Financial Econometrics A | Final Exam | February 10th, 2017

Solution Key

Question A:

Consider the following DAR-X model given by

$$x_t = \phi x_{t-1} + \varepsilon_t, \tag{A.1}$$

where the error term ε_t satisfies

$$\varepsilon_t = \sigma_t z_t, \quad z_t \sim i.i.d.N(0,1) \tag{A.2}$$
$$\sigma_t^2 = \omega + \alpha x_{t-1}^2 + \beta y_{t-1}^2. \tag{A.3}$$

Here y_t is some observed exogenous variable satisfying

$$y_t = \rho y_{t-1} + \eta_t, \quad \eta_t \sim i.i.d.N(0,1).$$

The model parameters $\theta = (\phi, \omega, \alpha, \beta, \rho)$ satisfy $\phi, \rho \in \mathbb{R}$, $\omega > 0$, and $\alpha, \beta \geq 0$. We assume that the processes (z_t) and (η_t) are independent.

Question A.1: Suppose that $\alpha = \beta = 0$. Under what conditions is x_t weakly mixing?

Solution: If $\alpha = \beta = 0$, $\varepsilon_t \sim i.i.d.N(0, \omega)$, and hence x_t is an AR(1) process. It is well-known from the lecture notes that x_t is weakly-mixing, if $|\phi| < 1$. The good answer should mention that this can be shown formally by verifying that x_t satisfies the drift criterion. Ideally, derivations should be included.

Question A.2: Consider the joint process $W_t = (x_t, y_t)$. Argue that W_t is a Markov chain.
It holds that the conditional density of W_t is

$$\begin{aligned} f(W_t | W_{t-1}) &= f(x_t, y_t | x_{t-1}, y_{t-1}) \\ &= f(x_t | x_{t-1}, y_{t-1}) f(y_t | y_{t-1}). \end{aligned}$$

Derive expressions for the conditional densities $f(x_t|x_{t-1}, y_{t-1})$ and $f(y_t|y_{t-1})$, and argue that $f(x_t, y_t|x_{t-1}, y_{t-1})$ is positive and continuous in (x_t, y_t) and (x_{t-1}, y_{t-1}).
Explain briefly what this insight can be used for.

It can be shown (but do not do so) that the Markov chain W_t satisfies the drift criterion with drift function $\delta(W_t) = 1 + \|W_t\|^2 = 1 + x_t^2 + y_t^2$ if $\max(\rho^2 + \beta, \phi^2 + \alpha) < 1$.

Solution: It holds that W_t is a Markov chain, as only (x_{t-1}, y_{t-1}) enters the dynamics of (x_t, y_t). Formally, $f(W_t|W_{t-1}, W_{t-2}, ...) = f(W_t|W_{t-1})$. Since, x_t is conditionally Gaussian with conditional mean ϕx_{t-1} and conditional variance σ_t^2 (conditional on x_{t-1} and y_{t-1}), it holds that

$$f(x_t|x_{t-1}, y_{t-1}) = \frac{1}{\sqrt{2\pi\sigma_t^2}} \exp\left\{-\frac{(x_t - \phi x_{t-1})^2}{2\sigma_t^2}\right\}.$$

Moreover, y_t is conditionally Gaussian with conditional mean ρy_{t-1} and conditional variance 1 (conditional on y_{t-1}). Hence,

$$f(y_t|y_{t-1}) = \frac{1}{\sqrt{2\pi}} \exp\left\{-\frac{(y_t - \rho y_{t-1})^2}{2}\right\}.$$

Based on the expressions for $f(x_t|x_{t-1}, y_{t-1})$ and $f(y_t|y_{t-1})$, it holds that $f(x_t|x_{t-1}, y_{t-1})$ and $f(y_t|y_{t-1})$ are positive and continuous in (x_t, y_t) and (x_{t-1}, y_{t-1}). Hence $f(W_t|W_{t-1})$ is positive and continuous in W_t and W_{t-1}. This property is useful when establishing that W_t satisfies the drift criterion.

Question A.3: Consider the OLS estimator for ϕ given by

$$\hat{\phi} = \frac{\sum_{t=1}^T x_{t-1}x_t}{\sum_{t=1}^T x_{t-1}^2}.$$

It holds that

$$\hat{\phi} - \phi = \frac{\sum_{t=1}^T x_{t-1}\varepsilon_t}{\sum_{t=1}^T x_{t-1}^2}.$$

Assume that $W_t = (x_t, y_t)$ is weakly mixing and satisfies the drift criterion such that $E[x_t^4] < \infty$ and $E[y_t^4] < \infty$.
Show that, as $T \to \infty$,

$$\frac{1}{\sqrt{T}} \sum_{t=1}^T x_{t-1}\varepsilon_t \xrightarrow{D} N(0, v),$$

with $v = E[\omega x_{t-1}^2 + \alpha x_{t-1}^4 + \beta y_{t-1}^2 x_{t-1}^2]$.
Explain briefly what this property can be used for.

Solution: The result is established by verifying the conditions of the CLT for weakly mixing processes (Theorem II.1 from the lecture notes). It holds that $\sum_{t=1}^T x_{t-1}\varepsilon_t = \sum_{t=1}^T f(x_t, y_t, x_{t-1}, y_{t-1})$, with

$$\begin{aligned} f(x_t, y_t, x_{t-1}, y_{t-1}) &= x_{t-1}(x_t - \phi x_{t-1}) \\ &= x_{t-1} z_t \sqrt{\omega + \alpha x_{t-1}^2 + \beta y_{t-1}^2} \end{aligned}$$

Hence the CLT is satisfied if $E[x_{t-1}z_t\sqrt{\omega+\alpha x_{t-1}^2+\beta y_{t-1}^2}|x_{t-1},y_{t-1}] = 0$ and $E[|x_{t-1}z_t\sqrt{\omega+\alpha x_{t-1}^2+\beta y_{t-1}^2}|^2] < \infty$. The first condition holds, since $E[z_t]=0$ and z_t and (x_{t-1},y_{t-1}) are independent. For the second condition, we have that

$$E[x_{t-1}^2 z_t^2(\omega+\alpha x_{t-1}^2+\beta y_{t-1}^2)] = E[\omega x_{t-1}^2+\alpha x_{t-1}^4+\beta y_{t-1}^2 x_{t-1}^2] = v,$$

where we have used that $E[z_t^2] = 1$. As $E[x_t^4] < \infty$ and $E[y_t^4] < \infty$, $E[y_{t-1}^2 x_{t-1}^2] < \infty$ by the Hölder (or Cauchy-Schwarz) inequality. We conclude that $v < \infty$. By the CLT,

$$\frac{1}{\sqrt{T}}\sum_{t=1}^{T} x_{t-1}\varepsilon_t \xrightarrow{D} N(0,v), \quad \text{as } T\to\infty.$$

This property is important for obtaining the asymptotic distribution of the OLS estimator. Specifically, using that (x_t,y_t) is weakly mixing and satisfies $E[x_t^4]<\infty$,

$$\frac{1}{T}\sum_{t=1}^{T} x_{t-1}^2 \xrightarrow{P} E[x_t^2],$$

as $T\to\infty$, by the LLN. Hence,

$$\sqrt{T}(\hat{\phi}-\phi) = \frac{\frac{1}{\sqrt{T}}\sum_{t=1}^{T} x_{t-1}\varepsilon_t}{\frac{1}{T}\sum_{t=1}^{T} x_{t-1}^2} \xrightarrow{D} N(0,(E[x_t^2])^{-2}v), \quad \text{as } T\to\infty.$$

The latter derivations are not required.

Question A.4: Instead of estimating only ϕ based on OLS, we may estimate all parameters $\theta = (\phi,\omega,\alpha,\beta,\rho)$ based on maximum likelihood estimation. For the model (A.1)-(A.3), the one-period VaR at risk level κ, $\text{VaR}_{T,1}^{\kappa}$, is

$$\text{VaR}_{T,1}^{\kappa} = -\phi x_T - \sigma_{T+1}\Phi^{-1}(\kappa), \quad \kappa\in(0,1),$$

where $\Phi^{-1}(\cdot)$ denotes the inverse cdf of the standard normal distribution. Explain briefly how you would compute an estimate of $\text{VaR}_{T,1}^{\kappa}$.

Solution: Given an estimate of $\theta = (\phi,\omega,\alpha,\beta,\rho)$, denoted $\hat{\theta} = (\hat{\phi},\hat{\omega},\hat{\alpha},\hat{\beta},\hat{\rho})$, obtained by maximum likelihood (or some other method), an estimate of σ_{T+1} is given by

$$\hat{\sigma}_{T+1} = \sqrt{\hat{\omega}+\hat{\alpha}x_T^2+\hat{\beta}y_T^2},$$

where x_T and y_T are contained in the data set. For given $\kappa\in(0,1)$, $\Phi^{-1}(\kappa)$ is known, since $\Phi^{-1}(\cdot)$ denotes the inverse cdf of the standard normal distribution. An estimate of $\text{VaR}_{T,1}^{\kappa}$ is thus computed as $-\hat{\phi}x_T - \hat{\sigma}_{T+1}\Phi^{-1}(\kappa)$.

Question B:

Suppose that the logarithm of the price of a share of stock is given by

$$p(t) = p(0) + \mu t + \sigma W(t), \quad t \in [0, T], \tag{B.1}$$

where $p(0) \in \mathbb{R}$ is some fixed initial value, $\mu \in \mathbb{R}$ and $\sigma > 0$ are constants, and $W(t)$ is a Brownian motion.

Recall here that the Brownian motion $W(t)$ has the properties

1. $W(0) = 0$.
2. W has independent increments, i.e. if $0 \leq r < s \leq t < u$, then

   $$W(u) - W(t) \text{ and } W(s) - W(r)$$

   are independent.
3. The increments are normally distributed, i.e.

   $$W(t) - W(s) \sim N(0, t - s)$$

   for all $0 \leq s \leq t$.

Suppose that we have observed the price $p(t)$ at $n + 1$ equidistant points

$$0 = t_0 < t_1 < \ldots < t_n = T,$$

with

$$t_i = \frac{i}{n}T, \quad i = 0, ..., n.$$

Based on these points we obtain n log-returns given by

$$r(t_i) = p(t_i) - p(t_{i-1}), \quad i = 1, ..., n.$$

Question B.1: Argue that $r(t_i)$ is normally distributed, i.e. show that

$$r(t_i) \sim N\left(\mu\frac{T}{n}, \sigma^2\frac{T}{n}\right).$$

Show that

$$\text{cov}(r(t_i), r(t_{i-1})) = 0.$$

Solution: The properties follow directly from the definition of $r(t_i)$ and the properties of the Brownian motion. Derivations should be included.

Question B.2: We now seek to estimate the model parameters (μ, σ^2) based on maximum likelihood. Given the n log-returns, the log-likelihood function is (up to a constant and a scaling factor)

$$L_n(\mu, \sigma^2) = \sum_{i=1}^{n} \left\{ -\log(\sigma^2 \frac{T}{n}) - \frac{\left[r(t_i) - \mu \frac{T}{n}\right]^2}{\sigma^2 \frac{T}{n}} \right\}.$$

Let $\hat{\mu}$ denote the maximum likelihood estimator of μ.
Show that

$$\hat{\mu} = \frac{1}{T} \sum_{i=1}^{n} r(t_i) = \frac{1}{T} \left[p(T) - p(0)\right].$$

Argue that the sampling frequency of the log-returns over the interval $[0, T]$ does not have any influence on the estimate of μ.

Solution: By solving the F.O.C. for maximization of $L_n(\mu, \sigma^2)$, that is solving

$$\frac{\partial L_n(\mu, \sigma^2)}{\partial \mu} = 0$$

for μ, yields the MLE

$$\hat{\mu} = \frac{1}{T} \sum_{i=1}^{n} r(t_i).$$

Derivations should be included. Moreover,

$$\sum_{i=1}^{n} r(t_i) = \sum_{i=1}^{n} p(t_i) - p(t_{i-1}) = p(t_n) - p(t_0) = p(T) - p(0),$$

by the definition of t_i. Hence, the MLE does not depend on n, i.e. the number of observations within the interval $[0, T]$.

Question B.3: Let $\hat{\mu}$ denote the maximum likelihood estimator derived in Question B.2.
Show that $\hat{\mu}$ is an unbiased estimator for μ, i.e. show that

$$E[\hat{\mu}] = \mu.$$

Moreover, show that the variance of the estimator is

$$\text{Var}(\hat{\mu}) = \frac{\sigma^2}{T}.$$

It can be shown (but do not do so) that these two properties ensure that $\hat{\mu} \xrightarrow{p} \mu$ as $T \to \infty$.

Solution: We have, using the dynamics for $p(t)$, that

$$\begin{aligned}\hat{\mu} &= \frac{1}{T}\left[p(T) - p(0)\right] \\ &= \frac{1}{T}[\mu T + \sigma W(T)] \\ &= \mu + \frac{1}{T}\sigma W(T).\end{aligned}$$

The expressions for $E[\hat{\mu}]$ and $\mathrm{Var}(\hat{\mu})$ then follow directly by the properties of the Brownian motion. Derivations should be included.

Question B.4: Assume now that $T = 1$, such that we have n observations of the log-returns over the time interval $[0, 1]$, which you may think of as the time interval over one trading day. Then the maximum likelihood estimator for σ^2 is given by

$$\hat{\sigma}^2 = \sum_{i=1}^{n}\left[r(t_i) - \frac{1}{n}\sum_{i=1}^{n} r(t_i)\right]^2.$$

Use that $r(t_i) = \frac{\mu}{n} + \frac{\sigma}{\sqrt{n}} z_i$, with $z_i \sim i.i.d. N(0, 1)$ in order to show that

$$\frac{1}{n}\sum_{i=1}^{n} r(t_i) \xrightarrow{p} 0 \quad \text{as } n \to \infty.$$

Explain briefly how $\hat{\sigma}^2$ is related to the Realized Volatility.

Solution: We have, that

$$\begin{aligned}\frac{1}{n}\sum_{i=1}^{n} r(t_i) &= \frac{1}{n}\sum_{i=1}^{n}\left(\frac{\mu}{n} + \frac{\sigma}{\sqrt{n}} z_i\right) \\ &= \frac{\mu}{n} + \frac{\sigma}{\sqrt{n}}\frac{1}{n}\sum_{i=1}^{n} z_i.\end{aligned}$$

For the first term, $\frac{\mu}{n} \to 0$ as $n \to \infty$. For the second term, $\frac{1}{n}\sum_{i=1}^{n} z_i \xrightarrow{p} E(z_i) = 0$ by the LLN for $i.i.d.$ processes. We conclude that $\frac{1}{n}\sum_{i=1}^{n} r(t_i) \xrightarrow{p} 0 \quad$ as $n \to \infty$.

The realized volatility (over the interval $[0, 1]$) is

$$\sum_{i=1}^{n} [r(t_i)]^2 .$$

Hence the realized volatility is obtained from $\hat{\sigma}^2 = \sum_{i=1}^{n} \left[r(t_i) - \frac{1}{n}\sum_{i=1}^{n} r(t_i)\right]^2$ by substiuting in the probability limit of $\frac{1}{n}\sum_{i=1}^{n} r(t_i)$ (that is equal to zero).

Question B.5: The following figure shows the daily log-returns of the S&P 500 index for the period January 4, 2010 to September 17, 2015.

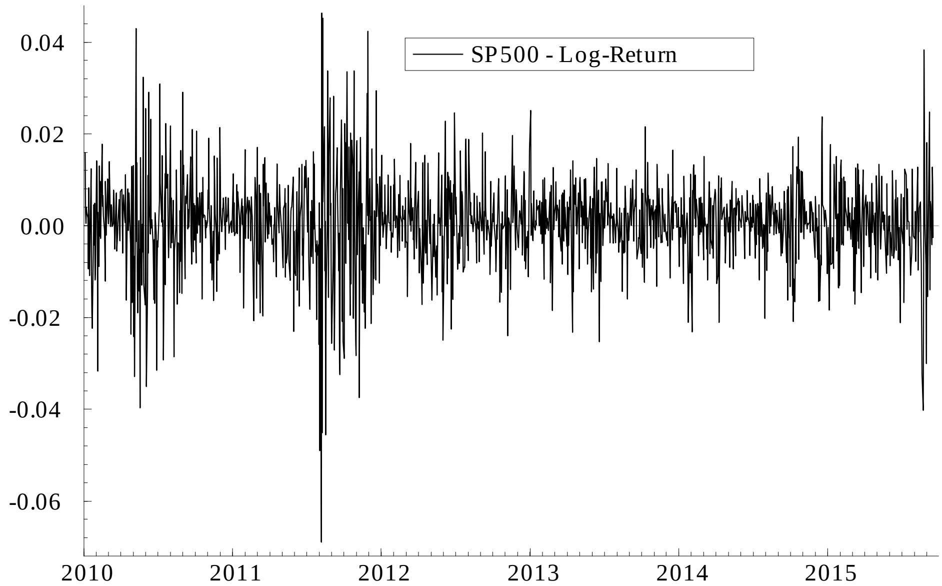

Discuss briefly whether the model in (B.1) is a reasonable model for the daily log returns of the S&P 500 index.

Solution: The model in (B.1) suggests that daily log-returns $r(t) = p(t) - p(t-1)$, $t = 1, 2, ...$, should be given by

$$\mu + \sigma(W(t) - W(t-1)).$$

By the properties of the Brownian motion, we would have that $r(t) \sim i.i.d.N(\mu, \sigma^2)$. I.e. the returns would be independent and Gaussian with constant mean and variance. By visual inspection of the series, it appears that the returns are heteroskedastic, and we know from the lectures that the returns are unconditionally heavy-tailed (i.e. non-Gaussian). This suggests that the model is not appropriate for modelling the main features of the daily return series. Ideally, a few derivations should be included.