

Suggested answers to Fall 2004 exam. MFE 2004.

The solutions below are not the only possible answers. The questions could be answered in a few possible ways. Take the answers below as a guideline (rather than a definite solution) of what was required in each question. Partial credit was awarded for partial solutions that demonstrated knowledge of the material.

Problem 1.

1. True. The GARCH is parsimonious. In general, parsimonious models perform well out-of-sample. Moreover, the GARCH process captures successfully an important feature of conditional volatility: persistence.

2. False. 5-minute stock returns are likely to be serially correlated because of microstructure effects (bid-ask spread, liquidity effects, order smoothing, etc.) Daily stock returns might still exhibit some serial correlation, but it would be much less significant than at the 5-minute frequency. The weekly serial correlation will be almost non-existent for most stocks. At monthly frequency, the serial correlation will be the weakest. Markets are efficient, if you give investors enough time to react to news.

3. False. The Variance Ratio test allows for more general serial correlations than a simple AR(1) model. The serial correlation model is not explicitly specified in the VR test.

4. True. The CAPM explains a large fraction of the cross sectional variation of returns. However, this does not mean that it explains all the cross-sectional variation. In fact, there is a lot of variation that is still unexplained, which leads to a rejection of the CAPM. However, the CAPM is a useful model against which all other models are benchmarked.

5. False. The Fama-French can be estimated only using time-series regressions, but models that do not express the factors as zero-cost (long-short) portfolios need to be run in two stages (see, Fama-MacBeth lecture, and Chen, Roll, and Ross article)

6. True, they both will lead to the estimate $\hat{\beta} = \sum x_t y_t / \sum x_t^2$ which is also the OLS estimate.

7. True, event studies often involve data that is, by its very nature, heterogeneous. Moreover, event studies are often conducted using small samples and the distributions might not be approximately normal.

8. True, this is precisely the difference between the two. Otherwise, they are both used to obtain a small sample distribution of the parameters of interest.

9. False. The exact opposite is true. A linear, parsimonious model is likely to do better out-of-sample. A non-linear model with more parameters will do better in-sample, but it might also lead to overfitting.

10. No, even in linear systems, the Kalman filter will do better. Actually, in a linear system, the Kalman filter is optimal, i.e. it produced the best linear forecasts. In other words, it will dominate any other estimation procedure, including the rolling window.

Problem 2.

1. No. If $P_t = \mu + P_{t-1} + \varepsilon_t$, then its variance, $V(P_t) = \sigma_\varepsilon t$ depends on t .

2. Yes, since $R_t = \Delta P_t$, it will be stationary. There are many ways of answering this question. For instance, R_t is obtained by applying the filter $(1 - L)$ to price and this filter produces stationary data. Or, in efficient markets, returns must be close to white noise (which is a stationary process). Or, from homework, we have estimated the process for returns $R_t = \varphi + \phi R_{t-1} + u_t$ and the estimate of ϕ was (for the market portfolio) very close to zero. Hence, the return process is stationary.

3. It is most likely that the estimated process is the yield of a bond. It cannot be a return, because returns are not that persistent, even at 5-minute intervals. It might be a price, with a downward biased estimate of ϕ . However, the bias needs to be quite large.

4. A process is stationary if its moments are independent of time.

5. If we have non-stationary time series, we run the risk of getting spuriously significant results. For instance, if we have the two non-stationary variables $Y_t = v + Y_{t-1} + u_t$ and $X_t = \eta + X_{t-1} + \varepsilon_t$ and the shocks are uncorrelated, running a regression of Y_t on X_t will produce statistically significant t-statistics and large R^2 . There is no good way of testing for spurious correlation.

6. One solution is to stationarize X_t by applying the filter $(1 - L)$ only to that variable. Alternatively, if we have a very strong prior that Y and X must enter (economically) into a linear relation $Y_t = \alpha + \gamma X_t + \varepsilon_t$, then we can estimate the regression without any filtering. This is the co-integration case. However, if there is no true underlying relation between the two variables, we will get a spuriously significant coefficient.

Problem 3.

1. The mean vector has dimensions 1000 by 1.

2. The covariance matrix has dimensions 1000 by 1000.

3. The unconditional mean is estimated by taking simple averages ($\frac{1}{T} \sum_t r_{i,t}$) of all stock returns and collecting them into a vector of estimates.

4. The covariance matrix has to be symmetric and positive semi-definite. To insure that it be symmetric, we will estimate the 1000 variances and only $(1000 * 999)/2$ covariances (off-diagonal terms). If we estimate each covariance independently, there is nothing that guarantees that the matrix will be positive semi-definite.

If we estimate each term of the covariance matrix in the same way, we will have

5. The Markowitz portfolio weights w_t will be proportional to the estimate of μ and inversely proportional to the estimate of Σ . Since neither of these estimates varies with time, the portfolio weights will also be the unconditional portfolio weights, i.e. time invariant.

6. The unconditional and conditional moments will differ if there is time variation in the process. Also, their estimates will differ if we take that time variation into account when estimating the moments.

7. Of course. Now, the Markowitz portfolio weights will be proportional to the estimate of $\mu_{t|t-1}$ and inversely proportional to the estimate of $\Sigma_{t|t-1}$. Since these estimates will vary from *period to period*, so would the portfolio weights.

8. There are many ways of attacking this problem. First, the options data can give us the implied vols of each stock at a point in time, which can be used in the estimation of $\Sigma_{t|t-1}$. You have to explain exactly how this is done for full credit. Second, the options market data can be used to infer the conditional risk premium of the stock, or $\mu_{t|t-1}$. This is the risk adjustment between the risk neutral and physical measures. We can get an estimate of that premium also at each point of time. You have to explain exactly how this is done for full credit. There are other possibilities.

9. To do that, one can run simple regressions $R_{i,t} = \eta_i + \delta_i X_{i,t-1} + u_{i,t}$ and the fitted values will be estimates of $\mu_{t|t-1}$. This regression can then be applied to out-of-sample forecasting. For full credit, you must discuss the distinctions between in- and out-of-sample forecasting.

10. Because *it* might be the case that $X_{i,t}$ is endogenous, or that $E(X_{i,t}u_{i,t}) \neq 0$. Hence, we won't be able to estimate δ_i consistently. Show calculations why?

11. This is related to question 4 above. One way is to decompose $\Sigma = CC'$ where the matrices C will be lower triangular with the standard deviations of the stock on the main diagonal. The estimates of the standard deviations can be obtained from the implied vols in the options market.

12. Yes, you can set up two sets of moments, one $E(X_{i,t-1}u_{i,t})$ for each stock and another $E(\text{vec}(C)) = 0$ where $\text{vec}()$ is a function that takes all elements of C and outputs a vector of elements. The matrix C is implicitly a function of option pricing data which will be used to extract the implied vols. Be more specific.