



Munich Personal RePEc Archive

Unbiased QML Estimation of Log-GARCH Models in the Presence of Zero Returns

Sucarrat, Genaro and Escribano, Alvaro

BI Norwegian Business School, Universidad Carlos III de Madrid

9 September 2013

Online at <https://mpra.ub.uni-muenchen.de/50699/>

MPRA Paper No. 50699, posted 16 Oct 2013 08:20 UTC

Unbiased QML Estimation of Log-GARCH Models in the Presence of Zero Returns

Genaro Sucarrat¹ and Alvaro Escribano²

9th September 2013

Abstract

A critique that has been directed towards the log-GARCH model is that its log-volatility specification does not exist in the presence of zero returns. A common “remedy” is to replace the zeros with a small (in the absolute sense) non-zero value. However, this renders Quasi Maximum Likelihood (QML) estimation asymptotically biased. Here, we propose a solution to the case where actual returns are equal to zero with probability zero, but zeros nevertheless are observed because of measurement error (due to missing values, discreteness approximation error, etc.). The solution treats zeros as missing values and handles these by combining QML estimation via the ARMA representation with the Expectation-maximisation (EM) algorithm. Monte Carlo simulations confirm that the solution corrects the bias, and several empirical applications illustrate that the bias-correcting estimator can make a substantial difference.

JEL Classification: C22, C58

Keywords: ARCH, exponential GARCH, log-GARCH, ARMA, Expectation-Maximisation (EM)

1	Introduction	2
2	Model, problem and solution	4
2.1	The log-GARCH model	4
2.2	The effect of zeros – a Monte Carlo study	5
2.3	A solution based on the EM-algorithm	7
3	Empirical illustrations	8
4	Conclusions	9
	References	12

¹Corresponding author. Department of Economics, BI Norwegian Business School, Nydalsveien 37, 0484 Oslo, Norway. Email genaro.sucarrat at bi.no, phone +47+46410779, fax +47+23264788. Webpage: <http://www.sucarrat.net/>.

²Department of Economics, Universidad Carlos III de Madrid (Spain), email: alvaroe@eco.uc3m.es.

1 Introduction

Models in the Autoregressive Conditional Heteroscedasticity (ARCH) class due to [Engle \(1982\)](#) have been extensively used to model the time-varying volatility of financial return (see [Francq and Zakoian \(2010\)](#) for a recent survey of ARCH models). In particular, the first-order Generalised ARCH model of [Bollerslev \(1986\)](#), i.e. the GARCH(1,1), has established itself as an almost unquestionable benchmark. [Pantula \(1986\)](#), [Geweke \(1986\)](#) and [Milhøj \(1987\)](#) independently proposed specifications within the log-ARCH class of models as an alternative to non-exponential ARCH models. Their main motivation was to ensure the positivity of fitted volatilities – this is not guaranteed in non-exponential ARCH models (in particular when additional exogenous or predetermined conditioning information is added), to allow for richer dynamics (e.g. negative ARCH parameters for cyclical or contrarian dynamics) and to enable tests for integrated log-variance via [Dickey and Fuller \(1979\)](#) tests for unit roots. [Engle and Bollerslev \(1986\)](#), however, argued against log-ARCH models because of the possibility of applying the log-operator on zero-values.³ This occurs whenever the return or de-meaned return equals zero. Subsequently [Nelson \(1991\)](#) proposed an alternative exponential ARCH specification, the EGARCH model, where the problem is sidestepped by replacing the problematic term with an expression that does not involve the log operator. This solution, however, comes at a considerable cost: Restrictive assumptions are needed to ensure that QML estimation provides consistent and asymptotically normal estimates ([Wintenberger \(2012\)](#)), and unconditional moments (e.g. the unconditional variance of returns) will generally not exist for t -distributed densities (see condition (A1.6) and the subsequent discussion in [Nelson \(1991, p. 365\)](#)).

Zero returns occur in two different types of situations. In the first the zero-probability of actual return is zero, but zeros are nevertheless observed due to, say, missing values, discreteness approximation error and other data issues. For example, in financial markets prices are usually quoted with a few digits only (typically two). Financial returns are thus often measured as zero even though the true returns are non-zero. One may thus argue that zeros should be treated as missing values instead of zeros. Similarly, missing quotes or transaction prices are typically replaced by the previous observation, which in many cases results in an observed zero return even though the actual one is non-zero. Finally, impulse dummies are sometimes used to de-mean returns in the conditional mean. This leads to de-meaned returns equal to zero. When the impulse dummies are intended to neutralise the effect of large

³Another critique that has been directed towards the log-GARCH (e.g. [Teräsvirta \(2009\)](#)) is that the first unconditional autocorrelations of the squared returns, a measure of volatility persistence, can be unreasonably high. But this only occurs in very specific cases: The log-GARCH class allows for a much larger range of autocorrelation patterns than ordinary GARCH models, since the autocorrelation pattern depends on the shape of the conditional density (the more fat-tailed, the lower correlations) in addition to the persistence parameters, see [Sucarrat et al. \(2013\)](#).

outliers or “jumps” – this is often the motivation in macroeconomics and finance, then one may argue that the zeros should be treated as missing observations of actual (de-meaned) returns. The second type of situation in which zero returns occur is when the zero-probability of actual return is truly non-zero. This type of situation is addressed in [Sucarrat \(2013\)](#). Here, the focus is on the first type of situation.

Two Quasi Maximum Likelihood Estimators (QMLEs) have been proposed for the log-GARCH model. [Sucarrat et al. \(2013\)](#) prove that QMLE via the ARMA representation provides consistent and asymptotically normal estimates by appropriately adjusting for the intercept bias in the log-volatility specification induced by the ARMA representation. [Francq et al. \(2013\)](#) prove consistency and asymptotic normality of the QMLE when estimation is undertaken via the density of conditional (de-meaned) return. Both estimation schemes are valid under mild assumptions, both rely on the assumption that the probability of a zero (de-meaned) return is zero and both estimators produce asymptotically biased estimates in the presence of zeros.

This paper makes two contributions. First, we quantify the bias produced by zeros in a Monte Carlo study. The results show that the downwards bias of volatility increases with the number of zeros, that the reaction to shocks is underestimated and that the empirical standard errors are larger in the presence of zeros. The extent of these features depend on the parameter values, on whether the conditional density is fat-tailed or not and on the type of estimator. Second, we propose an asymptotically unbiased QMLE procedure. The procedure treats zeros as missing observations, and combines the [Sucarrat et al. \(2013\)](#) QMLE with the Expectation-Maximisation (EM) algorithm to handle the missing observations. A similar procedure cannot be devised for the [Francq et al. \(2013\)](#) QMLE (see [Section 2.3](#)), since it does not make use of the ARMA representation. A Monte Carlo study confirms that the EM-based estimator corrects for the bias, and several empirical applications illustrate how much the parameter estimates and the fitted conditional standard deviations can differ in practice.

The rest of the paper is organised as follows. The next section, [Section 2](#), provides an overview of the log-GARCH model, studies the effect of zeros by means of a Monte Carlo study, and presents the unbiased QMLE procedure that we propose together with a Monte Carlo study of its properties. [Section 3](#) contains the empirical illustrations of how much the parameter estimates and the fitted conditional standard deviations can differ if zeros are not correctly accounted for. The illustrations are from a variety of very different markets. Finally, [Section 4](#) concludes. Tables and Figures are located at the end.

2 Model, problem and solution

2.1 The log-GARCH model

If ϵ_t denotes financial return (possibly de-meant), then the log-GARCH(P, Q) model is given by

$$\epsilon_t = \sigma_t z_t, \quad z_t \sim IID(0, 1), \quad Prob(z_t = 0) = 0, \quad \sigma_t > 0, \quad (1)$$

$$\ln \sigma_t^2 = \alpha_0 + \sum_{p=1}^P \alpha_p \ln \epsilon_{t-p}^2 + \sum_{q=1}^Q \beta_q \ln \sigma_{t-q}^2, \quad t \in \mathbb{Z}, \quad (2)$$

where P is the ARCH order and Q is the GARCH order. Denoting $P^* = \max\{P, Q\}$, if the roots of the lag polynomial $1 - (\alpha_1 + \beta_1)L - \dots - (\alpha_{P^*} + \beta_{P^*})L^{P^*}$ are all greater than 1 in modulus and if $|E(\ln z_t^2)| < \infty$, then $\ln \sigma_t^2$ is stable. In the context of log-GARCH models, the so-called inlier issue (see [Breidt and Carriquiry \(1996\)](#) for a discussion in a Stochastic Volatility (SV) context) amounts to whether $E(\ln z_t^2)$ exists. For the Student's t density and for the Generalised Error Distribution (GED), the two most common distributions in finance, $E(\ln z_t^2)$ generally exists. [Francq et al. \(2013\)](#) provide conditions for the existence of log-moments more generally.

If $|E(\ln z_t^2)| < \infty$, then the ARMA(P, Q) representation of the log-volatility specification (2) exists almost surely and is given by

$$\ln \epsilon_t^2 = \phi_0 + \sum_{p=1}^P \phi_p \ln \epsilon_{t-p}^2 + \sum_{q=1}^Q \theta_q u_{t-q} + u_t, \quad u_t \sim IID(0, \sigma_u^2), \quad t \in \mathbb{Z}, \quad (3)$$

where

$$\phi_0 = \alpha_0 + (1 - \sum_{q=1}^Q \beta_q) \cdot E(\ln z_t^2), \quad (4)$$

$$\phi_p = \alpha_p + \beta_p, \quad 1 \leq p \leq P, \quad (5)$$

$$\theta_q = -\beta_q, \quad 0 \leq q \leq Q, \quad (6)$$

$$u_t = \ln z_t^2 - E(\ln z_t^2). \quad (7)$$

In other words, consistent and asymptotically normal estimates of all the ARMA parameters – and hence all the log-GARCH parameters except the log-volatility intercept α_0 – are thus readily obtained via usual ARMA estimation methods (e.g. Gaussian QMLE) subject to appropriate assumptions, see e.g. [Brockwell and Davis \(2006\)](#). For a consistent estimate of α_0 , however, a consistent estimate of $E(\ln z_t^2)$ is needed. [Sucarrat et al. \(2013\)](#) prove that a simple estimator made up of the ARMA residuals \hat{u}_t provides a consistent and asymptotically normal estimate of $E(\ln z_t^2)$ under mild assumptions. As a consequence, they prove consistency and asymp-

otic normality of the log-GARCH(P, Q) model via the ARMA representation for a range of ARMA estimators, including the Gaussian QMLE. Additional terms, e.g. asymmetry/leverage terms, or exogenous or predetermined conditioning information (i.e. “X”), can also be added without affecting the relationship between the log-GARCH and ARMA parameters, nor the structure of the bias-correction procedure. So estimation via the ARMA representation and subsequent bias-correcting can be generalised to both univariate and multivariate log-GARCH-X models. In the (empirical) presence of zeros, however, the QMLE via the ARMA representation will be asymptotically biased if the zeros are replaced with non-zero values.

Francq et al. (2013) propose a slightly different version of the (symmetric) log-volatility specification. In their setup (2) is replaced by

$$\ln \sigma_t^2 = \alpha_0 + \sum_{p=1}^P \alpha_p I_{\{z_{t-p} \neq 0\}} \ln \epsilon_{t-p}^2 + \sum_{q=1}^Q \beta_q \ln \sigma_{t-q}^2, \quad \forall t \in \mathbb{Z}, \quad (8)$$

where $I_{\{z_{t-p} \neq 0\}}$ is an indicator function equal to 0 if $z_{t-p} = 0$, and 1 otherwise. Of course, theoretically (2) and (8) are equal almost surely. In empirical practice, however, the latter avoids the problem of possibly applying the natural logarithm operator on zero values. Nevertheless, since the Francq et al. (2013) QMLE (which is in ϵ_t rather than via the ARMA representation) also relies on the assumption $Prob(z_t = 0) = 0$, the empirical presence of zeros also leads to asymptotically biased estimates.

2.2 The effect of zeros – a Monte Carlo study

To study the effect of zeros a model of why they arise is needed. To this end we distinguish between the actual return ϵ_t , which is governed by (1)-(2), and observed return $\tilde{\epsilon}_t$ which is given by

$$\tilde{\epsilon}_t = \epsilon_t I_t, \quad I_t \in \{0, 1\}, \quad (9)$$

where I_t is an independent – but not necessarily identical – series with zero-probabilities $\pi_{0t} \geq 0, t \in \mathbb{Z}$. In other words, our model of observed return allows for a time-varying zero probability, π_{0t} , of zero occurrences that are independent. It is important that the zeros are determined independently of the process that determines volatility. Otherwise the zero process would itself be part of the DGP. However, the independence assumption can be relaxed to contemporaneous independence between z_t and I_t conditional on the past, by appropriately adapting the framework in Sucarrat (2013). In other words, volatilities and zero-probabilities can be a function of past zero-probabilities and volatilities in a mutually dependent manner. Zero returns that are the result of missing values, discreteness approximation error, impulse dummies in the mean specification and so on, are therefore viewed as occurring independently of each other, but possibly with a time-varying probability.

To shed light on the effect of observed zeros we undertake a Monte Carlo study for both the [Sucarrat et al. \(2013\)](#) QMLE and the [Francq et al. \(2013\)](#) QMLE. In the simulations the Data Generating Process (DGP) of return ϵ_t is given by the log-GARCH(1,1) specification

$$\epsilon_t = \sigma_t z_t, \quad z_t \sim IID(0, 1), \quad Prob(z_t = 0) = 0, \quad (10)$$

$$\ln \sigma_t^2 = \alpha_0 + \alpha_1 \ln \epsilon_{t-1}^2 + \beta_1 \ln \sigma_{t-1}^2 \quad (11)$$

for empirically relevant combinations of the parameters α_0 , α_1 and β_1 (these combinations are referred to as Experiments A, B and C), and with the zero probability being constant over time and equal to either 0, 0.05, 0.10 or 0.20. For the [Francq et al. \(2013\)](#) QMLE the zeros of observed return $\tilde{\epsilon}_t$ are simply not included in the recursion because of the indicator function in the log-volatility specification (8). For the [Sucarrat et al. \(2013\)](#) QMLE, however, estimation is undertaken by means of the zero-adjusted return

$$\tilde{\epsilon}_t = \begin{cases} \epsilon_t & \text{if } I_t = 1, \\ k & \text{if } I_t = 0, \end{cases} \quad (12)$$

where k is a non-zero real number. In other words, the log-volatility specification that is used for the recursions in the estimations is $\ln \sigma_t^2 = \alpha_0 + \alpha_1 \ln \tilde{\epsilon}_{t-1}^2 + \beta_1 \ln \sigma_{t-1}^2$. Clearly the choice of k will influence the results. In particular, the closer to zero, the larger the bias will be. In our simulations we set k equal to the minimum of the absolute non-zero values of observed return $\tilde{\epsilon}_t$.

Table 1 contains the results for the [Francq et al. \(2013\)](#) QMLE, whereas Table 2 contains the results for the [Sucarrat et al. \(2013\)](#) QMLE. For the first estimator the effect of zeros is straightforward: The higher zero probability, the greater bias, and the bias is higher when the conditional density is fat-tailed (i.e. $t(5)$). This is particularly clear from Figure 1, which represents the bias graphically. The log-volatility intercept α_0 is biased downwards, which means volatility will generally be biased downwards in the presence of zeros. The ARCH parameter α_1 , which controls the impact of shocks on volatility, is also biased downwards. The presence of zeros thus means volatility will be under-responsive to shocks. This effect is exacerbated by the upward bias of β_1 , since this parameter controls the effect of the long-term component of volatility. Finally, the empirical standard errors are higher in the presence of zeros, and increasing the zero probability generally increases the standard errors.

For the second estimator the biases are generally bigger compared with the [Francq et al. \(2013\)](#) QMLE, but not always as straightforward. This is most readily seen in Figure 2. Just as for the first estimator higher zero probability means larger negative bias for both α_0 and α_1 , although the bias is not always higher for $t(5)$. For β_1 , however, the effect of zeros is more complex since the bias can change sign. Finally, a non-zero probability means the empirical standard errors are always higher than when the zero probability is zero, although the increase is not always

monotone.

2.3 A solution based on the EM-algorithm

The actual return ϵ_t is correctly observed whenever $I_t = 1$ in (9). Whenever $I_t = 0$, however, then the actual return ϵ_t is incorrectly observed or “missing”. A common solution to missing observations is the Expectation-Maximisation (EM) algorithm, see Casella and Berger (2002). There is a voluminous literature on how missing observations can be handled in ARMA models in combination with the EM-algorithm, see e.g. Jones (1980), Kohn and Ansley (1986), Gomez and Maravall (1994), and Brockwell and Davis (2006). So missing observations are straightforwardly handled in the Sucarrat et al. (2013) QMLE, since estimation is via the ARMA representation. Specifically, in the ARMA recursion of (3) the missing values of $\ln \epsilon_t^2$ are replaced by the conditional expectation $E(\ln \epsilon_t^2 | \mathcal{I}_{t-1})$, i.e. the fitted value of the ARMA representation, where \mathcal{I}_{t-1} is the conditioning information. Next, after estimation of the ARMA representation, the ARMA residuals \hat{u}_t where returns are non-zero are used for the estimation of $E(\ln z_t^2)$. Consistent and asymptotically normal estimates of the log-GARCH parameters are then obtained via the formulas in (4)-(6). In the Francq et al. (2013) QMLE, by contrast, the EM-algorithm cannot be used, or at least not in a way that we are aware of. The reason for this is that an estimate of $\ln \epsilon_t^2$ is needed as a replacement for the missing observations in the recursion of the log-volatility specification (8), and this is not provided by the estimator when it is interpreted as a QMLE. In the specific case where the Francq et al. (2013) estimator is interpreted an *exact* MLE, however, then the EM-algorithm is available. In that case z_t is standard normal, so $E(\ln z_t^2) = -1.27$ and hence $E(\ln \epsilon_t^2 | \mathcal{I}_{t-1}) = \ln \sigma_t^2 - 1.27$.

To study the properties of the Sucarrat et al. (2013) QMLE in combination with the EM-algorithm we undertake a Monte Carlo experiment similar to the one in the previous subsection. Table 3 contains the results and Figure 3 represents the finite sample bias graphically. Compared with Figure 2 it is clear that the EM-algorithm corrects the Sucarrat et al. (2013) bias for all three parameters. Compared with the case where there are no zeros the finite sample biases increase slightly (and more so for $z_t \sim t(5)$) as the zero probability increases. But this is to be expected since observations are lost by treating zeros as missing values. The empirical standard errors are virtually unaffected as the the zero probability increases, which is in stark contrast to the QMLEs without the EM-algorithm. Finally, compared with the Francq et al. (2013) QMLE the finite sample bias is substantially smaller for the location-parameter α_0 , i.e. the most important parameter in determining the level of volatility.

3 Empirical illustrations

In this section we illustrate the difference in parameter estimates and fitted conditional standard deviations for five daily financial returns: The Apple stock, the EUR/USD exchange rate, the Standard and Poor’s 500 stock market index (SP500), the WTI oil price and the London gold price. This small selection of returns accounts for a variety of market characteristics. For example, whereas the EUR/USD is traded in a global market almost continuously 24-hours a day and seven days a week – possibly with thousands of trades per second, the London Gold price is only fixed twice a day, and presumably not on Bank holidays and in weekends. The sources of the data are Yahoo Finance (<http://finance.yahoo.com>) for the Apple and SP500 series, the European Central Bank (<http://www.ecb.int/>) for the EUR/USD series, the US Energy Information Agency (<http://www.eia.gov/>) for the WTI crude oil price (in USD) per barrel series, and Kitco (<http://www.kitco.com/>) for the London afternoon (i.e. PM) gold price series.

The sample dates and the descriptive statistics of the returns are contained in Table 4, whereas Figure 4 contains graphs of the returns. They confirm that the returns exhibit the usual properties of excess kurtosis compared with the normal, and ARCH as measured by first order serial correlation in the squared return. The number of zeros varies from only 2 observations (about 0.1% of the sample) for SP500 to 294 observations (about 4% of the sample) for Apple. The reasons for each zero are likely to differ substantially both within and across markets. We do not try to identify these reasons, since our main objective is to illustrate how the estimates and fitted conditional standard deviations differ according to estimation method.

Table 5 contains the estimates of the log-GARCH(1,1) specification

$$r_t = \epsilon_t, \quad \epsilon_t = \sigma_t z_t, \quad z_t \sim IID(0, 1), \quad (13)$$

$$\ln \sigma_t^2 = \alpha_0 + \alpha_1 \ln \epsilon_{t-1}^2 + \beta_1 \ln \sigma_{t-1}^2, \quad (14)$$

where r_t is the log-return in percent (i.e., the log-difference of the financial price multiplied by 100). In all cases one or more estimates differ already at the second decimal. This is the case even for SP500, where there are only two zeros. The smallest numerical differences are produced by oil, whereas the biggest are produced by the EUR/USD exchange rate. This is noteworthy because the proportion (about 1%) of oil zeros is slightly higher than for EUR/USD (about 0.9%), and much higher (about ten times) than for SP500. In other words, the number of zeros is not always the main source of the estimation bias. This is in accordance with the Monte Carlo studies, which revealed that parameter estimates and conditional density are sometimes more important. With respect to the estimate of the ARCH parameter α_1 , which controls the short-term impact of shocks or large (in absolute value) returns, the EM-estimates are substantially higher except for oil where they are only slightly higher. For the GARCH parameter β_1 by contrast, all the EM-

estimates are lower – sometimes substantially (e.g. EUR/USD).

Descriptive statistics and graphs of the fitted conditional standard deviations, their differences and their ratios, are contained in Table 6 and Figures 5-7. They clearly suggest that estimation method can matter a lot, both nominally and in relative terms. For example, for Apple the EM-estimates yield fitted conditional standard deviations that are at most 2.14 times higher, and the maximum nominal difference is 2.44. Such differences can make a huge difference in risk analysis and asset pricing. The Apple graphs also reveals what seems to be an inverse tendency. In the beginning of the sample the EM-estimates produce higher fitted conditional standard deviations. However, this is reversed in the second part of the sample. A possible reason is that there are fewer zeros in the second part of the sample. For most returns the average fitted conditional standard deviation is higher for the EM-estimates. This is most clearly seen in the graph of EUR/USD, where the fitted conditional standard deviations produced by the EM-estimates are clearly above almost everywhere. The only case where the average difference is not positive is oil. There, the average is approximately equal to zero. But the ratio graph clearly shows that, in relative terms, the EM-estimates occasionally produce values that are up to 66% higher. So all in all the comparison of fitted conditional standard deviations show that the EM-estimates generally produce higher values, and sometimes much higher.

4 Conclusions

We propose an asymptotically unbiased QML procedure for log-GARCH models in the presence of zero returns. The procedure combines the [Sucarrat et al. \(2013\)](#) QMLE with the Expectation-Maximisation (EM) algorithm, a procedure that is not available for the [Francq et al. \(2013\)](#) QMLE. The reason for this is that the former estimator is via the ARMA representation, whereas the latter is not. The QML procedure relies on the assumption that the actual return is zero with zero probability, but accommodates that observed return can be non-zero due to, say, missing values, discreteness approximation error, impulse dummies in the mean specification or other data issues. The zeros are assumed independent but not necessarily identically distributed, as the zero probability can be time-varying and conditionally dependent on the past. (The counterpart problem where actual return can be zero with non-zero probability is solved in a companion paper ([Sucarrat \(2013\)](#)).) Our Monte Carlo simulations and our empirical illustrations show that volatility is generally underestimated when zeros are present, and that the impact of shocks on volatility is underestimated in the presence of zeros. In practice this means that the fitted conditional standard deviations are generally underestimated – sometimes substantially.

The results in this paper can be extended in at least three ways. First, it is straightforward to devise unbiased QML procedures for univariate and multivari-

ate log-GARCH-X models by combining the EM-algorithm with the methods proposed in [Sucarrat et al. \(2013\)](#), since the relationship between the log-GARCH and (V)ARMA parameters are not affected by the addition of exogenous or predetermined conditioning information (leverage, volatility proxies, volume or other information arrival indicators, seasonality terms, etc.). Second, as a direct consequence of the first, unbiased QMLEs for univariate and multivariate log-MEM-X models can be devised, where MEM is short for Multiplicative Error Models, see [Brownlees et al. \(2012\)](#). MEM-models are particularly suited for non-negative financial data like volume, durations and trades, and because of its structure a QMLE for log-GARCH-X models is also a QMLE for log-MEM-X models. There are often zeros and/or missing values in volume, duration and trade data. When these zeros can be viewed as measurement error or as a result of missing values, then the methods in this paper can be used to adjust for the bias created by the zeros. Third, if zeros are the result of measurement error, then they also lead to biased ML-estimates for other ARCH models, e.g. the GARCH of [Bollerslev \(1986\)](#), the EGARCH of [Nelson \(1991\)](#) and the Beta-t-EGARCH model of [Harvey \(2013\)](#). Although QML estimation is not available for the latter due to its nature, exact ML estimation in combination with the EM-algorithm can be used in all three classes.

References

- Bollerslev, T. (1986). Generalized autoregressive conditional heteroscedasticity. *Journal of Econometrics* 31, 307–327.
- Breidt, F. J. and A. L. Carriquiry (1996). Improved Quasi-Maximum Likelihood Estimation for Stochastic Volatility Models. In J. Lee, W. Johnson, and A. Zellner (Eds.), *Modeling and Prediction: Honoring Seymour Geisser*. Berlin: Springer.
- Brockwell, P. J. and R. A. Davis (2006). *Time Series: Theory and Methods*. New York: Springer. 2nd. Edition.
- 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*. New Jersey: Wiley.
- Casella, G. and R. Berger (2002). *Statistical Inference*. Pacific Grove, CA: Thomson Learning. 2nd. Edition.
- Dickey, D. and W. Fuller (1979). Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association* 74, 426–431.
- Engle, R. (1982). Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica* 50, 987–1008.

- Engle, R. F. and T. Bollerslev (1986). Modelling the persistence of conditional variances. *Econometric Reviews* 5, 1–50.
- Francq, C., O. Wintenberger, and J.-M. Zakoïan (2013). GARCH Models Without Positivity Constraints: Exponential or Log-GARCH? Forthcoming in *Journal of Econometrics*, <http://dx.doi.org/10.1016/j.jeconom.2013.05.004>.
- Francq, C. and J.-M. Zakoïan (2010). *GARCH Models*. New York: Marcel Dekker.
- Geweke, J. (1986). Modelling the Persistence of Conditional Variance: A Comment. *Econometric Reviews* 5, 57–61.
- Gomez, V. and A. Maravall (1994). Estimation, Prediction, and Interpolation for Nonstationary Series with the Kalman Filter. *Journal of the American Statistical Association* 89, 611–624.
- Harvey, A. C. (2013). *Dynamic Models for Volatility and Heavy Tails*. New York: Cambridge University Press.
- Jones, R. (1980). Maximum likelihood fitting of ARMA models to time series with missing observations.
- Kohn, R. and C. Ansley (1986). Estimation, Prediction, and Interpolation for ARIMA Models With Missing Data. *Journal of the American Statistical Association* 81, 751–761.
- Ljung, G. and G. Box (1979). On a Measure of Lack of Fit in Time Series Models. *Biometrika* 66, 265–270.
- 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.
- R Core Team (2013). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Sucarrat, G. (2013). Models of Zero-Augmented Financial Return with Time-Varying Zero-Probabilities. Work in progress.
- Sucarrat, G., Á. Escribano, and S. Grønneberg (2013). Estimation and Inference in Univariate and Multivariate Log-GARCH-X Models When the Conditional Density is Unknown. <http://www.sucarrat.net/>.

Teräsvirta, T. (2009). An introduction to univariate GARCH models. In T. Mikosch, T. Kreiss, J.-P. Davis, R. Andersen, and T. Gustav (Eds.), *Handbook of Financial Time Series*. Berlin: Springer.

Wintenberger, O. (2012). QML estimation of the continuously invertible EGARCH(1,1) model. Unpublished working paper.

Table 1: [Francq et al. \(2013\)](#) QMLE of the log-GARCH(1,1)

z_t	π_0	DGP (ID: $\alpha_0, \alpha_1, \beta_1$)	$m(\hat{\alpha}_0)$	$se(\hat{\alpha}_0)$	$m(\hat{\alpha}_1)$	$se(\hat{\alpha}_1)$	$m(\hat{\beta}_1)$	$se(\hat{\beta}_1)$	$m(\hat{\pi}_0)$
$N(0, 1)$	0.00	A: 0, 0.10, 0.80	-0.001	0.009	0.100	0.004	0.800	0.010	0.000
		B: 0, 0.05, 0.90	-0.001	0.006	0.050	0.003	0.900	0.008	0.000
		C: 0, 0.03, 0.95	-0.001	0.005	0.030	0.002	0.950	0.004	0.000
	0.05	A: 0, 0.10, 0.80	-0.021	0.010	0.096	0.004	0.807	0.011	0.050
		B: 0, 0.05, 0.90	-0.011	0.007	0.048	0.003	0.904	0.009	0.050
		C: 0, 0.03, 0.95	-0.007	0.005	0.028	0.002	0.953	0.005	0.050
	0.10	A: 0, 0.10, 0.80	-0.040	0.011	0.093	0.005	0.814	0.011	0.099
		B: 0, 0.05, 0.90	-0.020	0.008	0.046	0.003	0.908	0.008	0.099
		C: 0, 0.03, 0.95	-0.013	0.006	0.027	0.002	0.955	0.005	0.099
	0.20	A: 0, 0.10, 0.80	-0.075	0.014	0.088	0.005	0.827	0.013	0.199
		B: 0, 0.05, 0.90	-0.037	0.010	0.043	0.004	0.915	0.010	0.199
		C: 0, 0.03, 0.95	-0.022	0.008	0.025	0.003	0.960	0.005	0.199
$t(5)$	0.00	A: 0, 0.10, 0.80	0.000	0.019	0.100	0.007	0.801	0.016	0.000
		B: 0, 0.05, 0.90	-0.001	0.013	0.050	0.006	0.899	0.013	0.000
		C: 0, 0.03, 0.95	-0.002	0.012	0.030	0.004	0.949	0.009	0.000
	0.05	A: 0, 0.10, 0.80	-0.025	0.020	0.095	0.008	0.808	0.018	0.050
		B: 0, 0.05, 0.90	-0.014	0.015	0.048	0.006	0.903	0.014	0.050
		C: 0, 0.03, 0.95	-0.010	0.013	0.028	0.004	0.952	0.009	0.050
	0.10	A: 0, 0.10, 0.80	-0.049	0.023	0.091	0.008	0.816	0.021	0.100
		B: 0, 0.05, 0.90	-0.027	0.017	0.046	0.006	0.907	0.016	0.100
		C: 0, 0.03, 0.95	-0.019	0.015	0.026	0.004	0.955	0.009	0.100
	0.20	A: 0, 0.10, 0.80	-0.093	0.030	0.084	0.009	0.827	0.025	0.201
		B: 0, 0.05, 0.90	-0.051	0.024	0.042	0.006	0.912	0.018	0.201
		C: 0, 0.03, 0.95	-0.033	0.021	0.023	0.004	0.959	0.011	0.201

DGP, $\ln \sigma_t^2 = \alpha_0 + \alpha_1 \ln \epsilon_{t-1}^2 + \beta_1 \ln \sigma_{t-1}^2$ with $z_t \sim IID(0, 1)$ and $T = 10000$. ID, experiment identifier (i.e. A, B or C). $m(\cdot)$, sample average of the Monte Carlo estimates. $se(\cdot)$, sample standard deviation of the Monte Carlo estimates (division by S , not by $S - 1$, where $S = 100$ is the number of Monte Carlo simulations). $\hat{\pi}$, proportion of zeros. Simulations in *R* version 3.0.0, see [R Core Team \(2013\)](#).

Table 2: [Sucarrat et al. \(2013\)](#) QMLE of the log-GARCH(1,1), with zeros replaced by the minimum of the absolute non-zero values

z_t	π_0	DGP (ID: $\alpha_0, \alpha_1, \beta_1$)	$m(\hat{\alpha}_0)$	$se(\hat{\alpha}_0)$	$m(\hat{\alpha}_1)$	$se(\hat{\alpha}_1)$	$m(\hat{\beta}_1)$	$se(\hat{\beta}_1)$	$m(\hat{\pi}_0)$
$N(0, 1)$	0.00	A: 0, 0.10, 0.80	-0.003	0.016	0.100	0.007	0.797	0.020	0.000
		B: 0, 0.05, 0.90	-0.003	0.011	0.050	0.005	0.899	0.013	0.000
		C: 0, 0.03, 0.95	-0.003	0.008	0.030	0.004	0.949	0.007	0.000
	0.05	A: 0, 0.10, 0.80	-0.071	0.048	0.030	0.008	0.860	0.047	0.050
		B: 0, 0.05, 0.90	-0.057	0.125	0.016	0.005	0.914	0.101	0.050
		C: 0, 0.03, 0.95	-0.031	0.031	0.010	0.004	0.963	0.020	0.050
	0.10	A: 0, 0.10, 0.80	-0.091	0.066	0.017	0.008	0.871	0.069	0.099
		B: 0, 0.05, 0.90	-0.069	0.123	0.009	0.005	0.918	0.100	0.099
		C: 0, 0.03, 0.95	-0.074	0.212	0.007	0.004	0.945	0.112	0.099
	0.20	A: 0, 0.10, 0.80	-0.157	0.227	0.010	0.008	0.843	0.188	0.199
		B: 0, 0.05, 0.90	-0.123	0.261	0.005	0.006	0.892	0.201	0.199
		C: 0, 0.03, 0.95	-0.107	0.186	0.004	0.005	0.937	0.095	0.199
$t(5)$	0.00	A: 0, 0.10, 0.80	-0.002	0.017	0.101	0.008	0.798	0.018	0.000
		B: 0, 0.05, 0.90	-0.003	0.012	0.051	0.006	0.897	0.014	0.000
		C: 0, 0.03, 0.95	-0.004	0.011	0.030	0.004	0.948	0.008	0.000
	0.05	A: 0, 0.10, 0.80	-0.090	0.055	0.031	0.008	0.860	0.043	0.050
		B: 0, 0.05, 0.90	-0.055	0.046	0.016	0.006	0.925	0.038	0.050
		C: 0, 0.03, 0.95	-0.034	0.031	0.010	0.003	0.966	0.017	0.050
	0.10	A: 0, 0.10, 0.80	-0.129	0.102	0.019	0.008	0.859	0.080	0.100
		B: 0, 0.05, 0.90	-0.086	0.143	0.010	0.006	0.918	0.101	0.100
		C: 0, 0.03, 0.95	-0.062	0.160	0.007	0.005	0.959	0.076	0.100
	0.20	A: 0, 0.10, 0.80	-0.210	0.372	0.009	0.009	0.838	0.239	0.201
		B: 0, 0.05, 0.90	-0.137	0.298	0.005	0.007	0.903	0.181	0.201
		C: 0, 0.03, 0.95	-0.317	0.680	0.004	0.007	0.863	0.275	0.201

DGP, $\ln \sigma_t^2 = \alpha_0 + \alpha_1 \ln \epsilon_{t-1}^2 + \beta_1 \ln \sigma_{t-1}^2$ with $z_t \sim IID(0, 1)$ and $T = 10000$. ID, experiment identifier (i.e. A, B or C). $m(\cdot)$, sample average of the Monte Carlo estimates. $se(\cdot)$, sample standard deviation of the Monte Carlo estimates (division by S , not by $S - 1$, where $S = 100$ is the number of Monte Carlo simulations). $\hat{\pi}$, proportion of zeros. Simulations in *R* version 3.0.0, see [R Core Team \(2013\)](#).

Table 3: [Sucarrat et al. \(2013\)](#) QMLE of the log-GARCH(1,1) in combination with the EM-algorithm

z_t	π_0	DGP (ID: $\alpha_0, \alpha_1, \beta_1$)	$m(\hat{\alpha}_0)$	$se(\hat{\alpha}_0)$	$m(\hat{\alpha}_1)$	$se(\hat{\alpha}_1)$	$m(\hat{\beta}_1)$	$se(\hat{\beta}_1)$	$m(\hat{\pi}_0)$
$N(0,1)$	0.05	A: 0, 0.10, 0.80	-0.002	0.017	0.101	0.008	0.795	0.021	0.050
		B: 0, 0.05, 0.90	-0.003	0.012	0.051	0.006	0.895	0.015	0.050
		C: 0, 0.03, 0.95	-0.005	0.010	0.031	0.004	0.945	0.009	0.050
	0.10	A: 0, 0.10, 0.80	0.000	0.017	0.103	0.008	0.793	0.020	0.099
		B: 0, 0.05, 0.90	-0.002	0.011	0.052	0.006	0.895	0.015	0.099
		C: 0, 0.03, 0.95	-0.004	0.009	0.032	0.005	0.945	0.009	0.099
	0.20	A: 0, 0.10, 0.80	0.005	0.017	0.105	0.008	0.790	0.021	0.199
		B: 0, 0.05, 0.90	0.000	0.011	0.053	0.007	0.893	0.015	0.199
		C: 0, 0.03, 0.95	-0.002	0.009	0.033	0.005	0.944	0.010	0.199
$t(5)$	0.05	A: 0, 0.10, 0.80	0.000	0.018	0.103	0.008	0.794	0.019	0.050
		B: 0, 0.05, 0.90	-0.003	0.013	0.052	0.006	0.893	0.015	0.050
		C: 0, 0.03, 0.95	-0.006	0.013	0.032	0.005	0.943	0.011	0.050
	0.10	A: 0, 0.10, 0.80	0.002	0.018	0.104	0.009	0.793	0.020	0.100
		B: 0, 0.05, 0.90	-0.002	0.013	0.053	0.007	0.893	0.016	0.100
		C: 0, 0.03, 0.95	-0.005	0.014	0.033	0.005	0.943	0.012	0.100
	0.20	A: 0, 0.10, 0.80	0.007	0.019	0.106	0.009	0.790	0.020	0.201
		B: 0, 0.05, 0.90	0.000	0.013	0.054	0.007	0.891	0.016	0.201
		C: 0, 0.03, 0.95	-0.004	0.014	0.034	0.005	0.942	0.012	0.201

DGP, $\ln \sigma_t^2 = \alpha_0 + \alpha_1 \ln \epsilon_{t-1}^2 + \beta_1 \ln \sigma_{t-1}^2$ with $z_t \sim IID(0,1)$ and $T = 10000$. ID, experiment identifier (i.e. A, B or C). $m(\cdot)$, sample average of the Monte Carlo estimates. $se(\cdot)$, sample standard deviation of the Monte Carlo estimates (division by S , not by $S - 1$, where $S = 100$ is the number of Monte Carlo simulations). $\hat{\pi}$, proportion of zeros. Simulations in *R* version 3.0.0, see [R Core Team \(2013\)](#).

Table 4: Descriptive statistics of financial returns

	s^2	s^4	$ARCH_1$ [p-val]	T	0s	$\hat{\pi}$
Apple (10 Sep. 1984 – 23 Aug. 2013)	9.25	55.03	7.12 [0.01]	7303	294	0.040
EUR/USD (5 Jan. 1999 – 23 Aug. 2013)	0.43	5.44	150.63 [0.00]	3751	32	0.009
SP500 (4 Jan. 1999 – 23 Aug. 2013)	1.73	10.30	143.10 [0.00]	3684	2	0.001
Oil (5 Apr. 1983 – 19 Aug. 2013)	5.72	18.80	160.60 [0.00]	7621	73	0.010
Gold (4 Jan. 2006 – 23 Aug. 2013)	1.85	7.29	10.94 [0.00]	1929	20	0.010

s^2 , sample variance. s^4 , sample kurtosis. $ARCH_1$, [Ljung and Box \(1979\)](#) test statistic of first-order serial correlation in the squared return. T , number of returns. 0s, number of zero returns in the sample. $\hat{\pi}$, proportion of zero returns in the sample.

Table 5: Empirical estimates of log-GARCH(1,1) specification for five daily financial returns

	Method	$\hat{\alpha}_0$	$\hat{\alpha}_1$	$se(\hat{\alpha}_1)$	$\hat{\beta}_1$	$se(\hat{\beta}_1)$
Apple:	0-adj	0.034	0.014	0.003	0.983	0.005
	EM	0.071	0.040	0.005	0.953	0.006
EUR/USD:	0-adj	0.022	0.019	0.004	0.976	0.005
	EM	0.066	0.048	0.007	0.901	0.016
SP500:	0-adj	0.070	0.045	0.006	0.946	0.008
	EM	0.092	0.056	0.006	0.931	0.009
Oil:	0-adj	0.074	0.043	0.004	0.951	0.005
	EM	0.074	0.045	0.004	0.948	0.005
Gold:	0-adj	0.055	0.029	0.006	0.959	0.009
	EM	0.090	0.047	0.007	0.932	0.011

[Sucarrat et al. \(2013\)](#) QML estimation of the log-GARCH(1,1) specification $\ln \sigma_t^2 = \alpha_0 + \alpha_1 \ln \epsilon_{t-1}^2 + \beta \ln \sigma_{t-1}^2$. 0-adj, zero returns replaced by the minimum of the absolute non-zero value before estimation. EM, estimation with the EM-algorithm (i.e. zeros not replaced before estimation). $se(\cdot)$, standard error of estimate. All computations in *R* version 3.0.0, see [R Core Team \(2013\)](#).

Table 6: Descriptive statistics of fitted conditional standard deviations

		Mean	s^2	Max	Min
Apple:	0-adj	2.970	0.753	6.346	1.511
	EM	3.001	0.808	6.442	1.040
	Diff	0.030	0.388	2.420	-1.428
	Ratio	1.029	0.059	2.141	0.579
EUR/USD:	0-adj	0.643	0.018	1.171	0.360
	EM	0.813	0.019	1.399	0.471
	Diff	0.171	0.006	0.459	-0.104
	Ratio	1.283	0.021	1.742	0.860
SP500:	0-adj	1.191	0.327	4.730	0.437
	EM	1.225	0.352	5.013	0.417
	Diff	0.034	0.004	0.391	-0.237
	Ratio	1.027	0.002	1.378	0.907
Oil:	0-adj	2.197	0.968	7.530	0.410
	EM	2.189	0.902	7.415	0.432
	Diff	-0.008	0.018	1.040	-0.212
	Ratio	1.006	0.006	1.660	0.938
Gold:	0-adj	1.317	0.096	2.580	0.723
	EM	1.375	0.126	3.017	0.663
	Diff	0.058	0.012	0.676	-0.149
	Ratio	1.041	0.006	1.330	0.855

Mean, sample average. s^2 , sample variance. Max, maximum value. Min, minimum value. Diff, the difference between fitted conditional standard deviations: $\hat{\sigma}_{t,EM} - \hat{\sigma}_{t,0-adj}$. Ratio, the ratio between fitted conditional standard deviations: $\hat{\sigma}_{t,EM}/\hat{\sigma}_{t,0-adj}$. All computations in *R* version 3.0.0, see [R Core Team \(2013\)](#).

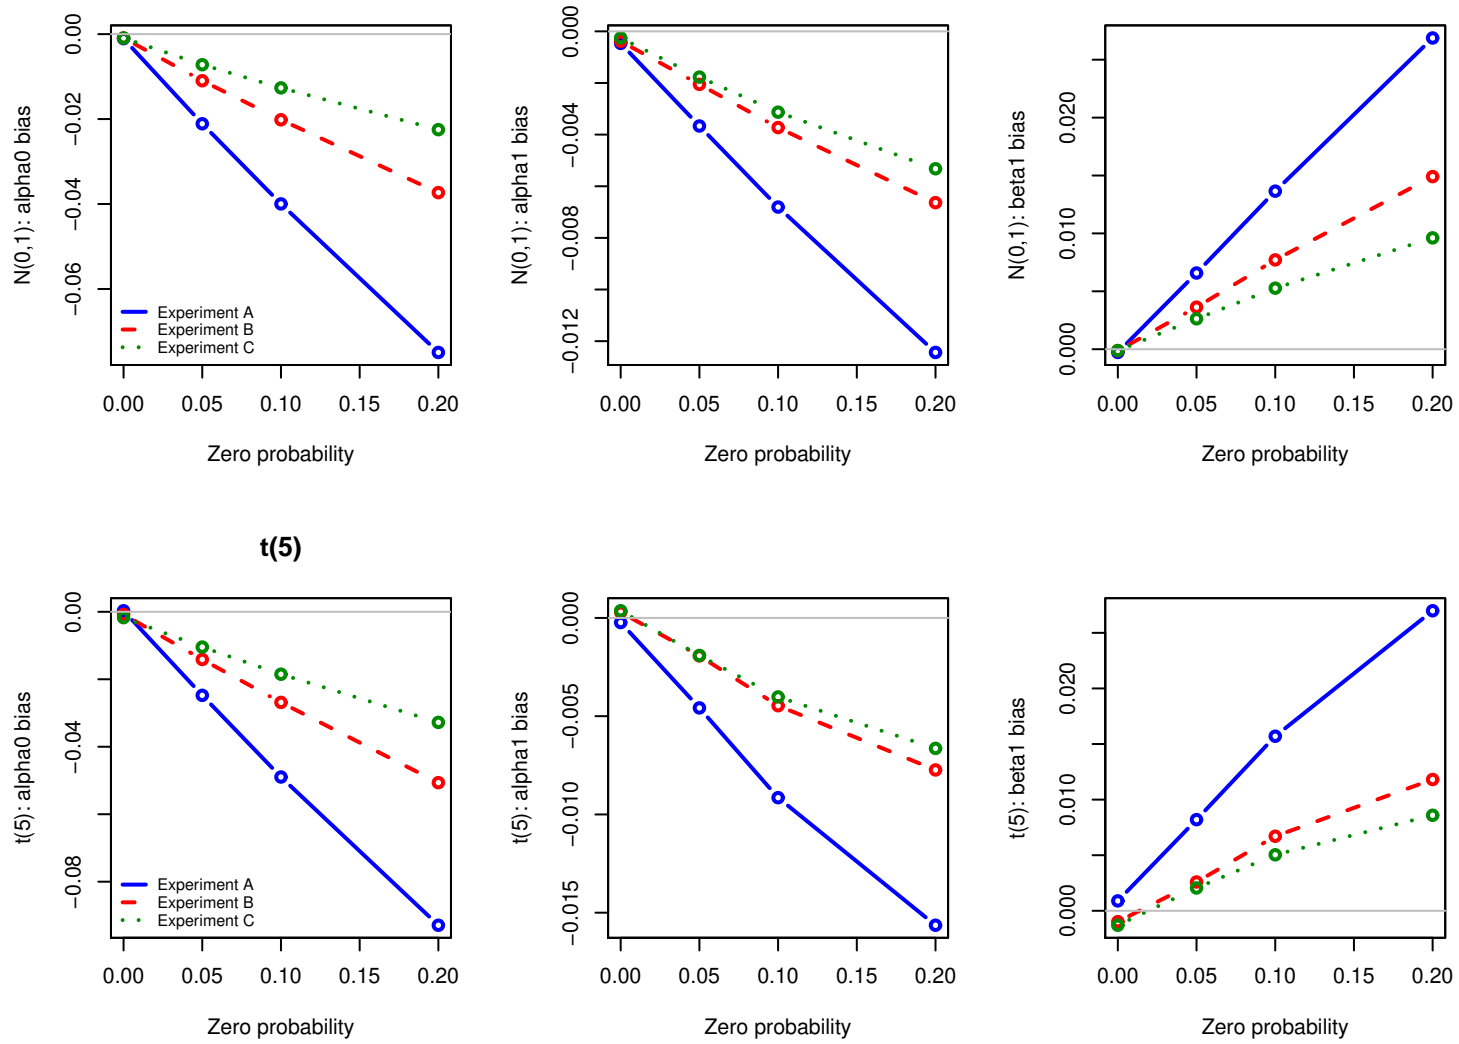


Figure 1: Bias ($estimate - true\ value$) of Francq et al. (2013) QMLE when $z_t \sim N(0, 1)$ (upper graphs) and $z_t \sim t(5)$ (lower graphs), see Table 1. Experiment A: Solid blue line, Experiment B: Dashed red line, Experiment C: Dotted green line

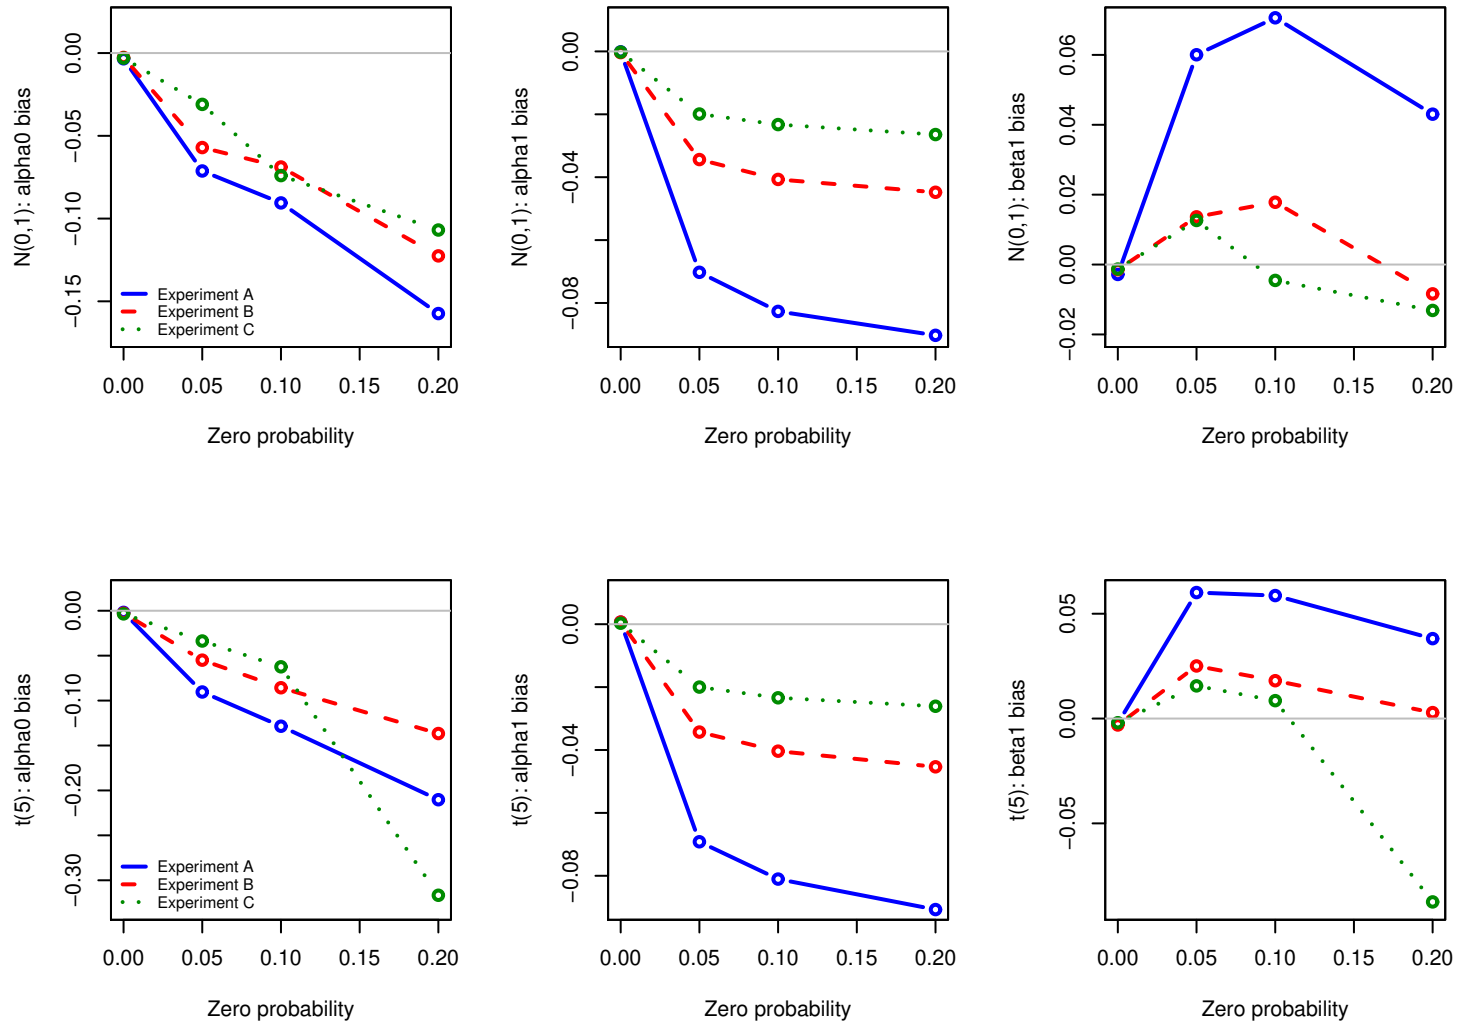


Figure 2: Bias ($estimate - true\ value$) of Sucarrat et al. (2013) QMLE when $z_t \sim N(0,1)$ (upper graphs) and $z_t \sim t(5)$ (lower graphs), see Table 2. Experiment A: Solid blue line, Experiment B: Dashed red line, Experiment C: Dotted green line

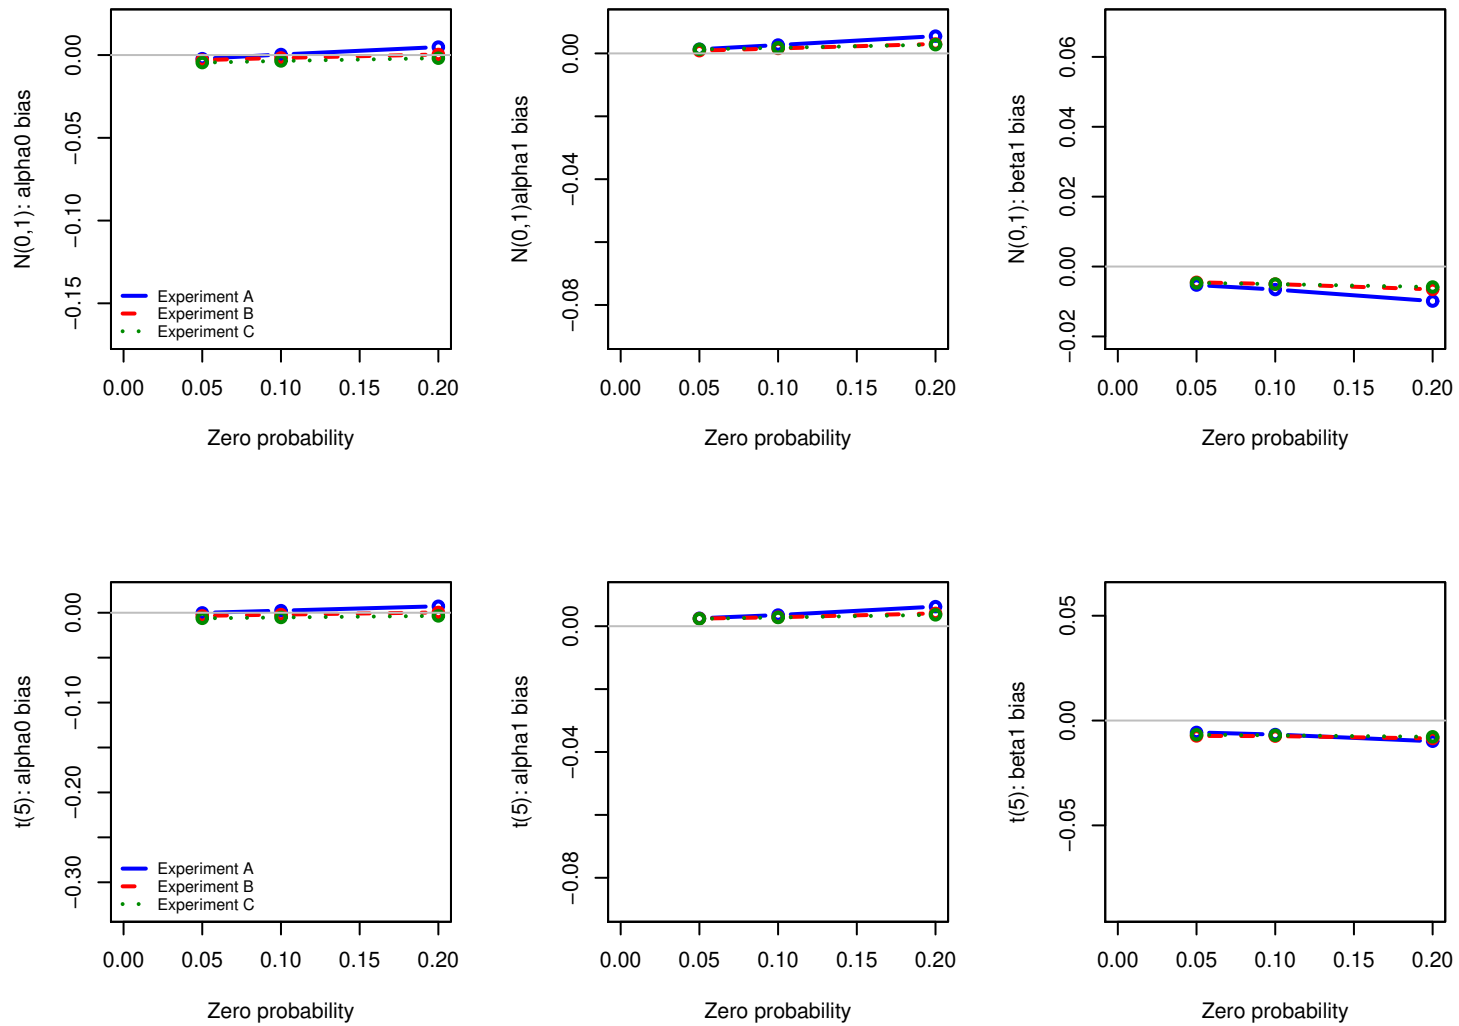


Figure 3: Bias ($estimate - true\ value$) of [Sucarrat et al. \(2013\)](#) QMLE combined with EM-algorithm for $z_t \sim N(0, 1)$ (upper graphs) and $z_t \sim t(5)$ (lower graphs), see Table 3. Experiment A: Solid blue line, Experiment B: Dashed red line, Experiment C: Dotted green line

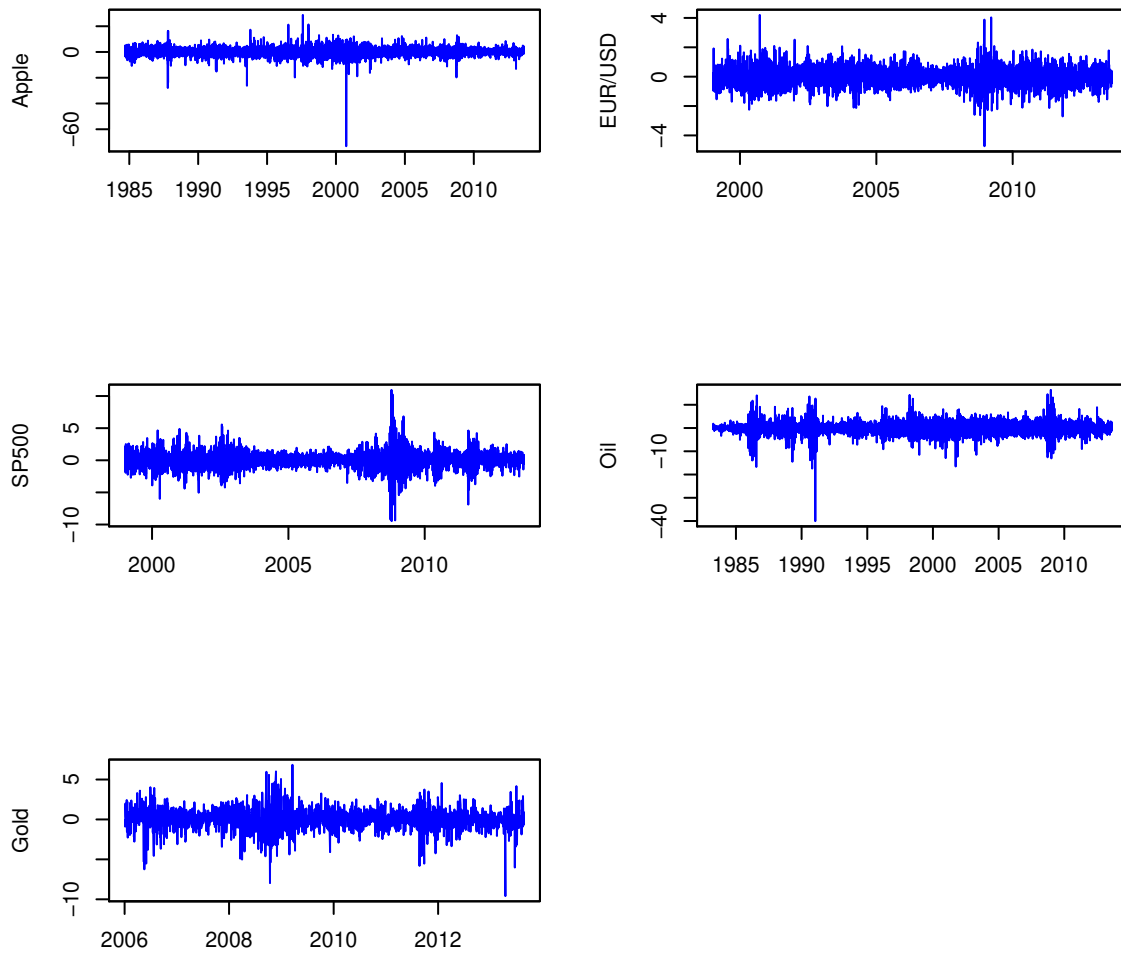


Figure 4: Daily financial log-returns in percent

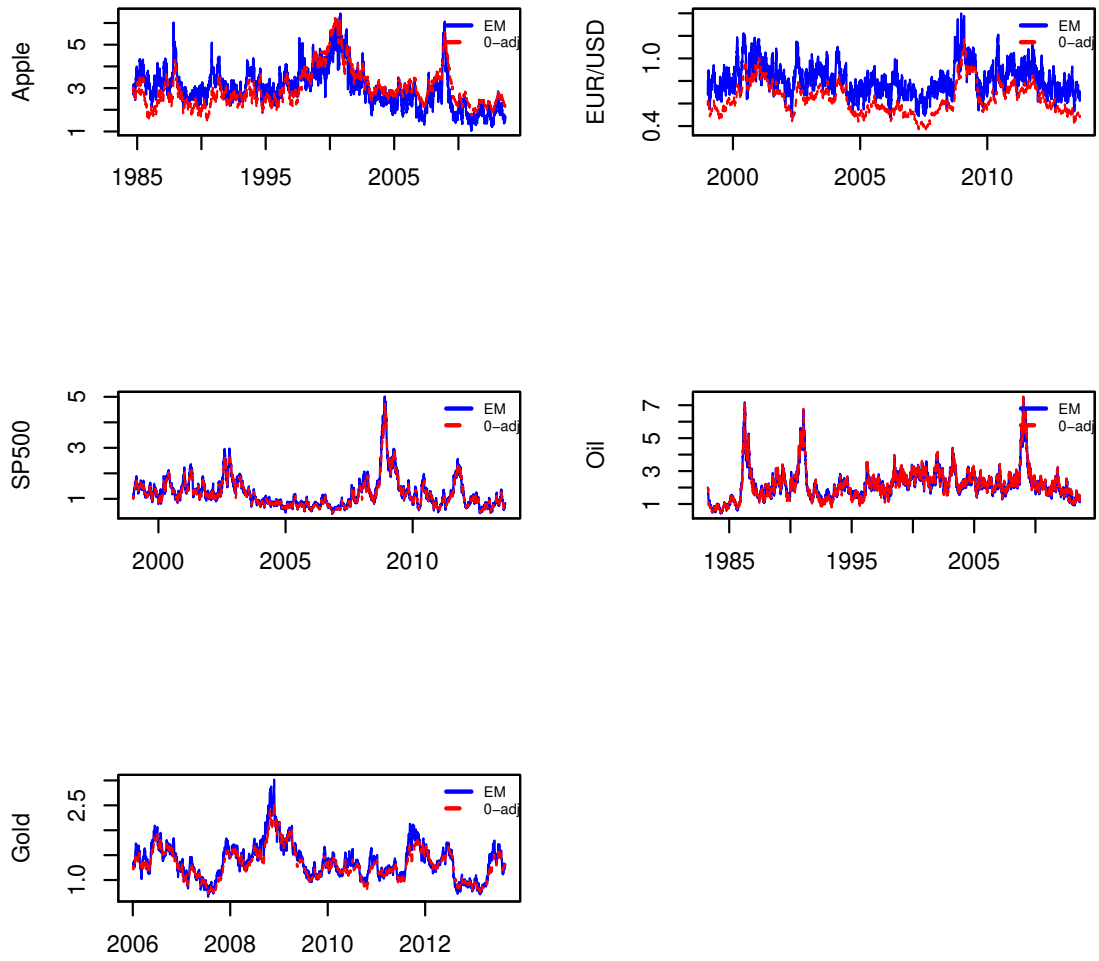


Figure 5: Fitted conditional standard deviations ($\hat{\sigma}_{t,EM}$ solid blue line, $\hat{\sigma}_{t,0\text{-adj}}$ dashed red line)

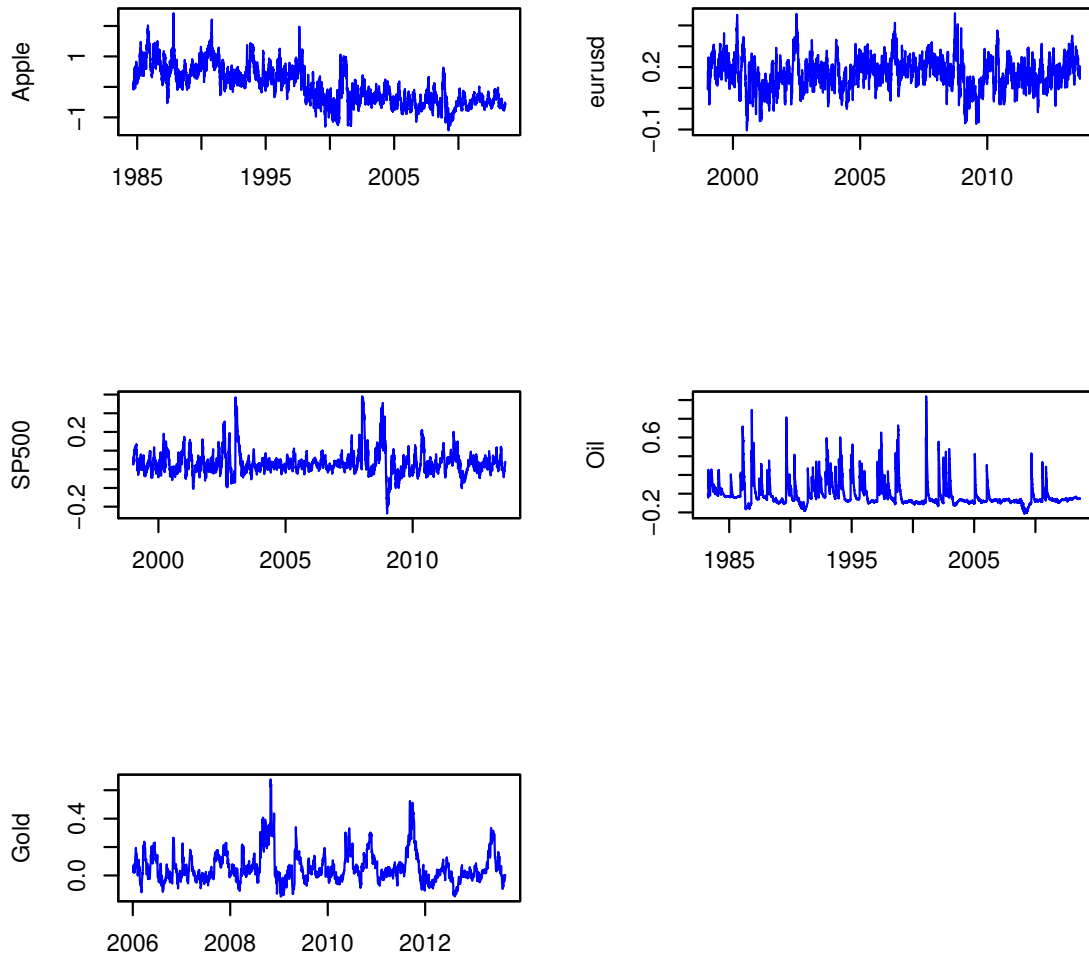


Figure 6: Difference of fitted conditional standard deviations (i.e. $\hat{\sigma}_{t,EM} - \hat{\sigma}_{t,0-adj}$)

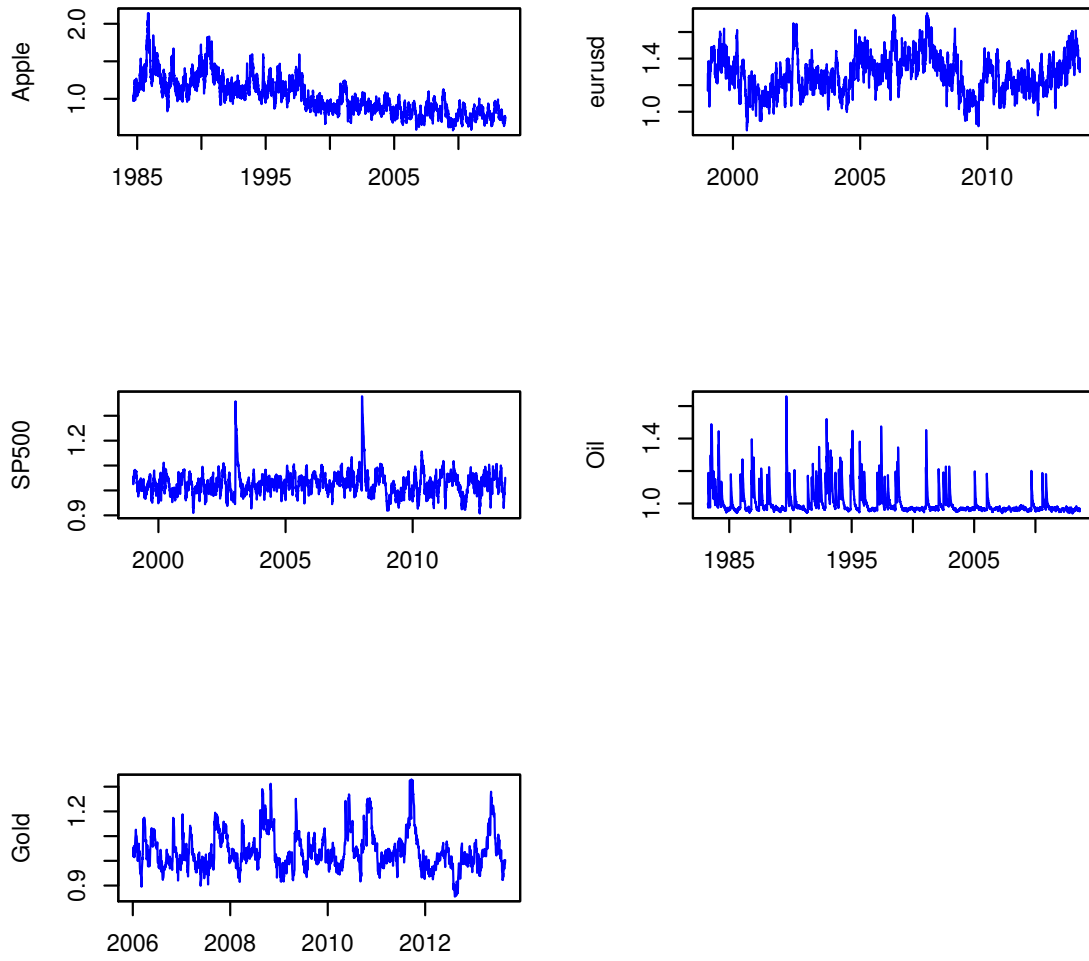


Figure 7: Ratios of fitted conditional standard deviations (i.e. $\hat{\sigma}_{t,EM}/\hat{\sigma}_{t,0-adj}$)