University of Washington
Department of Economics

Spring 2006
Eric Zivot

Economics 424

Midterm Exam Solutions

This is a closed book and closed note exam. However, you are allowed one page of notes (double-sided). Answer all questions and write all answers in a blue book or on separate sheets of paper. Time limit is 1 hours and 50 minutes. Total points = 100.

I. Return Calculations (25 pts, 5 points each)

1. Consider a one month investment in two Northwest stocks: Amazon and Costco. Suppose you buy Amazon and Costco at the end of March at $P_{A,t-1} = \$36.53,\ P_{C,t-1} = \$54.16$ and then sell at the end of April for $P_{A,t} = \$35.21,\ P_{C,t} = \$54.43$. (Note: these are actual closing prices for 2006 taken from Yahoo!)

a. What are the simple monthly returns for the two stocks?

```
> R.A = (p.A[2] - p.A[1])/p.A[1]
> R.A
[1] -0.03613
> R.C = (p.C[2] - p.C[1])/p.C[1]
> R.C
[1] 0.004985
```

b. What are the continuously compounded monthly returns for the two stocks?

```
> r.A = log(1 + R.A)
> r.A
[1] -0.0368
> r.C = log(1 + R.C)
> r.C
[1] 0.004973
```

c. Suppose Costco paid a \$1 per share cash dividend at the end of April. What is the monthly simple total return on Costco? What is the monthly dividend yield?

```
> div.C = 1
> R.total.C = (p.C[2] + div.C - p.C[1])/p.C[1]
> R.total.C
[1] 0.02345
> div.yield.C = div.C/p.C[1]
> div.yield.C
[1] 0.01846
```

d. Suppose the monthly returns on Amazon and Costco from question (a) above are the same every month for 1 year. Compute the simple annual returns as well as the continuously compounded annual returns for the two stocks.

```
> R.A.annual = (1 + R.A)^12 - 1
> R.C.annual = (1 + R.C)^12 - 1
> R.A.annual
[1] -0.357
> R.C.annual
[1] 0.06149
> r.A.annual = log(1 + R.A.annual)
> r.C.annual = log(1 + R.C.annual)
> r.A.annual
[1] -0.4416
> r.C.annual
[1] 0.05967
```

e. At the end of March, 2006, suppose you have $10,000 to invest in Amazon and Costco over the next month. Suppose you invest $6000 in Amazon and $4000 in Costco. Compute the monthly simple return and the monthly continuously compounded return on the portfolio. Assume that both stocks do not pay a dividend.

```
> x.A = 6000/10000
> x.C = 4000/10000
> x.A
[1] 0.6
> x.C
[1] 0.4
> R.port = x.A * R.A + x.C * R.C
> R.port
[1] -0.01969
> r.port = log(1 + R.port)
> r.port
[1] -0.01988
```

II. Probability Theory (25 points, 5 points each)

1. Consider an investment in a portfolio consisting of Starbucks and Amazon stock. Let R_{sbux} denote the monthly *simple* return on Starbucks and assume that $R_{sbux} \sim N(0.02,(0.20)^2)$. Let R_{amnz} denote the monthly *simple* return on Amazon and assume that $R_{amnz} \sim N(0.01,(0.10)^2)$. Further, assume that $corr(R_{sbux}, R_{amnz}) = 0.5$.

a) For a portfolio with 20% of wealth invested in Starbucks and the remaining 80% of wealth invested in Amazon, compute $E[R_p]$, $\text{var}(R_p)$, and $SD(R_p)$ where R_p denotes the simple return on the portfolio.

```
> x.sbux = 0.2
> x.amnz = 0.8
> mu.sbux = 0.02
> var.sbux = 0.2^2
> mu.amnz = 0.01
> var.amnz = 0.1^2
> rho.sa = 0.5
> cov.sa = rho.sa * sqrt(var.sbux * var.amnz)
> cov.sa
[1] 0.01
> mu.p = x.sbux * mu.sbux + x.amnz * mu.amnz
> mu.p
[1] 0.012
> var.p = x.sbux^2 * var.sbux + x.amnz^2 * var.amnz
          + 2 * x.sbux * x.amnz * cov.sa
> var.p
[1] 0.0112
> sd.p = sqrt(var.p)
> sd.p
[1] 0.1058
```

b) What is the probability distribution of R_p? Sketch this distribution, indicating the location of $E[R_p]$ and $E[R_p] \pm 2 \cdot SD(R_p)$.

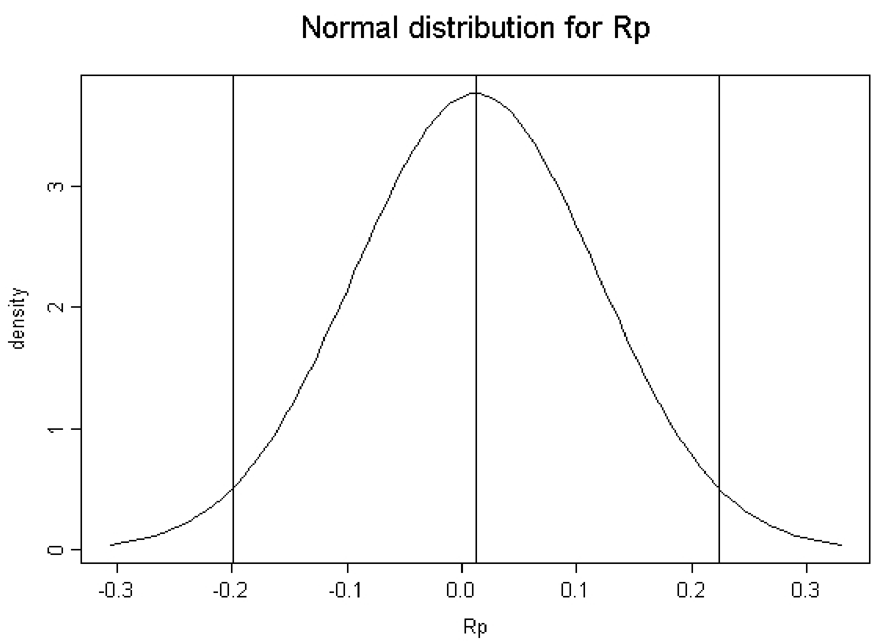

c)   Compute the 5% quantile of the distribution for R_p. Using this quantile, compute the monthly 5% value-at-risk ($VaR_{.05}$) of an initial \$1,000 investment. (Hint: the 5% quantile for a standard normal random variable is -1.645.)

```
> q.05 = mu.p + sd.p * (-1.645)
> q.05
[1] -0.1621
> VaR.05 = q.05 * 1000
> VaR.05
[1] -162.0904
```

2. Let $r_{sbux,t}$ denote the continuously compounded return on Starbucks stock in month t and assume that $r_{sbux,t} \sim iid\ N(0.02,(0.20)^2)$. Define the annual continuously compounded return as $r_t(12) = r_t + r_{t-1} + \cdots + r_{t-1} = \sum_{j=0}^{11} r_{t-j}$.

a) Compute $E[r_t(12)]$, $\text{var}(r_t(12))$, and $SD(r_t(12))$. Is the distribution of $r_t(12)$ normal?

```
> mu.a = 12 * mu.sbux
> mu.a
[1] 0.24
> var.a = 12 * var.sbux
> var.a
[1] 0.48
> sd.a = sqrt(var.a)
> sd.a
[1] 0.6928
```

The distribution of r(12) is normal.

b) Using the distribution of $r_t(12)$, compute the annual 5% value-at-risk ($VaR_{.05}$) of an initial $1,000 investment.

```
> q.05.a = mu.a + sd.a * (-1.645)
> q.05.a
[1] -0.8997
> VaR.05.a = (exp(q.05.a) - 1) * 1000
> VaR.05.a
[1] -593.3
```

III. Time Series Concepts (25 points)

1. Let $\{Y_t\}$ represent a stochastic process. Under what conditions is $\{Y_t\}$ covariance stationary? (5 points)

A covariance stationary process satisfies the following three conditions

$E[Y_t] = \mu$

$\text{var}(Y_t) = \sigma^2$

$\text{cov}(Y_t, Y_{t-j}) = \gamma_j$ (depends on j and not t)

2. The figures below give simulated observations from three different covariance stationary time series (labeled series 1, series 2, and series 3, respectively). Each series has mean zero. One series represents the Gaussian white noise process

$Y_t \sim iid\ N(0,1)$

One series represents the MA(1) process

$Y_t = \varepsilon_t + 0.75\varepsilon_{t-1},\ \varepsilon_t \sim iid\ N(0,1)$

One series represents the AR(1) process

$Y_t = 0.75Y_{t-1} + \varepsilon_t,\ \varepsilon_t \sim iid\ N(0,1)$

Series 1

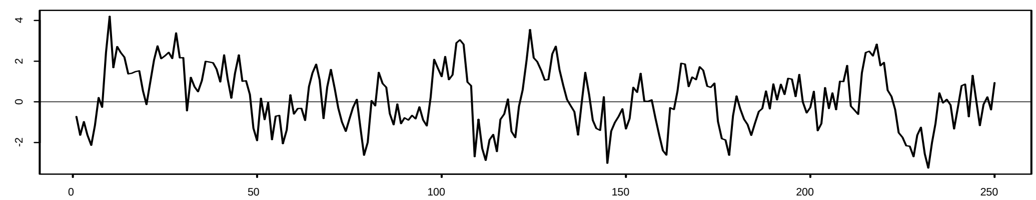

SACF: Series 1

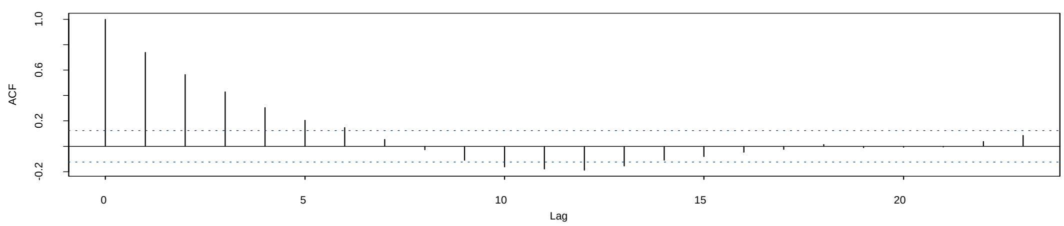

Series 2

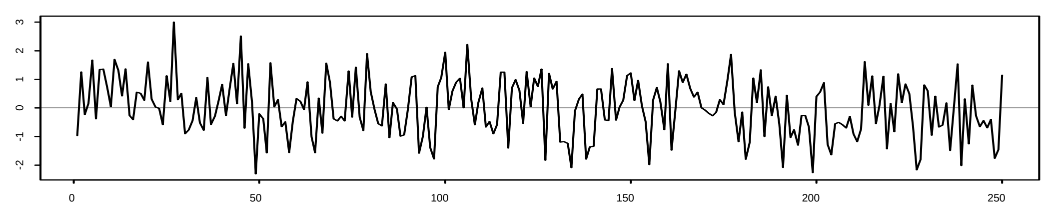

SACF: Series 2

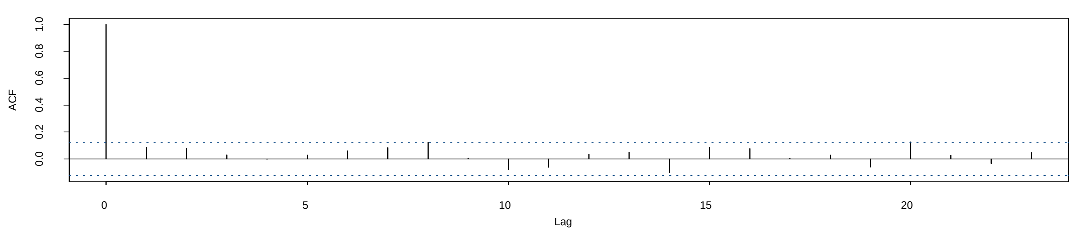

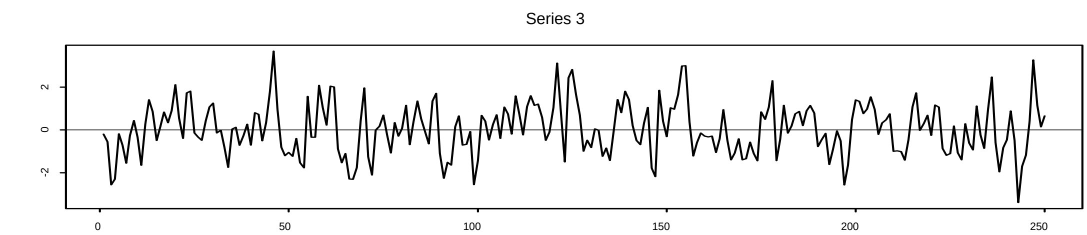

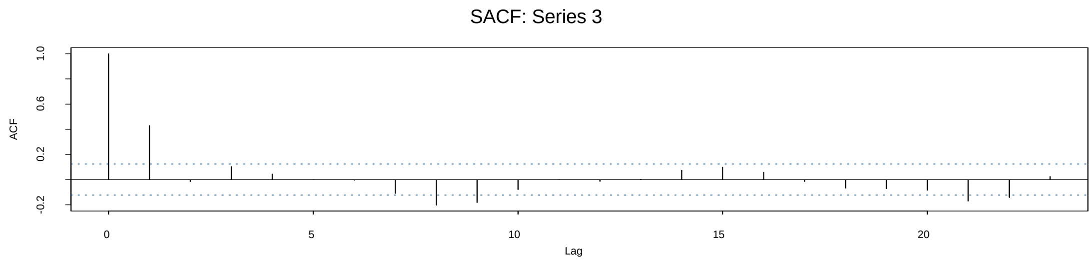

Which series is the Gaussian white noise, which series is the MA(1) process, and which series is the AR(1) process? Briefly justify your answers. (10 points)

Gaussian white noise: series 2 (SACF shows no autocorrelation)

MA(1) process: series 3 (SACF cuts off at lag 1)

AR(1) process: series 1 (SACF decays geometrically)

3. The figure below shows weekly (Friday) observations on the annualized US 3 month T-bill rate over the period 1/05/62 through 3/31/95 along with the sample ACF. (10 points)

a) Does the interest rate series look like a realization from a covariance stationary time series? Why or why not.
Interest rates don't really look like a stationary time series. The level of rates seems to change over time, and the volatility of the rates also seems to change over time.

b) Assume the interest rate series is covariance stationary. Based on the shape of the sample autocorrelation function, would an MA(1) process or an AR(1) process best describe the data? Briefly justify your answer.

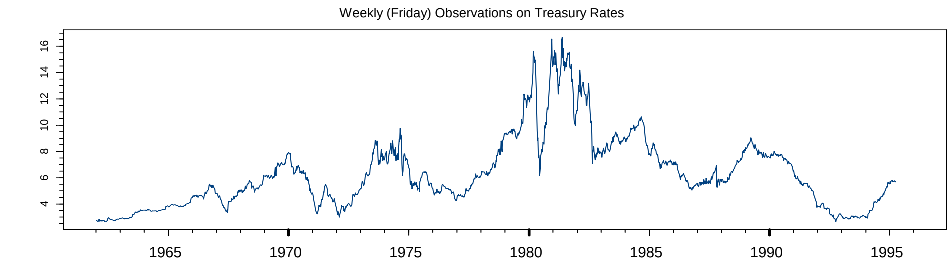

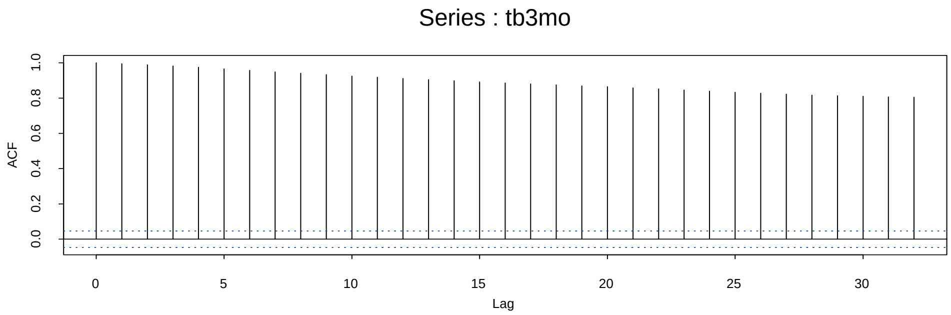

Only the first lag autocorrelation for an MA(1) process is non-zero. For interest rates, the autocorrelations are non-zero for up to 30 lags (weeks). This type of behavior can be captured by an AR(1) process with ϕ very close to 1.

IV. Descriptive Statistics (25 points, 5 points each)

1. Consider one year of *daily* continuously compounded (cc) returns on Microsoft stock computed using *daily* closing prices over the period May 2, 2005 – April 28, 2006. The drop in prices at the end of the sample is not a mistake. On April 28, 2006 Microsoft stock's price dropped 11.4% from the day before.

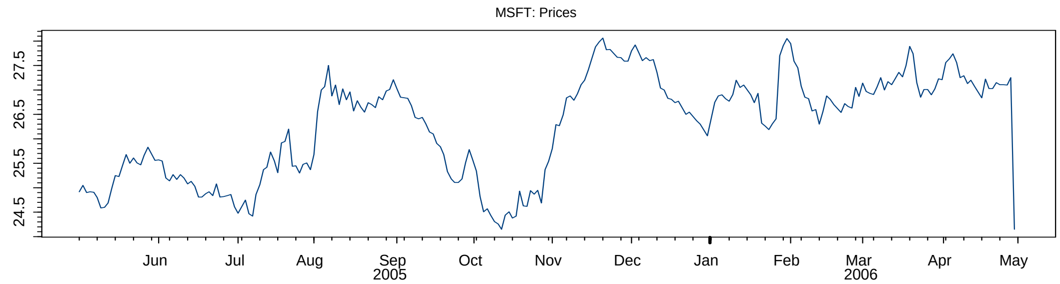

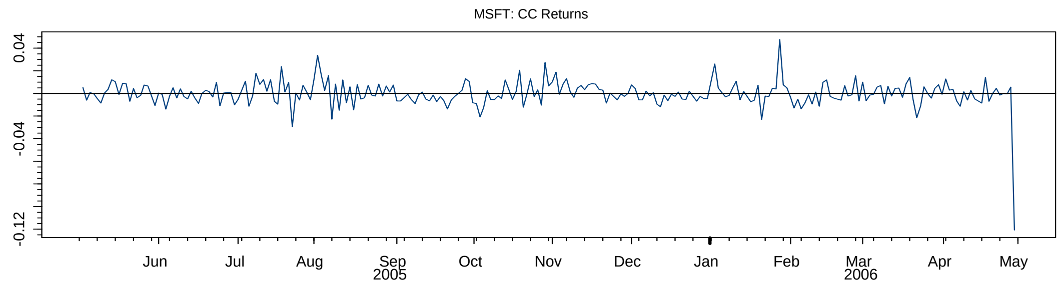

a. Do the daily cc returns appear to be a realization from a covariance stationary stochastic process? Briefly justify your answer.

Except for the last observation, the cc returns look to be reasonably covariance stationary. There is no obvious trend in the data. The volatility seems to vary a little bit with time.

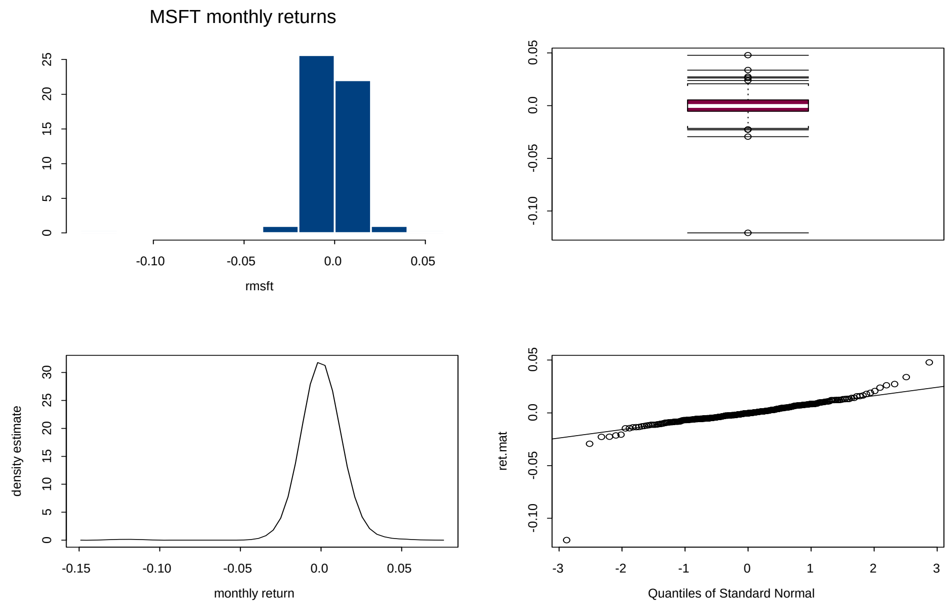

b. The figure above shows various graphical diagnostics regarding the empirical distribution of the daily cc returns on Microsoft. Based on these diagnostics, do you think that the normal distribution is a good model for the underlying probability distribution of the daily cc returns on Microsoft? Briefly justify your answer by commenting on each of the four plots.

The graphical diagnostics do not support the normal distribution for daily cc returns. The histogram is not too informative about the shape of the distribution. However, the smoothed histogram shows the impact of the large outlier in the long left tail and reveals some negative skewness. The boxplot and qqplot show strong departures from normal behavior. The boxplot shows several outliers besides the one at the end of the sample, the the qqplot shows a strong departure from linear behavior at the upper and lower quantiles.

c. Summary descriptive statistics, computed from S-PLUS, for the daily cc returns are given below. Which of these summary statistics indicate evidence for, or against, the normal distribution as the probability model for the daily cc returns?

```
> summaryStats(msftDailyReturns)

Sample Quantiles:
     min        1Q     median       3Q     max
 -0.1208 -0.005378 -0.0004037 0.005444 0.04769

Sample Moments:
       mean     std skewness kurtosis
 -0.0001255 0.01187   -3.812    45.29

Number of Observations:  250
```

Data from a normal distribution should have skewness close to zero and kurtosis close to three. The daily Microsoft cc returns has a large negative skewness and a very large kurtosis which is strong evidence against the normal distribution.

d. On April 28th, Microsoft's stock dropped 11.4% from the previous day. This resulted in over a $32 Billion loss in market value for Microsoft. Was this an unusual event if the cc returns on Microsoft follow a normal distribution? Briefly justify your answer.

If the Microsoft returns were normally distributed then there would be essentially a zero chance of seeing a -11.4% return (-12.08 cc return) in any given day. Recall, with a normal distribution 99% of the observations are between ± 3 standard deviations from the mean. The sample mean is -0.0001 and the sample SD is 0.01187. Therefore,

$$mean \pm 3SD \approx [-0.035, 0.035]$$

Hence we should see daily returns between -3.5% and 3.5% is the data were normally distributed. The observed -11.4% return is more than 10 values of the SD from the mean.

e) The empirical 1% and 5% quantiles from the daily cc returns are given below.

```
> quantile(msftDailyReturns, probs=c(0.01, 0.05))
          1%          5%
 -0.02285873 -0.01236073
```

Using these quantiles, compute the daily 1% and 5% value-at-risk (*VaR*) based on an investment of $100,000.

```
> q.msft = c(-0.02285873, -0.01236073)
> VaR.msft = (exp(q.msft) - 1) * 100000
> VaR.msft
[1] -2259.945 -1228.465
```

V. Constant Expected Return Model (25 points, 5 points each)

1. Consider the constant expected return model

$$R_{it} = \mu_i + \varepsilon_{it},\ \varepsilon_{it} \sim iid\ N(0, \sigma_i^2)$$
$$\operatorname{cov}(R_{it}, R_{jt}) = \sigma_{ij},\ corr(R_{it}, R_{jt}) = \rho_{ij}$$

for the monthly continuously compounded returns on Boeing and Microsoft (same data as lab 5) over the period July 1992 through October 2000. For this period there are 100 monthly observations.

a) Based on the S-PLUS output below, give the "plug-in principle" estimates for $\mu_i, \sigma_i^2, \sigma_i, \sigma_{ij}$ and ρ_{ij} for the two assets. Arrange these estimates nicely in a table.

```
        muhat.vals sigma2hat.vals sigmahat.vals
rboeing   0.012436      0.0057945      0.076121
  rmsft   0.027564      0.0114112      0.106823

                  covhat.vals rhohat.vals
rboeing,rmsft -0.000067409  -0.0082898
```

This is the "who's buried in Grant's tomb" question. The answer is given to you.

b) Using the above output, compute estimated standard errors for $\hat{\mu}_i,\ \hat{\sigma}_i,\ (i = boeing, microsoft)$ and $\hat{\rho}_{msft,boeing}$. Briefly comment on the precision of the estimates.

The analytic formulas for these estimated standard errors are

$$\widehat{SE}(\hat{\mu}) = \frac{\hat{\sigma}}{\sqrt{T}},\ \widehat{SE}(\hat{\sigma}) = \frac{\hat{\sigma}}{\sqrt{2T}},\ \widehat{SE}(\hat{\rho}) = \frac{(1-\hat{\rho}^2)}{\sqrt{T}}$$

Using the above formulas we get

```
> nobs
[1] 100
> se.muhat = sigmahat.vals/sqrt(nobs)
cbind(muhat.vals, se.muhat)
        muhat.vals se.muhat
rboeing    0.01244 0.007612
  rmsft    0.02756 0.010682

> se.sigmahat = sigmahat.vals/sqrt(2 * nobs)
> cbind(sigmahat.vals, se.sigmahat)
        sigmahat.vals se.sigmahat
rboeing       0.07612    0.005383
  rmsft       0.10682    0.007554

> se.rhohat = (1 - rhohat.vals^2)/sqrt(nobs)
> cbind(rhohat.vals, se.rhohat)
              rhohat.vals se.rhohat
rboeing,rmsft    -0.00829   0.09999
```

The SE values for the mean are one half the size of the mean, which is fairly large for returns. Therefore, the mean is not estimated very precisely. The SE values for the SD are about ten time smaller than the SD which is fairly small for returns. Therefore, the SD is estimated fairly precisely. The SE value, 0.1, for the correlation is larger than the correlation estimate. Since the correlation lies between -1 and 1 the uncertainty implied by the SE value is moderate.

c) For Boeing, compute a 95% confidence interval for μ. Also, compute a 95% confidence interval for $\rho_{msft,boeing}$. Briefly comment on the precision of the estimates. In particular, note if both positive and negative values are in the respective confidence intervals.

```
> ci.95.mu = c(muhat.vals[1] - 2 * se.muhat[1], muhat.vals[1]
             + 2 * se.muhat[1])
> names(ci.95.mu) = c("lower", "upper")
> ci.95.mu
     lower   upper
 -0.002789 0.02766

> ci.95.rho = c(rhohat.vals - 2 * se.rhohat, rhohat.vals
              + 2 * se.rhohat)
> names(ci.95.rho) = c("lower", "upper")
> ci.95.rho
   lower  upper
 -0.2083 0.1917
```

Both 95% confidence intervals contain positive and negative values. Because positive and negative mean returns have very different implications, this result indicates the data is not very informative about the true mean. For the correlation, the values in the *confidence interval indicate a weak linear association but the data is not clear about the direction of this association.*

The bootstrap can be used to compute an estimated standard error for the estimated 5% monthly value-at-risk, based on a $100,000 investment, computed using the formula

$$V\hat{a}R_{.05} = (e^{\hat{q}_{.05}} - 1)\cdot 100{,}000,\ \hat{q}_{.05} = \hat{\mu} + \hat{\sigma}\cdot(-1.645)$$

The following S-PLUS output is based on 1000 bootstrap replications of the above value-at-risk formula from the monthly cc returns for Boeing:

```
> VaRhat.boot
Call:
bootstrap(data = rboeing, statistic = Value.at.Risk)

Number of Replications: 1000

Summary Statistics:
        Observed  Bias   Mean   SE
rboeing   -10665 147.6 -10517 1325
```

d) What is the bootstrap estimate of the standard error for the value-at-risk estimate? Using this standard error, compute an approximate 95% confidence interval for the true 5% value-at-risk.

The bootstrap SE is

$$SE_{boot}(\widehat{VaR}_{0.05}) = 1325$$

and a 95% confidence interval is

$$\widehat{VaR}_{0.05} \pm 2 \times SE_{boot}(\widehat{VaR}_{0.05}) = [-\$13{,}315, -\$8{,}015]$$

e) The figure below (on the next page) shows a histogram and normal qq-plot of the bootstap value-at-risk values. Based on the shape of the bootstrap distribution, how good do you think is the 95% bootstrap confidence interval you computed in part d)?

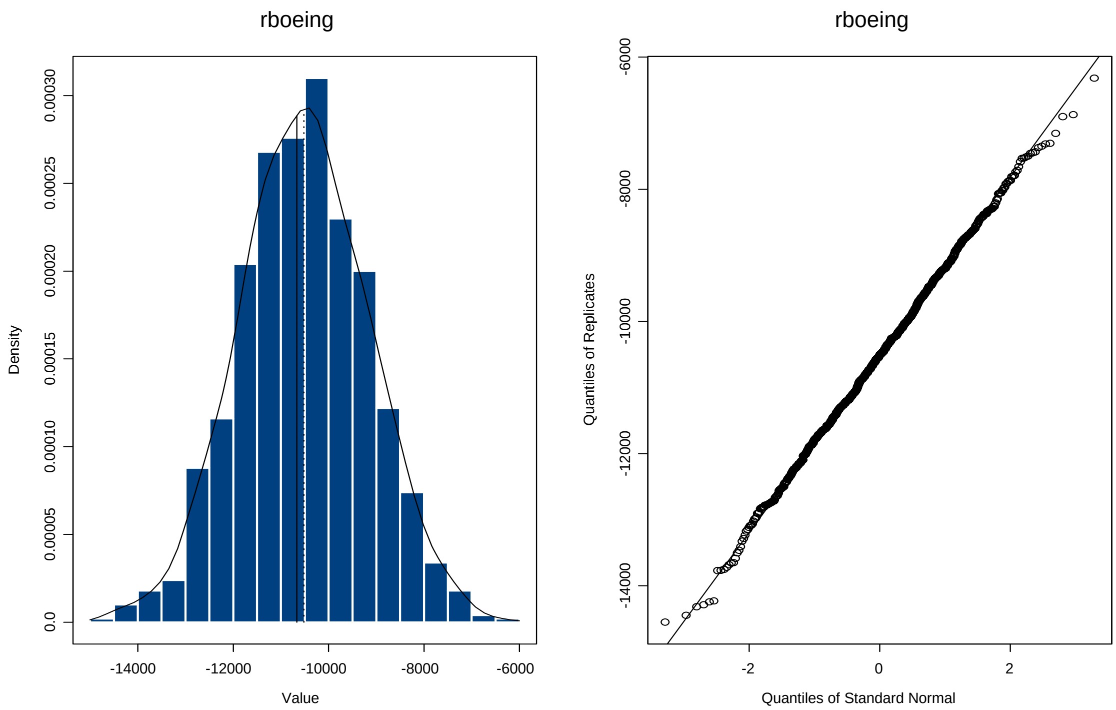

The bootstrap histogram looks pretty normal, and the qq-plot is mostly linear which also supports normality. Therefore, the boostrap estimates should be accurate.