



Munich Personal RePEc Archive

The Log-GARCH Model via ARMA Representations

Sucarrat, Genaro

BI Norwegian Business School

30 August 2018

Online at <https://mpra.ub.uni-muenchen.de/100386/>
MPRA Paper No. 100386, posted 15 May 2020 05:17 UTC

The Log-GARCH Model via ARMA Representations*

Genaro Sucarrat[†]

This version: 30th August 2018

[First version: 4th December 2017]

Abstract

The log-GARCH model provides a flexible framework for the modelling of economic uncertainty, financial volatility and other positively valued variables. Its exponential specification ensures fitted volatilities are positive, allows for flexible dynamics, simplifies inference when parameters are equal to zero under the null, and the log-transform makes the model robust to jumps or outliers. An additional advantage is that the model admits ARMA-like representations. This means log-GARCH models can readily be estimated by means of widely available software, and enables a vast range of well-known time-series results and methods. This chapter provides an overview of the log-GARCH model and its ARMA representation(s), and of how estimation can be implemented in practice. After the introduction, we delineate the univariate log-GARCH model with volatility asymmetry (“leverage”), and show how its (nonlinear) ARMA representation is obtained. Next, stationary covariates (“X”) are added, before a first-order specification with asymmetry is illustrated empirically. Then we turn our attention to multivariate log-GARCH-X models. We start by presenting the multivariate specification in its general form, but quickly turn our focus to specifications that can be estimated equation-by-equation – even in the presence of Dynamic Conditional Correlations (DCCs) of unknown form. Next, a multivariate non-stationary log-GARCH-X model is formulated, in which the X-covariates can be both stationary and/or nonstationary. A common critique directed towards the log-GARCH model is that its ARCH terms may not exist in the presence of inliers. An own Section is devoted to how this can be handled in practice. Next, the generalisation of log-GARCH models to logarithmic Multiplicative Error Models (MEMs) is made explicit. Finally, the chapter concludes.

JEL Classification: C22, C32, C51, C58

Keywords: Financial return, volatility, ARCH, exponential GARCH, log-GARCH, Multivariate GARCH

*I am grateful to the Bilel, a reviewer and Hamdi Raissi for useful comments, suggestions and questions. All errors are mine.

[†]Department of Economics, BI Norwegian Business School, Nydalsveien 37, 0484 Oslo, Norway. Email genaro.sucarrat@bi.no, phone +47+46410779, fax +47+23264788. Webpage: <http://www.sucarrat.net/>

Contents:

1	Introduction	2
2	Univariate log-GARCH models	5
2.1	The asymmetric log-GARCH	5
2.2	The ARMA representation	6
2.3	Adding stationary covariates (“X”)	7
2.4	Estimation of the coefficient covariance matrix	7
2.5	Empirical examples	9
3	Multivariate log-GARCH models	10
3.1	A multivariate asymmetric log-GARCH-X model	10
3.2	Equation-by-equation estimation	11
3.3	Non-stationary models	12
3.4	Dynamic Conditional Correlations (DCCs)	15
4	Handling zeros in practice	15
5	Modelling positively valued variables	16
6	Conclusions	17
	References	17

1 Introduction

The starting point of Engle’s (1982) Autoregressive Conditional Heteroscedasticity (ARCH) class of models is

$$y_t = \mu_t + \epsilon_t, \quad \epsilon_t = \sigma_t \eta_t, \quad \sigma_t > 0, \quad \eta_t \sim iid(0, 1),$$

where y_t denotes the variable of interest (e.g. financial return), μ_t is the mean specification (e.g. an AR-X model) or simply zero, ϵ_t is the error term or mean-corrected variable of interest, σ_t is the conditional standard deviation or volatility, and η_t is an innovation with mean zero and unit variance. Arguably, the most common specification of σ_t is the first-order Generalised ARCH (GARCH) model of Bollerslev (1986):

$$\sigma_t^2 = \omega + \alpha_1 \epsilon_{t-1}^2 + \beta_1 \sigma_{t-1}^2, \quad \omega > 0, \quad \alpha_1, \beta_1 \geq 0. \quad (1)$$

Usually, this model is referred to as the GARCH(1,1) model. The log-ARCH class of models was, independently, first proposed by Geweke (1986), Pantula (1986) and Milhøj (1987). However, the idea of modelling the log-variance goes at least back to Park (1966). The logarithmic counterpart of the GARCH(1,1) is the log-GARCH(1,1) model, which is given by

$$\ln \sigma_t^2 = \omega + \alpha_1 \ln \epsilon_{t-1}^2 + \beta_1 \ln \sigma_{t-1}^2, \quad \omega, \alpha_1, \beta_1 \in \mathbb{R}. \quad (2)$$

Just as in the GARCH model, the ω is the volatility intercept, α_1 is the ARCH-parameter and β_1 is the GARCH-parameter: ω controls the level of volatility (but in a multiplicative

way), α_1 controls the impact of shocks or news η_{t-1}^2 , whereas $\alpha_1 + \beta_1$ controls the degree to which volatility σ_t^2 is persistent. If ϵ_t is a (mean-corrected) daily financial return, then typical estimates of α_1 and β_1 lie around 0.05 and 0.9, respectively, both in the GARCH(1,1) and log-GARCH(1,1) cases. See Section 2.5 for an illustration of the latter. Let $\ln^+ x = \max\{x, 0\}$. If $Pr(\eta_t = 0) = 0$, $E|\ln \eta_t^2| < \infty$ and $E(\ln^+ |\ln \eta_t^2|) < \infty$, then a sufficient condition for strict stationarity and ergodicity of (2) is

$$|\alpha_1 + \beta_1| < 1, \quad (3)$$

see Francq et al. (2013, Theorem 2.1 on p. 36).

From a user-perspective, the log-GARCH model has many attractive features:

- *Fitted volatility is guaranteed to be positive due to the exponential specification.* This is particularly important in higher order specifications, which may be needed in daily data with weekly periodicity (e.g. electricity prices), and in quarterly and monthly data. For example, when Engle (1982) proposed his ARCH model, he had to substantially restrict his ARCH(4) specification of quarterly UK inflation uncertainty to ensure positivity of his conditional variance estimates: The specification was assumed to follow $\sigma_t^2 = \omega + \alpha_1(0.4\epsilon_{t-1}^2 + 0.3\epsilon_{t-2}^2 + 0.2\epsilon_{t-3}^2 + 0.1\epsilon_{t-4}^2)$ so that only ω and α_1 were estimated, see Engle (1982, p. 1002).
- *No non-negativity constraints on parameters.* In the GARCH model, parameters must satisfy non-negativity constraints. In the GARCH(1,1), for example, the constraints are $\omega > 0$ and $\alpha_1, \beta_1 \geq 0$. The first means σ_t^2 is bounded from below and hence cannot be smaller than ω . The second, i.e. $\alpha_1, \beta_1 \geq 0$, implies that autocorrelations of ϵ_t^2 are non-negative. The log-GARCH, by contrast, does not impose non-negativity constraints on the parameters. This means its volatility is bounded from below by 0 (and not ω as in the GARCH), and that it admits negative autocorrelations on ϵ_t^2 . This latter is useful in regular data (e.g. daily data with a weekly 5-day or 7-day cycle, and monthly and quarterly data), since there it is likely that one or more autocorrelations of ϵ_t^2 are negative, see e.g. Pretis et al. (2018, Section 5.4).
- *Standard inference valid under nullity of parameters.* Often it is of interest to formulate a test in which one or more parameters are 0 under the null, say, $H_0: \beta_1 = 0$. A test with this null can be carried out in a standard way (e.g. with a t -test) in the log-GARCH model, since the value of β_1 lies in the interior of the admissible parameter space under the null. In the GARCH model, by contrast, $\beta_1 = 0$ lies on the boundary of the admissible parameter space because of the non-negativity constraints. This means standard inference procedures are not valid. The practical implication of this is that it is usually easier to carry out hypothesis tests in the log-GARCH model when parameters are 0 under the null. Indeed, such tests can readily be carried out in standard software via the estimated ARMA representations.
- *No non-negativity constraints on covariates.* Often it is of interest to include covariates in the σ_t^2 specification. For example, does a large return yesterday in a different market (say, x_{t-1}^2) increase volatility σ_t^2 today? Numerous studies have been undertaken with covariates of this or a similar kind. For brevity, the inclusion

of covariates is sometimes indicated by “X”. In the log-GARCH-X model the covariates are not restricted to be non-negative. In the GARCH-X model, by contrast, the covariates must be non-negative to ensure that σ_t^2 is positive, see e.g. [Francq and Thieu \(2018\)](#). This limits the type of questions that can be answered within a GARCH-X model, and compounds the problem described above of inference under nullity of parameters. These restrictions and challenges do not characterise the log-GARCH.

- *Invariance to power-transformations.* Consider the δ th. power log-GARCH(1,1) specification

$$\ln \sigma_t^\delta = \omega_\delta + \alpha_1 \ln |\epsilon_{t-1}|^\delta + \beta_1 \ln \sigma_{t-1}^\delta, \quad \delta > 0. \quad (4)$$

This specification is of interest if the objective is to forecast, say, the conditional standard deviation σ_t (or any σ_t^δ with $\delta > 0$) rather than σ_t^2 . In contrast to the power GARCH counterpart, (4) can be re-written in terms of its 2nd. power as

$$\ln \sigma_t^2 = \omega + \alpha_1 \ln \epsilon_{t-1}^2 + \beta_1 \ln \sigma_{t-1}^2 \quad \text{with} \quad \omega = \frac{\omega_\delta}{\delta}. \quad (5)$$

In other words, an estimate of (4) for any power $\delta > 0$ is straightforwardly obtained via the estimate of (5). The power GARCH model, by contrast, is not characterised by this invariance to power transformations.

- *Robustness to outliers.* It is well-known that the GARCH model is fragile when outliers or large jumps in ϵ_t are present, see e.g. [Carnero et al. \(2007\)](#) and the references therein. Similarly, the unconditional variance in Nelson’s (1991) EGARCH may not exist if η_t is too fat-tailed, e.g. student’s t , see Nelson’s own discussion in the Appendix (same place). In the log-GARCH, by contrast, the effect of η_t^2 is dampened due to the log-transformation. This means the log-GARCH is much more robust to outliers or jumps and fat-tailedness of η_t . This can be illustrated by revisiting the daily Apple log-return series used by [Harvey and Sucarrat \(2014, pp. 320-321\)](#) to illustrate a similar robustness for the Beta-t-EGARCH model. On Thursday 28 September 2000 the firm Apple issued a profit warning after closing hours, which led its stock-value to fall from USD 26.75 to USD 12.88. Volatility, however, was not affected on the subsequent days. [Figure 1](#) contains a snapshot of the event and the surrounding days. The figure plots absolute returns, the fitted conditional standard deviations of a GARCH(1,1) specification, and the fitted conditional standard deviations of a log-GARCH(1,1). The GARCH forecasts (one-step-ahead) of standard deviations exceed absolute returns for almost two months after the event, a clear-cut example of forecast failure. The forecasts of the log-GARCH, by contrast, remain in the same range of variation as the absolute returns due to the log-transformation.¹ This provides an empirical example of the GARCH model being prone to forecast failure in the presence of large outliers or jumps.
- *Generality of specification.* Two common alternatives to the log-GARCH are the Stochastic Volatility (SV) class of models, and the EGARCH of [Nelson \(1991\)](#). In

¹Estimation in R ([R Core Team \(2018\)](#)). The GARCH model is estimated with the `garch` function from the `tseries` package of [Trapletti and Hornik \(2016\)](#). The log-GARCH model is estimated via its ARMA representation with the `lgarch` function from the `lgarch` package of [Sucarrat \(2015\)](#).

a sense, the log-GARCH is more general than these model-classes, since both admit a log-GARCH representation (but not necessarily vice-versa), see [Asai \(1998\)](#), and [Francq et al. \(2017\)](#).

- *Log-GARCH models admit ARMA representations.* This was already noted by [Pantula \(1986\)](#), and has been exploited in numerous subsequent works, see e.g. [Psaradakis and Tzavalis \(1999\)](#). The usefulness of this is that a vast number of results, methods and techniques from the time-series literature is available. In particular, widely available software provide routines for the estimation of ARMA and/or VARMA models can be applied, which means univariate and multivariate log-GARCH models can readily be estimated in practice via their (V)ARMA representation(s).

The focus of this chapter is the last of these features. In the next part, [Section 2](#), we provide an overview of univariate models. We start by outlining an asymmetric specification, before we turn to its ARMA representation. Next we add stochastic conditioning covariates (“X”), and then sketch how estimates of the coefficient-covariances can be obtained in numerical software. We complete the [Section](#) by empirical illustrations of the log-GARCH(1,1) model. [Section 3](#) provides an overview of multivariate models. Again, we start by outlining the asymmetric specification and its corresponding VARMA and VARMA-X representations. Next, we turn to specifications that are amenable to equation-by-equation estimation, both stationary and non-stationary versions, even in the presence of Dynamic Conditional Correlations (DCCs) of unknown form. The focus on multivariate specifications that can be estimated equation-by-equation is motivated by the fact estimation becomes infeasible in practice as the dimension grows too large. We end the section with a short note on how models of Dynamic Conditional Correlations (DCCs) can be estimated subsequently. [Section 4](#) provides some suggestions on how to handle zeros in practice, whereas [Section 5](#) outlines how log-GARCH models can be used to model positively valued variables. Finally, [Section 6](#) concludes and provides suggestions for further research.

2 Univariate log-GARCH models

2.1 The asymmetric log-GARCH

Financial returns are often more volatile after a fall in price compared to a rise. This is usually referred to as asymmetry or leverage. To accommodate this commonly found feature, [Francq et al. \(2013\)](#) proposed the asymmetric log-GARCH. If $Pr(\eta_t = 0) = 0$ for all t , then their asymmetric log-GARCH can be re-parametrised as

$$\ln \sigma_t^2 = \omega + \sum_{i=1}^p \alpha_i \ln \epsilon_{t-i}^2 + \sum_{j=1}^q \beta_j \ln \sigma_{t-j}^2 + \sum_{k=1}^r \gamma_k 1_{\{\epsilon_{t-k} < 0\}} \ln \epsilon_{t-k}^2, \quad (6)$$

where

$$1_{\{\epsilon_t < 0\}} \ln \epsilon_t^2 = \begin{cases} \ln \epsilon_t^2 & \text{if } \epsilon_t < 0 \\ 0 & \text{if } \epsilon_t > 0 \end{cases}$$

is the asymmetry or leverage term. The advantages of the re-parametrisation in [\(6\)](#) are that it is more straightforward to test for the presence of asymmetry in practice,

and that it closely resembles the most common asymmetric non-exponential GARCH-counterpart of [Glosten et al. \(1993\)](#). In the log-GARCH(1,1), for example, asymmetry can be tested by means of a simple t -test. The re-parametrisation implies, however, that the sufficient conditions for strict stationarity and ergodicity (i.e. Theorem 2.1 in [Francq et al. \(2013, p. 36\)](#)) also needs to be re-parametrised. For example, in the first order case (i.e. $p = q = r = 1$), the sufficient condition becomes

$$|\alpha_1 + \beta_1|^{Pr(\eta_t > 0)} \cdot |\alpha_1 + \beta_1 + \gamma_1|^{1-Pr(\eta_t > 0)} < 1.$$

In the absence of asymmetry we obtain the usual condition in (3), i.e. $|\alpha_1 + \beta_1| < 1$.

2.2 The ARMA representation

If $Pr(\eta_t = 0) = 0$ and $E|\ln \eta_t^2| < \infty$, then (6) admits, almost surely, a (nonlinear in variables) ARMA(p, q) representation. It is obtained in two steps. First, $\ln \eta_t^2$ is added to each side of (6). Second, $\sum_{j=1}^q \beta_j (\ln \eta_t^2 - E(\ln \eta_t^2)) - \sum_{j=1}^q \beta_j (\ln \eta_t^2 - E(\ln \eta_t^2))$ is added to the right-hand side. Re-organising gives the nonlinear ARMA representation

$$\ln \epsilon_t^2 = \omega^* + \sum_{i=1}^p \phi_i \ln \epsilon_{t-i}^2 + \sum_{j=1}^q \theta_j u_{t-j}^2 + \sum_{k=1}^r \gamma_k 1_{\{\epsilon_{t-k} < 0\}} \ln \epsilon_{t-k}^2 + u_t, \quad u_t \sim iid(0, \sigma_u^2), \quad (7)$$

where

$$\omega^* = \omega + (1 - \sum_{j=1}^q \beta_j) \cdot E(\ln \eta_t^2), \quad \phi_i = \alpha_i + \beta_i, \quad \theta_j = -\beta_j, \quad u_t = \ln \eta_t^2 - E(\ln \eta_t^2). \quad (8)$$

If, in addition, $E[(\ln \eta_t^2)^2] < \infty$, then $\sigma_u^2 < \infty$ with $\sigma_u^2 = E[(\ln \eta_t^2)^2] - E(\ln \eta_t^2)^2$. Note that the specification is a nonlinear (in variables) ARMA due to the asymmetry terms. The stationarity conditions of [Francq et al. \(2013\)](#) still apply, since $\ln \epsilon_t^2$ is simply a sum of the stationary variables $\ln \sigma_t^2$ and $\ln \eta_t^2$. The model is therefore amenable to estimation by well-known ARMA-methods and widely available software. All the ARCH and GARCH parameters are identified via the relations in (8), and inference – even under the null of zero parameters – is readily carried out via a suitable transformation of the estimated coefficient covariance matrix, see Section 2.4. However, to identify the volatility intercept ω an estimate of $E(\ln \eta_t^2)$ is needed, and $E(\ln \eta_t^2)$ depends on the distribution of η_t^2 . [Sucarrat et al. \(2016\)](#) show that, under mild and general assumptions,

$$-\ln \left[\frac{1}{T} \sum_{t=1}^T \exp(\hat{u}_t) \right] \xrightarrow{p} E(\ln \eta_t^2), \quad (9)$$

where T is the sample size and \hat{u}_t is the residual from the estimated ARMA representation. Note that the expression inside the square brackets of (9) is the smearing estimator of [Duan \(1983\)](#). The motivation behind this estimator is that, if $E(\eta_t^2) = 1$ and $E(\ln \eta_t^2) < \infty$, then the population counterpart is equal to $E(\ln \eta_t^2)$:

$$-\ln E[e^{u_t}] = -\ln E \left[e^{\ln \eta_t^2 - E(\ln \eta_t^2)} \right] = -\ln \left[\frac{1}{e^{E(\ln \eta_t^2)}} \cdot E(\eta_t^2) \right] = E(\ln \eta_t^2).$$

Subject to suitable assumptions, therefore, consistent estimation of the ARMA representation (7) and the log-moment $E(\ln \eta_t^2)$, leads to consistent estimation of all the log-GARCH parameters in (6).

Another notable property of the estimator in (9) is that it ensures the sample variance of the standardised residuals $\{\hat{\eta}_t\} = \{\epsilon_t/\hat{\sigma}_t\}$, where $\hat{\sigma}_t^2$ is the fitted value of σ_t^2 , is approximately equal to 1 in empirical applications. This is required for $\hat{\sigma}_t^2$ to be a valid estimate of the conditional variance σ_t^2 . To see that the estimator in (9) ensures that the sample variance of $\{\hat{\eta}_t\}$ is approximately equal to 1, let $\hat{\eta}_t^* = \epsilon_t/\hat{\sigma}_t^*$ denote the residual scaled by the square root of the fitted value of the exponentiated ARMA-representation: $\hat{\sigma}_t^{*2} = \exp(\ln \hat{\epsilon}_t^2)$, where $\ln \hat{\epsilon}_t^2$ is the fitted value of the ARMA-representation. Noting that we also have $\hat{\eta}_t^* = \exp(\hat{u}_t/2)$, it follows that

$$\frac{\hat{\eta}_t^*}{\sqrt{T^{-1} \sum_{t=1}^T \exp(\hat{u}_t)}} = \frac{\hat{\eta}_t^*}{\exp\left(\ln T^{-1} \sum_{t=1}^T \exp(\hat{u}_t)/2\right)} = \frac{\epsilon_t}{\exp\left(\ln \hat{\epsilon}_t^2/2 - \hat{E}(\ln \eta_t^2)/2\right)} = \hat{\eta}_t,$$

where $\hat{E}(\ln \eta_t^2)$ is the estimator in (9). In other words, the smearing estimate $T^{-1} \sum_{t=1}^T \exp(\hat{u}_t)$ is approximately equal to the sample variance of $\{\hat{\eta}_t^*\}$, thus ensuring the sample variance of $\{\hat{\eta}_t\}$ is always approximately equal to 1 in empirical applications.

2.3 Adding stationary covariates (“X”)

Let $\mathbf{x}_t = (x_{1t}, \dots, x_{st})'$ denote a vector of strictly stationary and ergodic covariates. The (asymmetric) log-GARCH-X model is given by

$$\ln \sigma_t^2 = \omega + \sum_{i=1}^p \alpha_i \ln \epsilon_{t-i}^2 + \sum_{j=1}^q \beta_j \ln \sigma_{t-j}^2 + \sum_{k=1}^r \gamma_k 1_{\{\epsilon_{t-k} < 0\}} \ln \epsilon_{t-k}^2 + \sum_{l=1}^s \lambda_l x_{l,t-1}. \quad (10)$$

A common example of a covariate is realised volatility, i.e. a volatility proxy, but another example is *extended* asymmetry. In other words, the extended asymmetric log-GARCH model of Francq et al. (2017) is nested in (10). The (nonlinear) ARMA-X representation is obtained in the same way as earlier (see above), and it is given by

$$\ln \epsilon_t^2 = \omega^* + \sum_{i=1}^p \phi_i \ln \epsilon_{t-i}^2 + \sum_{j=1}^q \theta_j u_{t-j}^2 + \sum_{k=1}^r \gamma_k 1_{\{\epsilon_{t-k} < 0\}} \ln \epsilon_{t-k}^2 + \sum_{l=1}^s \lambda_l x_{l,t-1} + u_t, \quad (11)$$

where the relations between the log-GARCH and ARMA parameters are exactly as before, i.e. they are given by (8). Also, as noted earlier, no non-negativity constraints on the parameters $(\lambda_1, \dots, \lambda_s)'$ nor on the covariates \mathbf{x}_t are needed. Accordingly, standard inference methods are available under the null of 0s on one or more of the $\lambda_1, \dots, \lambda_s$, i.e. that one or more covariate has no impact on volatility. To estimate $E(\ln \eta_t^2)$, the same formula as earlier, i.e. (9), can be used. Estimation of (10), therefore, can straightforwardly be undertaken in widely available software.

2.4 Estimation of the coefficient covariance matrix

For inference on the parameters an estimate of the coefficient covariance matrix is needed, and this expression depends on the estimator. The two most common estimators of

ARMA-models are Least Squares (LS) and Gaussian Maximum Likelihood (ML). Both provide consistent and asymptotically normal estimates under mild assumptions – even when the error u_t is non-Gaussian, and most of the asymptotic properties of the two estimators are identical, see e.g. [Brockwell and Davis \(2006\)](#). The LS and Gaussian ML estimators are asymptotically efficient when η_t is sufficiently fat-tailed or skewed (or both). If, however, η_t is Gaussian, then improved efficiency can be achieved with the exponential Chi-squared (Quasi) ML estimator proposed by [Francq and Sucarrat \(2018\)](#). Here, we outline the details of the LS estimator, but the approach is similar for both the Gaussian and Chi-squared ML estimators.

Let $\varphi = (\omega^*, \phi_1, \dots, \phi_p, \theta_1, \dots, \theta_q, \gamma_1, \dots, \gamma_r, \lambda_1, \dots, \lambda_s)'$ denote the parameter of the ARMA representation given by (11), and let

$$\hat{\varphi} = \arg \min_{\varphi} \frac{1}{T} \sum_{t=1}^T u_t^2 \quad (12)$$

denote its Least Squares (LS) estimate. Often, numerical software provide utility functions for the computation of the Hessian at the optimum. [Francq and Sucarrat \(2017, pp. 27-28\)](#) show that this can be used to build an estimate of the coefficient covariance matrix. Specifically, they show that an estimate of the asymptotic coefficient matrix is obtained as

$$\left(\frac{1}{T} \sum_{t=1}^T \hat{u}_t^2 \right) \cdot 2 \cdot \hat{\mathbf{S}}^{-1},$$

where \hat{u}_t is the residual of the estimated ARMA-representation and $\hat{\mathbf{S}}$ is the Hessian at $\hat{\varphi}$ based on (12). If LS estimation is implemented by minimising the sum instead of the average, i.e.

$$\hat{\varphi} = \arg \min_{\varphi} \sum_{t=1}^T u_t^2, \quad (13)$$

then the estimate of the asymptotic coefficient matrix is modified to

$$\left(\frac{1}{T} \sum_{t=1}^T \hat{u}_t^2 \right) \cdot 2T \cdot \hat{\mathbf{S}}^{-1},$$

where $\hat{\mathbf{S}}$ is now the Hessian at $\hat{\varphi}$ based on (13).

Let $\zeta = (\omega, \alpha_1, \dots, \alpha_p, \beta_1, \dots, \beta_q, \gamma_1, \dots, \gamma_r, \lambda_1, \dots, \lambda_s)'$ denote the parameter of the log-GARCH specification (10), and let $\hat{\zeta}$ denote its estimate. An estimate of its asymptotic coefficient matrix is available by using the relationships between the log-GARCH and ARMA-parameters given by (8). For example, if $\widehat{Var}(\hat{x})$ and $\widehat{Cov}(\hat{x}, \hat{y})$ denote the variance of the estimate \hat{x} and the covariance of the estimates \hat{x} and \hat{y} , respectively, then the variance of the ARCH-parameter $\hat{\alpha}_i$ is obtained as $\widehat{Var}(\hat{\alpha}_i) = \widehat{Var}(\hat{\phi}_i) + \widehat{Var}(\hat{\theta}_i) + 2\widehat{Cov}(\hat{\phi}_i, \hat{\theta}_i)$. Similarly, the variance of the GARCH-parameter $\hat{\beta}_i$ is obtained as $\widehat{Var}(\hat{\beta}_i) = \widehat{Var}(-\hat{\theta}_i) = \widehat{Var}(\hat{\theta}_i)$. All the variances and covariances are readily available in this way, apart from those associated with the estimate of the log-GARCH intercept $\hat{\omega}$. These computations are more involved and requires the use of the delta-method, see [Francq and Sucarrat \(2017, pp. 21-22\)](#).

2.5 Empirical examples

To provide an empirical illustration of the log-GARCH model, we re-visit six daily financial return series: The FTSE100 and SP100 indices (source: Bloomberg), the Apple stock price (source: Yahoo Finance, <https://yahoo.finance.com>), the USD/EUR exchange rate (source: The European Central Bank, <http://www.ecb.int/>), the Brent blend oil price (source: The US Energy Information Agency, <http://www.eia.gov/>) and the gold price (source: Kitco, <http://www.kitco.com/>). The first two return series were studied in Francq and Sucarrat (2017, Section 5.1), whereas the latter four return series were studied in Harvey and Sucarrat (2014, Section 6). Note that the Apple series is the same as the one used to illustrate the robustness to outliers of log-GARCH models in the introduction (Section 1).

Let P_t denote the index-value or price of the asset in question in day t . The return $y_t = \epsilon_t$ is computed as the log-return in percent, i.e. $\epsilon_t = (\ln P_t - \ln P_{t-1}) \cdot 100$. The sample periods and descriptive statistics of the returns are contained in the upper part of Table 1, whereas Figure 2 contains graphs of the return series. As commonly found, the returns exhibit excess kurtosis relative to the normal distribution, and first-order ARCH at 5% and higher significance levels according to a Ljung and Box (1979) test for first-order autocorrelation in ϵ_t^2 . Also, the plots in Figure 2 confirm that volatility is persistent in the sense that the returns are characterised by volatility clustering.

Arguably, the most common volatility model is the plain GARCH(1,1). The plain log-GARCH(1,1) counterpart is given by

$$\ln \sigma_t^2 = \omega + \alpha_1 \ln \epsilon_{t-1}^2 + \beta_1 \ln \sigma_{t-1}^2,$$

and estimates of this model are contained in the middle part of Table 1. Estimation is undertaken via the ARMA(1,1) representation with the `lgarch` function from the *R* package `lgarch`, see Sucarrat (2015). Usually, in ordinary GARCH(1,1) models, the estimate of the ARCH parameter α_1 lies around 0.05, and the estimate of the GARCH parameter β_1 lies around 0.95. The results show that this is also the case for the log-GARCH(1,1) models. When estimation is via the ARMA-representation, then an estimate of $E(\ln \eta_t^2)$ is needed in order to estimate ω . If $\eta_t \sim N(0, 1)$, then $E(\ln \eta_t^2) = -1.27$. In other words, the discrepancy from -1.27 can be viewed as a measure of departure from normality. For example, if η_t is a standardised t with 10 degrees of freedom (a “moderate” departure from normality), then $E(\ln \eta_t^2) = -1.39$. The estimates of $E(\ln \eta_t^2)$ range from -1.375 to -1.522 , which suggests η_t is non-normal, albeit not dramatically so.

Often daily financial return series exhibit volatility asymmetry, i.e. a negative return tends to increase the volatility on the subsequent day. For stocks, this is typically referred to as a leverage effect, since leverage is often cited as the reason for the effect. For other return series, the more generic label “asymmetry” may be more appropriate, since the effect can be positive instead of negative, and since the reason for asymmetry may not be leverage. For exchange rates, for example, the presence and sign of asymmetry will usually depend on the relative strength of the two currencies in question. In other words, asymmetry is unlikely to be present in the USD/EUR exchange rate, since both the USD and Euro currencies are considered as strong currencies in international money markets. To explore the presence of volatility asymmetry in the six return series we fit a log-GARCH(1,1) model with extended asymmetry, i.e.

$$\ln \sigma_t^2 = \omega + \alpha_1 \ln \epsilon_{t-1}^2 + \beta_1 \ln \sigma_{t-1}^2 + \gamma_1 1_{\{\epsilon_{t-1} < 0\}} \ln \epsilon_{t-1}^2 + \lambda_1 1_{\{\epsilon_{t-1} < 0\}}.$$

The $1_{\{\epsilon_{t-1} < 0\}} \ln \epsilon_{t-1}^2$ is the ordinary asymmetry term, and $1_{\{\epsilon_{t-1} < 0\}}$ is the extended asymmetry term. As noted by [Francq et al. \(2017\)](#), to ensure invariance to scale-transformations the extended asymmetry term is needed when the ordinary asymmetry term is present. If we use ± 2 as critical values in a two-sided t -test with zero as null, then both the ordinary and extended asymmetry terms are significant for the stock returns (i.e. FTSE100, SP100 and Apple). For the remaining returns, however, neither the ordinary nor the extended term is significant. In other words, the results suggest the stock returns tend to be more volatile on days subsequent to a negative return, but not the exchange rate, oilprice nor the gold return.

3 Multivariate log-GARCH models

Let $\mathbf{y}_t = (y_{1t}, \dots, y_{Mt})'$ denote an M -dimensional vector of variables (e.g. financial returns) at t . A generic model of \mathbf{y}_t can be written as (see e.g. [Engle \(2002\)](#))

$$\begin{aligned} \mathbf{y}_t &= \boldsymbol{\mu}_t + \boldsymbol{\epsilon}_t, \\ \boldsymbol{\epsilon}_t &= (\epsilon_{1t}, \dots, \epsilon_{Mt})', \quad \mathbf{H}_t = E_{t-1}(\boldsymbol{\epsilon}_t \boldsymbol{\epsilon}_t'), \quad \mathbf{D}_t^2 = \text{diag}(\mathbf{H}_t), \\ \boldsymbol{\eta}_t &= \mathbf{D}_t^{-1} \boldsymbol{\epsilon}_t, \quad \mathbf{R}_t = E_{t-1}(\boldsymbol{\eta}_t \boldsymbol{\eta}_t'), \end{aligned}$$

where $\boldsymbol{\mu}_t$ is, say, a VARMA-X model, $\boldsymbol{\epsilon}_t = (\epsilon_{1t}, \dots, \epsilon_{Mt})'$ is the error term, \mathbf{H}_t is an $M \times M$ covariance matrix conditional on the past information \mathcal{F}_{t-1} , $E_{t-1}(\cdot)$ is shorthand notation for $E(\cdot | \mathcal{F}_{t-1})$, \mathbf{D}_t^2 is a diagonal $M \times M$ matrix with the conditional variance or volatility $\boldsymbol{\sigma}_t^2 = (\sigma_{1t}^2, \dots, \sigma_{Mt}^2)'$ on the diagonal, $\boldsymbol{\eta}_t = (\eta_{1t}, \dots, \eta_{Mt})'$ is the standardised error, i.e. $E(\boldsymbol{\eta}_t) = \mathbf{0}$ and $\text{Var}(\boldsymbol{\eta}_t) = \mathbf{1}$ where $\mathbf{0}$ and $\mathbf{1}$ are $M \times 1$ vectors, \mathbf{D}_t^{-1} is a diagonal $M \times M$ matrix with $(1/\sigma_{1t}, \dots, 1/\sigma_{Mt})'$ on the diagonal and \mathbf{R}_t is the correlation matrix conditional on the past. The relationships between \mathbf{H}_t and \mathbf{R}_t are given by $\mathbf{H}_t = \mathbf{D}_t \mathbf{R}_t \mathbf{D}_t$ and $\mathbf{R}_t = \mathbf{D}_t^{-1} \mathbf{H}_t \mathbf{D}_t^{-1}$.

3.1 A multivariate asymmetric log-GARCH-X model

The multivariate asymmetric log-GARCH-X model is given by

$$\ln \boldsymbol{\sigma}_t^2 = \boldsymbol{\omega} + \sum_{i=1}^p \boldsymbol{\alpha}_i \ln \boldsymbol{\epsilon}_{t-i}^2 + \sum_{j=1}^q \boldsymbol{\beta}_j \ln \boldsymbol{\sigma}_{t-j}^2 + \sum_{k=1}^r \boldsymbol{\gamma}_k \mathbf{1}_{\{\epsilon_{t-k} < 0\}} \ln \boldsymbol{\epsilon}_{t-k}^2 + \boldsymbol{\lambda} \mathbf{x}_{t-1}, \quad (14)$$

where $\ln \boldsymbol{\sigma}_t^2 = (\ln \sigma_{1t}^2, \dots, \ln \sigma_{Mt}^2)'$, $\boldsymbol{\omega} = (\omega_1, \dots, \omega_M)'$, $\ln \boldsymbol{\epsilon}_{t-i}^2 = (\ln \epsilon_{1,t-i}^2, \dots, \ln \epsilon_{M,t-i}^2)'$, $\mathbf{1}_{\{\epsilon_{t-k} < 0\}} \ln \boldsymbol{\epsilon}_{t-k}^2 = (1_{\{\epsilon_{1,t-k} < 0\}} \ln \epsilon_{1,t-k}^2, \dots, 1_{\{\epsilon_{M,t-k} < 0\}} \ln \epsilon_{M,t-k}^2)'$ and $\mathbf{x}_{t-1} = (x_{1,t-1}, \dots, x_{M,t-1})'$ are all $M \times 1$ vectors, and where

$$\boldsymbol{\alpha}_i = \begin{pmatrix} \alpha_{11.i} & \cdots & \alpha_{1M.i} \\ \vdots & \ddots & \vdots \\ \alpha_{M1.i} & \cdots & \alpha_{MM.i} \end{pmatrix}, \quad \boldsymbol{\beta}_j = \begin{pmatrix} \beta_{11.j} & \cdots & \beta_{1M.j} \\ \vdots & \ddots & \vdots \\ \beta_{M1.j} & \cdots & \beta_{MM.j} \end{pmatrix}, \quad (15)$$

$$\boldsymbol{\gamma}_k = \begin{pmatrix} \gamma_{11.k} & \cdots & \gamma_{1M.k} \\ \vdots & \ddots & \vdots \\ \gamma_{M1.k} & \cdots & \gamma_{MM.k} \end{pmatrix}, \quad \boldsymbol{\lambda}_l = \begin{pmatrix} \lambda_{11.l} & \cdots & \lambda_{1M.l} \\ \vdots & \ddots & \vdots \\ \lambda_{M1.l} & \cdots & \lambda_{MM.l} \end{pmatrix} \quad (16)$$

are all $M \times M$ matrices. The (nonlinear) VARMA-X representation is obtained in the same way as in the univariate case, and it is given by

$$\ln \epsilon_t^2 = \omega^* + \sum_{i=1}^p \phi_i \ln \epsilon_{t-i}^2 + \sum_{j=1}^q \theta_j u_{t-j}^2 + \sum_{k=1}^r \gamma_k \mathbf{1}_{\{\epsilon_{t-k} < 0\}} \ln \epsilon_{t-k}^2 + \boldsymbol{\lambda} \mathbf{x}_{t-1} + \mathbf{u}_t, \quad (17)$$

where

$$\omega^* = \omega + \left(\mathbf{I} - \sum_{j=1}^q \beta_j \right) E(\ln \eta_t^2), \quad \phi_i = \alpha_i + \beta_i, \quad \theta_j = -\beta_j, \quad \mathbf{u}_t = \ln \eta_t^2 - E(\ln \eta_t^2). \quad (18)$$

Without asymmetry (i.e. $\gamma_1 = \dots = \gamma_r = 0$), (17) is simply a VARMA-X model. To conduct inference on the log-GARCH parameters, an approach similar to the one outlined in Section 2.4 can be used.

If η_t is iid, then the conditional correlation matrix \mathbf{R}_t is constant, so that (14) is a Constant Conditional Correlation (CCC) model. Under suitable stationarity and regularity conditions, the (nonlinear) VARMA-X representation (17) can then be estimated by common methods, e.g. multivariate Gaussian QML. If \mathbf{R}_t is time-varying (and stationary), then a reasonable conjecture is that estimates will still be consistent subject to suitable assumptions. However, the asymptotic properties of such an estimator are currently unknown.

3.2 Equation-by-equation estimation

Multivariate volatility models are plagued by the ‘‘curse of dimensionality’’: As the dimension grows, estimation becomes infeasible due to the large amount of parameters that are estimated. One solution, if available, is equation-by-equation estimation. For this to be possible the GARCH-matrices (i.e. β_1, \dots, β_q) must all be diagonal, and

$$\text{each } \eta_{mt}, m = 1, \dots, M, \text{ must be independent of the past information } \mathcal{F}_{t-1}, \quad (19)$$

see Francq and Zakoian (2016). Francq and Sucarrat (2017) propose a first order version (i.e. $p = q = 1$) of the multivariate log-GARCH-X that satisfies these properties, and which allows for certain types of Dynamic Conditional Correlations (DCCs) of unknown form. A generalisation of their model allows for higher orders and asymmetry, and the m th. equation in such a generalised model is given by

$$\ln \sigma_{mt}^2 = \omega_m + \sum_{i=1}^p \alpha_{m,i} \ln \epsilon_{t-i}^2 + \sum_{j=1}^q \beta_{mm,j} \ln \sigma_{m,t-j}^2 + \sum_{k=1}^r \gamma_{m,k} \mathbf{1}_{\{\epsilon_{t-k} < 0\}} \ln \epsilon_{t-k}^2 + \boldsymbol{\lambda}_m \mathbf{x}_{t-1}, \quad (20)$$

where $\alpha_{m,i}$, $\gamma_{m,k}$ and $\boldsymbol{\lambda}_m$ are $1 \times M$ vectors made up of the m th. row in the matrices ϕ_i , γ_k and $\boldsymbol{\lambda}$, respectively. The $\beta_{mm,j}$ is the m th. element of the m th. column in the diagonal matrix β_j .

The univariate ARMA-X representation of the m th. equation is

$$\begin{aligned} \ln \epsilon_{mt}^2 &= \omega_m^* + \sum_{i=1}^p \phi_{m,i} \ln \epsilon_{t-i}^2 + \sum_{j=1}^q \theta_{mm,j} u_{m,t-j} \\ &\quad + \sum_{k=1}^r \gamma_{m,k} \mathbf{1}_{\{\epsilon_{t-k} < 0\}} \ln \epsilon_{t-k}^2 + \boldsymbol{\lambda}_m \mathbf{x}_{t-1} + u_{mt}, \end{aligned} \quad (21)$$

where

$$\begin{aligned}\omega_m^* &= \omega_m + \left(1 - \sum_{j=1}^q \beta_{mm.j}\right) E(\ln \eta_{mt}^2), & \phi_{m.i} &= \alpha_{m.i} + \beta_{m.i}, \\ \theta_{mm.j} &= -\beta_{mm.j} & \text{and} & & u_{mt} &= \ln \eta_{mt}^2 - E(\ln \eta_{mt}^2).\end{aligned}\tag{22}$$

Under stationarity and suitable regularity conditions, (21) can be estimated consistently with standard software. Subsequently, the log-GARCH parameters can be identified via the relations in (22). To identify ω_m , the formula in (9) can be applied to the residuals of equation m to estimate $E(\ln \eta_{mt}^2)$. For inference on the log-GARCH parameters in equation m , an approach similar to the one outlined in Section 2.4 can be used. For inference that involves parameters from more than one equation, then the joint coefficient covariance is needed, see [Francq and Sucarrat \(2017\)](#).

For equation-by-equation estimation to be available the GARCH-matrices β_1, \dots, β_q must all be diagonal. To test whether this is indeed the case, a Lagrange-Multiplier (LM) test of equation m can be devised: Under the null all the elements of $\{\beta_{mi.j} : i \neq m, j = 1, \dots, q\}$ are equal to zero, whereas under the alternative one or more elements are non-zero. Formally, this has not been pursued yet in the theoretical log-GARCH literature.

3.3 Non-stationary models

A common approach to non-stationary volatility is to decompose σ_t^2 multiplicatively, see (amongst other) [Van Bellegem and Von Sachs \(2004\)](#), [Engle and Rangel \(2008\)](#), [Mazur and Pipien \(2012\)](#), and [Amado and Terasvirta \(2014a, 2014b\)](#). This means

$$\sigma_t^2 = \mathbf{g}_t \odot \mathbf{h}_t = (g_{1t}h_{1t}, \dots, g_{Mt}h_{Mt})',$$

where \mathbf{g}_t is the non-stationary component, \mathbf{h}_t is the stationary component (e.g. a GARCH-like process), and \odot is the elementwise (Hadamard) matrix product.² [Escribano and Sucarrat \(2018\)](#) propose a non-stationary multivariate log-GARCH-X specification that can be estimated equation-by-equation. Their motivation was the presence of non-stochastic periodicity in the intraday electricity price market. However, their idea applies more generally. The non-stationary component in their model is given by

$$\ln \mathbf{g}_t = \left(\ln g_1(\boldsymbol{\lambda}_1^f, \mathbf{x}_{1t}^f), \dots, \ln g_M(\boldsymbol{\lambda}_M^f, \mathbf{x}_{Mt}^f) \right)',$$

where $\ln g_1, \dots, \ln g_M$ are known functions (linear or nonlinear), $\mathbf{x}_{1t}^f, \dots, \mathbf{x}_{Mt}^f$ are known, non-stochastic or fixed (hence the superscript f) regressors, and $\boldsymbol{\lambda}_1^f, \dots, \boldsymbol{\lambda}_M^f$ are unknown parameters to be estimated. Neither the \mathbf{x}_{mt}^f 's nor the $\ln g_m$'s are restricted to be equal across equations, and the $\ln g_m$'s can assume a variety of shapes. In the simplest case the $\ln g_m$'s are linear functions made up of time dummies (e.g. calendar effects), but it can also take the shape of an exponential spline as in [Engle and Rangel \(2008\)](#), the Fourier Flexible Form (FFF) as in [Mazur and Pipien \(2012\)](#), or smooth threshold models as in [Amado](#)

²For example, if \mathbf{a} and \mathbf{b} are two equally sized $M \times 1$ vectors, say, $\mathbf{a} = (a_1, \dots, a_M)'$ and $\mathbf{b} = (b_1, \dots, b_M)'$, then $\mathbf{a} \odot \mathbf{b} = (a_1 b_1, \dots, a_M b_M)'$.

and Terasvirta (2014a, 2014b). The functions may also be estimated nonparametrically, as in Van Bellegem and Von Sachs (2004).

If we for notational simplicity exclude asymmetry and covariates, then the stationary component is given by

$$\ln \mathbf{h}_t = \boldsymbol{\omega} + \sum_{i=1}^p \boldsymbol{\alpha}_i \ln \tilde{\boldsymbol{\epsilon}}_{t-i}^2 + \sum_{j=1}^q \boldsymbol{\beta}_j \ln \mathbf{h}_{t-j}, \quad (23)$$

where $\ln \mathbf{h}_t = \ln \boldsymbol{\sigma}_t^2 - \ln \mathbf{g}_t = (\ln h_{1,t}, \dots, \ln h_{M,t})'$, $\boldsymbol{\omega} = (\omega_1, \dots, \omega_M)'$, $\ln \tilde{\boldsymbol{\epsilon}}_t^2 = (\ln \boldsymbol{\epsilon}_t^2 - \ln \mathbf{g}_t) = (\ln h_{1t}\eta_{1t}^2, \dots, \ln h_{Mt}\eta_{Mt}^2)'$, and $\boldsymbol{\alpha}_i$ and $\boldsymbol{\beta}_j$ are both $M \times M$ matrices as in (15). The matrices $\boldsymbol{\beta}_j$ need not be diagonal. However, we will impose this restriction to enable an equation-by-equation estimation scheme. The m th. log-volatility equation thus becomes

$$\ln \sigma_{mt}^2 = \ln g_{mt} + \ln h_{mt}, \quad (24)$$

$$\ln g_{mt} = \ln g_m(\boldsymbol{\lambda}_m^f, \mathbf{x}_{mt}^f), \quad (25)$$

$$\ln h_{mt} = \omega_m + \sum_{i=1}^p \boldsymbol{\alpha}_{m,i} \ln \tilde{\boldsymbol{\epsilon}}_{t-i}^2 + \sum_{j=1}^q \beta_{mm,j} \ln h_{m,t-j}^2, \quad (26)$$

where $\boldsymbol{\alpha}_{m,i}$ is the m th. row of $\boldsymbol{\alpha}_i$, i.e. $\boldsymbol{\alpha}_{m,i} = (\alpha_{m1,i}, \dots, \alpha_{mM,i})$. Let λ_{m0}^f denote the unconditional mean of $\ln \tilde{\boldsymbol{\epsilon}}_{mt}^2$, i.e. $\lambda_{m0}^f = E(\ln \tilde{\boldsymbol{\epsilon}}_{mt}^2)$ with $E|\ln \tilde{\boldsymbol{\epsilon}}_{mt}^2| < \infty$. If we add $\ln \eta_{mt}^2$ to each side of (24), and then $\lambda_{m0}^f - \lambda_{m0}^f$ to the right-hand side, we obtain

$$\ln \epsilon_{mt}^2 = \lambda_{m0}^f + \ln g_m(\boldsymbol{\lambda}_m^f, \mathbf{x}_{mt}^f) + w_{mt}, \quad w_{mt} = (\ln \tilde{\boldsymbol{\epsilon}}_{mt}^2 - \lambda_{m0}^f).$$

This is simply a regression with a fixed or non-stochastic part, i.e. $\lambda_{m0}^f + \ln g_m(\boldsymbol{\lambda}_m^f, \mathbf{x}_{mt}^f)$, and a zero-mean stationary error governed by the mean-corrected ARMA model

$$w_{mt} = \sum_{i=1}^p \phi_{m,i} w_{t-i} + \sum_{j=1}^q \theta_{mm,j} u_{m,t-j} + u_{mt}, \quad (27)$$

where $w_{mt} = \ln \tilde{\boldsymbol{\epsilon}}_{mt}^2 - E(\ln \tilde{\boldsymbol{\epsilon}}_{mt}^2)$ and $\mathbf{w}_t = (w_{1t}, \dots, w_{Mt})'$. This means the m th. equation can be estimated in three steps:

1. Estimate λ_{m0}^f and $\boldsymbol{\lambda}_m^f$ via the auxiliary regression

$$\ln \epsilon_{mt}^2 = \lambda_{m0}^f + \ln g_m(\boldsymbol{\lambda}_m^f, \mathbf{x}_{mt}^f) + w_{mt},$$

where λ_{m0}^f is the intercept and w_{mt} is a zero-mean stationary error-term governed by (27). If $\boldsymbol{\lambda}_m^f$ enters linearly in $\ln g_m$, then the parameters can simply be estimated by OLS.

2. Fit an ARMA model to the residuals \hat{w}_{mt} from the first step. The relation between the parameters of the log-GARCH model and the parameters of the mean-corrected ARMA-representation are the same as in the case where the ARMA-representation is not mean-corrected, i.e. (22). So this provides an estimate of all the log-GARCH parameters apart from the intercept ω_m . An estimate of ω_m , however, is not needed if the aim is to estimate σ_{mt}^2 . The reason for this is that the fitted values from the

first two steps provide estimates of $E(\ln \tilde{\epsilon}_{mt}^2) + \ln g_{mt}$ and $E_{t-1}(y_{mt})$, respectively. Adding these gives

$$\begin{aligned} E(\ln \tilde{\epsilon}_{mt}^2) + \ln g_{mt} + E_{t-1}(w_{mt}) &= \ln g_{mt} + E_{t-1}(\ln \tilde{\epsilon}_{mt}^2) \\ &= \ln g_{mt} + \ln h_t + E(\ln \eta_{mt}^2), \end{aligned}$$

since $\ln \tilde{\epsilon}_{mt}^2 = \ln h_{mt} + \ln \eta_{mt}^2$. So only an estimate of $E(\ln \eta_{mt}^2)$ is needed to complete the estimate of σ_{mt}^2 .

3. Estimate the log-moment $E(\ln \eta_{mt}^2)$ needed to complete the estimate of σ_{mt}^2 . Again, we can use the residuals from Step 2 in combination with (9).

Summarised, then, the estimate of σ_{mt}^2 is given by

$$\hat{\sigma}_{mt}^2 = \exp \left(\underbrace{\hat{E}(\ln \tilde{\epsilon}_{mt}^2) + \ln \hat{g}_{mt}}_{\text{Step 1}} + \underbrace{\hat{E}_{t-1}(w_{mt})}_{\text{Step 2}} - \underbrace{\hat{E}(\ln \eta_{mt}^2)}_{\text{Step 3}} \right),$$

where $\hat{E}(\ln \tilde{\epsilon}_{mt}^2) + \ln \hat{g}_{mt}$ is the fitted value of the auxiliary regression in Step 1, $\hat{E}_{t-1}(y_{mt})$ is the fitted value of the mean-corrected ARMA representation in Step 2, and $\hat{E}(\ln \eta_{mt}^2)$ is the estimate of $E(\ln \eta_{mt}^2)$ in Step 3. Note that the three-step procedure can in fact be reduced to two steps if the centred exponential Chi-squared QMLE of [Francq and Sucarrat \(2018\)](#) is used in the second step, since $E(\ln \eta_{mt}^2)$ enters explicitly as a parameter to be estimated in the centred exponential Chi-squared density. This will also be more efficient if η_{mt} is normal or close to normal.

An estimate of ω_m requires estimation of the other equations, in addition to equation m . This is because the expression for $E(\ln \tilde{\epsilon}_{mt}^2)$, which can be written as $E(\ln \tilde{\epsilon}_{mt}^2) = \omega_m^* + \sum_{i=1}^p \phi_{m,i} E(\ln \tilde{\epsilon}_t^2)$, depends on the unconditional expectations of the other equations. Recalling, from (22), that $\omega_m^* = \omega_m + (1 - \sum_{j=1}^q \beta_{mm,j}) E(\ln \eta_{mt}^2)$ when the GARCH-matrices are diagonal, solving for ω_m in the expression for $E(\ln \tilde{\epsilon}_{mt}^2)$ gives

$$\omega_m = (1 - \sum_{j=1}^q \beta_{mm,j}) E(\ln \tilde{\epsilon}_{mt}^2) - \sum_{i=1}^p \alpha_{m,i} E(\ln \tilde{\epsilon}_t^2) - (1 - \sum_{j=1}^q \beta_{mm,j}) E(\ln \eta_{mt}^2), \quad (28)$$

where we have used that $\sum_{i=1}^p \phi_{m,i} E(\ln \tilde{\epsilon}_t^2) = \sum_{i=1}^p \alpha_{m,i} E(\ln \tilde{\epsilon}_t^2) + \sum_{j=1}^q \beta_{mm,j} E(\ln \tilde{\epsilon}_{mt}^2)$. It should be noted that only the elements in $E(\ln \tilde{\epsilon}_t^2)$, apart from the m th. entry, comes from the other equations. In other words, if there is no feedback effects (i.e. all entries in the $\alpha_{m,i}$'s apart from the m th. entry are zero), then there is no need to estimate the other equations in order to estimate ω_m .

Asymmetry and stochastic covariates (“X”) can be added without affecting the estimation procedure just sketched. The only caveat is that they need to be mean corrected. Specifically, if \mathbf{x}_{t-1} is a $(r+s) \times 1$ vector that collects all the asymmetry terms and conditioning covariates of the stationary part, then they need to enter as $(\mathbf{x}_{t-1} - \bar{\mathbf{x}})$ in the ARMA representation, where $\bar{\mathbf{x}} = (\bar{x}_1, \dots, \bar{x}_M)'$ are the sample means of the stationary covariates. The stationary component is thus

$$\ln h_t^2 = \omega + \sum_{i=1}^p \alpha_i \ln \tilde{\epsilon}_{t-i}^2 + \sum_{j=1}^q \beta_j \ln h_{t-j}^2 + \delta(\mathbf{x}_{t-1} - \bar{\mathbf{x}}),$$

where $\boldsymbol{\delta}$ is a parameter-matrix of appropriate size, and the mean-corrected ARMA representation of equation m is

$$w_{mt} = \sum_{i=1}^p \phi_{m,i} \mathbf{w}_{t-i} + \sum_{j=1}^q \theta_{m,m,j} u_{m,t-j} + \boldsymbol{\delta}_m (\mathbf{x}_{t-1} - \bar{\mathbf{x}}) + u_{mt}, \quad (29)$$

where w_{mt} , \mathbf{w}_t and u_{mt} are defined as earlier, and $\boldsymbol{\delta}_m$ is the m th. row of $\boldsymbol{\delta}$. The practical consequence of this is that the three step estimation procedure described above only requires one minor modification: Estimate (29) instead of (27) in Step 2. The other steps are unchanged, and if an estimate of ω_m is needed, then formula (28) can still be used.

The asymptotic theory of non-stationary log-GARCH models has not been formally developed yet. Nevertheless, approximate inference procedures are readily available. For the stationary ARMA-representation a procedure similar to the one outlined in Section 2.4 can be used for inference within a single equation. The unknown is whether, or to what extent, this procedure is affected by the prior estimation of the non-stationary part. For inference that involves parameters from more than one equation, then an approximate joint coefficient covariance can be obtained along the lines of Francq and Sucarrat (2017). For inference regarding the parameters in the non-stationary part, then an approximate coefficient covariance can be computed by classical methods. For example, if the parameters of the non-stationary part in equation m are estimated by OLS, and if \mathbf{X}_m denotes the $T \times k$ regressor matrix of the OLS estimator, then an approximate expression is obtained as

$$(\mathbf{X}'_m \mathbf{X}_m)^{-1} \mathbf{X}'_m \widehat{\boldsymbol{\Omega}}_m \mathbf{X}_m (\mathbf{X}'_m \mathbf{X}_m)^{-1},$$

where $\widehat{\boldsymbol{\Omega}}_m$ is an estimate of the autocovariance matrix of w_{m1}, \dots, w_{mT} . The estimation results of the stationary part can be used to compute $\widehat{\boldsymbol{\Omega}}_m$. Indeed, if the stationary part is an ARMA, then this procedure is already available in a number of softwares.

3.4 Dynamic Conditional Correlations (DCCs)

Assumption (19) allow for certain types of DCCs when a multivariate log-GARCH is estimated equation-by-equation. The estimation procedures described above, however, do not provide estimates of the DCCs. Nevertheless, they can – if needed – be estimated in a subsequent step. The estimates $\widehat{\sigma}_{1t}^2, \dots, \widehat{\sigma}_{Mt}^2$ lead to the standardised residuals $\widehat{\boldsymbol{\eta}}_t = (\widehat{\eta}_{1t}, \dots, \widehat{\eta}_{Mt})'$, where $\widehat{\eta}_{mt} = \epsilon_{mt} / \widehat{\sigma}_{mt}$. These residuals can be used to estimate a DCC specification of $\mathbf{R}_t = E(\boldsymbol{\eta}_t \boldsymbol{\eta}'_t | \mathcal{F}_{t-1})$. An example is the DCC of Engle (2002), or alternatively the corrected version of Aielli (2013), see e.g. the empirical section of Francq and Sucarrat (2017). Another option is the robust (to spikes) DCC model proposed for electricity prices by Dupuis (2017).

4 Handling zeros in practice

Throughout we have relied on the theoretical assumption $Pr(\eta_t = 0) = 0$. In practice, however, if no conditional mean equation is fitted (i.e. we set $\mu_t = 0$ for all t), we may experience that $\epsilon_t = 0$ for some t . The most straightforward solution to this is to fit a specification μ_t , e.g. an intercept. This is not only justifiable in most contexts, it is also recommendable in order to ensure that the fitted values of ϵ_t are centred about zero.

A second solution consists of replacing zeros with some non-zero value c . One such value is a number very close to zero, say, the machine epsilon (e.g. $2.22e - 16$) of the software used. This is probably the worst possible choice! The reason for this is that $\ln c^2$ will usually be much smaller than any empirical non-zero value of $\ln \epsilon_t^2$. Accordingly, this will induce a large ARCH shock (or inlier) at each zero location. A more sensible solution is to set c equal to a value informed by the economic application in question. If there is no (obvious) economic motivation to inform the choice of c , then one may choose a certain quantile of the non-zero values of ϵ_t^2 (e.g. 10%), or the sample average of $\ln \epsilon_t^2$ (zeros excluded), or simply the value 1. The latter is very neat and justified in the log-GARCH(1,1) when the estimates of α_1 and β_1 are typical, i.e. about 0.05 and 0.9, respectively. Setting $c = 1$ thus means $\ln c^2 = 0$, so that all the weight (in predicting $\ln \sigma_t^2$) is shifted on to the GARCH term, i.e. $\ln \sigma_{t-1}^2$. If β_1 is large (e.g. about 0.9), then this is a very sensible solution.

A third solution consists of estimating the replacement value. This is the solution proposed by [Sucarrat and Escribano \(2018\)](#). Specifically, they propose to treat zeros as missing values, and to impute the missing values by the estimate of $E_{t-1}(\ln \epsilon_t^2)$ at each missing location. This means an optimal replacement value is inserted at each missing location in the ARMA representation, where “optimal” means the conditional (squared) forecast error is minimised, and/or that the likelihood is maximised. Arguably, treating zeros as missing values is the most appealing solution if no conditional mean is fitted. However, implementing the solution usually requires more of the user, and consistent parameter estimates are not guaranteed – in particular if the proportion of zeros is large. The freely available *R* package `lgarch` ([Sucarrat \(2015\)](#)) implements the missing value approach.

A fourth solution consists of adding a non-zero value c to *all* the squared observations $\epsilon_1^2, \dots, \epsilon_T^2$. This leads to a new series $\{\epsilon_t^{*2}\}$ with $\epsilon_t^{*2} = \epsilon_t^2 + c$ and $\epsilon_t^* = \sigma_t^* \eta_t^*$, such that σ_t^{*2} is approximately equal to $\sigma_t^2 + c$. In other words, approximate forecasts of σ_t^2 can be obtained by using the estimates of σ_t^{*2} , and noting that $\sigma_t^2 \approx \sigma_t^{*2} - c$. If the values of ϵ_t^2 are sufficiently large compared with c , then adding c will not alter the dynamics of ϵ_t^2 in a notable manner. An example is the case where ϵ_t^2 is interpreted as volume, i.e. a positively valued variable (see Section 5). In this case ϵ_t^2 will usually be much larger than, say, $c = 1$.

5 Modelling positively valued variables

[Engle and Russell \(1998\)](#) noted that ϵ_t^2 could be interpreted as positively valued variable, and hence showed that σ_t^2 can be interpreted as the conditional expectation of the positively valued variable. Put differently, [Engle and Russell \(1998\)](#) showed that the ARCH-class of models can be used to model positively valued variables like duration, volume, price-spread or realised volatility. This spurred the Multiplicative Error Model (MEM) literature, see [Brownlees et al. \(2012\)](#) for an overview. A particularly useful characteristic of the MEM interpretation is that, in practice, an ARCH estimation routine can be used to estimate a MEM. For example, suppose y_t denotes the positively valued variable in question. By providing the software in question with $\sqrt{y_t}$, then the software will return estimates of the MEM.

Formally, the MEM-class is given by

$$y_t = \psi_t \zeta_t, \quad \zeta_t \sim iid(1, \sigma_\zeta^2), \quad \psi_t, \zeta_t \geq 0,$$

where $\psi_t = E_{t-1}(y_t)$ is interpreted as the expectation of y_t conditional on past information. Bauwens and Giot (2000) were the first to propose a logarithmic version of the MEM, i.e. a log-MEM. They noted that the MEM-interpretation associated with ordinary GARCH models was restricted by the non-negativity constraints on the parameters, and by the non-negativity constraints on the conditioning variables. Instead, they proposed a first order log-MEM with a single (stochastic) conditioning covariate. More generally, however, the univariate log-MEM(p, q) with stochastic conditioning variables can be written as

$$\ln \psi_t = \omega + \sum_{i=1}^p \alpha_i \ln y_{t-i} + \sum_{j=1}^q \beta_j \ln \psi_{t-j} + \sum_{l=1}^s \lambda_l x_{l,t-1}.$$

That is, a univariate log-GARCH-X without asymmetry (ordinary and extended asymmetry-terms are not meaningful in MEMs). If $Pr(\zeta_t = 0) = 0$ and $E|\ln \zeta_t| < \infty$, then the model admits – as earlier – an ARMA-representation almost surely given by

$$\ln y_t = \omega^* + \sum_{i=1}^p \phi_i \ln y_{t-i} + \sum_{j=1}^q \theta_j u_{t-j} + \sum_{l=1}^s \lambda_l x_{l,t-1} + u_t,$$

and the relations between the log-MEM and ARMA parameters are exactly the same as before, i.e. they are given by (22). The existence of an ARMA-representation of the log-MEM was, to the best of my knowledge, first pointed out by Allen et al. (2008). Similarly, the model can be extended to the multivariate case as in Section 3, and non-stationarities can be introduced in the same way as in Section 3.3. Finally, fractionally integrated extensions of the log-MEM are considered in Beran et al. (2015), and in Feng and Zhou (2015).

6 Conclusions

The log-GARCH model provides a very flexible framework for the modelling of economic uncertainty, financial volatility and a range of other positively valued variables. In this chapter we have outlined how univariate and multivariate log-GARCH models can be represented by (V)ARMA-like representations, and – as a consequence – how well-known (V)ARMA results can be used to estimate univariate and multivariate log-GARCH models. Nevertheless, there is still a large, unexploited potential. There exists a wide range of well-known results on time-varying coefficients, non-stationarities, missing data and efficient estimation, amongst other, that can potentially be used to shed light on and further extend the log-GARCH class of models via its (V)ARMA-like representations.

References

- Aielli, G. P. (2013). Dynamic Conditional Correlations: On Properties and Estimation. *Journal of Business and Economic Statistics* 31, 282–299. <http://dx.doi.org/10.1080/07350015.2013.771027>.

- Allen, D., F. Chan, M. McAleer, and S. Peiris (2008). Finite sample properties of the QMLE for the log-acd model: Application to australian stocks. *Journal of Econometrics* 147, 163–185.
- Amado, C. and T. Terasvirta (2014a). Modelling Changes in the unconditional variance of long stock return series. *Journal of Empirical Finance* 25, 15–35.
- Amado, C. and T. Terasvirta (2014b). Modelling volatility by variance decomposition. *Journal of Econometrics* 175, 142–153.
- Asai, M. (1998). A new method to estimate stochastic volatility models: A log-GARCH approach. *Journal of the Japanese Statistical Society* 28, 101–114.
- Bauwens, L. and P. Giot (2000). The logarithmic ACD model: an application to the bid-ask quote process of three NYSE stocks. *Annales d'Economie et de Statistique* 60, 117–149.
- Beran, J., Y. Feng, and S. Gosh (2015). Modelling long-range dependence and trends in duration series: an approach based on EFARIMA and ESEMIFAR models. *Statistical Papers* 56, 431–451.
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroscedasticity. *Journal of Econometrics* 31, 307–327.
- Brockwell, P. J. and R. A. Davis (2006). *Time Series: Theory and Methods*. New York: Springer. 2nd. Edition, first published in 1991.
- Brownlees, C., F. Cipollini, and G. Gallo (2012). Multiplicative Error Models. In L. Bauwens, C. Hafner, and S. Laurent (Eds.), *Handbook of Volatility Models and Their Applications*, pp. 223–247. New Jersey: Wiley.
- Carnero, M. A., D. Pena, and E. Ruiz (2007). Effects of outliers on the identification and estimation of GARCH models. *Journal of Time Series Analysis* 28, 471–497.
- Duan, N. (1983). Smearing Estimate: A Nonparametric Re-transformation Method. *Journal of the American Statistical Association* 78, pp. 605–610.
- Dupuis, D. J. (2017). Electricity price dependence in new york state zones: A robust detrended correlation approach. *Annals of Applied Statistics* 11, 248–273.
- Engle, R. (1982). Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica* 50, 987–1008.
- Engle, R. (2002). Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models. *Journal of Business and Economic Statistics* 20, 339–350.
- Engle, R. F. and J. G. Rangel (2008). The Spline GARCH Model for Low Frequency Volatility and its Global Macroeconomic Causes. *Review of Financial Studies* 21, 1187–1222.

- Engle, R. F. and J. R. Russell (1998). Autoregressive Conditional Duration: A New Model of Irregularly Spaced Transaction Data. *Econometrica* 66, 1127–1162.
- Escribano, Á. and G. Sucarrat (2018). Equation-by-Equation Estimation of Multivariate Periodic Electricity Price Volatility. *Energy Economics*. In press. Working paper version: MPRA Paper No. 72736 at <https://mpra.ub.uni-muenchen.de/72736/>.
- Feng, Y. and C. Zhou (2015). Forecasting financial market activity using a semiparametric fractionally integrated Log-ACD. *International Journal of Forecasting* 31, 349–363.
- Francq, C. and G. Sucarrat (2017). An Equation-by-Equation Estimator of a Multivariate Log-GARCH-X Model of Financial Returns. *Journal of Multivariate Analysis* 153, 16–32.
- Francq, C. and G. Sucarrat (2018). An Exponential Chi-Squared QMLE for Log-GARCH Models Via the ARMA Representation. *Journal of Financial Econometrics* 16, 129–154. Working Paper version: <http://mpra.ub.uni-muenchen.de/51783/>.
- Francq, C. and L. Q. Thieu (2018). Qml inference for volatility models with covariates. *Econometric Theory*. <https://doi.org/10.1017/S0266466617000512>.
- Francq, C., O. Wintenberger, and J.-M. Zakoïan (2013). GARCH Models Without Positivity Constraints: Exponential or Log-GARCH? *Journal of Econometrics* 177, 34–36.
- Francq, C., O. Wintenberger, and J.-M. Zakoïan (2017). Goodness-of-fit tests for log and exponential GARCH models. *TEST*. <http://dx.doi.org/10.1007/s11749-016-0506-2>.
- Francq, C. and J.-M. Zakoïan (2016). Estimating multivariate volatility models equation by equation. *The Journal of the Royal Statistical Society. Series B* 78, 613–635. Working paper version: MPRA Paper No. 54250: <http://mpra.ub.uni-muenchen.de/54250/>.
- Geweke, J. (1986). Modelling the Persistence of Conditional Variance: A Comment. *Econometric Reviews* 5, 57–61.
- Glosten, L. R., R. Jagannathan, and D. E. Runkle (1993). On the Relation between the Expected Value and the Volatility of the Nominal Excess Return on Stocks. *Journal of Finance* 48, 1779–1801.
- Harvey, A. C. and G. Sucarrat (2014). EGARCH models with fat tails, skewness and leverage. *Computational Statistics and Data Analysis* 76, 320–338.
- Ljung, G. and G. Box (1979). On a Measure of Lack of Fit in Time Series Models. *Biometrika* 66, 265–270.
- Mazur, B. and M. Pipien (2012). On the Empirical Importance of Periodicity in the Volatility of Financial Returns – Time Varying GARCH as a Second Order APC(2) Process. *Central European Journal of Economic Modelling and Econometrics* 4, 95–116.
- Milhøj, A. (1987). A Multiplicative Parametrization of ARCH Models. Research Report 101, University of Copenhagen: Institute of Statistics.

- Nelson, D. B. (1991). Conditional Heteroskedasticity in Asset Returns: A New Approach. *Econometrica* 59, 347–370.
- Pantula, S. (1986). Modelling the Persistence of Conditional Variance: A Comment. *Econometric Reviews* 5, 71–73.
- Park, R. (1966). Estimation with Heteroscedastic Error Terms. *Econometrica* 34, 888–888.
- Pretis, F., J. Reade, and G. Sucarrat (2018). Automated General-to-Specific (GETS) Regression Modeling and Indicator Saturation for Outliers and Structural Breaks. *Journal of Statistical Software* 86, 1–44.
- Psaradakis, Z. and E. Tzavalis (1999). On regression-based tests for persistence in logarithmic volatility models. *Econometric Reviews* 18, 441–448.
- R Core Team (2014). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- R Core Team (2018). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Sucarrat, G. (2015). *lgarch: Simulation and Estimation of Log-GARCH Models*. R package version 0.6-2.
- Sucarrat, G. and Á. Escibano (2018). Estimation of Log-GARCH Models in the Presence of Zero Returns. *European Journal of Finance* 24, 809–827. <http://dx.doi.org/10.1080/1351847X.2017.1336452>.
- Sucarrat, G., S. Grønneberg, and Á. Escibano (2016). Estimation and Inference in Univariate and Multivariate Log-GARCH-X Models When the Conditional Density is Unknown. *Computational Statistics and Data Analysis* 100, 582–594.
- Trapletti, A. and K. Hornik (2016). *tseries: Time Series Analysis and Computational Finance*. R package version 0.10-35.
- Van Bellegem, S. and R. Von Sachs (2004). Forecasting economic time-series with unconditional time-varying variance. *International Journal of Forecasting* 20, 611–627.

Table 1: Empirical examples of log-GARCH models (see Section 2.5)

Descriptive statistics:							
	Sample	T	s^2	s^4	ARCH [p -val]		
FTSE100	5/1/1998–2/6/2015	4397	1.51	8.59	97.9 [0.00]		
SP100	5/1/1998–1/6/2015	4379	1.60	10.10	220.1 [0.00]		
Apple	10/9/1984–12/10/2011	6835	9.63	53.85	5.90 [0.02]		
USD/EUR	5/1/1999–12/10/2011	3274	0.45	5.45	205.9 [0.00]		
Oilprice	21/5/1987–4/10/2011	6190	5.59	17.47	78.6 [0.00]		
Gold	4/1/2006–12/10/2011	1448	1.96	6.18	6.76 [0.01]		
Plain log-GARCH(1,1) models:							
		$\hat{\omega}$	$\hat{\alpha}_1$ (<i>s.e.</i>)	$\hat{\beta}_1$ (<i>s.e.</i>)	$\hat{E}(\ln \eta_t^2)$		
FTSE100		0.067	0.047 (0.006)	0.942 (0.008)	−1.415		
SP100		0.072	0.047 (0.006)	0.945 (0.008)	−1.513		
Apple		0.055	0.032 (0.005)	0.963 (0.007)	−1.375		
USD/EUR		0.025	0.022 (0.005)	0.971 (0.007)	−1.380		
Oilprice		0.069	0.037 (0.005)	0.952 (0.007)	−1.401		
Gold		0.054	0.032 (0.007)	0.958 (0.010)	−1.522		
Asymmetric log-GARCH(1,1) models:							
		$\hat{\omega}$	$\hat{\alpha}_1$ (<i>s.e.</i>)	$\hat{\beta}_1$ (<i>s.e.</i>)	$\hat{\gamma}_1$ (<i>s.e.</i>)	$\hat{\lambda}_1$ (<i>s.e.</i>)	$\hat{E}(\ln \eta_t^2)$
FTSE100		−0.126	0.003 (0.006)	0.945 (0.006)	0.069 (0.010)	0.362 (0.035)	−1.365
SP100		−0.112	0.011 (0.007)	0.934 (0.008)	0.070 (0.012)	0.374 (0.040)	−1.453
Apple		0.046	0.028 (0.007)	0.954 (0.009)	0.016 (0.007)	0.053 (0.019)	−1.361
USD/EUR		0.013	0.024 (0.007)	0.970 (0.007)	−0.003 (0.010)	0.021 (0.027)	−1.374
Oilprice		0.058	0.030 (0.006)	0.952 (0.007)	0.014 (0.008)	0.024 (0.016)	−1.396
Gold		0.063	0.036 (0.009)	0.961 (0.011)	−0.012 (0.014)	−0.025 (0.049)	−1.517

T , number of non-missing returns. s^2 , sample variance. s^4 , sample kurtosis. ARCH, Ljung and Box (1979) test statistic of first-order serial correlation in the squared return. p -val, the p -value of the test-statistic. *s.e.*, approximate standard errors (obtained via the numerically estimated Hessian) of estimate. All computations in R (R Core Team (2014)), estimation with the `lgarch` function (Sucarrat (2015)).

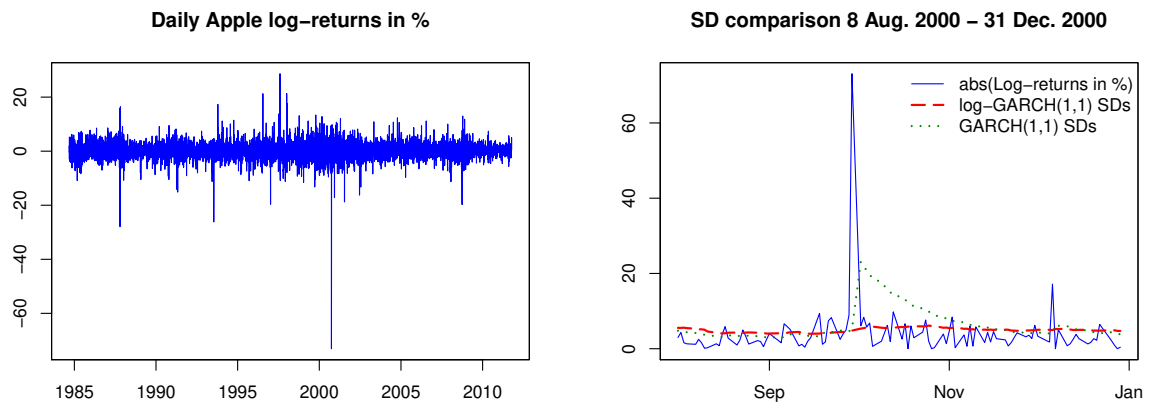


Figure 1: Daily Apple log-return in % (left graph) and a snapshot (right graph) of the period before and after a profit warning 28 September 2000, see Section 1. Datasource: Yahoo Finance

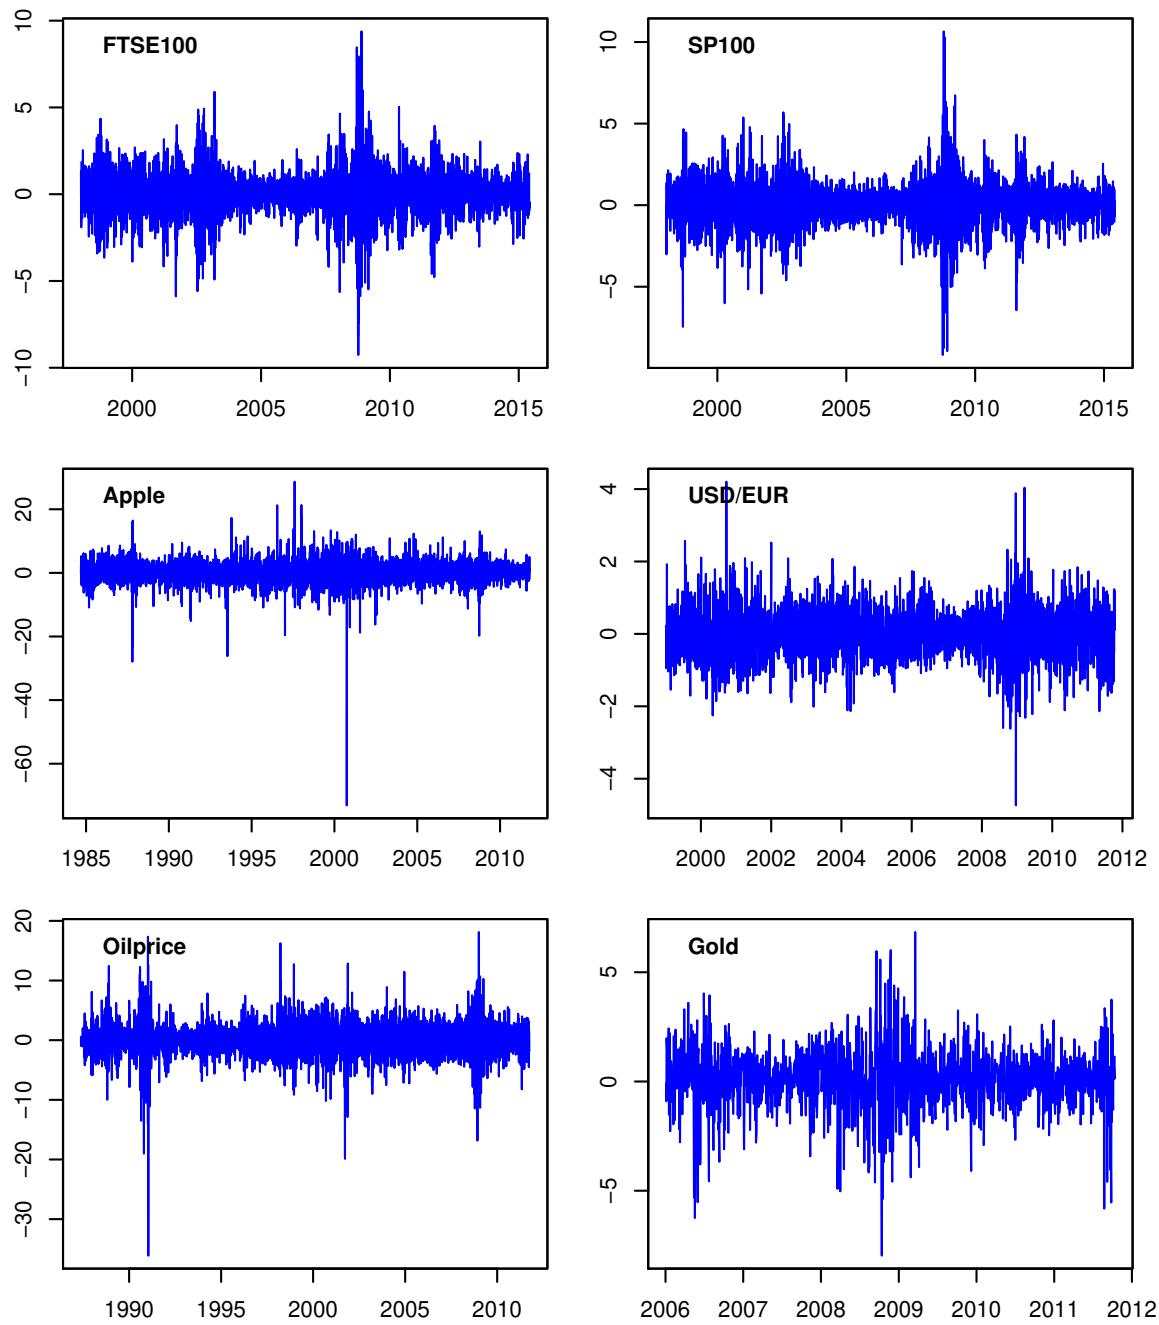


Figure 2: Daily financial log-returns in %, see Section 2.5 and Table 1. Datasources: Bloomberg, Yahoo Finance, European Central Bank, US Energy Information Agency and Kitco