

Value-at-Risk Prediction in R with the GAS Package

by David Ardia, Kris Boudt and Leopoldo Catania

Abstract GAS models have been recently proposed in time-series econometrics as valuable tools for signal extraction and prediction. This paper details how financial risk managers can use GAS models for Value-at-Risk (VaR) prediction using the novel **GAS** package for R. Details and code snippets for prediction, comparison and backtesting with GAS models are presented. An empirical application considering Dow Jones Index constituents investigates the VaR forecasting performance of GAS models.

Introduction

Following the regulatory process of the Basel Accords (currently the Basel III Accords), banks and financial institutions are required to meet capital requirements, and must rely on state-of-the-art risk systems. In particular, they must assess the uncertainty about the future values of their portfolios, and estimate the extent and the likelihood of potential losses. This is typically achieved in two steps. First, the distribution of future profit & loss (*i.e.*, future portfolios' or assets' returns) is modelled. Second, financial risk is measured from the distribution; nowadays, the Value-at-Risk (VaR) risk measure is the standard (Jorion, 1997). This metric gives, for a given time horizon, the asset's loss (or return) that is expected to be exceeded with a given probability level α (referred to as the *risk level*, and which is typically set to one or five percent, *i.e.*, $\alpha \in \{0.01, 0.05\}$). Hence, the VaR is nothing else than a given percentile of the returns distribution. The popularity of VaR mostly relies on: (i) the simple rationale behind it, (ii) the ease of computation, and (iii) its role in the financial regulation (see Basel Committee, 2010).

Formally, assuming a continuous cumulative density function (*cdf*) with time-varying parameters $\theta_t \in \mathbb{R}^d$ and additional static parameters $\psi \in \mathbb{R}^q$, $F(\cdot; \theta_t, \psi)$, for the asset log-return at time t , $r_t \in \mathbb{R}$, the $VaR_t(\alpha)$ is given by:

$$VaR_t(\alpha) \equiv F^{-1}(\alpha; \theta_t, \psi),$$

where $F^{-1}(\cdot)$ denotes the inverse of the *cdf*, *i.e.*, the quantile function. It follows that $VaR_t(\alpha)$ is nothing more than the α -quantile of the return distribution at time t .¹ Evidently, remaining within the fully parametric framework considered here, the crucial point is the determination of $F(\cdot)$ and its parameters θ_t and ψ . Indeed, the majority of the financial econometrics literature focuses on this aspect; see *e.g.*, McNeil et al. (2015).

Recently, the new class of Score Driven (SD) models has been introduced by Creal et al. (2013) and Harvey (2013) offering an alternative to the Generalized AutoRegressive Heteroscedasticity (GARCH) framework pioneered by Bollerslev (1986) to model the conditional variance of financial returns. SD models are also referred to as Generalized Autoregressive Score (GAS) and Dynamic Conditional Score (DCS) models. In this paper, we follow the GAS nomenclature.

Formally, in GAS models the vector of time-varying parameters, θ_t , is updated through a dynamic equation based on the score of the conditional² probability density function of r_t , $f(\cdot; \theta_t, \psi)$, *i.e.*:

$$\theta_{t+1} \equiv \kappa + \mathbf{A}\mathbf{s}_t + \mathbf{B}\theta_t, \quad (1)$$

where \mathbf{s}_t is the scaled score of $f(\cdot; \theta_t, \psi)$ with respect to θ_t , evaluated in r_t ; see Creal et al. (2013). The coefficients κ , \mathbf{A} and \mathbf{B} control for the evolution of θ_t and need to be estimated along with ψ from the data, usually by Maximum Likelihood. GAS models have been employed for a variety of applications in financial econometrics, mostly because they offer a link between the two most common frameworks to model volatility, namely: GARCH models and stochastic volatility (SV) models (Taylor, 1986). Indeed, while resorting on straightforward estimation procedures as for GARCH, they update models parameters accounting for the whole shape of the conditional distribution of the data, as in the SV context; see Koopman et al. (2016). Hence, they offer a good trade-off between ease of estimation and flexibility.

¹Sometimes VaR is defined with respect to the the loss variable $l_t = -r_t$, *i.e.*, the negative of the return. Clearly, all the arguments of this paper can be easily adapted to this case.

²The conditioning is intended with respect to the past observations r_{t-s} ($s > 0$), however, for notational purposes, this is not always reported.

Within the **R** statistical environment, the **GAS** package of [Catania et al. \(2016\)](#) allows practitioners and researchers to easily deal with GAS models. **GAS** implements routines to estimate, predict, and simulate with GAS models; see [Ardia et al. \(2016\)](#).

The aim of this paper is to show how the **GAS** package can be used for VaR evaluation, prediction and backtesting. We do this by detailing the functionality of the **GAS** package devoted to these aspects as well as reporting an empirical application using financial time series. We focus on the three major steps a practitioner involved in risk management faces during his job: (i) prediction of future VaR levels, (ii) backtesting, and (iii) comparison with alternative models. The empirical part of the article deals with these three points from an applied perspective while the computational part details the **GAS** functionalities devoted to VaR.

The article is organized in the following manner. We first introduce a general GAS model suited for financial returns. Then we review the methodology for assessing and comparing VaR predictions in the context of GAS models. **GAS** functionalities devoted to VaR analysis are subsequently illustrated. Finally, we report an empirical application using a moderately large set of financial returns.

A skew–Student–*t* GAS model with time–varying volatility

Financial returns exhibit several stylized facts that need to be taken into consideration in order to produce reliable risk forecasts. Empirically, the distribution of returns is (left–)skewed and fat–tailed, and its variance is time–varying (*i.e.*, returns exhibit the so–called *volatility clustering*); see, *e.g.*, [McNeil et al. \(2015\)](#).

To account for these features, we consider a very flexible specification, in which we assume that the log–return at time t , r_t , is distributed conditionally on past observations as follows:

$$r_t | \mathcal{I}_{t-1} \sim \mathcal{SKST}(r_t; \mu, \sigma_t, \xi, \nu),$$

where \mathcal{I}_{t-1} is the information set up to time $t-1$, and $\mathcal{SKST}(r_t; \cdot)$ denotes the skew–Student–*t* distribution of [Fernández and Steel \(1998\)](#) with location $\mu \in \mathbb{R}$, time–varying scale $\sigma_t > 0$, and shape and skewness parameters $\nu > 2$ and $\xi > 0$, respectively. We parametrise the \mathcal{SKST} distribution as in [Bauwens and Laurent \(2005\)](#) such that $\mathbb{E}_{t-1}[r_t] = \mu$ and $\text{Var}_{t-1}[r_t] = \sigma_t^2$. The log–density of the \mathcal{SKST} distribution evaluated in r_t is given by:

$$\log f_{\mathcal{SKST}}(r_t; \mu, \sigma_t, \xi, \nu) = \log g + \log s + c - \frac{\nu+1}{2} \log \left[1 + \frac{\left[\left(\frac{r_t - \mu}{\sigma_t} \right) s + m \right]^2}{(\nu-2) (\xi^*)^2} \right],$$

where:

$$\begin{aligned} m &\equiv \mu_1 \left(\xi - \frac{1}{\xi} \right) \\ s &\equiv \sqrt{(1 - \mu_1^2) \left(\xi^2 + \frac{1}{\xi^2} \right) + 2\mu_1^2 - 1} \\ g &\equiv \frac{2}{\xi + \frac{1}{\xi}} \\ c &\equiv \frac{1}{2} [\log \nu - \log(\nu-2) - \log \pi + \log \nu] + \log \Gamma \left(\frac{\nu+1}{2} \right) - \log \Gamma \left(\frac{\nu}{2} \right), \end{aligned}$$

with:

$$\mu_1 \equiv \frac{2\sqrt{\nu-2}}{(\nu-1)} \frac{\Gamma(\frac{\nu+1}{2})}{\Gamma(\frac{\nu}{2})\Gamma(\frac{1}{2})},$$

and:

$$\xi_t^* \equiv \begin{cases} 1 & \text{if } z_t = 0 \\ \frac{1}{\xi} & \text{if } z_t < 0, \\ \xi & \text{if } z_t > 0 \end{cases},$$

where $z_t \equiv \left(\frac{r_t - \mu}{\sigma_t} \right) s + m$. It follows that, similar to [Laurent et al. \(2016\)](#), the time–varying GAS parameter θ_t in (1) is given by $\theta_t \equiv \theta_t \equiv \log \sigma_t$ and that its time–variation is also determined by the

static parameters $\boldsymbol{\psi} \equiv (\mu, \nu, \xi)'$. The score of $SKST(r_t; \cdot)$ with respect to θ_t, s_t , is given by:

$$s_t \equiv \left(\frac{sz_t (\nu + 1) (r_t - \mu)}{(\xi_t^*)^2 (\nu - 2)} - 1 \right),$$

and enters linearly in the updating equation (1). Starting from the general $SKST$ distribution, we recover as special cases:

- (i) The Student- t distribution, \mathcal{ST} , imposing $\xi = 1$;
- (ii) The Normal distribution, \mathcal{N} , imposing $\nu = \infty$ and $\xi = 1$.

Given a series of T log-returns, r_1, \dots, r_T , the model parameters are estimated by maximizing the log-likelihood function; see Blasques et al. (2014). Prediction with GAS models is straightforward thanks to the recursive nature of the updating equation (1). Specifically, the one-step ahead predictive distribution $F(\cdot; \hat{\theta}_{T+1}, \boldsymbol{\psi})$ is available in closed form whereas the h -step ahead distribution ($h > 1$) needs to be simulated; see Blasques et al. (2016). As a direct consequence, at time t , $VaR_{t+1}(\alpha)$ is directly available, whereas $VaR_{t+h}(\alpha)$ ($h > 1$) has to be evaluated as the quantile of simulated values.

Backtesting and comparing the VaR

The recursive method of forecasting (Marcellino et al., 2006) is usually employed in order to backtest the adequacy of a statistical model as well as to perform models comparison in terms of VaR predictions. The objective of a backtesting analysis is to verify the precision of the prediction by separating the estimation window and the evaluation period. Differently, models comparison usually order models according to a loss function. To this end, the full sample of T returns is divided in an in-sample period of length S , and an out-of-sample period of length H . Model parameters are firstly estimated over the in-sample period, subsequently the h -step ahead prediction of the return distribution at time $S + h$ is generated along with the corresponding VaR level. These steps are repeated augmenting the in-sample period with new observations in a recursive way until we reach the end of the series, T . If during the data augmentation step, past observations are eliminated, we are considering a rolling window, otherwise we have an expanding window. In this paper, we use the rolling window configuration for our analysis and we set $h = 1$.

Once a series of VaR predictions is available, forecasts adequacy is assessed through backtesting procedures. VaR backtesting procedures usually checks the correct coverage of the unconditional and conditional left-tail of log-returns distribution. Correct unconditional coverage (UC) was first considered by Kupiec (1995), while correct conditional coverage (CC) by Christoffersen (1998). The main difference between UC and CC concerns the distribution we are focusing on. For instance, UC considers correct coverage of the left-tail of the marginal return distribution, $f(r_t)$, while CC deals with the conditional density $f(r_t | \mathcal{I}_{t-1})$. From an inferential perspective, UC looks at the ratio between the expected number of VaR violations implied by the chosen confidence level, α , during the forecast period, *i.e.*, αH , and the realized VaR violations observed from the data. Formally, the Actual over Expected (AE) ratio is defined as:

$$AE \equiv \frac{\sum_{s=1}^H d_{t+s}}{\alpha H}, \quad (2)$$

where $d_{t+s} \equiv \mathbf{1}_{\{r_{t+s} < VaR_{t+s}(\alpha)\}}$ and $\mathbf{1}_{\{\cdot\}}$ is the indicator function equal to one if the condition holds, and zero otherwise. Hence, if $AE > 1$ the model underestimate risk, while if $AE < 1$ the model is too conservative. Clearly, both situations are problematic and imply capital losses, the former due to unexpected negative returns, the latter due to overestimation of capital requirements.

In order to investigate CC, Christoffersen (1998) proposed a test on the series of VaR exceedance $\{d_{t+s}, s = 1, \dots, H\}$, usually referred to as "hitting series". Specifically, if correct conditional coverage is achieved by the model, VaR exceedances should be independently distributed over time.

Recently, a general testing procedure for dynamic quantile models, also suited for VaR backtesting, has been proposed by Engle and Manganelli (2004). The dynamic quantile test (DQ) jointly tests for UC and CC and has more power than previous alternatives under some form of model misspecification. The series of interest is defined as $\{d_{t+s} - \alpha, s = 1, \dots, H\}$. Under correct model specification, we have the following moment conditions: $\mathbb{E}[d_{t+s} - \alpha] = 0$, $\mathbb{E}_{t+s-1}[d_{t+s} - \alpha] = 0$, $\mathbb{E}[(d_{t+s} - \alpha)(d_{t+j} - \alpha)] = 0$ for $j \neq s$; see Engle and Manganelli (2004).

Real world applications generally consider several models for VaR prediction. If correct unconditional/conditional coverage is achieved by more than one model, the practitioner faces the problem of not being able to choose between different alternatives. In this situation, model comparison

techniques are used in order to choose the best performing model. Models ranking is achieved thanks to the definition of a loss function. Among several available loss functions for quantile prediction (McAleer and Da Veiga, 2008), the Quantile Loss (QL) used for quantile regressions (Koenker and Bassett, 1978) is one the most frequent choice in the VaR context; see González-Rivera et al. (2004). Formally, given a VaR prediction at confidence level α for time $t + 1$, the associated quantile loss, $QL_{t+1}(\alpha)$, is defined as:

$$QL_{t+1}(\alpha) \equiv (\alpha - d_{t+1})(y_{t+1} - VaR_{t+1}(\alpha)).$$

Evidently, QL is an asymmetric loss function that penalizes more heavily with weight $(1 - \alpha)$ the observations for which we observe returns VaR exceedance. Quantile losses are then averaged over the forecasting period, models with lower averages are preferred. Outperformance of model \mathcal{A} versus model \mathcal{B} is finally assessed looking at the ratio between the average QLs, associated with the two models, *i.e.*, if $QL_{\mathcal{A}}/QL_{\mathcal{B}} < 1$ then model \mathcal{A} outperforms model \mathcal{B} and vice versa.

Summarizing, given a set of available models, a typical VaR forecasting exercise consists of the three major steps:

- (i) perform rolling forecast during the out-of-sample period,
- (ii) statistical backtest of VaR predictions using UC, CC and DQ tests,
- (iii) model comparison looking at the average QL of each model.

The next section is devoted to detail the implementation of these steps with the **GAS** package.

VaR prediction with the GAS package

We first briefly review how to make predictions with GAS models using the **GAS** package. We refer the reader to Ardia et al. (2016) where additional details and examples are presented. The rest of the section is devoted to VaR backtest and model comparison through the DQ test and the QL measure previously introduced.

One-step ahead rolling forecasts applications are straightforward to implement with the **GAS** package. Specifically, the user just need to specify the model through the `UniGASSpec()` function, and then perform rolling predictions with the `UniGASRoll()` function. In the **GAS** package, models are specified through the definition of the conditional distribution assumed for the data, `Dist`, and the list of time-varying parameters, `GASpar`.

`Dist` is a character equal to the label of the distribution. For instance, \mathcal{SKST} is identified as "sstd", \mathcal{ST} as "std" and \mathcal{N} as "norm"; see Table 1 of Ardia et al. (2016) for the list of distributions and associated labels available in the **GAS** package.

`GASpar` is a list with named boolean elements. Entries name are: `location`, `scale`, `skewness` and `shape` and indicate whether the associated distribution parameters are time-varying or not. By default we have `GASpar = list(location = FALSE, scale = TRUE, skewness = FALSE, shape = FALSE)`, *i.e.*, only volatility is time-varying. For instance, in order to specify the three GAS models: $\text{GAS-}\mathcal{SKST}$, $\text{GAS-}\mathcal{ST}$ and $\text{GAS-}\mathcal{N}$, we need to execute the following lines:

```
> library("GAS")
> GASSpec_sstd <- UniGASSpec(Dist = "sstd", GASpar = list(scale = TRUE))
> GASSpec_std <- UniGASSpec(Dist = "std", GASpar = list(scale = TRUE))
> GASSpec_norm <- UniGASSpec(Dist = "norm", GASpar = list(scale = TRUE))
```

`UniGASSpec()` delivers an object of the class `uGASSpec` which comes with several methods; see `help("UniGASSpec")`.

The `UniGASRoll()` function accepts an object of the class `uGASSpec`, `GASSpec`, a numeric vector for the series of returns, `data`, and other arguments such as:

- the length of the out-of-sample period: `ForecastLength`,
- the type of the rolling window used to update the data: `RefitWindow`,
- the number of observations within each model re-estimation: `RefitEvery`,

among others; see `help("UniGASRoll")`. `ForecastLength` and `RefitEvery` are numeric elements while `RefitWindow` is a character equal to "moving" (the default) for a fixed window scheme or "recursive" for expanding window. As previously mentioned, in this paper we consider the case `RefitWindow = "moving"`. For reproducibility purposes, the `dji30ret` data set which consists of the Dow Jones Index 30 Constituents closing value log-returns is included in the **GAS** package. For our examples, we consider the last $T = 2500$ observations of `dji30ret`:

```
> data("dji30ret", package = "GAS")
> dji30ret = tail(dji30ret, 2500)
```

We fix the length of the out-of-sample period to 1000 ($H = 1000$), and re-estimate the model parameters each time a new observation arrives. One-step ahead rolling predictions for the first series of returns using the GAS- \mathcal{N} model are then computed as:

```
> library("parallel")
> cluster = makeCluster(8)
> H = 1000
> Roll_1N = UniGASRoll(dji30ret[, 1], GASSpec_norm, ForecastLength = H,
                      RefitEvery = 1, cluster = cluster)
```

where we have also made use of parallel processing through the definition of a `cluster` object exploiting the `parallel` package included in R since version 2.14.0.

The output of `UniGASRoll()` is an object of the class `uGASRoll` which comes with several methods; see `help("UniGASRoll")`. VaR forecasts at confidence level $\alpha = 0.01$ can be computed from `Roll_1N` using the `quantile` method, *i.e.*:

```
> alpha = 0.01
> VaR_1N = quantile(Roll_1N, probs = alpha)
```

where `VaR_1N` is a matrix of dimension 1000×1 containing the VaR forecasts.³ VaR backtest procedures are implemented through the `BacktestVaR()` function. This function accepts the following arguments:

- `data`, `numeric` containing the out-of-sample data;
- `VaR`, `numeric` containing the series of VaR forecasts;
- `alpha`, the VaR confidence level α ;
- `Lags`, the number of lags used in the DQ test, by default `Lags = 4`; see Engle and Manganelli (2004).

The function returns a `list` with named entries:

- `LRuc`, the test statistic and associated p -value for the UC test of Kupiec (1995);
- `LRcc`, the test statistic and associated p -value for the CC test of Christoffersen (1998);
- `DQ`, the test statistic and associated p -value for the DQ test of Engle and Manganelli (2004);
- `Loss`, the average QL and QL series used by González-Rivera et al. (2004),
- `AD`, the mean and max VaR Absolute Deviation (AD) used by McAleer and Da Veiga (2008);
- `AE`, the AE ratio reported in (2).

For instance, in order to compute the VaR backtest measures defined above on the forecast series `VaR_1N`, we use:

```
> VaRBacktest_1N = BacktestVaR(data = tail(dji30ret[, 1], H), VaR = VaR_1N, alpha = alpha)
```

Hence, the DQ test statistic and its associated p -value can be extracted as:

```
> VaRBacktest_1N$DQ
$stat
      [,1]
[1,] 42.0586

$pvalue
      [,1]
[1,] 1.79043e-07
```

which, in this case, is against the null of correct model specification for the 1% VaR level.

Now, if we evaluate VaR forecasts using the GAS- \mathcal{ST} model, we obtain, say, the `VaRBacktest_1ST` object. For GAS- \mathcal{ST} the DQ test reports:

³The `probs` argument in `quantile` can also be a `numeric` vector of p VaR levels. In this case, `VaR_1N` would be a $1000 \times p$ matrix.

```
R> VaRBacktest_1ST$DQ
$stat
      [,1]
[1,] 7.26818

$pvalue
      [,1]
[1,] 0.2967566
```

The large p -value indicates that the null of correct model specification for the 1% VaR cannot be rejected at the usual levels of significance. Models comparison in terms of average QL also favours GAS- \mathcal{ST} . Indeed the ratio between the QL of GAS- \mathcal{ST} and GAS- \mathcal{N} :

```
> round(VaRBacktest_1ST$Loss$Loss / VaRBacktest_1N$Loss$Loss, 2)
[1] 0.89
```

This indicates that GAS- \mathcal{ST} outperforms GAS- \mathcal{N} by 11% in terms of average QL.

Empirical application: US Industrial firms

We report now a short empirical application in order to illustrate the benefits of using GAS models for VaR predictions.

We predict the one-step ahead VaR level with the GAS- \mathcal{N} , GAS- \mathcal{ST} and GAS- \mathcal{SKST} models using the same configuration of the previous section, *i.e.*, we set, $T = 2500$, $H = 1000$, and $\text{RefitEvery} = 1$. The analysis considers all the 30 constituents of the Dow Jones Index (DJI) available in the `dji30ret` data set previously detailed. The out-of-sample period starts on 14, February 2005, and includes the recent Global Financial Crisis of 2007–2008.

We consider two VaR confidence levels $\alpha = 1\%$ and $\alpha = 5\%$. The code used for this application is available in the GitHub GAS repository: <https://github.com/LeopoldoCatania/GAS>.

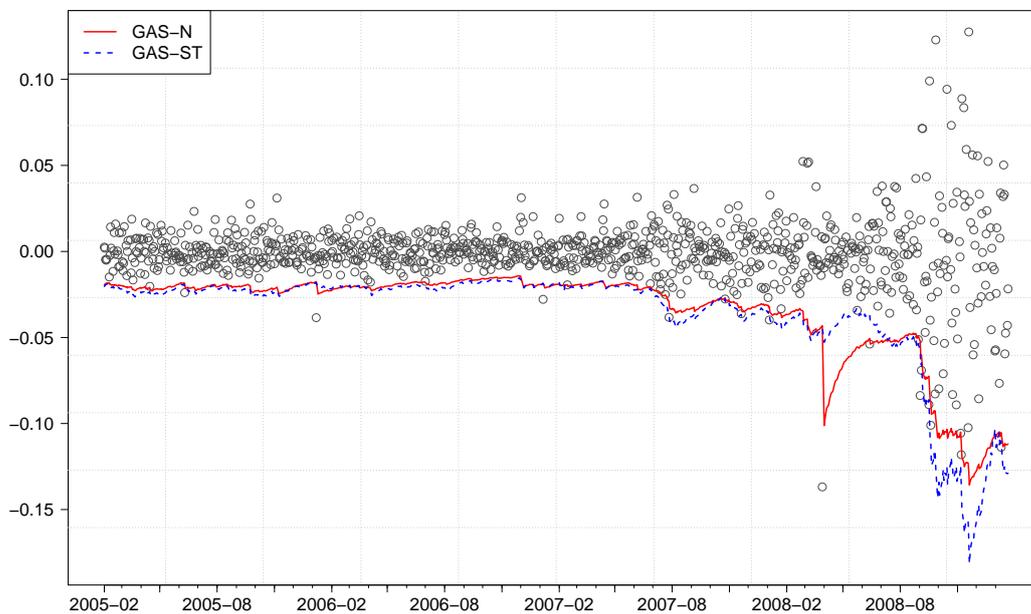


Figure 1: One-step ahead VaR forecasts for General Electric (GE) at the $\alpha = 1\%$ confidence level for the GAS- \mathcal{N} (*solid*) and GAS- \mathcal{ST} (*dotted*) models. Gray points indicate realized log-returns calculated as the differences between the logarithm of two consecutive prices. The forecast period spans from 14, February 2005, to February 3rd, 2009 for a total of 1,000 predictions.

Figure 1 depicts the one-step ahead VaR forecasts for General Electric (GE) at the $\alpha = 1\%$ confidence level for the GAS- \mathcal{N} (*solid*) and GAS- \mathcal{ST} (*dotted*) specifications. We clearly see the impact of the recent Global Financial Crisis on the volatility of the series. Indeed, the last part of

the figure presents much more variability than the initial one. What is also clearly evident from Figure 1, is the robustness of the GAS- ST model to extreme observations compared to GAS- \mathcal{N} . Indeed, on 11, April 2008, General Electric reported an unexpected net income drop of 6%, which in turn translated to a fall of about 12% of its market value. The signal capture by GAS- \mathcal{N} was that of an abrupt increase in volatility, with the consequence of large VaR level predictions. In contrast, the GAS- ST model slightly increased the volatility level and continued to predict reasonable VaR levels. What happened is that, GAS- ST treated the 12% negative return as realization from the fat-tailed Student- t distribution, hence tapering its impact on the conditional volatility level. On the contrary, GAS- \mathcal{N} treated the negative return as a realization from the Normal distribution, which is a clear signal of increase of volatility.

Table 1 reports the DQ test p -values for the three model specifications and the two VaR confidence levels.

Asset	$\alpha = 1\%$			$\alpha = 5\%$		
	GAS- \mathcal{N}	GAS- ST	GAS- $SKST$	GAS- \mathcal{N}	GAS- ST	GAS- $SKST$
AA	0.00	0.29	0.29	0.00	0.00	0.00
AIG	0.00	0.00	0.00	0.00	0.00	0.00
AXP	0.08	0.99	0.67	0.17	0.18	0.08
BA	0.21	0.92	0.20	0.09	0.36	0.40
BAC	0.00	0.00	0.08	0.00	0.02	0.06
C	0.00	0.00	0.00	0.01	0.00	0.01
CAT	0.00	0.15	0.09	0.00	0.51	0.48
CVX	0.00	0.45	0.20	0.02	0.02	0.02
DD	0.03	0.31	0.24	0.26	0.22	0.11
DIS	0.00	0.98	0.23	0.00	0.28	0.12
GE	0.00	0.21	0.25	0.17	0.33	0.03
GM	0.00	0.04	0.01	0.02	0.19	0.11
HD	0.02	0.01	0.06	0.37	0.89	0.67
HPQ	0.00	0.01	0.02	0.00	0.06	0.17
IBM	0.00	0.02	0.02	0.00	0.02	0.02
INTC	0.02	0.04	0.05	0.02	0.04	0.05
JNJ	1.00	1.00	0.96	0.01	0.00	0.00
JPM	0.24	0.93	0.27	0.29	0.42	0.17
KO	0.06	0.03	0.20	0.00	0.12	0.22
MCD	0.81	0.98	1.00	0.27	0.18	0.40
MMM	0.00	0.29	0.00	0.00	0.13	0.07
MRK	0.00	0.06	0.06	0.00	0.08	0.03
MSFT	0.07	0.20	0.31	0.30	0.36	0.44
PFE	0.13	0.28	0.22	0.09	0.78	0.76
PG	0.00	0.12	0.08	0.16	0.10	0.19
T	0.26	0.27	0.31	0.02	0.01	0.04
UTX	0.00	0.31	0.04	0.17	0.82	0.93
VZ	0.01	0.99	0.98	0.89	0.81	0.90
WMT	0.01	0.00	0.01	0.83	0.21	0.10
XOM	0.00	0.00	0.00	0.08	0.02	0.04

Table 1: DQ test statistic p -values for DJI constituents one-step ahead VaR forecasts at the two confidence levels $\alpha = 1\%$ and $\alpha = 5\%$. Under the null hypothesis we have correct model specification for the chosen quantile level. Light gray cells indicate p -values lower than 1%. The out-of-sample period spans from 14, February 2005, to February 3rd, 2009 for a total of 1,000 observations.

Under the null hypothesis we have correct model specification for the α -quantile level. Our results suggest that, GAS- ST and GAS- $SKST$ perform similar in terms of correct unconditional and conditional coverage. Hence, the inclusion of skewness does not seem to increase the performance of VaR predictions for the considered series. Indeed, sometimes results are even worse for GAS- $SKST$ compared to GAS- ST , indicating that the estimation error for the additional skewness parameter could worsen VaR predictions. This is the case for example for General Electric (GE) when GAS- $SKST$ reject the null both for $\alpha = 1\%$ and $\alpha = 5\%$, whereas GAS- ST does never reject.

GAS- \mathcal{N} is suboptimal with respect to GAS- ST and GAS- $SKST$. This result is somehow expected since GAS- ST and GAS- $SKST$ exhibit excess kurtosis and deliver more robust updates for the volatility parameter than GAS- \mathcal{N} .

Asset	$\alpha = 1\%$			$\alpha = 5\%$		
	GAS- \mathcal{N}	GAS- \mathcal{ST}	GAS- \mathcal{SKST}	GAS- \mathcal{N}	GAS- \mathcal{ST}	GAS- \mathcal{SKST}
AA	1.00	0.92	0.93	1.00	0.98	0.98
AIG	1.00	0.98	1.00	1.00	1.02	1.02
AXP	1.00	0.93	0.96	1.00	0.99	1.00
BA	1.00	0.99	0.99	1.00	0.99	0.99
BAC	1.00	0.83	0.85	1.00	1.01	1.02
C	1.00	0.93	0.94	1.00	1.01	1.01
CAT	1.00	0.84	0.85	1.00	0.94	0.94
CVX	1.00	1.00	1.00	1.00	1.01	1.01
DD	1.00	0.95	0.96	1.00	0.98	0.98
DIS	1.00	0.97	0.97	1.00	1.01	1.01
GE	1.00	0.94	0.97	1.00	0.97	0.98
GM	1.00	0.89	0.91	1.00	0.95	0.96
HD	1.00	1.02	1.02	1.00	1.00	1.00
HPQ	1.00	0.96	0.97	1.00	0.95	0.95
IBM	1.00	1.00	1.00	1.00	0.94	0.94
INTC	1.00	0.96	0.95	1.00	0.95	0.95
JNJ	1.00	0.99	1.00	1.00	0.99	1.00
JPM	1.00	0.92	0.94	1.00	0.99	0.99
KO	1.00	0.99	0.98	1.00	0.97	0.97
MCD	1.00	1.00	1.00	1.00	0.99	0.99
MMM	1.00	0.85	0.86	1.00	0.92	0.91
MRK	1.00	0.84	0.84	1.00	0.89	0.89
MSFT	1.00	0.88	0.89	1.00	0.93	0.92
PFE	1.00	0.85	0.85	1.00	0.92	0.92
PG	1.00	0.88	0.89	1.00	0.97	0.97
T	1.00	0.97	0.96	1.00	1.01	1.01
UTX	1.00	0.98	0.99	1.00	1.00	1.00
VZ	1.00	0.91	0.90	1.00	0.98	0.98
WMT	1.00	0.95	0.96	1.00	0.99	0.99
XOM	1.00	1.02	0.99	1.00	1.02	1.01

Table 2: QL ratios for GAS- \mathcal{ST} and GAS- \mathcal{SKST} with respect to GAS- \mathcal{N} for the two VaR confidence levels $\alpha = 1\%$ and $\alpha = 5\%$. Values greater than 1 indicate outperformance of GAS- \mathcal{N} and vice versa. Light gray cells indicate ratios higher than 1, *i.e.*, when the GAS- \mathcal{N} performs better. The out-of-sample period spans from 14, February 2005, to February 3rd, 2009 for a total of 1,000 observations.

Table 2 reports the ratios between the average QL of the considered models over the one delivered by GAS- \mathcal{N} for $\alpha = 1\%$ and $\alpha = 5\%$. Values greater than one indicate outperformance of GAS- \mathcal{N} , and vice versa. Consistently with the DQ test results, we find that GAS- \mathcal{ST} and GAS- \mathcal{SKST} perform similar and are preferred to GAS- \mathcal{N} . Indeed, GAS- \mathcal{ST} outperforms GAS- \mathcal{N} up to 16% for $\alpha = 1\%$ and 11% for $\alpha = 5\%$ in terms of average QL.

Conclusion

This article detailed how to use GAS models for VaR predictions in R using the GAS package of Catania et al. (2016). We briefly review the four steps that compose a typical VaR analysis which are: (i) models specifications, (ii) VaR predictions, (iii) backtesting and (iv) model comparison. We illustrate the use of the GAS package for VaR analysis using DJI constituents returns series. Our results indicate that the use of a fat-tailed conditional distribution is required by the considered series of returns.

Bibliography

- D. Ardia, K. Boudt, and L. Catania. Generalized autoregressive score models in R: The **GAS** package, Sept. 2016. Working paper. [p2, 4]
- Basel Committee. Basel III: A global regulatory framework for more resilient banks and banking systems. Technical report, Bank for International Settlements, 2010. [p1]
- L. Bauwens and S. Laurent. A new class of multivariate skew densities, with application to generalized autoregressive conditional heteroscedasticity models. *Journal of Business & Economic Statistics*, 2005. [p2]
- F. Blasques, S. J. Koopman, and A. Lucas. Maximum likelihood estimation for correctly specified generalized autoregressive score models: Feedback effects, contraction conditions and asymptotic properties. techreport TI 14-074/III, Tinbergen Institute, 2014. URL <http://www.tinbergen.nl/discussionpaper/?paper=2332>. [p3]
- F. Blasques, S. J. Koopman, K. Lasak, and A. Lucas. In-sample confidence bands and out-of-sample forecast bands for time-varying parameters in observation-driven models. *International Journal of Forecasting*, 32(3):875–887, 2016. [p3]
- T. Bollerslev. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307–327, 1986. [p1]
- L. Catania, K. Boudt, and D. Ardia. **GAS: Generalised Autoregressive Score Models**, 2016. URL <https://cran.r-project.org/package=GAS>. R package version 0.1.4. [p2, 8]
- P. F. Christoffersen. Evaluating interval forecasts. *International Economic Review*, 39(4):841–862, 1998. [p3, 5]
- D. Creal, S. J. Koopman, and A. Lucas. Generalized autoregressive score models with applications. *Journal of Applied Econometrics*, 28(5):777–795, 2013. [p1]
- R. F. Engle and S. Manganelli. CAViaR: Conditional autoregressive Value at Risk by regression quantiles. *Journal of Business & Economic Statistics*, 22(4):367–381, 2004. [p3, 5]
- C. Fernández and M. F. Steel. On Bayesian modeling of fat tails and skewness. *Journal of the American Statistical Association*, 93(441):359–371, 1998. [p2]
- G. González-Rivera, T.-H. Lee, and S. Mishra. Forecasting volatility: A reality check based on option pricing, utility function, Value-at-Risk, and predictive likelihood. *International Journal of Forecasting*, 20(4):629 – 645, 2004. [p4, 5]
- A. C. Harvey. *Dynamic Models for Volatility and Heavy Tails: With Applications to Financial and Economic Time Series*. Cambridge University Press, 2013. [p1]
- P. Jorion. *Value at risk*. McGraw-Hill, New York, 1997. [p1]
- R. Koenker and G. Bassett. Regression quantiles. *Econometrica*, 46(1):33–50, 1978. [p4]
- S. J. Koopman, A. Lucas, and M. Scharth. Predicting time-varying parameters with parameter-driven and observation-driven models. *Review of Economics and Statistics*, 98(1):97–110, 2016. [p1]
- P. H. Kupiec. Techniques for verifying the accuracy of risk measurement models. *The Journal of Derivatives*, 3(2), 1995. [p3, 5]
- S. Laurent, C. Lecourt, and F. C. Palm. Testing for jumps in conditionally Gaussian ARMA-GARCH models, a robust approach. *Computational Statistics & Data Analysis*, 100:383 – 400, 2016. [p2]
- M. Marcellino, J. H. Stock, and M. W. Watson. A comparison of direct and iterated multistep AR methods for forecasting macroeconomic time series. *Journal of Econometrics*, 135(1):499–526, 2006. [p3]
- M. McAleer and B. Da Veiga. Single-index and portfolio models for forecasting Value-at-Risk thresholds. *Journal of Forecasting*, 27(3):217–235, 2008. [p4, 5]
- A. J. McNeil, R. Frey, and P. Embrechts. *Quantitative Risk Management: Concepts, Techniques and Tools*. Princeton university press, Princeton, 2015. [p1, 2]

S. Taylor. *Modelling Financial Time Series*. Wiley, Chichester, UK, 1986. [p1]

David Ardia
Institute of Financial Analysis
University of Neuchâtel, Switzerland
and
Department of Finance, Insurance and Real Estate
Laval University, Canada
david.ardia@unine.ch

Kris Boudt
Vrije Universiteit Brussel, Belgium
and
Vrije Universiteit Amsterdam, The Netherlands
kris.boudt@vub.ac.be

Leopoldo Catania (corresponding author)
Department of Economics and Finance
University of Rome, "Tor Vergata", Italy
leopoldo.catania@uniroma2.it