

Introduction to the rugarch package. (Version 1.2-8)

Alexios Ghalanos

November 10, 2013

Contents

1	Introduction	3
2	Model Specification	3
2.1	Univariate ARFIMAX Models	4
2.2	Univariate GARCH Models	5
2.2.1	The standard GARCH model ('sGARCH')	6
2.2.2	The integrated GARCH model ('iGARCH')	7
2.2.3	The exponential GARCH model	7
2.2.4	The GJR-GARCH model ('gjrGARCH')	7
2.2.5	The asymmetric power ARCH model ('apARCH')	8
2.2.6	The family GARCH model ('fGARCH')	9
2.2.7	The Component sGARCH model ('csGARCH')	11
2.2.8	The Multiplicative Component sGARCH model ('mcsGARCH')	12
2.3	Conditional Distributions	13
2.3.1	The Normal Distribution	14
2.3.2	The Student Distribution	14
2.3.3	The Generalized Error Distribution	15
2.3.4	Skewed Distributions by Inverse Scale Factors	16
2.3.5	The Generalized Hyperbolic Distribution and Sub-Families	16
2.3.6	The Generalized Hyperbolic Skew Student Distribution	20
2.3.7	Johnson's Reparametrized SU Distribution	20
3	Fitting	20
3.1	Fit Diagnostics	23
4	Filtering	27
5	Forecasting and the GARCH Bootstrap	28
6	Simulation	31
7	Rolling Estimation	31
8	Simulated Parameter Distribution and RMSE	34
9	The ARFIMAX Model with constant variance	38

10 Mispecification and Other Tests	38
10.1 The GMM Orthogonality Test	38
10.2 Parametric and Non-Parametric Density Tests	38
10.3 Directional Accuracy Tests	40
10.4 VaR and Expected Shortfall Tests	41
11 Future Development	42
12 FAQs and Guidelines	42

1 Introduction

The pioneering work of Box et al. (1994) in the area of autoregressive moving average models paved the way for related work in the area of volatility modelling with the introduction of ARCH and then GARCH models by Engle (1982) and Bollerslev (1986), respectively. In terms of the statistical framework, these models provide motion dynamics for the dependency in the conditional time variation of the distributional parameters of the mean and variance, in an attempt to capture such phenomena as autocorrelation in returns and squared returns. Extensions to these models have included more sophisticated dynamics such as threshold models to capture the asymmetry in the news impact, as well as distributions other than the normal to account for the skewness and excess kurtosis observed in practice. In a further extension, Hansen (1994) generalized the GARCH models to capture time variation in the full density parameters, with the Autoregressive Conditional Density Model¹, relaxing the assumption that the conditional distribution of the standardized innovations is independent of the conditioning information.

The `rugarch` package aims to provide for a comprehensive set of methods for modelling univariate GARCH processes, including fitting, filtering, forecasting, simulation as well as diagnostic tools including plots and various tests. Additional methods such as rolling estimation, bootstrap forecasting and simulated parameter density to evaluate model uncertainty provide a rich environment for the modelling of these processes. This document discusses the finer details of the included models and conditional distributions and how they are implemented in the package with numerous examples.

The `rugarch` package is available on CRAN (<http://cran.r-project.org/web/packages/rugarch/index.html>) and the development version on Google code (<http://code.google.com/p/rugarch>). Some online examples and demos are available on my website (<http://www.unstarched.net>).

The package is provided AS IS, without any implied warranty as to its accuracy or suitability. A lot of time and effort has gone into the development of this package, and it is offered under the GPL-3 license in the spirit of open knowledge sharing and dissemination. If you do use the model in published work DO remember to cite the package and author (type `citation("rugarch")` for the appropriate BibTeX entry) , and if you have used it and found it useful, drop me a note and let me know.

USE THE R-SIG-FINANCE MAILING LIST FOR QUESTIONS.

A section on FAQ is included at the end of this document.

2 Model Specification

This section discusses the key step in the modelling process, namely that of the specification. This is defined via a call to the `ugarchspec` function,

```
> args(ugarchspec)
```

```
function (variance.model = list(model = "sGARCH", garchOrder = c(1,
  1), submodel = NULL, external.regressors = NULL, variance.targeting = FALSE),
  mean.model = list(armaOrder = c(1, 1), include.mean = TRUE,
    archm = FALSE, archpow = 1, arfima = FALSE, external.regressors = NULL,
    archex = FALSE), distribution.model = "norm", start.pars = list(),
  fixed.pars = list(), ...)
```

Thus a model, in the `rugarch` package, may be described by the dynamics of the conditional mean and variance, and the distribution to which they belong, which determines any additional

¹This is in the works and will be released as a companion package.

parameters. The following sub-sections will outline the background and details of the dynamics and distributions implemented in the package.

2.1 Univariate ARFIMAX Models

The univariate GARCH specification allows to define dynamics for the conditional mean from the general ARFIMAX model with the addition of ARCH-in-mean effects introduced in Engle et al. (1987). The ARFIMAX-ARCH-in-mean specification may be formally defined as,

$$\Phi(L)(1-L)^d(y_t - \mu_t) = \Theta(L)\varepsilon_t, \quad (1)$$

with the left hand side denoting the Fractional AR specification on the demeaned data and the right hand side the MA specification on the residuals. (L) is the lag operator, $(1-L)^d$ the long memory fractional process with $0 < d < 1$, and equivalent to the Hurst Exponent $H - 0.5$, and μ_t defined as,

$$\mu_t = \mu + \sum_{i=1}^{m-n} \delta_i x_{i,t} + \sum_{i=m-n+1}^m \delta_i x_{i,t} \sigma_t + \xi \sigma_t^k, \quad (2)$$

where we allow for m external regressors x of which n (last n of m) may optionally be multiplied by the conditional standard deviation σ_t , and ARCH-in-mean on either the conditional standard deviation, $k = 1$ or conditional variance $k = 2$. These options can all be passed via the arguments in the `mean.model` list in the `ugarchspec` function,

- `armaOrder` (default = (1,1). The order of the ARMA model.)
- `include.mean` (default = TRUE. Whether the mean is modelled.)
- `archm` (default = FALSE. The ARCH-in-mean parameter.)
- `archpow` (default = 1 for standard deviation, else 2 for variance.)
- `arfima` (default = FALSE. Whether to use fractional differencing.)
- `external.regressors` (default = NULL. A matrix of external regressors of the same length as the data.)
- `archex` (default = FALSE. Either FALSE or integer denoting the number of external regressors from the end of the matrix to multiply by the conditional standard deviation.).

Since the specification allows for both fixed and starting parameters to be passed, it is useful to provide the naming convention for these here,

- AR parameters are `'ar1'`, `'ar2'`, ...,
- MA parameters are `'ma1'`, `'ma2'`, ...,
- mean parameter is `'mu'`
- `archm` parameter is `'archm'`
- the `arfima` parameter is `'arfima'`
- the external regressor parameters are `'mxreg1'`, `'mxreg2'`, ...,

Note that estimation of the mean and variance equations in the maximization of the likelihood is carried out jointly in a single step. While it is perfectly possible and consistent to perform a 2-step estimation, the one step approach results in greater efficiency, particularly for smaller datasets.

2.2 Univariate GARCH Models

In GARCH models, the density function is usually written in terms of the location and scale parameters, normalized to give zero mean and unit variance,

$$\alpha_t = (\mu_t, \sigma_t, \omega), \quad (3)$$

where the conditional mean is given by

$$\mu_t = \mu(\theta, x_t) = E(y_t | x_t), \quad (4)$$

and the conditional variance is,

$$\sigma_t^2 = \sigma^2(\theta, x_t) = E((y_t - \mu_t)^2 | x_t), \quad (5)$$

with $\omega = \omega(\theta, x_t)$ denoting the remaining parameters of the distribution, perhaps a shape and skew parameter. The conditional mean and variance are used to scale the innovations,

$$z_t(\theta) = \frac{y_t - \mu(\theta, x_t)}{\sigma(\theta, x_t)}, \quad (6)$$

having conditional density which may be written as,

$$g(z|\omega) = \frac{d}{dz} P(z_t < z|\omega), \quad (7)$$

and related to $f(y|\alpha)$ by,

$$f(y_t | \mu_t, \sigma_t^2, \omega) = \frac{1}{\sigma_t} g(z_t | \omega). \quad (8)$$

The **rugarch** package implements a rich set of univariate GARCH models and allows for the inclusion of external regressors in the variance equation as well as the possibility of using variance targeting as in Engle and Mezrich (1995). These options can all be passed via the arguments in the **variance.model** list in the **ugarchspec** function,

- **model** (default = 'sGARCH' (vanilla GARCH). Valid models are 'iGARCH', 'gjrGARCH', 'eGARCH', 'apARCH', 'fGARCH', 'csGARCH' and 'mcsGARCH').
- **garchOrder** (default = c(1,1). The order of the GARCH model.)
- **submodel** (default = NULL. In the case of the 'fGARCH' omnibus model, valid choices are 'GARCH', 'TGARCH', 'GJRGARCH', 'AVGARCH', 'NGARCH', 'NAGARCH', 'APARCH' and 'ALLGARCH').
- **external.regressors** (default = NULL. A matrix of external regressors of the same length as the data).
- **variance.targeting** (default = FALSE. Whether to include variance targeting. It is also possible to pass a numeric value instead of a logical, in which case it is used for the calculation instead of the variance of the conditional mean equation residuals).

The rest of this section discusses the various flavors of GARCH implemented in the package, while Section 2.3 discusses the distributions implemented and their standardization for use in GARCH processes.

2.2.1 The standard GARCH model ('sGARCH')

The standard GARCH model (Bollerslev (1986)) may be written as:

$$\sigma_t^2 = \left(\omega + \sum_{j=1}^m \zeta_j v_{jt} \right) + \sum_{j=1}^q \alpha_j \varepsilon_{t-j}^2 + \sum_{j=1}^p \beta_j \sigma_{t-j}^2, \quad (9)$$

with σ_t^2 denoting the conditional variance, ω the intercept and ε_t^2 the residuals from the mean filtration process discussed previously. The GARCH order is defined by (q, p) (ARCH, GARCH), with possibly m external regressors v_j which are passed *pre-lagged*. If variance targeting is used, then ω is replaced by,

$$\bar{\sigma}^2 (1 - \hat{P}) - \sum_{j=1}^m \zeta_j \bar{v}_j \quad (10)$$

where $\bar{\sigma}^2$ is the unconditional variance of ε^2 which is consistently estimated by its sample counterpart at every iteration of the solver following the mean equation filtration, and \bar{v}_j represents the sample mean of the j^{th} external regressors in the variance equation (assuming stationarity), and \hat{P} is the persistence and defined below. If a numeric value was provided to the *variance.targeting* option in the specification (instead of logical), this will be used instead of $\bar{\sigma}^2$ for the calculation.² One of the key features of the observed behavior of financial data which GARCH models capture is volatility clustering which may be quantified in the persistence parameter \hat{P} . For the 'sGARCH' model this may be calculated as,

$$\hat{P} = \sum_{j=1}^q \alpha_j + \sum_{j=1}^p \beta_j. \quad (11)$$

Related to this measure is the 'half-life' (call it $h2l$) defined as the number of days it takes for half of the expected reversion back towards $E(\sigma^2)$ to occur,

$$h2l = \frac{-\log_e 2}{\log_e \hat{P}}. \quad (12)$$

Finally, the unconditional variance of the model $\hat{\sigma}^2$, and related to its persistence, is,

$$\hat{\sigma}^2 = \frac{\hat{\omega}}{1 - \hat{P}}, \quad (13)$$

where $\hat{\omega}$ is the estimated value of the intercept from the GARCH model. The naming conventions for passing fixed or starting parameters for this model are:

- ARCH(q) parameters are 'alpha1', 'alpha2', ...,
- GARCH(p) parameters are 'beta1', 'beta2', ...,
- variance intercept parameter is 'omega'
- the external regressor parameters are 'vxreg1', 'vxreg2', ...,

²Note that this should represent a value related to the variance in the plain vanilla GARCH model. In more general models such as the APARCH, this is a value related to σ^δ , which may not be obvious since δ is not known prior to estimation, and therefore care should be taken in those cases. Finally, if scaling is used in the estimation (via the *fit.control* option), this value will also be automatically scale adjusted by the routine.

2.2.2 The integrated GARCH model ('iGARCH')

The integrated GARCH model (see Engle and Bollerslev (1986)) assumes that the persistence $\hat{P} = 1$, and imposes this during the estimation procedure. Because of unit persistence, none of the other results can be calculated (i.e. unconditional variance, half life etc). The stationarity of the model has been established in the literature, but one should investigate the possibility of omitted structural breaks before adopting the iGARCH as the model of choice. The way the package enforces the sum of the ARCH and GARCH parameters to be 1, is by subtracting $1 - \sum_{i=1}^q \alpha_i - \sum_{i>1}^p \beta_i$, so that the last beta is never estimated but instead calculated.

2.2.3 The exponential GARCH model

The exponential model of Nelson (1991) is defined as,

$$\log_e (\sigma_t^2) = \left(\omega + \sum_{j=1}^m \zeta_j v_{jt} \right) + \sum_{j=1}^q (\alpha_j z_{t-j} + \gamma_j (|z_{t-j}| - E|z_{t-j}|)) + \sum_{j=1}^p \beta_j \log_e (\sigma_{t-j}^2) \quad (14)$$

where the coefficient α_j captures the sign effect and γ_j the size effect. The expected value of the absolute standardized innovation, z_t is,

$$E|z_t| = \int_{-\infty}^{\infty} |z| f(z, 0, 1, \dots) dz \quad (15)$$

The persistence \hat{P} is given by,

$$\hat{P} = \sum_{j=1}^p \beta_j. \quad (16)$$

If variance targeting is used, then ω is replaced by,

$$\log_e (\bar{\sigma}^2) (1 - \hat{P}) - \sum_{j=1}^m \zeta_j \bar{v}_j \quad (17)$$

The unconditional variance and half life follow from the persistence parameter and are calculated as in Section 2.2.1.

2.2.4 The GJR-GARCH model ('gjrGARCH')

The GJR GARCH model of Glosten et al. (1993) models positive and negative shocks on the conditional variance asymmetrically via the use of the indicator function I ,

$$\sigma_t^2 = \left(\omega + \sum_{j=1}^m \zeta_j v_{jt} \right) + \sum_{j=1}^q (\alpha_j \varepsilon_{t-j}^2 + \gamma_j I_{t-j} \varepsilon_{t-j}^2) + \sum_{j=1}^p \beta_j \sigma_{t-j}^2, \quad (18)$$

where γ_j now represents the 'leverage' term. The indicator function I takes on value of 1 for $\varepsilon \leq 0$ and 0 otherwise. Because of the presence of the indicator function, the persistence of the model now crucially depends on the asymmetry of the conditional distribution used. The persistence of the model \hat{P} is,

$$\hat{P} = \sum_{j=1}^q \alpha_j + \sum_{j=1}^p \beta_j + \sum_{j=1}^q \gamma_j \kappa, \quad (19)$$

where κ is the expected value of the standardized residuals z_t below zero (effectively the probability of being below zero),

$$\kappa = E [I_{t-j} z_{t-j}^2] = \int_{-\infty}^0 f(z, 0, 1, \dots) dz \quad (20)$$

where f is the standardized conditional density with any additional skew and shape parameters (...). In the case of symmetric distributions the value of κ is simply equal to 0.5. The variance targeting, half-life and unconditional variance follow from the persistence parameter and are calculated as in Section 2.2.1. The naming conventions for passing fixed or starting parameters for this model are:

- ARCH(q) parameters are 'alpha1', 'alpha2', ...,
- Leverage(q) parameters are 'gamma1', 'gamma2', ...,
- GARCH(p) parameters are 'beta1', 'beta2', ...,
- variance intercept parameter is 'omega'
- the external regressor parameters are 'vxreg1', 'vxreg2', ...,

Note that the Leverage parameter follows the order of the ARCH parameter.

2.2.5 The asymmetric power ARCH model ('apARCH')

The asymmetric power ARCH model of Ding et al. (1993) allows for both leverage and the Taylor effect, named after Taylor (1986) who observed that the sample autocorrelation of absolute returns was usually larger than that of squared returns.

$$\sigma_t^\delta = \left(\omega + \sum_{j=1}^m \zeta_j v_{jt} \right) + \sum_{j=1}^q \alpha_j (|\varepsilon_{t-j}| - \gamma_j \varepsilon_{t-j})^\delta + \sum_{j=1}^p \beta_j \sigma_{t-j}^\delta \quad (21)$$

where $\delta \in \mathbb{R}^+$, being a Box-Cox transformation of σ_t , and γ_j the coefficient in the leverage term. Various submodels arise from this model:

- The simple GARCH model of Bollerslev (1986) when $\delta = 2$ and $\gamma_j = 0$.
- The Absolute Value GARCH (AVGARCH) model of Taylor (1986) and Schwert (1990) when $\delta = 1$ and $\gamma_j = 0$.
- The GJR GARCH (GJRGARCH) model of Glosten et al. (1993) when $\delta = 2$.
- The Threshold GARCH (TGARCH) model of Zakoian (1994) when $\delta = 1$.
- The Nonlinear ARCH model of Higgins et al. (1992) when $\gamma_j = 0$ and $\beta_j = 0$.
- The Log ARCH model of Geweke (1986) and Pantula (1986) when $\delta \rightarrow 0$.

The persistence of the model is given by,

$$\hat{P} = \sum_{j=1}^p \beta_j + \sum_{j=1}^q \alpha_j \kappa_j \quad (22)$$

where κ_j is the expected value of the standardized residuals z_t under the Box-Cox transformation of the term which includes the leverage coefficient γ_j ,

$$\kappa_j = E(|z| - \gamma_j z)^\delta = \int_{-\infty}^{\infty} (|z| - \gamma_j z)^\delta f(z, 0, 1, \dots) dz \quad (23)$$

If variance targeting is used, then ω is replaced by,

$$\bar{\sigma}^\delta (1 - \hat{P}) - \sum_{j=1}^m \zeta_j \bar{v}_j. \quad (24)$$

Finally, the unconditional variance of the model $\hat{\sigma}^2$ is,

$$\hat{\sigma}^2 = \left(\frac{\hat{\omega}}{1 - \hat{P}} \right)^{2/\delta} \quad (25)$$

where $\hat{\omega}$ is the estimated value of the intercept from the GARCH model. The half-life follows from the persistence parameter and is calculated as in Section 2.2.1. The naming conventions for passing fixed or starting parameters for this model are:

- ARCH(q) parameters are 'alpha1', 'alpha2', ...,
- Leverage(q) parameters are 'gamma1', 'gamma2', ...,
- Power parameter is 'delta',
- GARCH(p) parameters are 'beta1', 'beta2', ...,
- variance intercept parameter is 'omega'
- the external regressor parameters are 'vxreg1', 'vxreg2', ...,

In particular, to obtain any of the submodels simply pass the appropriate parameters as fixed.

2.2.6 The family GARCH model ('fGARCH')

The family GARCH model of Hentschel (1995) is another omnibus model which subsumes some of the most popular GARCH models. It is similar to the apARCH model, but more general since it allows the decomposition of the residuals in the conditional variance equation to be driven by different powers for z_t and σ_t and also allowing for both shifts and rotations in the news impact curve, where the shift is the main source of asymmetry for small shocks while rotation drives large shocks.

$$\sigma_t^\lambda = \left(\omega + \sum_{j=1}^m \zeta_j v_{jt} \right) + \sum_{j=1}^q \alpha_j \sigma_{t-j}^\lambda (|z_{t-j} - \eta_{2j}| - \eta_{1j} (z_{t-j} - \eta_{2j}))^\delta + \sum_{j=1}^p \beta_j \sigma_{t-j}^\lambda \quad (26)$$

which is a Box-Cox transformation for the conditional standard deviation whose shape is determined by λ , and the parameter δ transforms the absolute value function which it subject to rotations and shifts through the η_{1j} and η_{2j} parameters respectively. Various submodels arise from this model, and are passed to the `ugarchspec` 'variance.model' list via the submodel option,

- The simple GARCH model of Bollerslev (1986) when $\lambda = \delta = 2$ and $\eta_{1j} = \eta_{2j} = 0$ (submodel = 'GARCH').

- The Absolute Value GARCH (AVGARCH) model of Taylor (1986) and Schwert (1990) when $\lambda = \delta = 1$ and $|\eta_{1j}| \leq 1$ (submodel = 'AVGARCH').
- The GJR GARCH (GJRGARCH) model of Glosten et al. (1993) when $\lambda = \delta = 2$ and $\eta_{2j} = 0$ (submodel = 'GJRGARCH').
- The Threshold GARCH (TGARCH) model of Zakoian (1994) when $\lambda = \delta = 1$, $\eta_{2j} = 0$ and $|\eta_{1j}| \leq 1$ (submodel = 'TGARCH').
- The Nonlinear ARCH model of Higgins et al. (1992) when $\delta = \lambda$ and $\eta_{1j} = \eta_{2j} = 0$ (submodel = 'NGARCH').
- The Nonlinear Asymmetric GARCH model of Engle and Ng (1993) when $\delta = \lambda = 2$ and $\eta_{1j} = 0$ (submodel = 'NAGARCH').
- The Asymmetric Power ARCH model of Ding et al. (1993) when $\delta = \lambda$, $\eta_{2j} = 0$ and $|\eta_{1j}| \leq 1$ (submodel = 'APARCH').
- The Exponential GARCH model of Nelson (1991) when $\delta = 1$, $\lambda = 0$ and $\eta_{2j} = 0$ (not implemented as a submodel of fGARCH).
- The Full fGARCH model of Hentschel (1995) when $\delta = \lambda$ (submodel = 'ALLGARCH').

The persistence of the model is given by,

$$\hat{P} = \sum_{j=1}^p \beta_j + \sum_{j=1}^q \alpha_j \kappa_j \quad (27)$$

where κ_j is the expected value of the standardized residuals z_t under the Box-Cox transformation of the absolute value asymmetry term,

$$\kappa_j = E(|z_{t-j} - \eta_{2j}| - \eta_{1j}(z_{t-j} - \eta_{2j}))^\delta = \int_{-\infty}^{\infty} (|z - \eta_{2j}| - \eta_{1j}(z - \eta_{2j}))^\delta f(z, 0, 1, \dots) dz \quad (28)$$

If variance targeting is used, then ω is replaced by,

$$\bar{\sigma}^\lambda (1 - \hat{P}) - \sum_{j=1}^m \zeta_j \bar{v}_j \quad (29)$$

Finally, the unconditional variance of the model $\hat{\sigma}^2$ is,

$$\hat{\sigma}^2 = \left(\frac{\hat{\omega}}{1 - \hat{P}} \right)^{2/\lambda} \quad (30)$$

where $\hat{\omega}$ is the estimated value of the intercept from the GARCH model. The half-life follows from the persistence parameter and is calculated as in Section 2.2.1. The naming conventions for passing fixed or starting parameters for this model are:

- ARCH(q) parameters are 'alpha1', 'alpha2', ...,
- Asymmetry1(q) - rotation - parameters are 'eta11', 'eta12', ...,
- Asymmetry2(q) - shift - parameters are 'eta21', 'eta22', ...,
- Asymmetry Power parameter is 'delta',

- Conditional Sigma Power parameter is 'lambda',
- GARCH(p) parameters are 'beta1', 'beta2', ...,
- variance intercept parameter is 'omega'
- the external regressor parameters are 'vxreg1', 'vxreg2', ...,

2.2.7 The Component sGARCH model ('csGARCH')

The model of Lee and Engle (1999) decomposes the conditional variance into a permanent and transitory component so as to investigate the long- and short-run movements of volatility affecting securities. Letting q_t represent the permanent component of the conditional variance, the component model can then be written as:

$$\begin{aligned}\sigma_t^2 &= q_t + \sum_{j=1}^q \alpha_j (\varepsilon_{t-j}^2 - q_{t-j}) + \sum_{j=1}^p \beta_j (\sigma_{t-j}^2 - q_{t-j}) \\ q_t &= \omega + \rho q_{t-1} + \phi (\varepsilon_{t-1}^2 - \sigma_{t-1}^2)\end{aligned}\tag{31}$$

where effectively the intercept of the GARCH model is now time-varying following first order autoregressive type dynamics. The difference between the conditional variance and its trend, $\sigma_{t-j}^2 - q_{t-j}$ is the transitory component of the conditional variance. The conditions for the non-negativity of the conditional variance are given in Lee and Engle (1999) and imposed during estimation by the stationarity option in the fit.control list of the `ugarchfit` method, and related to the stationarity conditions that the sum of the (α, β) coefficients be less than 1 and that $\rho < 1$ (effectively the persistence of the transitory and permanent components).

The multistep, $n > 1$ ahead forecast of the conditional variance proceeds as follows:

$$\begin{aligned}E_{t-1}(\sigma_{t+n}^2) &= E_{t-1}(q_{t+n}) + \sum_{j=1}^q \alpha_j (\varepsilon_{t+n-j}^2 - q_{t+n-j}) + \sum_{j=1}^p \beta_j (\sigma_{t+n-j}^2 - q_{t+n-j}) \\ E_{t-1}(\sigma_{t+n}^2) &= E_{t-1}(q_{t+n}) + \sum_{j=1}^q \alpha_j E_{t-1}[\varepsilon_{t+n-j}^2 - q_{t+n-j}] + \sum_{j=1}^p \beta_j E_{t-1}[\sigma_{t+n-j}^2 - q_{t+n-j}]\end{aligned}\tag{32}$$

However, $E_{t-1}[\varepsilon_{t+n-j}^2] = E_{t-1}[\sigma_{t+n-j}^2]$, therefore:

$$\begin{aligned}E_{t-1}(\sigma_{t+n}^2) &= E_{t-1}(q_{t+n}) + \sum_{j=1}^q \alpha_j E_{t-1}[\sigma_{t+n-j}^2 - q_{t+n-j}] + \sum_{j=1}^p \beta_j E_{t-1}[\sigma_{t+n-j}^2 - q_{t+n-j}] \\ E_{t-1}(\sigma_{t+n}^2) &= E_{t-1}(q_{t+n}) + \left[\sum_{j=1}^{\max(p,q)} (\alpha_j + \beta_j) \right]^n (\sigma_t^2 - q_t)\end{aligned}\tag{33}$$

The permanent component forecast can be represented as:

$$E_{t-1} [q_{t+n}] = \omega + \rho E_{t-1} [q_{t+n-1}] + \phi E_{t-1} [\varepsilon_{t+n-j}^2 - \sigma_{t+n-j}^2] \quad (34)$$

$$= \omega + \rho E_{t-1} [q_{t+n-1}] \quad (35)$$

$$= \omega + \rho [\omega + \rho E_{t-1} [q_{t+n-2}]] \quad (36)$$

$$= \dots \quad (37)$$

$$= (1 + \rho + \dots + \rho^{n-1}) \omega + \rho^n q_t \quad (38)$$

$$= \frac{1 - \rho^n}{1 - \rho} \omega + \rho^n q_t \quad (39)$$

$$(40)$$

As $n \rightarrow \infty$ the unconditional variance is:

$$E_{t-1} [\sigma_{t+n}^2] = E_{t-1} [q_{t+n}] = \frac{\omega}{1 - \rho} \quad (41)$$

In the **rugarch** package, the parameters ρ and ϕ are represented by η_{11} ('eta11') and η_{21} ('eta21') respectively.

2.2.8 The Multiplicative Component sGARCH model ('mcsGARCH')

A key problem with using GARCH models for intraday data is the seasonality in the absolute returns observed at the beginning and end of the trading session. For regularly sampled time intervals (1-min, 5-min etc.), a number of models have tried to either 'de-seasonalize' the residuals and then fit the GARCH model by using for instance a Flexible Fourier method as in Andersen and Bollerslev (1997) or incorporate seasonality directly into the model as in the Periodic GARCH model of Bollerslev and Ghysels (1996). A rather simple and parsimonious approach to the de-seasonalization was recently presented in Engle and Sokalska (2012) (henceforth ES2012). Consider the continuously compounded return $r_{t,i}$, where t denotes the day and i the regularly spaced time interval at which the return was calculated. Under this model, the conditional variance is a multiplicative product of daily, diurnal and stochastic (intraday) components, so that the return process may be represented as:

$$\begin{aligned} r_{t,i} &= \mu_{t,i} + \varepsilon_{t,i} \\ \varepsilon_{t,i} &= (q_{t,i} \sigma_t s_i) z_{t,i} \end{aligned} \quad (42)$$

where $q_{t,i}$ is the stochastic intraday volatility, σ_t a daily exogenously determined forecast volatility, s_i the diurnal volatility in each regularly spaced interval i , $z_{t,i}$ the i.i.d (0,1) standardized innovation which conditional follows some appropriately chosen distribution. In ES2012, the forecast volatility σ_t is derived from a multifactor risk model externally, but it is just as possible to generate such forecasts from a daily GARCH model. The seasonal (diurnal) part of the process is defined as:

$$s_i = \frac{1}{T} \sum_{t=1}^T \left(\varepsilon_{t,i}^2 / \sigma_t^2 \right). \quad (43)$$

Dividing the residuals by the diurnal and daily volatility gives the normalized residuals ($\bar{\varepsilon}$):

$$\bar{\varepsilon}_{t,i} = \varepsilon_{t,i} / (\sigma_t s_i) \quad (44)$$

which may then be used to generate the stochastic component of volatility $q_{t,i}$ which assuming a simple vanilla GARCH model has the following motion dynamics:

$$q_{t,i}^2 = \left(\omega + \sum_{j=1}^m \zeta_j v_{jt} \right) + \sum_{j=1}^p \alpha_j \bar{\varepsilon}_{t-j}^2 + \sum_{j=1}^q \beta_j q_{t-j}^2 \quad (45)$$

see Section 2.2.1 for details. In the `rugarch` package, unlike the paper of ES2012, the conditional mean and variance equations (and hence the diurnal component on the residuals from the conditional mean filtration) are estimated jointly. Furthermore, and unlike ES2012, it is possible to include ARMAX dynamics in the conditional mean, though because of the complexity of the model and its use of time indices, ARCH-m is not currently allowed, but may be included once the `xts` package is fully translated to Rcpp. Finally, as an additional point of departure from ES2012, the diurnal component calculation uses the median instead of the mean which was found to provide for a more robust alternative, particularly given the type and size of datasets typically used [changed in version 1.2-3]. As currently stands, the model has methods for estimation (`ugarchfit`), filtering (`ugarchfilter`), forecast from fit (`ugarchforecast`) but not from spec (secondary dispatch method), simulation from fit (`ugarchsim`) but not from spec (`ugarchpath`), rolling estimation (`ugarchroll`) but not the bootstrap (`ugarchboot`). Some of the plots, which depend on the `xts` package will not render nicely since `plot.xts` does not play well with intraday data. Some plots however such as `VaRplot` have been amended to properly display intraday data coming from an `xts` object, and more may be added as time allows. An extensive example of the model may be found on <http://www.unstarched.net>. The paper by ES2012 is currently freely available here: <http://jfec.oxfordjournals.org/content/10/1/54.full>.

2.3 Conditional Distributions

The `rugarch` package supports a range of univariate distributions including the Normal (`'norm'`), Generalized Error (`'ged'`), Student (`'std'`) and their skew variants (`'snorm'`, `'sged'` and `'sstd'`) based on the transformations described in Fernandez and Steel (1998) and Ferreira and Steel (2006).³ Additionally, the Generalized Hyperbolic (`'ghyp'`), Normal Inverse Gaussian (`'nig'`) and GH Skew-Student (`'ghst'`)⁴ distributions are also implemented as is Johnson's reparametrized SU (`'jsu'`) distribution⁵ The choice of distribution is entered via the `'distribution.model'` option of the `ugarchspec` method. The package also implements a set of functions to work with the parameters of these distributions. These are:

- `ddist(distribution = "norm", y, mu = 0, sigma = 1, lambda = -0.5, skew = 1, shape = 5)`. The density (d^*) function.
- `pdist(distribution = "norm", q, mu = 0, sigma = 1, lambda = -0.5, skew = 1, shape = 5)`. The distribution (p^*) function.
- `qdist(distribution = "norm", p, mu = 0, sigma = 1, lambda = -0.5, skew = 1, shape = 5)`. The quantile (q^*) function.
- `rdist(distribution = "norm", n, mu = 0, sigma = 1, lambda = -0.5, skew = 1, shape = 5)`. The sampling (q^*) function.
- `fitdist(distribution = "norm", x, control = list())`. A function for fitting data using any of the included distributions.
- `dskewness(distribution = "norm", skew = 1, shape = 5, lambda = -0.5)`. The distribution skewness (analytical where possible else by quadrature integration).
- `dkurtosis(distribution = "norm", skew = 1, shape = 5, lambda = -0.5)`. The distribution excess kurtosis (analytical where it exists else by quadrature integration).

³These were originally taken from the `fBasics` package but have been adapted and re-written in C for the likelihood estimation.

⁴Since version 1.0-8.

⁵From the `gamlss` package.

This section provides a dry but comprehensive exposition of the required standardization of these distributions for use in GARCH modelling.

The conditional distribution in GARCH processes should be self-decomposable which is a key requirement for any autoregressive type process, while possessing the linear transformation property is required to center $(x_t - \mu_t)$ and scale (ε_t/σ_t) the innovations, after which the modelling is carried out directly using the zero-mean, unit variance, distribution of the standardized variable z_t which is a scaled version of the same conditional distribution of x_t , as described in Equations 6, 7 and 8.

2.3.1 The Normal Distribution

The Normal Distribution is a spherical distribution described completely by its first two moments, the mean and variance. Formally, the random variable x is said to be normally distributed with mean μ and variance σ^2 (both of which may be time varying), with density given by,

$$f(x) = \frac{e^{\frac{-0.5(x-\mu)^2}{\sigma^2}}}{\sigma\sqrt{2\pi}}. \quad (46)$$

Following a mean filtration or whitening process, the residuals ε , standardized by σ yield the standard normal density given by,

$$f\left(\frac{x-\mu}{\sigma}\right) = \frac{1}{\sigma}f(z) = \frac{1}{\sigma}\left(\frac{e^{-0.5z^2}}{\sqrt{2\pi}}\right). \quad (47)$$

To obtain the conditional likelihood of the GARCH process at each point in time (LL_t), the conditional standard deviation σ_t from the GARCH motion dynamics, acts as a scaling factor on the density, so that:

$$LL_t(z_t; \sigma_t) = \frac{1}{\sigma_t}f(z_t) \quad (48)$$

which illustrates the importance of the scaling property. Finally, the normal distribution has zero skewness and zero excess kurtosis.

2.3.2 The Student Distribution

The GARCH-Student model was first used described in Bollerslev (1987) as an alternative to the Normal distribution for fitting the standardized innovations. It is described completely by a shape parameter ν , but for standardization we proceed by using its 3 parameter representation as follows:

$$f(x) = \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\sqrt{\beta\nu\pi}\Gamma\left(\frac{\nu}{2}\right)}\left(1 + \frac{(x-\alpha)^2}{\beta\nu}\right)^{-\left(\frac{\nu+1}{2}\right)} \quad (49)$$

where α , β , and ν are the location, scale⁶ and shape parameters respectively, and Γ is the Gamma function. Similar to the GED distribution described later, this is a unimodal and symmetric distribution where the location parameter α is the mean (and mode) of the distribution while the variance is:

$$Var(x) = \frac{\beta\nu}{(\nu-2)}. \quad (50)$$

⁶In some representations, mostly Bayesian, this is represented in its inverse form to denote the precision.

For the purposes of standardization we require that:

$$\begin{aligned} \text{Var}(x) &= \frac{\beta\nu}{(\nu-2)} = 1 \\ \therefore \beta &= \frac{\nu-2}{\nu} \end{aligned} \quad (51)$$

Substituting $\frac{(\nu-2)}{\nu}$ into 49 we obtain the standardized Student's distribution:

$$f\left(\frac{x-\mu}{\sigma}\right) = \frac{1}{\sigma} f(z) = \frac{1}{\sigma} \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\sqrt{(\nu-2)} \pi \Gamma\left(\frac{\nu}{2}\right)} \left(1 + \frac{z^2}{(\nu-2)}\right)^{-\left(\frac{\nu+1}{2}\right)}. \quad (52)$$

In terms of R's standard implementation of the Student density ('dt'), and including a scaling by the standard deviation, this can be represented as:

$$\frac{dt\left(\frac{\varepsilon_t}{\sigma\sqrt{(\nu-2)/\nu}}, \nu\right)}{\sigma\sqrt{(\nu-2)/\nu}} \quad (53)$$

The Student distribution has zero skewness and excess kurtosis equal to $6/(\nu-4)$ for $\nu > 4$.

2.3.3 The Generalized Error Distribution

The Generalized Error Distribution (GED) is a 3 parameter distribution belonging to the exponential family with conditional density given by,

$$f(x) = \frac{\kappa e^{-0.5\left|\frac{x-\alpha}{\beta}\right|^\kappa}}{2^{1+\kappa^{-1}} \beta \Gamma(\kappa^{-1})} \quad (54)$$

with α , β and κ representing the location, scale and shape parameters. Since the distribution is symmetric and unimodal the location parameter is also the mode, median and mean of the distribution (i.e. μ). By symmetry, all odd moments beyond the mean are zero. The variance and kurtosis are given by,

$$\begin{aligned} \text{Var}(x) &= \beta^2 2^{2/\kappa} \frac{\Gamma(3\kappa^{-1})}{\Gamma(\kappa^{-1})} \\ \text{Ku}(x) &= \frac{\Gamma(5\kappa^{-1}) \Gamma(\kappa^{-1})}{\Gamma(3\kappa^{-1}) \Gamma(3\kappa^{-1})} \end{aligned} \quad (55)$$

As κ decreases the density gets flatter and flatter while in the limit as $\kappa \rightarrow \infty$, the distribution tends towards the uniform. Special cases are the Normal when $\kappa = 2$, the Laplace when $\kappa = 1$. Standardization is simple and involves rescaling the density to have unit standard deviation:

$$\begin{aligned} \text{Var}(x) &= \beta^2 2^{2/\kappa} \frac{\Gamma(3\kappa^{-1})}{\Gamma(\kappa^{-1})} = 1 \\ \therefore \beta &= \sqrt{2^{-2/\kappa} \frac{\Gamma(\kappa^{-1})}{\Gamma(3\kappa^{-1})}} \end{aligned} \quad (56)$$

Finally, substituting into the scaled density of z :

$$f\left(\frac{x-\mu}{\sigma}\right) = \frac{1}{\sigma} f(z) = \frac{1}{\sigma} \frac{\kappa e^{-0.5\left|\sqrt{2^{-2/\kappa} \frac{\Gamma(\kappa^{-1})}{\Gamma(3\kappa^{-1})}} z\right|^\kappa}}{\sqrt{2^{-2/\kappa} \frac{\Gamma(\kappa^{-1})}{\Gamma(3\kappa^{-1})}} 2^{1+\kappa^{-1}} \Gamma(\kappa^{-1})} \quad (57)$$

2.3.4 Skewed Distributions by Inverse Scale Factors

Fernandez and Steel (1998) proposed introducing skewness into unimodal and symmetric distributions by introducing inverse scale factors in the positive and negative real half lines. Given a skew parameter, ξ^7 , the density of a random variable z can be represented as:

$$f(z|\xi) = \frac{2}{\xi + \xi^{-1}} [f(\xi z) H(-z) + f(\xi^{-1} z) H(z)] \quad (58)$$

where $\xi \in \mathbb{R}^+$ and $H(\cdot)$ is the Heaviside function. The absolute moments, required for deriving the central moments, are generated from the following function:

$$M_r = 2 \int_0^\infty z^r f(z) dz. \quad (59)$$

The mean and variance are then defined as:

$$\begin{aligned} E(z) &= M_1 (\xi - \xi^{-1}) \\ Var(z) &= (M_2 - M_1^2) (\xi^2 + \xi^{-2}) + 2M_1^2 - M_2 \end{aligned} \quad (60)$$

The Normal, Student and GED distributions have skew variants which have been standardized to zero mean, unit variance by making use of the moment conditions given above.

2.3.5 The Generalized Hyperbolic Distribution and Sub-Families

In distributions where the expected moments are functions of all the parameters, it is not immediately obvious how to perform such a transformation. In the case of the GHYP distribution, because of the existence of location and scale invariant parameterizations and the possibility of expressing the variance in terms of one of those parametrization, namely the (ζ, ρ) , the task of standardizing and estimating the density can be broken down to one of estimating those 2 parameters, representing a combination of shape and skewness, followed by a series of transformation steps to demean, scale and then translate the parameters into the $(\alpha, \beta, \delta, \mu)$ parametrization for which standard formulae exist for the likelihood function. The (ξ, χ) parametrization, which is a simple transformation of the (ζ, ρ) , could also be used in the first step and then transformed into the latter before proceeding further. The only difference is the kind of 'immediate' inference one can make from the different parameterizations, each providing a different direct insight into the kind of dynamics produced and their place in the overall GHYP family particularly with regards to the limit cases.

The `rugarch` package performs estimation using the (ζ, ρ) parametrization⁸, after which a series of steps transform those parameters into the $(\alpha, \beta, \delta, \mu)$ while at the same time including the necessary recursive substitution of parameters in order to standardize the resulting distribution.

Proof 1 *The Standardized Generalized Hyperbolic Distribution.* Let ε_t be a r.v. with mean (0) and variance (σ^2) distributed as $GHYP(\zeta, \rho)$, and let z be a scaled version of the r.v. ε with variance (1) and also distributed as $GHYP(\zeta, \rho)$.⁹ The density $f(\cdot)$ of z can be expressed as

$$f\left(\frac{\varepsilon_t}{\sigma}; \zeta, \rho\right) = \frac{1}{\sigma} f_t(z; \zeta, \rho) = \frac{1}{\sigma} f_t(z; \tilde{\alpha}, \tilde{\beta}, \tilde{\delta}, \tilde{\mu}), \quad (61)$$

⁷When $\xi = 1$, the distribution is symmetric.

⁸Credit is due to Diethelm Wurtz for his original implementation in the fBasics package of the transformation and standardization function.

⁹The parameters ζ and ρ do not change as a result of being location and scale invariant

where we make use of the $(\alpha, \beta, \delta, \mu)$ parametrization since we can only naturally express the density in that parametrization. The steps to transforming from the (ζ, ρ) to the $(\alpha, \beta, \delta, \mu)$ parametrization, while at the same time standardizing for zero mean and unit variance are given henceforth.

Let

$$\zeta = \delta \sqrt{\alpha^2 - \beta^2} \quad (62)$$

$$\rho = \frac{\beta}{\alpha}, \quad (63)$$

which after some substitution may be also written in terms of α and β as,

$$\alpha = \frac{\zeta}{\delta \sqrt{(1 - \rho^2)}}, \quad (64)$$

$$\beta = \alpha \rho. \quad (65)$$

For standardization we require that,

$$\begin{aligned} E(X) &= \mu + \frac{\beta \delta}{\sqrt{\alpha^2 - \beta^2}} \frac{K_{\lambda+1}(\zeta)}{K_{\lambda}(\zeta)} = \mu + \frac{\beta \delta^2}{\zeta} \frac{K_{\lambda+1}(\zeta)}{K_{\lambda}(\zeta)} = 0 \\ \therefore \mu &= -\frac{\beta \delta^2}{\zeta} \frac{K_{\lambda+1}(\zeta)}{K_{\lambda}(\zeta)} \end{aligned} \quad (66)$$

$$\begin{aligned} Var(X) &= \delta^2 \left(\frac{K_{\lambda+1}(\zeta)}{\zeta K_{\lambda}(\zeta)} + \frac{\beta^2}{\alpha^2 - \beta^2} \left(\frac{K_{\lambda+2}(\zeta)}{K_{\lambda}(\zeta)} - \left(\frac{K_{\lambda+1}(\zeta)}{K_{\lambda}(\zeta)} \right)^2 \right) \right) = 1 \\ \therefore \delta &= \left(\frac{K_{\lambda+1}(\zeta)}{\zeta K_{\lambda}(\zeta)} + \frac{\beta^2}{\alpha^2 - \beta^2} \left(\frac{K_{\lambda+2}(\zeta)}{K_{\lambda}(\zeta)} - \left(\frac{K_{\lambda+1}(\zeta)}{K_{\lambda}(\zeta)} \right)^2 \right) \right)^{-0.5} \end{aligned} \quad (67)$$

Since we can express, $\beta^2 / (\alpha^2 - \beta^2)$ as,

$$\frac{\beta^2}{\alpha^2 - \beta^2} = \frac{\alpha^2 \rho^2}{\alpha^2 - \alpha^2 \rho^2} = \frac{\alpha^2 \rho^2}{\alpha^2 (1 - \rho^2)} = \frac{\rho^2}{(1 - \rho^2)}, \quad (68)$$

then we can re-write the formula for δ in terms of the estimated parameters $\hat{\zeta}$ and $\hat{\rho}$ as,

$$\delta = \left(\frac{K_{\lambda+1}(\hat{\zeta})}{\hat{\zeta} K_{\lambda}(\hat{\zeta})} + \frac{\hat{\rho}^2}{(1 - \hat{\rho}^2)} \left(\frac{K_{\lambda+2}(\hat{\zeta})}{K_{\lambda}(\hat{\zeta})} - \left(\frac{K_{\lambda+1}(\hat{\zeta})}{K_{\lambda}(\hat{\zeta})} \right)^2 \right) \right)^{-0.5} \quad (69)$$

Transforming into the $(\tilde{\alpha}, \tilde{\beta}, \tilde{\delta}, \tilde{\mu})$ parametrization proceeds by first substituting 69 into 64 and

simplifying,

$$\begin{aligned}
\tilde{\alpha} &= \frac{\hat{\zeta} \left(\frac{K_{\lambda+1}(\hat{\zeta})}{\hat{\zeta} K_{\lambda}(\hat{\zeta})} + \frac{\hat{\rho}^2 \left(\frac{K_{\lambda+2}(\hat{\zeta})}{K_{\lambda}(\hat{\zeta})} - \frac{(K_{\lambda+1}(\hat{\zeta}))^2}{(K_{\lambda}(\hat{\zeta}))^2} \right)}{(1-\hat{\rho}^2)} \right)^{0.5}}{\sqrt{(1-\hat{\rho}^2)}} , \\
&= \frac{\left(\frac{\hat{\zeta} K_{\lambda+1}(\hat{\zeta})}{K_{\lambda}(\hat{\zeta})} + \frac{\hat{\zeta}^2 \hat{\rho}^2 \left(\frac{K_{\lambda+2}(\hat{\zeta})}{K_{\lambda}(\hat{\zeta})} - \frac{(K_{\lambda+1}(\hat{\zeta}))^2}{(K_{\lambda}(\hat{\zeta}))^2} \right)}{(1-\hat{\rho}^2)} \right)^{0.5}}{\sqrt{(1-\hat{\rho}^2)}} , \\
&= \left(\frac{\frac{\hat{\zeta} K_{\lambda+1}(\hat{\zeta})}{K_{\lambda}(\hat{\zeta})}}{(1-\hat{\rho}^2)} + \frac{\hat{\zeta}^2 \hat{\rho}^2 \left(\frac{K_{\lambda+2}(\hat{\zeta})}{K_{\lambda+1}(\hat{\zeta})} \frac{K_{\lambda+1}(\hat{\zeta})}{K_{\lambda}(\hat{\zeta})} - \frac{(K_{\lambda+1}(\hat{\zeta}))^2}{(K_{\lambda}(\hat{\zeta}))^2} \right)}{(1-\hat{\rho}^2)^2} \right)^{0.5} , \\
&= \left(\frac{\frac{\hat{\zeta} K_{\lambda+1}(\hat{\zeta})}{K_{\lambda}(\hat{\zeta})}}{(1-\hat{\rho}^2)} \left(1 + \frac{\hat{\zeta} \hat{\rho}^2 \left(\frac{K_{\lambda+2}(\hat{\zeta})}{K_{\lambda+1}(\hat{\zeta})} - \frac{K_{\lambda+1}(\hat{\zeta})}{K_{\lambda}(\hat{\zeta})} \right)}{(1-\hat{\rho}^2)} \right) \right)^{0.5} . \tag{70}
\end{aligned}$$

Finally, the rest of the parameters are derived recursively from $\tilde{\alpha}$ and the previous results,

$$\tilde{\beta} = \tilde{\alpha} \hat{\rho}, \tag{71}$$

$$\tilde{\delta} = \frac{\hat{\zeta}}{\tilde{\alpha} \sqrt{1-\hat{\rho}^2}}, \tag{72}$$


$$\tilde{\mu} = \frac{-\tilde{\beta} \tilde{\delta}^2 K_{\lambda+1}(\hat{\zeta})}{\hat{\zeta} K_{\lambda}(\hat{\zeta})}. \tag{73}$$

For the use of the (ξ, χ) parametrization in estimation, the additional preliminary steps of converting to the (ζ, ρ) are,

$$\zeta = \frac{1}{\hat{\xi}^2} - 1, \tag{74}$$

$$\rho = \frac{\hat{\chi}}{\hat{\xi}}. \tag{75}$$

Particular care should be exercised when choosing the GH distribution in GARCH models since allowing the GIG λ parameter to vary is quite troublesome in practice and may lead to identification problems since different combinations of the 2 shape (λ, ζ) and 1 skew (ρ) parameters may lead to the same or close likelihood. In addition, large sections of the likelihood surface for some combinations of the distribution parameters is quite flat. Figure 1 shows the skewness, kurtosis and 2 quantiles surfaces for different combinations of the (ρ, ζ) parameters for two popular choices of λ .



(a) $\lambda = 1$ (HYP)

(b) $\lambda = -0.5$ (NIG)

Figure 1: GH Distribution Skewness, Kurtosis and Quantile Surfaces

2.3.6 The Generalized Hyperbolic Skew Student Distribution

The GH Skew-Student distribution was popularized by Aas and Haff (2006) because of its uniqueness in the GH family in having one tail with polynomial and one with exponential behavior. This distribution is a limiting case of the GH when $\alpha \rightarrow |\beta|$ and $\lambda = -\nu/2$, where ν is the shape parameter of the Student distribution. The domain of variation of the parameters is $\beta \in \mathbb{R}$ and $\nu > 0$, but for the variance to be finite $\nu > 4$, while for the existence of skewness and kurtosis, $\nu > 6$ and $\nu > 8$ respectively. The density of the random variable x is then given by:

$$f(x) = \frac{2^{(1-\nu)/2} \delta^\nu |\beta|^{(\nu+1)/2} K_{(\nu+1)/2} \left(\sqrt{\beta^2 (\delta^2 + (x - \mu)^2)} \right) \exp(\beta(x - \mu))}{\Gamma(\nu/2) \sqrt{\pi} \left(\sqrt{\delta^2 + (x - \mu)^2} \right)^{(\nu+1)/2}} \quad (76)$$

To standardize the distribution to have zero mean and unit variance, I make use of the first two moment conditions for the distribution which are:

$$\begin{aligned} E(x) &= \mu + \frac{\beta \delta^2}{\nu - 2} \\ Var(x) &= \frac{2\beta^2 \delta^4}{(\nu - 2)^2 (\nu - 4)} + \frac{\delta^2}{\nu - 2} \end{aligned} \quad (77)$$

We require that $Var(x) = 1$, thus:

$$\delta = \left(\frac{2\bar{\beta}^2}{(\nu - 2)^2 (\nu - 4)} + \frac{1}{\nu - 2} \right)^{-1/2} \quad (78)$$

where I have made use of the 4th parametrization of the GH distribution given in Prause (1999) where $\hat{\beta} = \beta\delta$. The location parameter is then rescaled by substituting into the first moment formula δ so that it has zero mean:

$$\bar{\mu} = -\frac{\beta \delta^2}{\nu - 2} \quad (79)$$

Therefore, we model the GH Skew-Student using the location-scale invariant parametrization $(\bar{\beta}, \nu)$ and then translate the parameters into the usual GH distribution's $(\alpha, \beta, \delta, \mu)$, setting $\alpha = abs(\bar{\beta}) + 1e - 12$. As of version 1.2-8, the quantile function (via **qdist**) is calculated using the SkewHyperbolic package of Scott and Grimson using the spline method (for speed), as is the distribution function (via **pdist**).

2.3.7 Johnson's Reparametrized SU Distribution

The reparameterized Johnson SU distribution, discussed in Rigby and Stasinopoulos (2005), is a four parameter distribution denoted by $JSU(\mu, \sigma, \nu, \tau)$, with mean μ and standard deviation σ for all values of the skew and shape parameters ν and τ respectively. The implementation is taken from the GAMLSS package of Stasinopoulos et al. (2009) and the reader is referred there for further details.

3 Fitting

Once a **uGARCHspec** has been defined, the **ugarchfit** method takes the following arguments:

```
> args(ugarchfit)
```

```
function (spec, data, out.sample = 0, solver = "solnp", solver.control = list(),
        fit.control = list(stationarity = 1, fixed.se = 0, scale = 0,
        rec.init = "all"), ...)
```

The `out.sample` option controls how many data points from the end to keep for out of sample forecasting, while the `solver.control` and `fit.control` provide additional options to the fitting routine. Importantly, the *stationarity* option controls whether to impose a stationarity constraint during estimation, which is usually closely tied to the persistence of the process. The *fixed.se* controls whether, for those values which are fixed, numerical standard errors should be calculated. The *scale* option controls whether the data should be scaled prior to estimation by its standard deviation (scaling sometimes facilitates the estimation process). The option *rec.init*, introduced in version 1.0-14 allows to set the type of method for the conditional recursion initialization, with default value 'all' indicating that all the data is used to calculate the mean of the squared residuals from the conditional mean filtration. To use the first 'n' points for the calculation, a positive integer greater than or equal to one (and less than the total estimation datapoints) can instead be provided. If instead a positive numeric value less than 1 is provided, this is taken as the weighting in an exponential smoothing backcast method for calculating the initial recursion value.

Currently, 5 solvers¹⁰ are supported, with the main one being the augmented Lagrange solver `solnp` of Ye (1997) implemented in R by Ghalanos and Theussl (2011). The main functionality, namely the GARCH dynamics and conditional likelihood calculations are done in C for speed. For reference, there is a benchmark routine called `ugarchbench` which provides a comparison of `rugarch` against 2 published GARCH models with analytic standard errors, and a small scale comparison with a commercial GARCH implementation. The fitted object is of class `uGARCHfit` which can be passed to a variety of other methods such as `show` (summary), `plot`, `ugarchsim`, `ugarchforecast` etc. The following example illustrates its use, but the interested reader should consult the documentation on the methods available for the returned class.

```
> spec = ugarchspec()
> data(sp500ret)
> fit = ugarchfit(spec = spec, data = sp500ret)
> show(fit)
```

```
*-----*
*          GARCH Model Fit          *
*-----*
```

Conditional Variance Dynamics

```
-----
GARCH Model      : sGARCH(1,1)
Mean Model       : ARFIMA(1,0,1)
Distribution      : norm
```

Optimal Parameters

```
-----
      Estimate Std. Error t value Pr(>|t|)
mu      0.000521   0.000087   6.0113      0
ar1      0.872797   0.070239  12.4261      0
```

¹⁰Since version 1.0 – 8 the 'nlopt' solver of Johnson (interfaced to R by Jelmer Ypma in the 'nloptr' package) has been added, greatly expanding the range of possibilities available via its numerous subsolver options - see documentation.

ma1	-0.899910	0.062590	-14.3779	0
omega	0.000001	0.000000	5.2449	0
alpha1	0.087817	0.007715	11.3819	0
beta1	0.904886	0.008458	106.9853	0

Robust Standard Errors:

	Estimate	Std. Error	t value	Pr(> t)
mu	0.000521	0.000098	5.3010	0.000000
ar1	0.872797	0.083728	10.4243	0.000000
ma1	-0.899910	0.074752	-12.0387	0.000000
omega	0.000001	0.000001	2.1508	0.031488
alpha1	0.087817	0.029804	2.9465	0.003214
beta1	0.904886	0.029183	31.0074	0.000000

LogLikelihood : 17902.65

Information Criteria

Akaike	-6.4808
Bayes	-6.4736
Shibata	-6.4808
Hannan-Quinn	-6.4783

Q-Statistics on Standardized Residuals

	statistic	p-value
Lag[1]	5.533	0.01866
Lag[p+q+1][3]	6.047	0.01393
Lag[p+q+5][7]	8.071	0.15239
d.o.f=2		
H0 : No serial correlation		

Q-Statistics on Standardized Squared Residuals

	statistic	p-value
Lag[1]	1.092	0.2959
Lag[p+q+1][3]	1.507	0.2196
Lag[p+q+5][7]	2.499	0.7766
d.o.f=2		

ARCH LM Tests

	Statistic	DoF	P-Value
ARCH Lag[2]	1.467	2	0.4803
ARCH Lag[5]	1.709	5	0.8878
ARCH Lag[10]	3.028	10	0.9808

Nyblom stability test

Joint Statistic: 174.9123

Individual Statistics:

mu 0.2123
ar1 0.1525
ma1 0.1091
omega 21.3662
alpha1 0.1344
beta1 0.1124

Asymptotic Critical Values (10% 5% 1%)

Joint Statistic: 1.49 1.68 2.12

Individual Statistic: 0.35 0.47 0.75

Sign Bias Test

```
-----  
                t-value      prob sig  
Sign Bias      0.4084 6.830e-01  
Negative Sign Bias 2.9508 3.183e-03 ***  
Positive Sign Bias 2.4061 1.616e-02 **  
Joint Effect    28.9261 2.321e-06 ***
```

Adjusted Pearson Goodness-of-Fit Test:

```
-----  
group statistic p-value(g-1)  
1 20 181.0 1.923e-28  
2 30 189.4 1.800e-25  
3 40 219.5 5.196e-27  
4 50 229.7 3.459e-25
```

Elapsed time : 1.421875

3.1 Fit Diagnostics

The summary method for the `uGARCHfit` object provides the parameters and their standard errors (and a robust version), together with a variety of tests which can also be called individually. The robust standard errors are based on the method of White (1982) which produces asymptotically valid confidence intervals by calculating the covariance (V) of the parameters (θ) as:

$$\hat{V} = (-A)^{-1}B(-A)^{-1} \quad (80)$$

where,

$$A = L''(\hat{\theta})$$
$$B = \sum_{i=1}^n g_i(x_i|\hat{\theta})^T g_i(x_i|\hat{\theta}) \quad (81)$$

which is the Hessian and covariance of the scores at the optimum. The robust standard errors are the square roots of the diagonal of V .

The `inforcriteria` method on a fitted or filtered object returns the Akaike (AIC), Bayesian

(BIC), Hannan-Quinn (HQIC) and Shibata (SIC) information criteria to enable model selection by penalizing overfitting at different rates. Formally, they may be defined as:

$$\begin{aligned}
AIC &= \frac{-2LL}{N} + \frac{2m}{N} \\
BIC &= \frac{-2LL}{N} + \frac{m \log_e(N)}{N} \\
HQIC &= \frac{-2LL}{N} + \frac{(2m \log_e(\log_e(N)))}{N} \\
SIC &= \frac{-2LL}{N} + \log_e \left(\frac{(N + 2m)}{N} \right)
\end{aligned} \tag{82}$$

were any parameters fixed during estimation are excluded from the calculation. The Q-Statistics are the test statistics from the test of Ljung and Box (1978) on the standardized residuals with $(1, p + q + 1, p + q + 5)$ lags (where p = AR order and q = MA order) and d.o.f the number of the AR and MA parameters, and squared standardized residuals with $(1, p + q + 1, p + q + 5)$ (where q = ARCH order and p = GARCH order) lags and d.o.f the number of ARCH and GARCH parameters (q, p) . Looking at the summary report, the high p-values for the standardized squared residuals indicates that there is little chance of serial correlation at the lags tested. The evidence for the standardized residuals is not as convincing but one should consider other factors, particularly when it comes to forecasting models.¹¹ The ARCH LM test of Engle (1982) tests the presence of ARCH effects by regressing the squared residuals of a series against its own lags. Since the Null is of No ARCH effects, a high p-value, as evidenced by the summary indicates that the GARCH model used was adequate to remove any such effects present prior to fitting (i.e. it is a good idea to test the series prior to fitting a GARCH model!). The **signbias** calculates the Sign Bias Test of Engle and Ng (1993), and is also displayed in the summary. This tests the presence of leverage effects in the standardized residuals (to capture possible misspecification of the GARCH model), by regressing the squared standardized residuals on lagged negative and positive shocks as follows:

$$\hat{z}_t^2 = c_0 + c_1 I_{\hat{z}_{t-1} < 0} + c_2 I_{\hat{z}_{t-1} < 0} \hat{z}_{t-1} + c_3 I_{\hat{z}_{t-1} \geq 0} \hat{z}_{t-1} + u_t \tag{83}$$

where I is the indicator function and \hat{z}_t the estimated standardized residuals from the GARCH process. The Null Hypotheses are $H_0 : c_i = 0$ (for $i = 1, 2, 3$), and that jointly $H_0 : c_1 = c_2 = c_3 = 0$. As can be inferred from the summary of the previous fit, there is significant Negative and Positive reaction to shocks. Using instead a model such as the apARCH would likely alleviate these effects.

The **gof** calculates the chi-squared goodness of fit test, which compares the empirical distribution of the standardized residuals with the theoretical ones from the chosen density. The implementation is based on the test of Palm (1996) which adjusts the tests in the presence on non-i.i.d. observations by reclassifying the standardized residuals not according to their value (as in the standard test), but instead on their magnitude, calculating the probability of observing a value smaller than the standardized residual, which should be identically standard uniform distributed. The function must take 2 arguments, the fitted object as well as the number of bins to classify the values. In the summary to the fit, a choice of (20, 30, 40, 50) bins is used, and from the summary of the previous example it is clear that the Normal distribution does not adequately capture the empirical distribution based on this test.

The **nyblom** test calculates the parameter stability test of Nyblom (1989), as well as the joint test. Critical values against which to compare the results are displayed, but this is not available for the joint test in the case of more than 20 parameters.

¹¹The WeightedPortTest based on Fisher and Gallagher (2012), and found on CRAN implements a number of more sophisticated set of weighted Portmanteau type tests, which the interested reader might like to consult.

Finally, some informative plots can be drawn either interactively(which = 'ask'), individually (which = 1:12) else all at once (which = 'all') as in Figure 2.

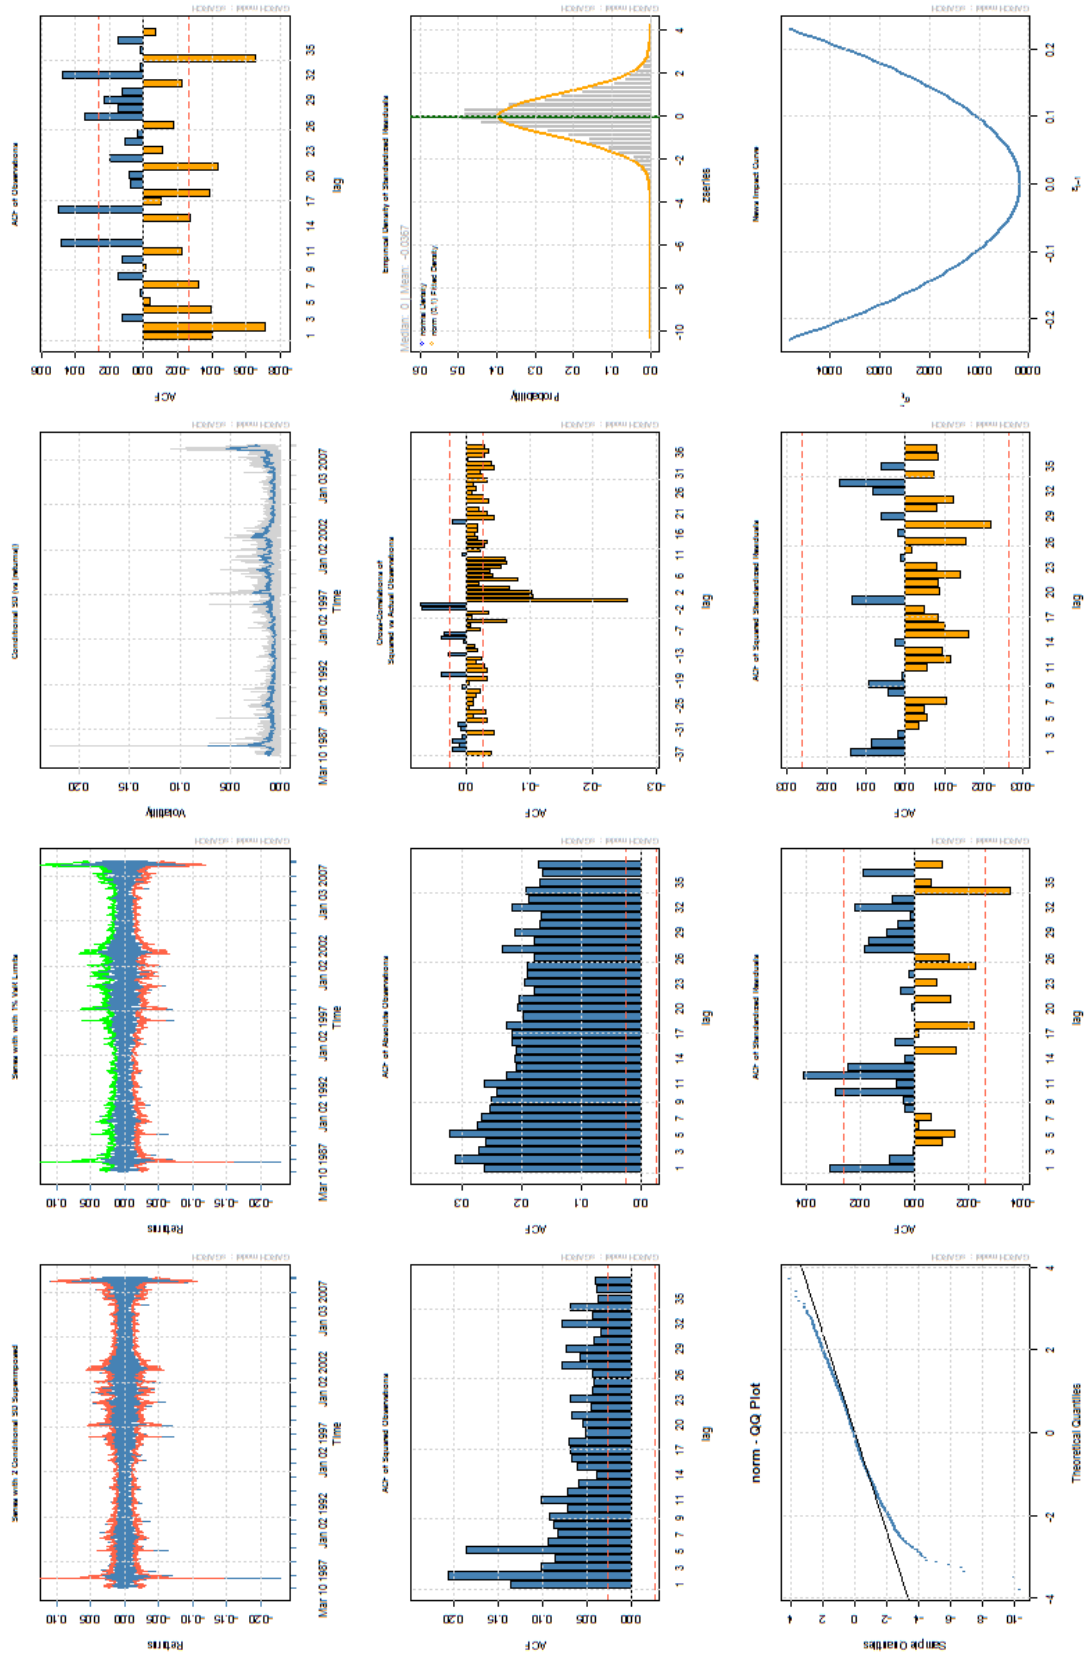


Figure 2: uGARCHfit Plots

4 Filtering

Sometimes it is desirable to simply filter a set of data with a predefined set of parameters. This may for example be the case when new data has arrived and one might not wish to re-fit. The `ugarchfilter` method does exactly that, taking a `uGARCHspec` object with fixed parameters. Setting fixed or starting parameters on the GARCH spec object may be done either through the `ugarchspec` function when it is called via the `fixed.pars` arguments to the function, else by using the `setfixed<-` method on the spec object. The example which follows explains how:

```
> data(sp500ret)
> spec = ugarchspec(variance.model = list(model = "apARCH"), distribution.model = "std")
> setfixed(spec) <- list(mu = 0.01, ma1 = 0.2, ar1 = 0.5, omega = 1e-05,
+   alpha1 = 0.03, beta1 = 0.9, gamma1 = 0.01, delta = 1, shape = 5)
> filt = ugarchfilter(spec = spec, data = sp500ret)
> show(filt)
```

```
*-----*
*           GARCH Model Filter           *
*-----*
```

Conditional Variance Dynamics

```
-----
GARCH Model      : apARCH(1,1)
Mean Model       : ARFIMA(1,0,1)
Distribution      : std
```

Filter Parameters

```
-----
mu      1e-02
ar1     5e-01
ma1     2e-01
omega   1e-05
alpha1  3e-02
beta1   9e-01
gamma1  1e-02
delta   1e+00
shape   5e+00
```

LogLikelihood : 5627.404

Information Criteria

```
-----
Akaike      -2.0378
Bayes       -2.0378
Shibata     -2.0378
Hannan-Quinn -2.0378
```

Q-Statistics on Standardized Residuals

	statistic	p-value
Lag[1]	1178	0
Lag[p+q+1][3]	1220	0
Lag[p+q+5][7]	1222	0

d.o.f=2
H0 : No serial correlation

Q-Statistics on Standardized Squared Residuals

	statistic	p-value
Lag[1]	170.4	0
Lag[p+q+1][3]	175.6	0
Lag[p+q+5][7]	178.1	0

d.o.f=2

ARCH LM Tests

	Statistic	DoF	P-Value
ARCH Lag[2]	171.5	2	0
ARCH Lag[5]	176.4	5	0
ARCH Lag[10]	177.9	10	0

Sign Bias Test

	t-value	prob	sig
Sign Bias	8.2990	1.308e-16	***
Negative Sign Bias	5.6952	1.296e-08	***
Positive Sign Bias	0.5734	5.664e-01	
Joint Effect	95.5847	1.382e-20	***

Adjusted Pearson Goodness-of-Fit Test:

group	statistic	p-value(g-1)
1	20	86296
2	30	128891
3	40	169014
4	50	207506

The returned object is of class `uGARCHfilter` and shares many of the methods as the `uGARCHfit` class. Additional arguments to the function are explained in the documentation. Note that the information criteria shown here are based on zero estimated parameters (they are all fixed), and the same goes for the *infocriteria* method on a `uGARCHfilter` object.

5 Forecasting and the GARCH Bootstrap

There are 2 types of forecasts available with the package. A rolling method, whereby consecutive 1-ahead forecasts are created based on the `out.sample` option set in the fitting routine, and an unconditional method for $n > 1$ ahead forecasts. (and it is also possible to combine the 2 creating

a rather complicated object). In the latter case, it is also possible to make use of the GARCH bootstrap, described in Pascual et al. (2006) and implemented in the function `ugarchboot`, with the added innovation of an optional extra step of fitting either a kernel or semi-parametric density (SPD) to the standardized residuals prior to sampling in order to provide for (possibly) more robustness in the presence of limited data. To understand what the GARCH bootstrap does, consider that there are two main sources of uncertainty about n -ahead forecasting from GARCH models: that arising from the form of the predictive density and that due to parameter uncertainty. The bootstrap method in the `rugarch` package is based on resampling standardized residuals from the empirical distribution of the fitted model to generate future realizations of the series and sigma. Two methods are implemented: one takes into account parameter uncertainty by building a simulated distribution of the parameters through simulation and refitting, and one which only considers distributional uncertainty and hence avoids the expensive and lengthy parameter distribution estimation. In the latter case, prediction intervals for the 1-ahead sigma forecast will not be available since only the parameter uncertainty is relevant in GARCH type models in this case. The following example provides for a brief look at the partial method, but the interested reader should consult the more comprehensive examples in the `inst` folder of the package.

```
> data(sp500ret)
> spec = ugarchspec(variance.model=list(model="csGARCH"), distribution="std")
> fit = ugarchfit(spec, sp500ret)
> bootp = ugarchboot(fit, method = c("Partial", "Full")[1],
+ n.ahead = 500, n.bootpred = 500)
> show(bootp)
```

```
*-----*
*      GARCH Bootstrap Forecast      *
*-----*
```

```
Model : csGARCH
n.ahead : 500
Bootstrap method: partial
Date (T[0]): 2009-01-30
```

Series (summary):

	min	q.25	mean	q.75	max	forecast[analytic]
t+1	-0.10855	-0.013343	0.000668	0.016285	0.090454	0.001944
t+2	-0.11365	-0.010796	0.001632	0.015721	0.085783	0.001707
t+3	-0.28139	-0.013203	-0.000378	0.015496	0.082250	0.001512
t+4	-0.10459	-0.014830	0.000346	0.015602	0.109223	0.001352
t+5	-0.21915	-0.012494	0.001196	0.016627	0.098003	0.001220
t+6	-0.11029	-0.012119	0.001008	0.015000	0.083469	0.001112
t+7	-0.22818	-0.013280	0.000398	0.015250	0.094184	0.001023
t+8	-0.25722	-0.014854	-0.001401	0.016074	0.088067	0.000949
t+9	-0.34629	-0.017681	-0.004484	0.012847	0.154058	0.000889
t+10	-0.11328	-0.013566	0.000957	0.018291	0.140734	0.000840
.....						

Sigma (summary):

	min	q0.25	mean	q0.75	max	forecast[analytic]
t+1	0.026387	0.026387	0.026387	0.026387	0.026387	0.026387
t+2	0.025518	0.025564	0.026345	0.026493	0.038492	0.026577




Figure 3: GARCH Bootstrap Forecast Plots

t+3	0.024698	0.025021	0.026332	0.026768	0.039903	0.026614
t+4	0.023925	0.024682	0.026440	0.027087	0.077525	0.026649
t+5	0.023259	0.024456	0.026474	0.027507	0.074919	0.026682
t+6	0.022622	0.024173	0.026498	0.027687	0.072876	0.026714
t+7	0.021988	0.023855	0.026463	0.027700	0.069719	0.026744
t+8	0.021503	0.023684	0.026424	0.027785	0.070750	0.026772
t+9	0.021044	0.023677	0.026690	0.028065	0.071725	0.026799
t+10	0.020560	0.023589	0.027050	0.028363	0.095243	0.026824
.....						

The 'recipe' for the full GARCH bootstrap is summarized below:

1. Extract the standardized residuals from the estimated object. If it is a specification with fixed parameters, first filter using the supplied dataset and then extract the standardized residuals from the filtered object.
2. Sample $n.bootfit$ sets of size N (original dataset less any out of sample periods) from either the raw standardized residuals, using the `spd` or kernel based methods.
3. Simulate $n.bootfit$ paths of size N , using as innovations the sampled standardized residuals. The simulation is initiated with the last values from the dataset at point N (T_0 in simulation time).
4. The $n.bootfit$ simulated series are then estimated with the same specification used by the originally supplied object in order to generate a set of coefficients representing the parameter uncertainty.
5. Filter the original dataset with the $n.bootfit$ set of estimated coefficients.
6. Use the last values of the filtered conditional sigma (and if it is the `csGARCH` model, then also the permanent component q) and residuals from the previous step to initialize a

new simulation with horizon `n.ahead` and `m.sim=n.bootpred`, using again the standardized residuals sampled as in step 2 and the new set of estimated coefficients. The simulation now contains uncertainty about the conditional `n-ahead` density as well as parameter uncertainty.

6 Simulation

Simulation may be carried out either directly on a fitted object (`ugarchsim`) else on a GARCH spec with fixed parameters (`ugarchpath`). The `ugarchsim` method takes the following arguments:

```
> args(ugarchsim)

function (fit, n.sim = 1000, n.start = 0, m.sim = 1, startMethod = c("unconditional",
  "sample"), presigma = NA, prereturns = NA, preresiduals = NA,
  rseed = NA, custom.dist = list(name = NA, distfit = NA),
  mexsimdata = NULL, vexsimdata = NULL, ...)
```

where the `n.sim` indicates the length of the simulation while `m.sim` the number of independent simulations. For reasons of speed, when `n.sim` is large relative to `m.sim`, the simulation code is executed in C, while for large `m.sim` a special purpose C++ code (using Rcpp and RcppArmadillo) is used which was found to lead to significant speed increase. Key to replicating results is the `rseed` argument which is used to pass a user seed to initialize the random number generator, else one will be assigned by the program. In any case, the returned object, of class `uGARCHsim` (or `uGARCHpath`) contains a slot with the seed(s) used.

7 Rolling Estimation

The `ugarchroll` method allows to perform a rolling estimation and forecasting of a model/dataset combination, optionally returning the VaR at specified levels. More importantly, it returns the distributional forecast parameters necessary to calculate any required measure on the forecasted density. The following example illustrates the use of the method where use is also made of the parallel functionality and run on 10 cores.¹² Figure 4 is generated by calling the plot function on the returned `uGARCHroll` object. Additional methods, and more importantly extractor functions can be found in the documentation. **Note that only `n.ahead=1` is allowed at present (more complicated rolling forecasts can be created by the user with the `ugarchfit` and `ugarchforecast` functions).** Finally, there is a new method called *resume* which allows resumption of estimation of an object which had non-converged windows, optionally supplying a different solver and solver control combination.

```
> data(sp500ret)
> library(parallel)
> cl = makePSOCKcluster(10)
> spec = ugarchspec(variance.model = list(model = "eGARCH"), distribution.model = "jsu")
> roll = ugarchroll(spec, sp500ret, n.start = 1000, refit.every = 100,
  refit.window = "moving", solver = "hybrid", calculate.VaR = TRUE,
  VaR.alpha = c(0.01, 0.05), cluster = cl, keep.coef = TRUE)
```

¹²Since version 1.0-14 the parallel functionality is based on the `parallel` package and it is upto the user to initialize a cluster object and pass it to the function, and then terminate it once it is no longer required. Eventually, this approach to the parallel usage will filter through to all the functions in `rugarch` and `rmgarch`.

```
>show(roll)
>stopCluster(cl)
```

```
*-----*
*                GARCH Roll                *
*-----*
```

```
No.Refits      : 46
Refit Horizon   : 100
No.Forecasts    : 4523
GARCH Model     : eGARCH(1,1)
Distribution     : jsu
```

Forecast Density:

	Mu	Sigma	Skew	Shape	Shape(GIG)	Realized
1991-02-21	4e-04	0.0102	-0.2586	1.5065	0	-0.0005
1991-02-22	2e-04	0.0099	-0.2586	1.5065	0	0.0019
1991-02-25	4e-04	0.0095	-0.2586	1.5065	0	0.0044
1991-02-26	3e-04	0.0093	-0.2586	1.5065	0	-0.0122
1991-02-27	1e-04	0.0101	-0.2586	1.5065	0	0.0135
1991-02-28	7e-04	0.0099	-0.2586	1.5065	0	-0.0018

.....

	Mu	Sigma	Skew	Shape	Shape(GIG)	Realized
2009-01-23	0.0015	0.0259	-0.87	2.133	0	0.0054
2009-01-26	0.0005	0.0243	-0.87	2.133	0	0.0055
2009-01-27	-0.0002	0.0228	-0.87	2.133	0	0.0109
2009-01-28	-0.0011	0.0212	-0.87	2.133	0	0.0330
2009-01-29	-0.0039	0.0191	-0.87	2.133	0	-0.0337
2009-01-30	0.0009	0.0220	-0.87	2.133	0	-0.0231

Elapsed: 12.97949 secs

```
> report(roll, type = "VaR", VaR.alpha = 0.05, conf.level = 0.95)
```

VaR Backtest Report

```
=====
```

```
Model:                eGARCH-jsu
Backtest Length:      4523
```

```
=====
```

```
alpha:                5%
Expected Exceed:      226.2
Actual VaR Exceed:    253
Actual %:              5.6%
```

Unconditional Coverage (Kupiec)

```
Null-Hypothesis:      Correct Exceedances
LR.uc Statistic:      0
LR.uc Critical:        3.841
LR.uc p-value:        1
Reject Null:          NO
```





Figure 4: eGARCH Rolling Forecast Plots

Conditional Coverage (Christoffersen)

Null-Hypothesis: Correct Exceedances and
Independence of Failures

LR.cc Statistic: 0
LR.cc Critical: 5.991
LR.cc p-value: 1
Reject Null: NO

8 Simulated Parameter Distribution and RMSE

It is sometimes instructive to be able to investigate the underlying density of the estimated parameters under different models. The `ugarchdistribution` method performs a monte carlo experiment by simulating and fitting a model multiple times and for different 'window' sizes. This allows to obtain some insight on the consistency of the parameter estimates as the data window increases by looking at the rate of decrease of the Root Mean Squared Error and whether we have \sqrt{N} consistency. This is a computationally expensive exercise and as such should only be undertaken in the presence of ample computing power and RAM. As in other functions, parallel functionality is enabled if available. The example which follows illustrates an instance of this test on one model and one set of parameters. Figures 5 and 6 complete this example.

```
> spec = ugarchspec(variance.model = list(model = "gjrGARCH"),
+   distribution.model = "ged")
> print(persistence(pars = unlist(list(mu = 0.001, ar1 = 0.4, ma1 = -0.1,
+   omega = 1e-06, alpha1 = 0.05, beta1 = 0.9, gamma1 = 0.05,
+   shape = 1.5))), distribution = "ged", model = "gjrGARCH"))

persistence
      0.975

> library(parallel)
> cl = makePSOCKcluster(10)
> setfixed(spec) <- list(mu = 0.001, ar1 = 0.4, ma1 = -0.1, omega = 1e-06,
+   alpha1 = 0.05, beta1 = 0.9, gamma1 = 0.05, shape = 1.5)
> dist = ugarchdistribution(fitORspec = spec, n.sim = 2000, n.start = 1,
+   m.sim = 100, recursive = TRUE, recursive.length = 6000, recursive.window = 1000,
+   rseed = 1066, solver = "solnp", solver.control = list(trace = 0),
+   cluster = cl)
> stopCluster(cl)
> show(dist)

*-----*
*   GARCH Parameter Distribution   *
*-----*
Model : gjrGARCH
No. Paths (m.sim) : 100
Length of Paths (n.sim) : 2000
Recursive : TRUE
Recursive Length : 6000
Recursive Window : 1000

Coefficients: True vs Simulation Mean (Window-n)
      mu      ar1      ma1      omega      alpha1      beta1      gamma1      shape
true-coef 0.00100000 0.40000 -0.100000 1.0000e-06 0.050000 0.90000 0.050000 1.5000
window-2000 0.00097099 0.39679 -0.097585 1.0411e-06 0.046161 0.90177 0.051709 1.4945
window-3000 0.00101520 0.38907 -0.089494 9.9344e-07 0.046741 0.90192 0.052902 1.4957
window-4000 0.00099876 0.39329 -0.094523 9.9556e-07 0.049385 0.90179 0.047515 1.4924
window-5000 0.00099491 0.40223 -0.103767 9.7545e-07 0.048651 0.90220 0.049442 1.4900
window-6000 0.00098876 0.39540 -0.096411 9.8975e-07 0.048753 0.90096 0.050907 1.4901
```

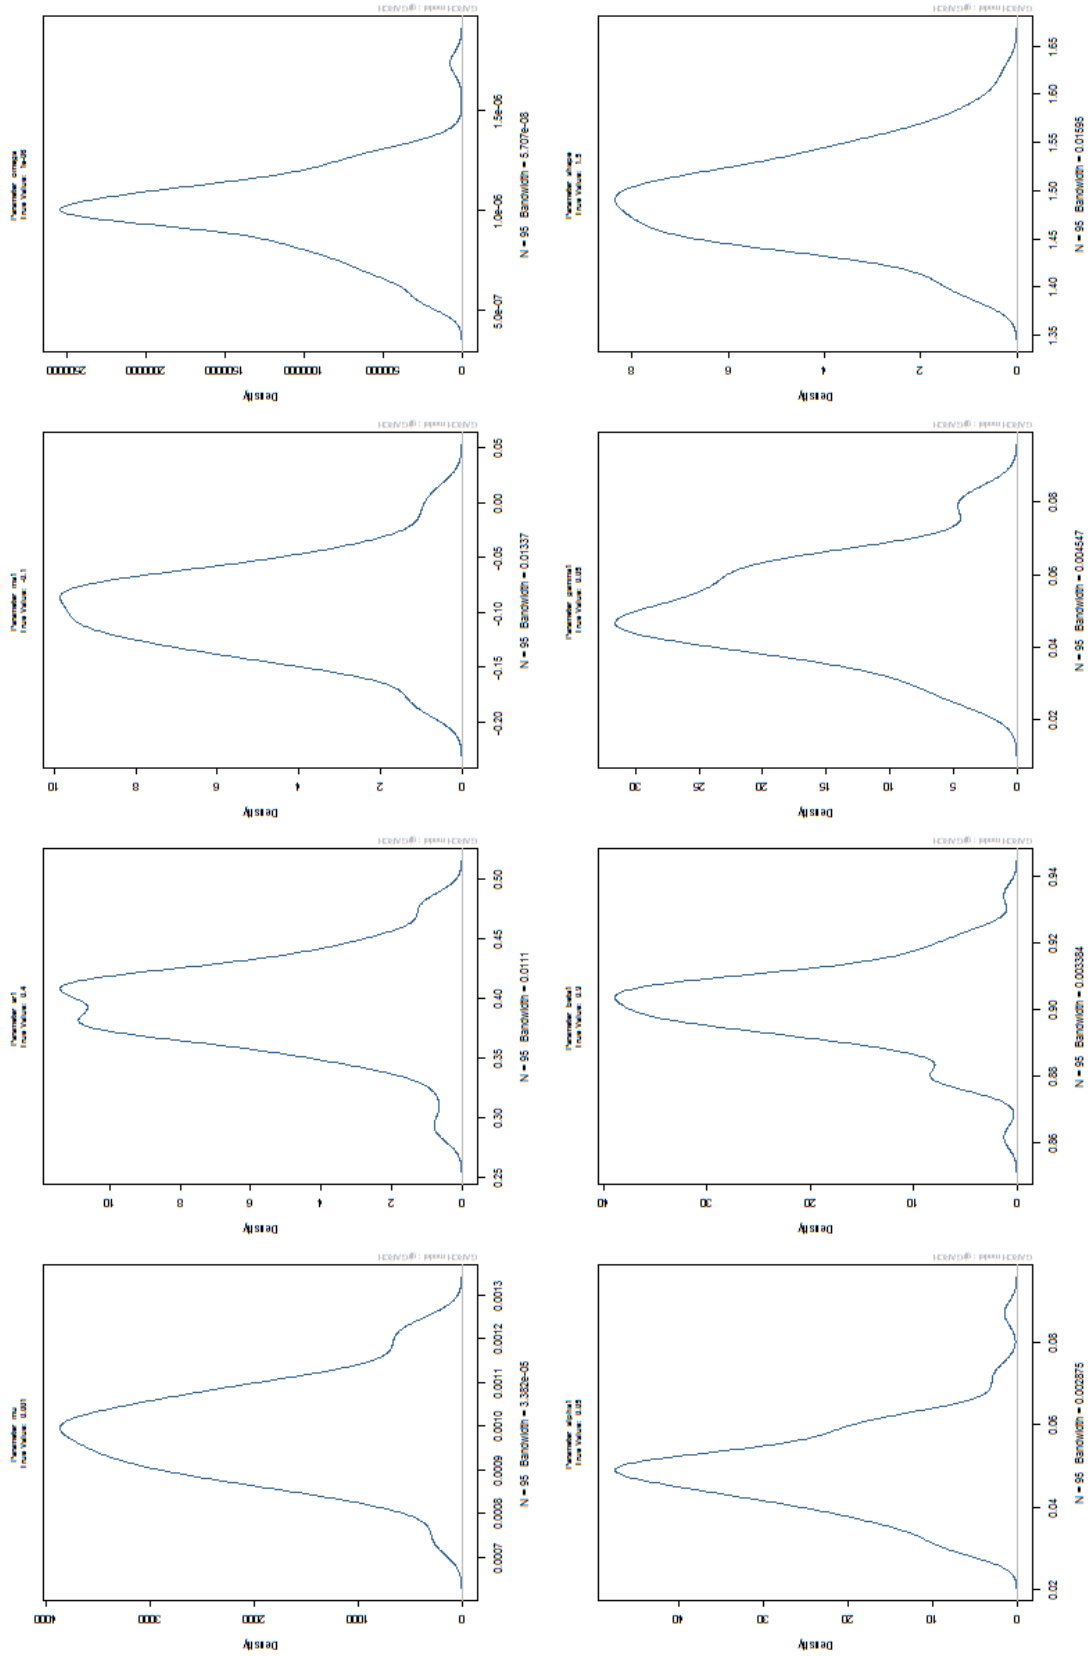


Figure 5: Simulated Parameter Density





Figure 6: RMSE Rate of Change



(b) GARCH Stat Plots

(a) Bivariate Parameter Plots

Figure 7: GARCH Simulated Parameters Density

9 The ARFIMAX Model with constant variance

The `rugarch` package implements an additional set of methods and classes, mirroring those of the GARCH specification, for modelling ARFIMAX processes with constant variance via Maximum Likelihood. With the exception of plots, the functionality is very similar to that covered so far for GARCH methods. The main functions are `arfimaspec`, `arfimafit`, `arfimaforecast`, `arfimasim`, `arfimapath`, `arfimadistribution` and `arfimaroll`. The usual extractor, inference and summary methods are replicated for all the ARFIMA classes and the user should consult the documentation for further details.

10 Misspecification and Other Tests

Apart from the usual tests presented in the summary to the fit object, a number of other interesting and useful tests are uniquely implemented in the `rugarch` package, and described in this section.

10.1 The GMM Orthogonality Test

The GMM type moment (orthogonality) tests of Hansen (1982) have been applied to test the adequacy of model in a variety of setups. Under a correctly specified model, certain population moment conditions should be satisfied and hold in the sample using the standardized residuals. The moment conditions can be tested both individually using a t-test or jointly using a Wald test. Formally, the following moment conditions are tested:

$$\begin{aligned}
 M_1 & E[z_t] = 0 \\
 M_2 & E[z_t^2 - 1] = 0 \\
 M_3 & E[z_t^3] = 0 \\
 M_4 & E[z_t^4 - 3] = 0 \\
 Q_2 & E\left[(z_t^2 - 1)\left(z_{t-j}^2 - 1\right)\right] = 0 \\
 Q_3 & E\left[z_t^3\left(z_{t-j}^3\right)\right] = 0 \\
 Q_4 & E\left[(z_t^4 - 3)\left(z_{t-j}^4 - 3\right)\right] = 0
 \end{aligned} \tag{84}$$

where $j = 1 \dots p$ is the lag (defaults to 4 in the function), M_1 to M_4 denotes the individual moment conditions (t-test), and Q_2 to Q_4 the joint conditional moment conditions (variance, skewness and kurtosis) which are distributed χ^2 with p d.o.f. All the moment conditions can also be tested jointly using a Wald test distributed χ^2 with $3p+4$ d.o.f. The test is implemented under the name: `GMMTest`.

10.2 Parametric and Non-Parametric Density Tests

A novel method to analyze how well a conditional density fits the underlying data is through the probability integral transformation (*PIT*) discussed in Rosenblatt (1952) and defined as:

$$x_t = \int_{-\infty}^{y_t} \hat{f}(u) du = \hat{F}(y_t) \tag{85}$$

which transforms the data y_t , using the estimated distribution \hat{F} into i.i.d. $U(0,1)$ under the correctly specified model. Based on this transformation, Tay et al. (1998) provide for a visual assessment test, while Berkowitz (2001) provides a more formal test, implemented in the package

under the name **BerkowitzTest**. Because of the difficulty in testing a $U(0, 1)$ sequence, the PIT data is transformed into $N(0, 1)$ by Berkowitz using the normal quantile function, and tested using a Lagrange Multiplier (*LM*) test for any residual autocorrelation given a specified number of lags. In addition, a tail test based on the censored Normal is also provided, under the Null that the standardized tail data has mean zero and unit variance. More recently, Hong and Li (2005) introduced a nonparametric portmanteau test, building on the work of Ait-Sahalia (1996), which tests the joint hypothesis of i.i.d *AND* $U(0, 1)$ for the sequence x_t . As noted by the authors, testing x_t using a standard goodness-of-fit test (such as the Kolmogorov-Smirnov) would only check the $U(0, 1)$ assumption under i.i.d. and not the joint assumption of $U(0, 1)$ and i.i.d. Their approach is to compare a kernel estimator $\hat{g}_j(x_1, x_2)$ for the joint density $g_j(x_1, x_2)$ of the pair $\{x_t, x_{t-j}\}$ (where j is the lag order) with unity, the product of two $U(0, 1)$ densities. Given a sample size n and lag order $j > 0$, their joint density estimator is:

$$\hat{g}_j(x_1, x_2) \equiv (n-j)^{-1} \sum_{t=j+1}^n K_h(x_1, \hat{X}_t) K_h(x_2, \hat{X}_{t-j}) \quad (86)$$

where $\hat{X}_t = X_t(\hat{\theta})$, and $\hat{\theta}$ is a \sqrt{n} consistent estimator of θ_0 . The function K_h is a boundary modified kernel defined as:

$$K_h(x, y) \equiv \begin{cases} h^{-1}k\left(\frac{x-y}{h}\right) / \int_{-(x/h)}^1 k(u) du, & \text{if } x \in [0, h), \\ h^{-1}k\left(\frac{x-y}{h}\right), & \text{if } x \in [h, 1-h], \\ h^{-1}k\left(\frac{x-y}{h}\right) / \int_{-1}^{(1-x)/h} k(u) du, & \text{if } x \in (1-h, 1], \end{cases} \quad (87)$$

where $h \equiv h(n)$ is a bandwidth such that $h \rightarrow 0$ as $n \rightarrow \infty$, and the kernel $k(\cdot)$ is a pre-specified symmetric probability density, which is implemented as suggested by the authors using a quartic kernel,

$$k(u) = \frac{15}{16}(1-u^2)^2 \mathbf{1}(|u| \leq 1), \quad (88)$$

where $\mathbf{1}(\cdot)$ is the indicator function. Their portmanteau test statistic is defined as:

$$\hat{W}(p) \equiv p^{-1/2} \sum_{j=1}^p \hat{Q}(j), \quad (89)$$

where

$$\hat{Q}(j) \equiv \left[(n-j)h\hat{M}(j) - A_h^0 \right] / V_0^{1/2}, \quad (90)$$

and

$$\hat{M}(j) \equiv \int_0^1 \int_0^1 [\hat{g}_j(x_1, x_2) - 1]^2 dx_1 dx_2. \quad (91)$$

The centering and scaling factors A_h^0 and V_0 are defined as:

$$\begin{aligned} A_h^0 &\equiv \left[(h^{-1} - 2) \int_{-1}^1 k^2(u) du + 2 \int_0^1 \int_{-1}^b k_b^2(u) dudb \right]^2 - 1 \\ V_0 &\equiv 2 \left[\int_{-1}^1 \left[\int_{-1}^1 k(u+v) k(v) dv \right]^2 du \right]^2 \end{aligned} \quad (92)$$

where,

$$k_b(\cdot) \equiv k(\cdot) / \int_{-1}^b k(v) dv. \quad (93)$$

Under the correct model specification, the authors show that $\hat{W}(p) \rightarrow N(0, 1)$ in distribution. Because negative values of the test statistic only occur under the Null Hypothesis of a correctly

specified model, the authors indicate that only upper tail critical values need be considered. The test is quite robust to model misspecification as parameter uncertainty has no impact on the asymptotic distribution of the test statistic as long as the parameters are \sqrt{n} consistent. Finally, in order to explore possible causes of misspecification when the statistic rejects a model, the authors develop the following test statistic:

$$M(m, l) \equiv \left[\sum_{j=1}^{n-1} w^2(j/p) (n-j) \hat{\rho}_{ml}^2(j) - \sum_{j=1}^{n-1} w^2(j/p) \right] / \left[2 \sum_{j=1}^{n-2} w^4(j/p) \right]^{1/2} \quad (94)$$

where $\hat{\rho}_{ml}(j)$ is the sample cross-correlation between \hat{X}_t^m and $\hat{X}_{t-|j|}^l$, and $w(\cdot)$ is a weighting function of lag order j , and as suggested by the authors implemented as the Bartlett kernel. As in the $\hat{W}(p)$ statistic, the asymptotic distribution of $M(m, l)$ is $N(0, 1)$ and upper critical values should be considered. As an experiment, Table 1 considers the cost of fitting a GARCH-Normal model when the true model is GARCH-Student, using the **HLTest** on simulated data using the **ugarchsim** function. The results are clear: At low levels of the shape parameter ν , representing a very high excess kurtosis, the model is overwhelmingly rejected by the test, and as that parameter increases to the point where the Student approximates the Normal, the rejections begin to reverse. Also of interest, but not surprising, the strength of the rejection is somewhat weaker for smaller datasets ($N = 500, 1000$). For example, in the case of using only 500 data points and a shape parameter of 4.1 (representing an excess kurtosis of 60!), 5% of the time, in this simulation, the test failed to reject the GARCH-Normal.

Table 1: GARCH-Student: Misspecification Exercise.

	$\nu[4.1]$	$\nu[4.5]$	$\nu[5]$	$\nu[5.5]$	$\nu[6]$	$\nu[6.5]$	$\nu[7]$	$\nu[7.5]$	$\nu[8]$	$\nu[10]$	$\nu[15]$	$\nu[20]$	$\nu[25]$	$\nu[30]$
N500														
stat	10.10	6.70	5.08	4.07	2.64	2.22	1.47	1.46	1.05	0.19	-0.34	-0.36	-0.54	-0.71
%reject	95	89	82	76	59	54	42	41	34	22	12	13	6	8
N1000														
stat	18.54	13.46	9.46	7.64	6.16	5.14	4.17	2.95	3.03	1.31	0.28	-0.15	-0.48	-0.47
%reject	100	100	98	97	90	86	79	64	69	39	24	11	7	12
N2000														
stat	32.99	26.46	19.41	15.53	12.41	10.35	7.76	6.79	5.79	3.20	0.87	0.09	0.03	-0.21
%reject	100	100	100	100	100	99	95	94	92	71	32	22	22	16
N3000														
stat	47.87	37.03	27.38	21.67	17.85	14.22	11.46	9.73	7.99	5.12	1.60	0.35	0.10	-0.09
%reject	100	100	100	100	100	100	100	99	96	85	46	27	22	15

Note: The table presents the average test statistic of Hong and Li (2005) and number of rejections at the 95% confidence level for fitting a GARCH(1,1)-Normal model to a GARCH(1,1)-Student model for different values of the shape parameter ν , and sample size (N). For each sample of size N , 250 simulated series were created from a GARCH student model with parameters $(\mu, \omega, \alpha, \beta) = (5.2334e - 04, 4.3655e - 06, 5.898e - 02, 9.2348e - 01)$, and ν in the range of $[4.1, 30]$, and fitted using a GARCH(1,1)-Normal model. The standardized residuals of the fitted model were then transformed via the normal distribution function into $U(0, 1)$ series and evaluated using the test of Hong and Li (2005).

10.3 Directional Accuracy Tests

High and significant Directional Accuracy (*DA*) could imply either an ability to predict the sign of the mean forecast or could merely be the result of volatility dependence in the absence of mean predictability as argued by Christoffersen and Diebold (2006). In either case, the function **DACTest** provides 2 tests for determining the significance of sign predictability and mean predictability. The Directional Accuracy (*DA*) Test of Pesaran and Timmermann (1992) and the Excess Profitability (*EP*) Test of Anatolyev and Gerko (2005), both of which are Hausmann type tests. The EP test statistic is formally defined as:

$$EP = \frac{A_T - B_T}{\sqrt{\hat{V}_{EP}}} \quad (95)$$

with,

$$\begin{aligned} A_T &= \frac{1}{T} \sum_t r_t \\ B_T &= \left(\frac{1}{T} \sum_t \text{sgn}(\hat{y}_t) \right) \left(\frac{1}{T} \sum_t y_t \right) \end{aligned} \quad (96)$$

with \hat{y} being the forecast of y and $r_t = \text{sgn}(\hat{y}_t)(y_t)$. According to the authors of the test, the estimated variance of EP , \hat{V}_{EP} may be estimated as:

$$\hat{V}_{EP} = \frac{4}{T^2} \hat{p}_{\hat{y}} (1 - \hat{p}_{\hat{y}}) \sum (y_t - \bar{y})^2 \quad (97)$$

where $\hat{p}_{\hat{y}} = \frac{1}{2} \left(1 + \frac{1}{T} \sum_t \text{sgn}(\hat{y}_t) \right)$. The EP statistic is asymptotically distributed as $N(0, 1)$.

For the DA test the interested reader can consult the relevant literature for more details.

10.4 VaR and Expected Shortfall Tests

The unconditional coverage, or proportion of failures, test of Kupiec (1995) allows to test whether the observed frequency of VaR exceedances is consistent with the expected exceedances, given the chosen quantile and a confidence level. Under the Null hypothesis of a correctly specified model, the number of exceedances X follows a binomial distribution. A probability below a given significance level leads to a rejection of the Null hypothesis. The test is usually conducted as a likelihood ratio test, with the statistic taking the form,

$$LR_{uc} = -2 \ln \left(\frac{(1-p)^{N-X} p^X}{\left(1 - \frac{X}{N}\right)^{N-X} \left(\frac{X}{N}\right)^X} \right) \quad (98)$$

where p is the probability of an exceedance for the chosen confidence level and N is the sample size. Under the Null the test statistic is asymptotically distributed as a χ^2 with 1 degree of freedom. The test does not consider any potential violation of the assumption of the independence of the number of exceedances. The conditional coverage test of Christoffersen et al. (2001) corrects this by jointly testing the frequency as well as the independence of exceedances, assuming that the VaR violation is modelled with a first order Markov chain. The test is a likelihood ratio, asymptotically distributed as χ^2 with 2 degrees of freedom, where the Null is that the conditional and unconditional coverage are equal to α . The test is implemented under the name **VaRTest**.

In a further paper, Christoffersen and Pelletier (2004) considers the duration between VaR violations as a stronger test of the adequacy of a risk model. The duration of time between VaR violations (no-hits) should ideally be independent and not cluster. Under the Null hypothesis of a correctly specified risk model, the no-hit duration should have no memory. Since the only continuous distribution which is memory free is the exponential, the test can be conducted on any distribution which embeds the exponential as a restricted case, and a likelihood ratio test then conducted to see whether the restriction holds. Following Christoffersen and Pelletier (2004), the Weibull distribution is used with parameter $b = 1$ representing the case of the exponential. The test is implemented under the name **VaRDurTest**.

Because VaR tests deal with the occurrences of hits, they are by definition rather crude measures to compare how well one model has done versus another, particularly with short data sets. The expected shortfall test of McNeil and Frey (2000) measures the mean of the shortfall violations which should be zero under the Null of a correctly specified risk model. The test is implemented in the function **ESTest** which also provides for the option of bootstrapping the distribution of

the p-value, hence avoiding any strong assumptions about the underlying distribution of the excess shortfall residuals.

Finally, it is understood that these tests are applied to out-of-sample forecasts and NOT insample, for which no correction to the tests have been made to account for parameter uncertainty.

11 Future Development

Any future extensions will likely be 'add-on' packages released in the google code repository of the package.

12 FAQs and Guidelines

This section provides for answers to some Frequently Asked Questions (Q) as well as Guidelines (G) for the use of the `rugarch` package.

Q: What is the format of the data accepted by the package?

Since version 1.01-3, only xts data is supported, or data which can be coerced to this. This is meant to simplify maintenance of the package whilst at the same time use what is a very popular and widely adopted 'format/wrapper'. Some of the extractor functions will now also return an xts formatted object.

Q: Where can I found out about changes to the package?

Read the changelog for version specific changes.

Q: Does the package support parallel computation?

Yes. Since version 1.0-14, `rugarch` makes exclusive use of the `parallel` package for all parallel computations. Certain functions take as input a user supplied cluster object (created by calling `parallel::makeCluster`), which is then used for parallel computations. It is then up to the user to terminate that cluster once it is no longer needed. Allowing a cluster object to be provided in this way was deemed the most flexible approach to the parallel computation problem across different architectures and resources.

Q: My model does not converge, what can I do?

There are several avenues to consider here. The package offers 4 different solvers, namely 'solnp', 'gosolnp', 'nlminb' and 'L-BGFS-U' (from optim). Each solver has its own merits, and control parameters which may, and should be passed, via the solver.control list in the fitting routines, depending on your particular data. For problems where neither 'solnp' nor 'nlminb' seem to work, try the 'gosolnp' solver which does a search of the parameter space based on a truncated normal distribution for the parameters and then initializes multiple restarts of the 'solnp' solver based on the best identified candidates. The numbers of randomly generated parameters (n.sim) and solver restarts (n.restarts) can be passed via the solver.control list. Additionally, in the fit.control list of the fitting routines, the option to perform scaling of the data prior to fitting usually helps, although it is not available under some setups. Finally, consider the amount of data you are using for modelling GARCH processes, which leads to another FAQ below.

Q: How much data should I use to model GARCH processes with confidence?

The distribution of the parameters varies by model, and is left to the reader to consult relevant literature on this. However, using 100 data points to try and fit a model is unlikely to be a sound approach as you are unlikely to get very efficient parameter estimates. The `rugarch` package does provide a method (`ugarchdistribution`) for simulating from a pre-specified model, data of different sizes, fitting the model to the data, and inferring the distribution of the parameters as well as the RMSE rate of change as the data length increases. This is a very computationally expensive way to examine the distribution of the parameters (but the only way in the non-Bayesian world), and as such should be used with care and in the presence of ample computing power.

Q: Where can one find more examples?

The package has a folder called `'rugarch.tests'` which contains many tests which I use for debugging and checking. The files in the folder should be `'sourced'` by the user, and the `'runtests.R'` file contains some wrapper functions which describe what each test does, and optionally runs chosen tests. The output will be a combination of text files (`.txt`) and figures (either `.eps` or `.png`) in an output directory which the user can define in the arguments to the wrapper function `'rugarch.runtests'`. It is quite instructive to read and understand what each test is doing prior to running it. There are also online examples which you can find by typing `rugarch` in a search engine.

Q: What to do if I find an error or have questions related to the package?

Please use the R-SIG-Finance mailing list to post your questions. If you do mail me directly, do consider carefully your email, debug information you submit, and correct email etiquette (i.e. do not send me a 1 MB `.csv` file of your data and at no time send me an Excel file).

References

- K. Aas and I.H. Haff. The generalized hyperbolic skew student's t-distribution. *Journal of Financial Econometrics*, 4(2):275–309, 2006.
- Y. Ait-Sahalia. Testing continuous-time models of the spot interest rate. *Review of Financial Studies*, 9(2):385–426, 1996.
- S. Anatolyev and A. Gerko. A trading approach to testing for predictability. *Journal of Business and Economic Statistics*, 23(4):455–461, 2005.
- Torben G Andersen and Tim Bollerslev. Intraday periodicity and volatility persistence in financial markets. *Journal of Empirical Finance*, 4(2):115–158, 1997.
- J. Berkowitz. Testing density forecasts, with applications to risk management. *Journal of Business and Economic Statistics*, 19(4):465–474, 2001.
- T. Bollerslev. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31:307–327, 1986.
- T. Bollerslev. A conditionally heteroskedastic time series model for speculative prices and rates of return. *The Review of Economics and Statistics*, 69(3):542–547, 1987.
- Tim Bollerslev and Eric Ghysels. Periodic autoregressive conditional heteroscedasticity. *Journal of Business & Economic Statistics*, 14(2):139–151, 1996.
- G.E.P. Box, G.M. Jenkins, and G.C. Reinsel. *Time series analysis: Forecasting and control*. Prentice Hall, 1994.
- P. Christoffersen and D. Pelletier. Backtesting value-at-risk: A duration-based approach. *Journal of Financial Econometrics*, 2(1):84–108, 2004.
- P. Christoffersen, J. Hahn, and A. Inoue. Testing and comparing value-at-risk measures. *Journal of Empirical Finance*, 8(3):325–342, 2001.
- P.F. Christoffersen and F.X. Diebold. Financial asset returns, direction-of-change forecasting, and volatility dynamics. *Management Science*, 52(8):1273–1287, 2006.
- Z. Ding, C.W.J. Granger, and R.F. Engle. A long memory property of stock market returns and a new model. *Journal of Empirical Finance*, 1(1):83–106, 1993.
- R.F. Engle. Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica*, 50(4):987–1007, 1982.
- R.F. Engle and T. Bollerslev. Modelling the persistence of conditional variances. *Econometric Reviews*, 5(1):1–50, 1986.
- R.F. Engle and J. Mezrich. Grappling with garch. *Risk*, 8(9):112–117, 1995.
- R.F. Engle and V.K. Ng. Measuring and testing the impact of news on volatility. *Journal of Finance*, 48(5):1749–1778, 1993.
- R.F. Engle, D.M. Lilien, and R.P. Robins. Estimating time varying risk premia in the term structure: The arch-m model. *Econometrica: Journal of the Econometric Society*, 55(2):391–407, 1987.

- Robert F. Engle and Magdalena E. Sokalska. Forecasting intraday volatility in the us equity market. multiplicative component garch. *Journal of Financial Econometrics*, 10(1):54–83, 2012. doi: 10.1093/jjfinec/nbr005. URL <http://jfec.oxfordjournals.org/content/10/1/54.abstract>.
- C. Fernandez and M.F. Steel. On bayesian modeling of fat tails and skewness. *Journal of the American Statistical Association*, 93(441):359–371, 1998.
- J.T. Ferreira and M.F. Steel. A constructive representation of univariate skewed distributions. *Journal of the American Statistical Association*, 101(474):823–829, 2006.
- T.J. Fisher and C.M. Gallagher. New weighted portmanteau statistics for time series goodness of fit testing. *Journal of the American Statistical Association*, 107(498):777–787, 2012.
- J.Y.P. Geweke. Modelling the persistence of conditional variances: A comment. *Econometric Reviews*, 5:57–61, 1986.
- A. Ghalanos and S. Theussl. *Rsolnp: General non-linear optimization using augmented Lagrange multiplier method.*, 1.11 edition, 2011.
- L.R. Glosten, R. Jagannathan, and D.E. Runkle. On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance*, 48(5):1779–1801, 1993.
- B.E. Hansen. Autoregressive conditional density estimation. *International Economic Review*, 35:705–730, 1994.
- L.P. Hansen. Large sample properties of generalized method of moments estimators. *Econometrica*, 50(4):1029–1054, 1982.
- L. Hentschel. All in the family nesting symmetric and asymmetric garch models. *Journal of Financial Economics*, 39(1):71–104, 1995.
- M.L. Higgins, A.K. Bera, et al. A class of nonlinear arch models. *International Economic Review*, 33(1):137–158, 1992.
- Y. Hong and H. Li. Nonparametric specification testing for continuous-time models with applications to term structure of interest rates. *Review of Financial Studies*, 18(1):37–84, 2005.
- P.H. Kupiec. Techniques for verifying the accuracy of risk measurement models. *The Journal of Derivatives*, 3(2):73–84, 1995. ISSN 1074-1240.
- G.J. Lee and R.F. Engle. A permanent and transitory component model of stock return volatility. In *Cointegration Causality and Forecasting A Festschrift in Honor of Clive WJ Granger*, pages 475–497. Oxford University Press, 1999.
- G.M. Ljung and G.E.P. Box. On a measure of lack of fit in time series models. *Biometrika*, 65(2):297–303, 1978.
- A.J. McNeil and R. Frey. Estimation of tail-related risk measures for heteroscedastic financial time series: an extreme value approach. *Journal of Empirical Finance*, 7(3-4):271–300, 2000.
- D.B. Nelson. Conditional heteroskedasticity in asset returns: A new approach. *Econometrica*, 59(2):347–70, 1991.
- J. Nyblom. Testing for the constancy of parameters over time. *Journal of the American Statistical Association*, 84(405):223–230, 1989.

- F.C. Palm. Garch models of volatility. *Handbook of Statistics*, 14:209–240, 1996.
- S.G. Pantula. Comment: Modelling the persistence of conditional variances. *Econometric Reviews*, 5(1):71–74, 1986.
- L. Pascual, J. Romo, and E. Ruiz. Bootstrap prediction for returns and volatilities in garch models. *Computational Statistics and Data Analysis*, 50(9):2293–2312, 2006.
- M.H. Pesaran and A. Timmermann. A simple nonparametric test of predictive performance. *Journal of Business and Economic Statistics*, 10(4):461–465, 1992.
- K. Prause. *The generalized hyperbolic model: Estimation, financial derivatives, and risk measures*. PhD thesis, University of Freiburg, 1999.
- R.A. Rigby and D.M. Stasinopoulos. Generalized additive models for location, scale and shape. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 54(3):507–554, 2005.
- M. Rosenblatt. Remarks on a multivariate transformation. *The Annals of Mathematical Statistics*, 23(3):470–472, 1952.
- G.W. Schwert. Stock volatility and the crash of ’87. *Review of Financial Studies*, 3(1):77, 1990.
- David Scott and Fiona Grimson. *SkewHyperbolic: The Skew Hyperbolic Student t-Distribution*, version 0.3-3 edition. URL <http://cran.r-project.org/web/packages/SkewHyperbolic/index.html>.
- D.M. Stasinopoulos, B.A. Rigby, and C. Akantziliotou. *gamlss: Generalized additive models for location, scale and shape.*, 1.11 edition, 2009.
- A. Tay, F.X. Diebold, and T.A. Gunther. Evaluating density forecasts: With applications to financial risk management. *International Economic Review*, 39(4):863–883, 1998.
- S.J. Taylor. *Modelling financial time series*. Wiley, 1986.
- H. White. Maximum likelihood estimation of misspecified models. *Econometrica: Journal of the Econometric Society*, 50(1):1–25, 1982.
- Y. Ye. *Interior point algorithms: Theory and analysis*, volume 44. Wiley-Interscience, 1997.
- J.M. Zakoian. Threshold heteroskedastic models. *Journal of Economic Dynamics and Control*, 18(5):931–955, 1994.