

Exercise Session for Financial Data Analysis
Summer term 2011
Problem Set 3

Write to `haas@stat.uni-muenchen.de` if you want to present. You can also indicate multiple exercises (ordered according to your preference) in case your most preferred problem has already been assigned.

Problem 1 Suppose that the return r_t of your portfolio is generated by

$$\begin{aligned}r_t &= 0.025 + \epsilon_t \\ \epsilon_t &= \eta_t \sigma_t, \quad \eta_t \stackrel{iid}{\sim} N(0, 1) \\ \sigma_t^2 &= 0.025 + 0.075\epsilon_{t-1}^2 + 0.9\sigma_{t-1}^2.\end{aligned}$$

Your current estimate for σ_t^2 is its unconditional expectation. Unfortunately, however, due to unpredictable adverse market conditions, your portfolio suffers from an unusually large negative shock, so that $r_t = -4.75$. Calculate the 1% Value-at-Risk for period $t + 1$.

Problem 2

Consider the following GARCH(1,1) model for returns r_t ,

$$r_t = c + \epsilon_t \tag{1}$$

$$\epsilon_t = \sigma_t \eta_t, \quad \eta_t \stackrel{iid}{\sim} \text{Normal}(0, 1) \tag{2}$$

$$\sigma_t^2 = \omega + \alpha \epsilon_{t-1}^2 + \beta \sigma_{t-1}^2 \tag{3}$$

$$\omega > 0, \quad \alpha \geq 0, \quad \beta \geq 0 \tag{4}$$

$$\alpha + \beta < 1 \tag{5}$$

$$3\alpha^2 + 2\alpha\beta + \beta^2 < 1. \tag{6}$$

- (a) What is the interpretation of σ_t^2 ?
- (b) What is the reason for the parameter restrictions (4) and (5)?
- (c) Derive the unconditional variance of r_t defined by (1)–(5).

- (d) Derive the autocorrelation of the squared process implied by model (2)–(6), i.e., $\text{Corr}(\epsilon_t^2, \epsilon_{t-\tau}^2)$, $\tau = 1, 2, \dots$

Hint: Without derivation, you can use the result that, for a stationary ARMA(1,1) process,

$$Y_t = \phi Y_{t-1} + \theta u_{t-1} + u_t,$$

where $\{u_t\}$ is white noise, the autocorrelation function is given by

$$\text{Corr}(Y_t, Y_{t-\tau}) = \phi^{\tau-1} \frac{(1 + \phi\theta)(\phi + \theta)}{1 + 2\phi\theta + \theta^2}, \quad \tau = 1, 2, \dots$$

Problem 3

Model (1)–(5) has been fitted to the daily percentage log-returns¹ of the German stock market index DAX 30 from January 2000 to March 2010 ($T = 2671$ daily observations).

- (i) Estimation results are reported in Table 1 on the next page. Would you say that the estimates imply a high persistence of volatility?
- (ii) As an alternative model with the same number of parameters, an ARCH(2) process is considered for the DAX returns, where (3)–(5) are replaced by

$$\sigma_t^2 = \omega + \alpha_1 \epsilon_{t-1}^2 + \alpha_2 \epsilon_{t-2}^2, \quad \omega > 0, \quad \alpha_1, \alpha_2 \geq 0, \quad \alpha_1 + \alpha_2 < 1. \quad (7)$$

Figure 2 shows, both for the ARCH(2) and the GARCH(1,1), the autocorrelation function (ACF) of the squares of the *standardized residuals*

$$\hat{\eta}_t = \frac{\hat{\epsilon}_t}{\hat{\sigma}_t}, \quad (8)$$

where $\hat{\epsilon}_t$ and $\hat{\sigma}_t$ are estimates of ϵ_t and σ_t , as implied by the fitted models, $t = 1, \dots, T$.

On the basis of Figure 2, which of the two models appears to provide the more appropriate specification for the volatility dynamics of the DAX returns? Could the outcome of this comparison have been expected from the properties of the GARCH(1,1) model, as compared to low-order ARCH specifications? Explain.

- (iii) For the standardized residuals (8), Table 1 also reports the Jarque–Bera test statistic (denoted by JB) along with its p -value underneath. Does this test indicate that there is a problem with model (1)–(5)? If so, what could be done to improve the model?

Table 1: GARCH(1,1) parameter estimates for the DAX returns

| \hat{c} | $\hat{\omega}$ | $\hat{\alpha}$ | $\hat{\beta}$ | JB* |
|--------------------|--------------------|--------------------|--------------------|-------------------|
| 0.0668 (0.0213) | 0.0211 (0.0053) | 0.0942 (0.0104) | 0.8985 (0.0103) | 162.3 (0.0000) |

The first four columns report parameter estimates for the GARCH(1,1) model (1)–(5) fitted to the DAX returns, with standard errors given in parentheses.

*JB is the value of the Jarque–Bera test statistic applied to the standardized residuals $\hat{\eta}_t = \hat{\epsilon}_t/\hat{\sigma}_t$, as defined in (8), with p -value underneath.

Problem 4 We have seen that the volatility forecasts of a covariance stationary GARCH model converge to the long-term (or unconditional) volatility as the forecast horizon increases. Alexander (2008)² argues that “the long term volatility forecast is very sensitive to small changes in the estimated values of the GARCH parameters”, so that it makes sense that the analyst imposes a personal view for long-term volatility before the parameters driving the volatility dynamics are estimated, “and then use the GARCH model to fill in the forecasts of volatility over the next day, week, month etc. that are consistent with this view.” This is also referred to as *volatility targeting*.

- (a) The long-term variance is supposed to be $\bar{\sigma}^2$. How can we reparameterize the standard GARCH(1,1) model $\sigma_t^2 = \omega + \alpha\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2$ such that, when estimated, we get a long-term volatility of $\bar{\sigma}^2$?
- (b) Consider the exponentially weighted moving average (EWMA) volatility model (also known as RiskMetrics[®] model³),

$$\sigma_t^2 = (1 - \lambda)\epsilon_{t-1}^2 + \lambda\sigma_{t-1}^2 = (1 - \lambda) \sum_{i=1}^{\infty} \lambda^{i-1} \epsilon_{t-i}^2, \quad 0 < \lambda < 1. \quad (9)$$

Find the multi-step-ahead variance forecasts for the process (9) and compare them with those of a covariance stationary GARCH(1,1) model from last week’s problem set.

¹That is, if I_t is the index level at day t , then $r_t = 100 \times \log(I_t/I_{t-1})$.

²*Market Risk Analysis Vol. 2: Practical Financial Econometrics*, Wiley, p. 144.

³Cf. JP Morgan: RiskMetrics—Technical Document, New York, 1996; see <http://www.riskmetrics.com/system/files/private/td4e.pdf>.

Problem 5 Assume that the return of your portfolio, r_p , is normally distributed with mean μ and variance σ^2 , i.e., its density function is given by

$$f_{\text{Normal}}(r_p; \mu, \sigma^2) = \frac{1}{\sqrt{2\pi}\sigma} \exp \left\{ -\frac{(r_p - \mu)^2}{2\sigma^2} \right\}, \quad r_p \in \mathbb{R}. \quad (10)$$

- (a) Assume that $\mu = 2.5$ and $\sigma^2 = 2.25$. Your portfolio is worth 1 \$. Compute the 1% Value-at-Risk.
- (b) Now assume that the portfolio return has a *Laplace* (or *double exponential*)⁴ distribution with mean μ and variance σ^2 , i.e., its density function is

$$f_{\text{Laplace}}(r_p; \mu, \sigma^2) = \frac{1}{\sqrt{2}\sigma} \exp \left\{ -\sqrt{2} \left| \frac{r_p - \mu}{\sigma} \right| \right\}, \quad r_p \in \mathbb{R}. \quad (11)$$

The Laplace distribution has sometimes been found to be useful for modeling stock return distributions,⁵ see also Figure 10 for the standardized densities with $\mu = 0$ and $\sigma^2 = 1$.

Assume that, in (11), $\mu = 2.5$ and $\sigma^2 = 2.25$. Find the 1% Value-at-Risk for a portfolio which is worth 1 \$. Compare the result with what has been obtained in part (a).

Problem 6 For the monthly (demeaned) returns of the MSCI Germany index from January 1970 to June 2008 ($T = 462$ observations), the regression

$$\hat{\epsilon}_t^2 = b_0 + \sum_{i=1}^{10} b_i \hat{\epsilon}_{t-i}^2 + u_t,$$

where $\hat{\epsilon}$ is the demeaned return, gives a coefficient of determination of $R^2 = 0.0444$

- (i) Explain how you can use this information for conducting a test for ARCH effects.
- (ii) What is the result of the test? What can we say about its p -value?

⁴This is a special case of the GED distribution with $p = 1$.

⁵E.g., C. W. J. Granger and Zhuanxin Ding (1995): Some Properties of Absolute Return, An Alternative Measure of Risk, *Annales D'économie et de Statistique*, 40, 67–91; and Mikael Linden (2001): A Model for Stock Return Distribution, *International Journal of Finance and Economics*, 6, 159–169.

Table 2: Quantiles of the χ^2 distribution.

| ν | $z_{0.9}$ | $z_{0.95}$ | $z_{0.975}$ | $z_{0.99}$ |
|-------|-----------|------------|-------------|------------|
| 1 | 2.7055 | 3.8415 | 5.0239 | 6.6349 |
| 2 | 4.6052 | 5.9915 | 7.3778 | 9.2103 |
| 3 | 6.2514 | 7.8147 | 9.3484 | 11.3449 |
| 4 | 7.7794 | 9.4877 | 11.1433 | 13.2767 |
| 5 | 9.2364 | 11.0705 | 12.8325 | 15.0863 |
| 6 | 10.6446 | 12.5916 | 14.4494 | 16.8119 |
| 7 | 12.0170 | 14.0671 | 16.0128 | 18.4753 |
| 8 | 13.3616 | 15.5073 | 17.5345 | 20.0902 |
| 9 | 14.6837 | 16.9190 | 19.0228 | 21.6660 |
| 10 | 15.9872 | 18.3070 | 20.4832 | 23.2093 |

ν denotes the degrees of freedom of the χ^2 distribution, and z_α is the α -Quantile, that is, z_α is such that

$$\int_0^{z_\alpha} \chi^2(z; \nu) dz = \alpha, \quad (12)$$

where $\chi^2(z; \nu)$ is the density function of a χ^2 random variable with ν degrees of freedom.

Problem 7 Consider the AR(1)–GARCH(1,1) process

$$\begin{aligned} r_t &= c + \phi r_{t-1} + \epsilon_t, \quad \epsilon_t = \eta_t \sigma_t, \quad \eta_t \stackrel{iid}{\sim} N(0, 1), \\ \sigma_t^2 &= \omega + \alpha \epsilon_{t-1}^2 + \beta \sigma_{t-1}^2, \\ \omega &> 0, \quad \alpha, \beta \geq 0, \quad \alpha + \beta < 1. \end{aligned}$$

For a sample r_1, r_2, \dots, r_T , construct the conditional log-likelihood function. What if η_t is assumed to have a Student's t distribution with ν degrees of freedom instead of a standard normal distribution?

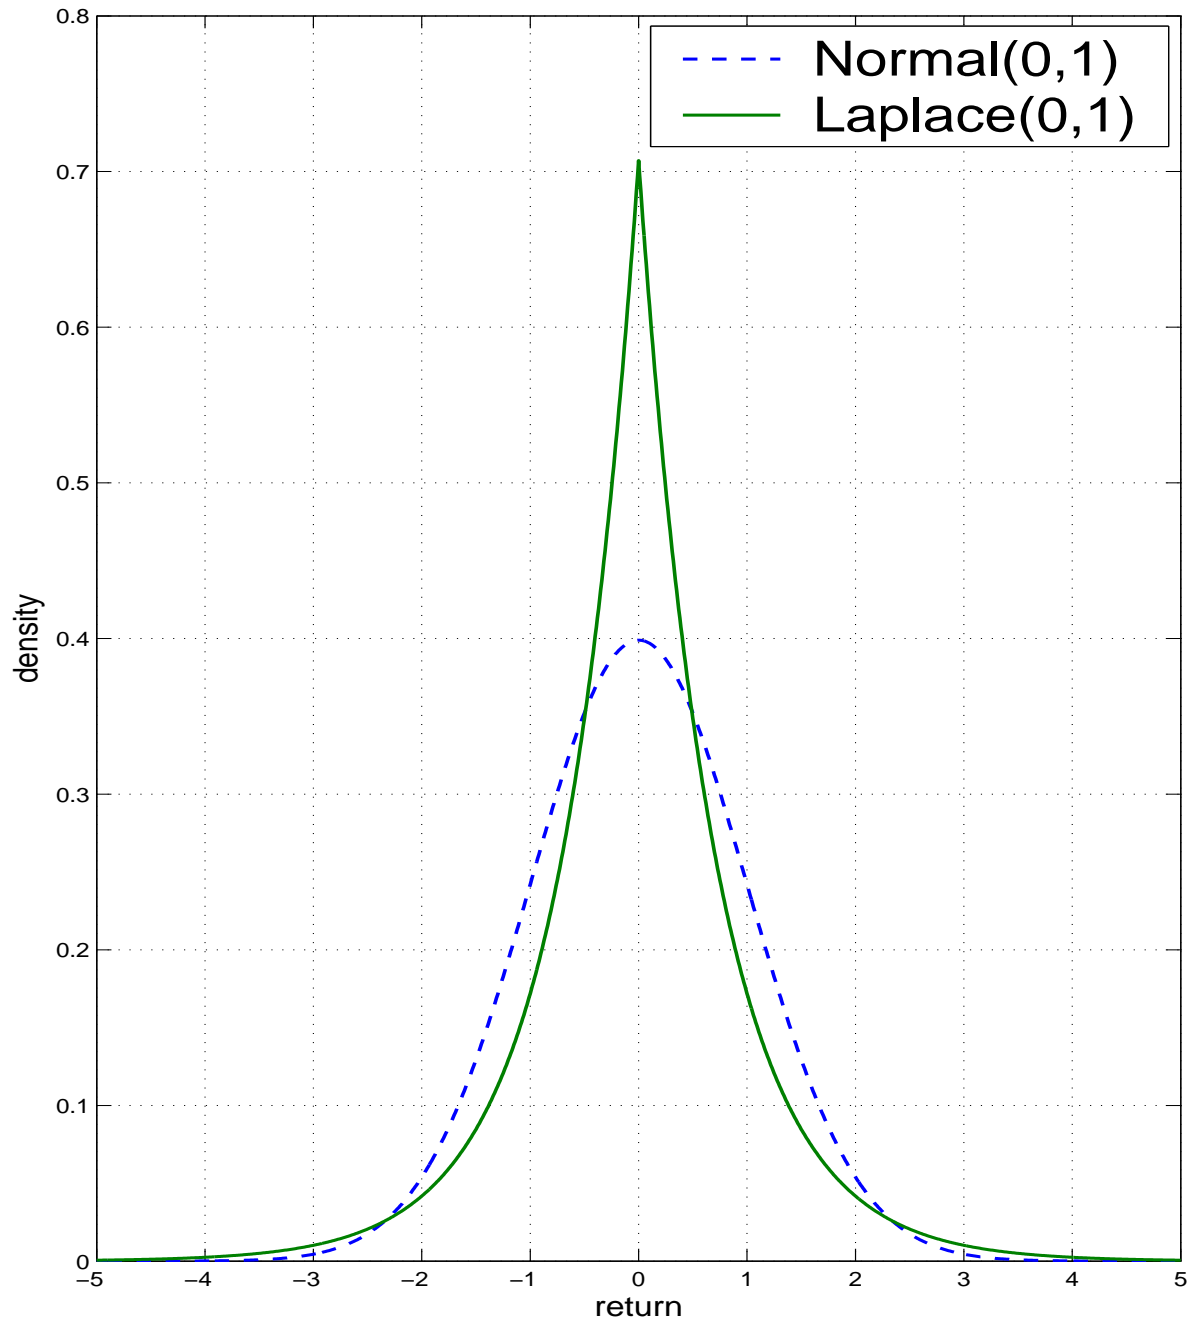


Figure 1: Standard (i.e., zero mean and unit variance) normal (dashed) and Laplace (solid) densities, given by (10) and (11), respectively, with $\mu = 0$ and $\sigma^2 = 1$.

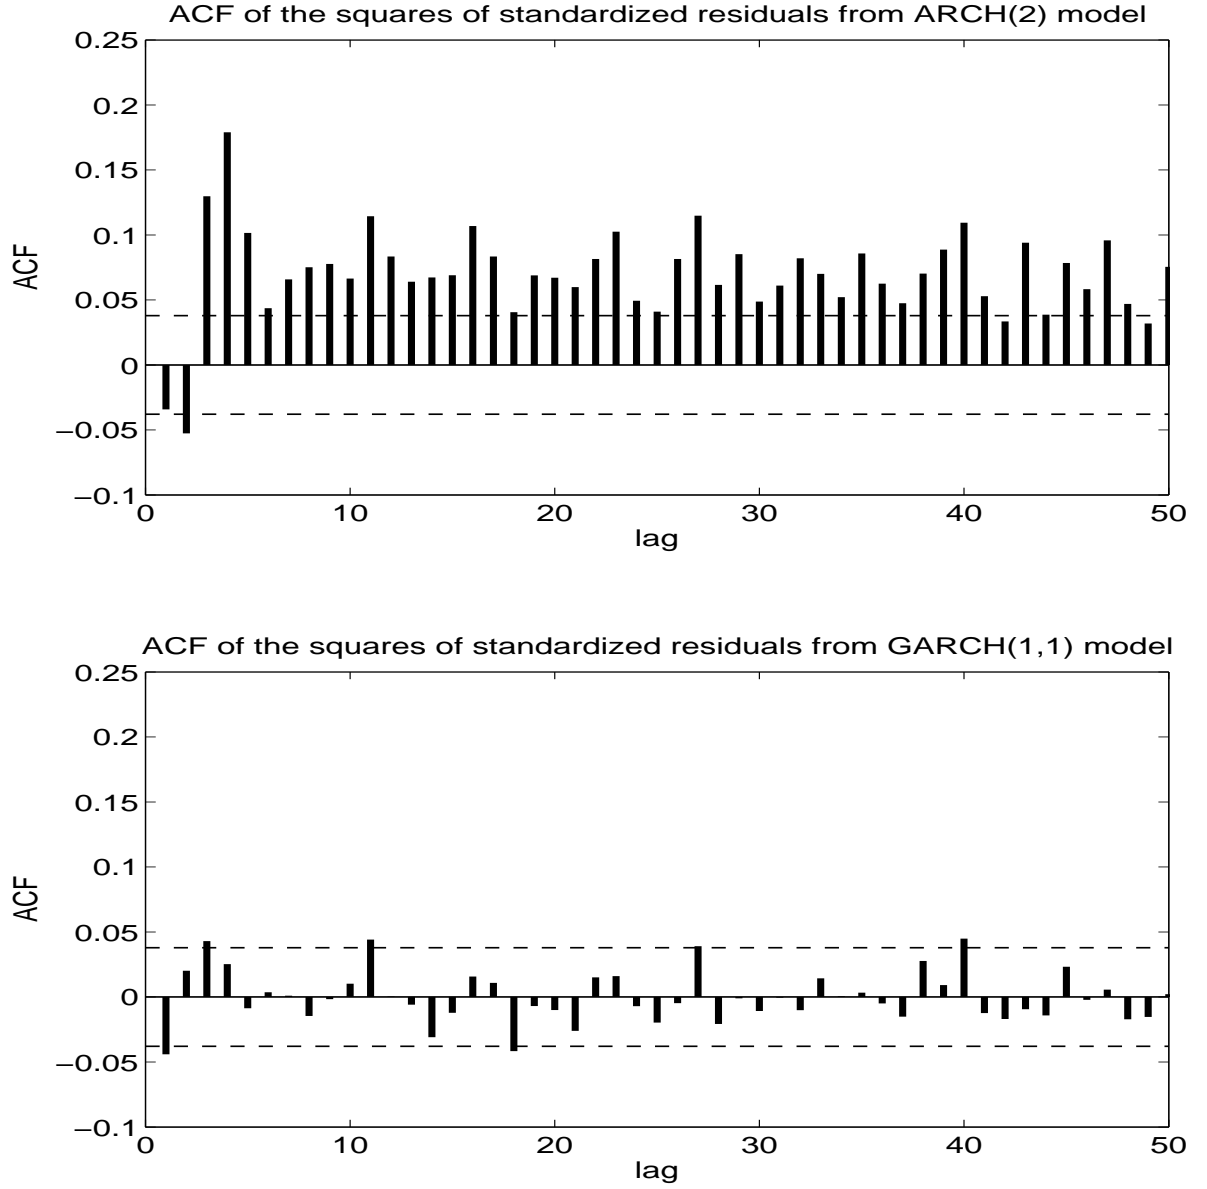


Figure 2: Shown are, for the fitted ARCH(2) (top plot) and GARCH(1,1) (bottom plot) processes, the autocorrelation functions (ACF) of the squares of the standardized residuals (8), i.e., $\hat{\eta}_t = \hat{\epsilon}_t / \hat{\sigma}_t$. Dashed lines represent approximate 95% confidence intervals.