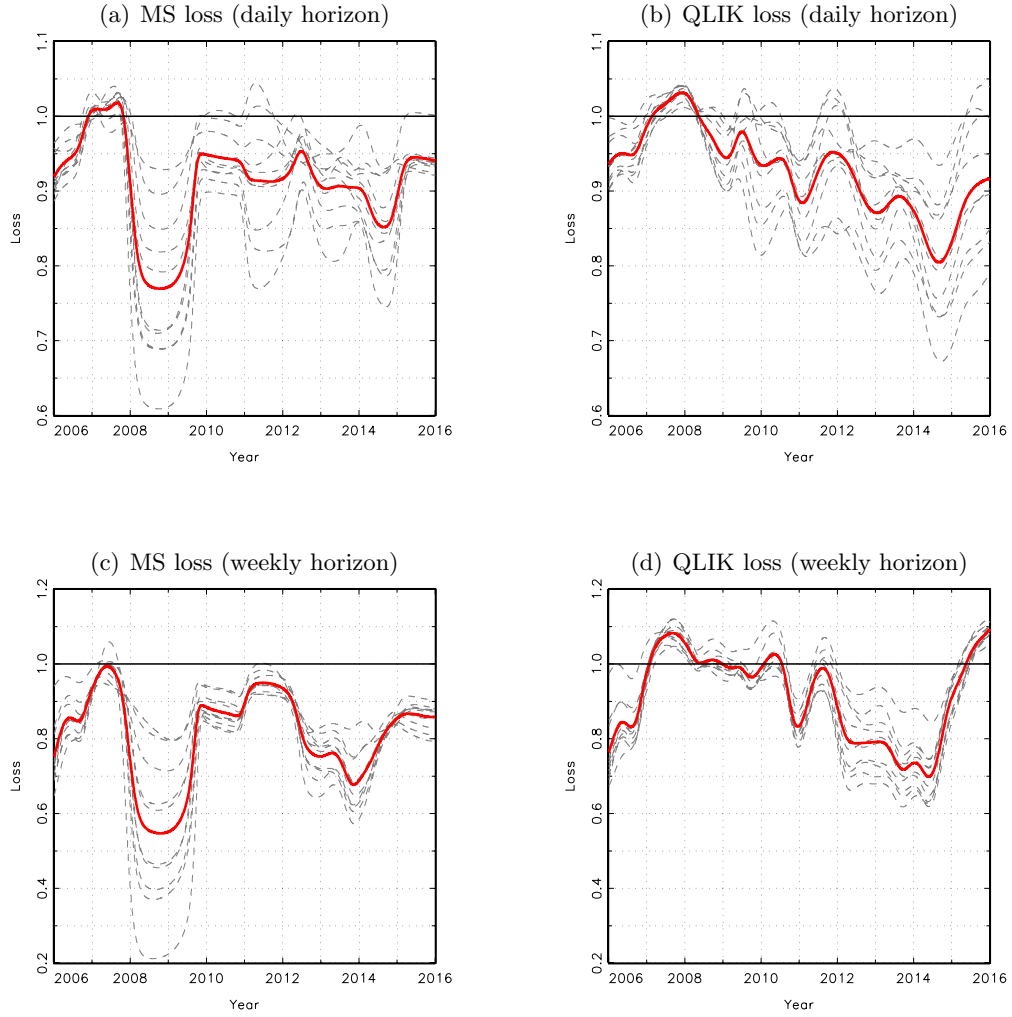


Figure 2 – Forecast losses over time

This figure contains plots of the time-varying mean forecast losses for quartic root transformed HAR models relative to the HAR model (both using IW estimation). Solid lines represent the relative mean forecast losses across all realised variance measures, and the dashed lines are the individual relative mean forecast losses for each realised variance measure.