University of Washington | Fall 2006
Department of Economics | Eric Zivot

Economics 424

Midterm Exam Suggested 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 = 120.

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

1. Consider a one year investment in two Northwest stocks: Amazon and Costco. Suppose you buy Amazon and Costco at the end of November 2005 at $P_{A,t-1} = \$48.46,\ P_{C,t-1} = \$49.59$ and then sell at the end of November 2006 for $P_{A,t} = \$37.56,\ P_{C,t} = \$52.91$. (Note: these are actual closing prices taken from Yahoo!)

```
> pa.1 = 48.46
> pa.2 = 37.56
> pc.1 = 49.59
> pc.2 = 52.91
```

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

```
> ra = (pa.2 - pa.1)/pa.1
> rc = (pc.2 - pc.1)/pc.1
> ra
[1] -0.2249
> rc
[1] 0.06695
```

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

```
> log(1 + ra)
[1] -0.2548
> log(1 + rc)
[1] 0.0648
```

c. Costco paid the following dividends between November 2005 and November 2006: \$0.11 per share cash dividend in February, \$0.11 per share cash dividend in May, and \$0.13 per share cash dividend in July. What is the annual simple total return on Costco? What is the annual dividend yield?

```
> rc.total = (pc.2 + 0.11 + 0.11 + 0.13 - pc.1)/pc.1
> div.y = (0.11 + 0.11 + 0.13)/pc.1
```

```
> rc.total
[1] 0.07401
> div.y
[1] 0.007058
> # total return = cap gain + div yeild
> rc + div.y
[1] 0.07401
```

d. The inflation rate between November 2005 and November 2006 was about 3%. Using this information, determine the simple and continuously compounded real annual returns on Amazon and Costco.

```
> inflat = 0.03
> # simple real returns
> ra.real = (1+ra)/(1+inflat) - 1
> rc.real = (1+rc)/(1+inflat) - 1
> ra.real
[1] -0.2475
> rc.real
[1] 0.03587
> # cc real returns
> log(1+ra.real)
[1] -0.2844
> log(1+rc.real)
[1] 0.03524
```

e. At the end of November, 2005, suppose you have \$100,000 to invest in Amazon and Costco over the next year. Suppose you sell short \$60,000 in Amazon and use the proceeds to buy \$160,000 in Costco. Using the results from part a, compute the annual simple return on the portfolio. Assume that both stocks do not pay a dividend.

```
> xa = -60000/100000
> xc = 160000/100000
> xa
[1] -0.6
> xc
[1] 1.6
> rp = xa*ra + xc*rc
> rp
[1] 0.2421
```

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

1. Suppose you currently hold \$100,000 in Starbucks stock. That is, your initial wealth at the beginning of the month is $W_0 = \$100,000$. Let R_{sbux} denote the monthly *simple* return

on Starbucks stock, and assume that $R_{sbux} \sim N(0.03,(0.20)^2)$. Let $W_1 = W_0(1+R_{SBUX})$ be a random variable representing your wealth at the end of the month.

a) Explain why the normal distribution may not be appropriate for describing the distribution of simple returns.

The normal distribution is defined from $-\infty$ to ∞. *Simple returns cannot be smaller than -1. Also, multi-period simple returns are multiplicative (geometric average). That is, the 2 period return is a geometric average (multiplicative) of two 1 period returns. If the 1 period returns are normally distributed then the 2 period return will not be normal.*

b) Compute $E[W_1]$, $\text{var}(W_1)$ and $SD(W_1)$

$$E[W_1] = W_0 + W_0 E[R_{SBUX}]$$
$$= \$100,000 + \$100,000 \times (0.03) = \$103,000$$
$$\text{var}(W_1) = W_0^2\, \text{var}(R_{SBUX})$$
$$= (100,000)^2 (0.20)^2 = 4 \times \text{e+08}$$
$$SD(W) = W_0 SD(R_{SBUX})$$
$$\$100,000 \times (0.20) = \$20,000$$

c) What is the probability distribution of W_1? Sketch the distribution, indicating the location of $E[W_1]$ and $E[W_1] \pm 2 \cdot SD(W_1)$.

Since R is normally distributed and W1 is a linear function of R, W1 is also normally distributed. The mean and variance of W1 is given in part b above.

d) Compute the 5% value-at-risk over the month. Hint: $q^z_{.05} = -1.645$ is the 5% quantile from the standard normal distribution.

$$q^R_{.05} = 0.03 + (0.20)(-1.645) = -0.299$$
$$VaR_{.05} = q^R_{.05} W_0 = -\$29,900$$

e) Now assume that R_{sbux} denote the monthly *continuously compounded* return on Starbucks stock. Compute the 5% value-at-risk over the month.

$$q^r_{.05} = 0.03 + (0.20)(-1.645) = -0.299$$
$$VaR_{.05} = (e^{q^r_{.05}} - 1)W_0 = -\$25.844$$

III. Time Series Concepts (20 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

$$E[Y_t] = \mu$$
$$\mathrm{var}(Y_t) = \sigma^2$$
$$\mathrm{cov}(Y_t, Y_{t-j}) = \gamma_j \text{ (depends on } j \text{ and not } t\text{)}$$

2. Consider the random walk model

$$Y_t = Y_{t-1} + \varepsilon_t = Y_0 + \sum_{j=1}^{t} \varepsilon_j,\ Y_0 = \text{constant.}$$
$$\varepsilon_t \sim iid\ N(0, \sigma^2)$$

Is $\{Y_t\}$ a covariance stationary stochastic process? Why or why not? (5 points)

No. The random walk process is not stationary. The variance of the random walk process depends on time:

$$\mathrm{var}(Y_t) = \sum_{j=1}^{t} \mathrm{var}(\varepsilon_j) = \sum_{j=1}^{t} \sigma^2 = t\sigma^2$$

3. The figure below shows annual observations on the dividend yield of the S&P 500 index over the period 1871 through 2000 along with the sample ACF. (10 points)

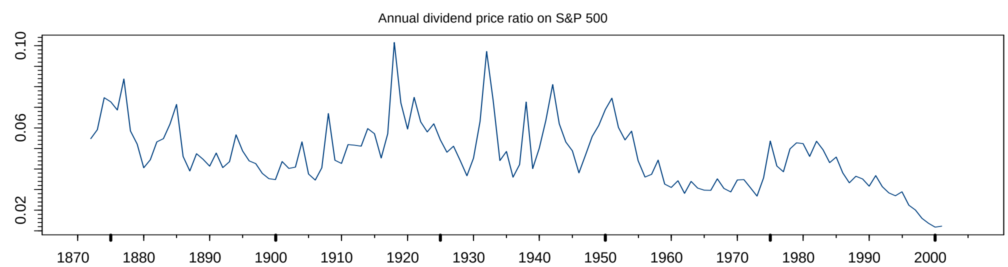

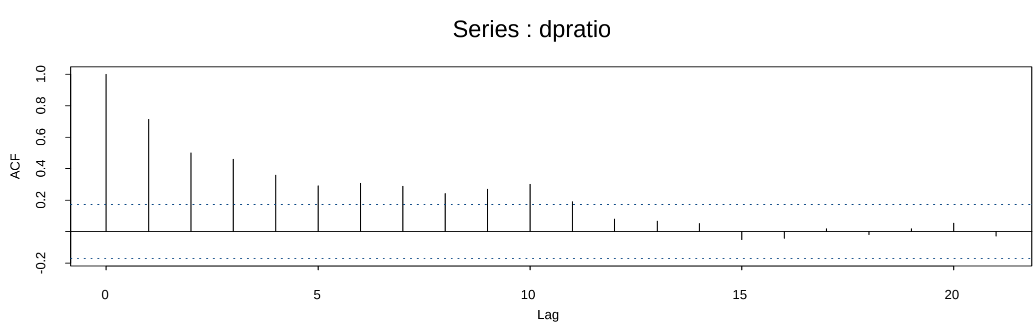

a) Does the dividend yield look like a realization from a covariance stationary time series? Why or why not.

There is no obvious trend in the dividend yield data. However, the average level looks a bit lower at the end of the sample (3%) than at the beginning of the sample (6%). It could be from a covariance stationary process, or it could be from a process with a mean value that depends on time.

b) Assume the dividend yield 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.

The SACF decays toward zero and does not cut off at lag 1. Therefore, it looks more like an AR(1) process than an MA(1) process.

VI. Constant Expected Return Model (50 points, 5 points each)

1. Consider the constant expected return model

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

for the *monthly* continuously compounded returns on the Vanguard extended market index (vexmx) and the Vanguard long-term bond index (vbltx) (subset of class project data) over the period September 2001 through September 2006. For this period there are T=60 monthly observations. The data are shown in the graph below.

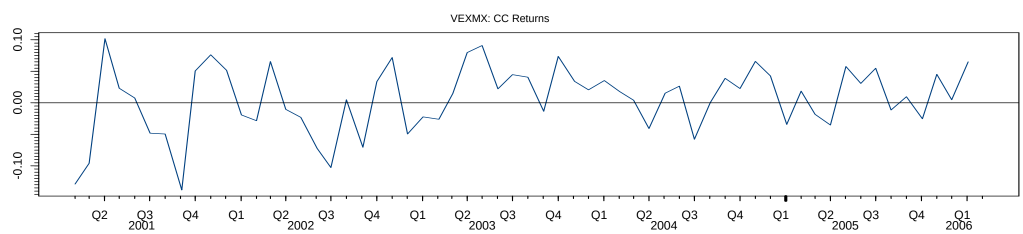

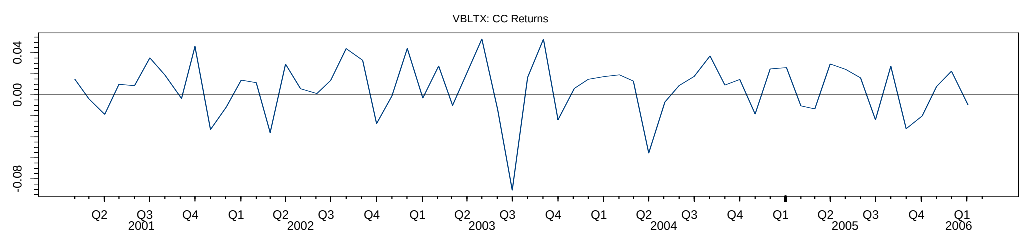

a) Do the monthly continuously compounded return data look like they come from a covariance stationary stochastic process? Why or why not?

There is no trend in the returns – they seem to fluctuate around a constant level. Also, the variability of the returns looks reasonably constant for vbltx; however, the variability for vexmx is lower toward the end of the sample than it is in the beginning of the sample. A covariance stationary process seems like a reasonable assumption for vbltx but not as reasonable for vexmx since the variance of vexmx appears to change over time.

b) What are the formulas used to compute estimates for μ_i, σ_i^2, and σ_{ij}?

$$\hat{\mu}_i = \frac{1}{T}\sum_{t=1}^{T} r_t,\ \hat{\sigma}_i^2 = \frac{1}{T-1}\sum_{t=1}^{T}(r_t - \hat{\mu})^2$$

$$\hat{\sigma}_{ij} = \frac{1}{T}\sum_{t=1}^{T}(r_{it} - \hat{\mu}_i)(r_{jt} - \hat{\mu}_j)$$

c) The following S-PLUS output gives the estimates of μ_i, σ_i, and σ_{ij} for vexmx and vbltx from the 5 years of monthly data.

```
> cbind(muhat.vals,sigmahat.vals)
      muhat.vals sigmahat.vals
vexmx   0.005770       0.05308
vbltx   0.006143       0.02648

> rhohat.vals
 vexmx,vbltx
     -0.2038
```

Briefly discuss these estimates in light of what we have learned about the CER model so far.

Both monthly mean estimates are positive and close to zero. The annualized average returns are about 6.9% for vexmx and 7.4% for vbltx. Over this period an investment in bonds did better than an investment in stocks. The SD estimates are also fairly small, with the SD of vbltx smaller than vexmx. The annualized SD values are about 18% for vexmx and about 9% for bonds. This is expected since bonds are generally thought to be safer than stocks. Since vexmx is a well diversified portfolio, its small SD relative to individual stocks is due to the diversification of risk. Interesting, the correlation between vbltx and vexmx is negative. Typically, stock and bond returns are slightly negatively correlated or uncorrelated. This is one reason to hold both stocks and bonds in your portfolio.

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

```
> nobs = numRows(projectReturns.ts)
> se.muhat = sigmahat.vals/sqrt(nobs)
> se.sigma = sigmahat.vals/sqrt(2*nobs)
```

```
> se.rho = sqrt( (1 - rhohat.vals^2)/nobs )
> se.muhat
    vexmx    vbltx
 0.006853 0.003419
> se.sigma
    vexmx    vbltx
 0.004846 0.002417
> se.rho
 vexmx,fmagx
      0.1264
```

For both vbltx and vexmx, the SE values for the mean estimates are about half the size of the mean estimates. This reflects a fair bit of uncertainty about the true value of the mean. Similarly, the SE value for the estimated correlation is about half the size of the estimated correlation. This implies a fairly imprecise estimate of correlation.

e) The analytic formulas for $SE(\hat{\sigma})$ and $SE(\hat{\rho})$ are approximations based on the Central Limit Theorem and may not be accurate for small sample sizes. The bootstrap provides an alternative means of computing these standard errors. Briefly describe how you would use the bootstrap to compute a numerical standard error for $\hat{\sigma}$. (Note: I do not want you to tell me how to do this in S-PLUS – I want you to describe the bootstrap algorithm)

The bootstrap works as follows.

Step 1: create B bootstrap samples by sampling with replacement from the original data. Each bootstrap sample has the same number of observations as the original sample

Step 2: From each bootstrap sample, compute an estimate of σ. This gives B bootstrap estimates $(\hat{\sigma}_1^*, \ldots, \hat{\sigma}_B^*)$

Step 3: Approximate the standard error for $\hat{\sigma}$ *by computing the sample standard deviation of* $(\hat{\sigma}_1^*, \ldots, \hat{\sigma}_B^*)$

$$SE_{boot}(\hat{\sigma}) = \left(\frac{1}{B-1} \sum_{i=1}^{B} (\sigma_i^{*2} - \bar{\sigma}^{*2})^2 \right)^{1/2}$$

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

```
> # 95% ci for mu
> upper = muhat.vals["vexmx"]+2*se.muhat["vexmx"]
> lower = muhat.vals["vexmx"]-2*se.muhat["vexmx"]
```

```
> cbind(lower,upper)
          lower   upper
vexmx -0.007935 0.01948

> # 95% ci for rho
> upper = rhohat.vals+2*se.rho
> lower = rhohat.vals-2*se.rho
> upper = rhohat.vals+2*se.rho
> lower = rhohat.vals-2*se.rho
> cbind(lower,upper)
              lower   upper
vexmx,vbltx -0.4566 0.04893
```

The 95% CI for the mean of vexmx contains both small positive and negative numbers, so it is not clear if the expected value is positive. Similarly, the 95% CI for the correlation between vbltx and vexmx contains both positive and negative values so it is not clear if the true correlation is negative.

g) Test the following hypotheses using a 5% significance level:
$H_0 : \mu_{vexmx} = 0$ vs. $H_1 : \mu_{vexmx} \neq 0$; $H_0 : \rho_{vexmx,vbltx} = 0$ vs. $H_1 : \rho_{vexmx,vbltx} \neq 0$.

You can do the hypotheses tests in two ways: (1) reject the null hypothesis if the value under the null hypothesis is not in the 95% confidence interval; (2) compute a t-statistic and reject the null hypothesis of the absolute value of the t-statistic is greater than 2.

Using the confidence interval method, we see that we cannot reject the null hypotheses that $\mu_{vexmx} = 0$ and $\rho_{vexmx,vbltx} = 0$ *at the 5% level because zero is in both 95% confidence intervals.*

For the t-statistic approach, we compute the following t-statistics

```
> t.stat.mu0 = muhat.vals["vexmx"]/se.muhat["vexmx"]
> t.stat.rho0 = rhohat.vals/se.rho
> abs(t.stat.mu0)
 vexmx
 0.842
> abs(t.stat.rho0)
 vexmx,vbltx
       1.613
```

Since the absolute value of both t-stats are less than 2, we do not reject the null hypotheses that $\mu_{vexmx} = 0$ and $\rho_{vexmx,vbltx} = 0$ *at the 5% level.*

h) The figure below gives some graphical diagnostics of the return distribution for vexmx. Also, estimated values of the skewness and excess kurtosis for vexmx are

```
> kurtosis(projectReturns.ts[,"vexmx"])
[1] 0.2538
> skewness(projectReturns.ts[,"vexmx"])
[1] -0.6503
```

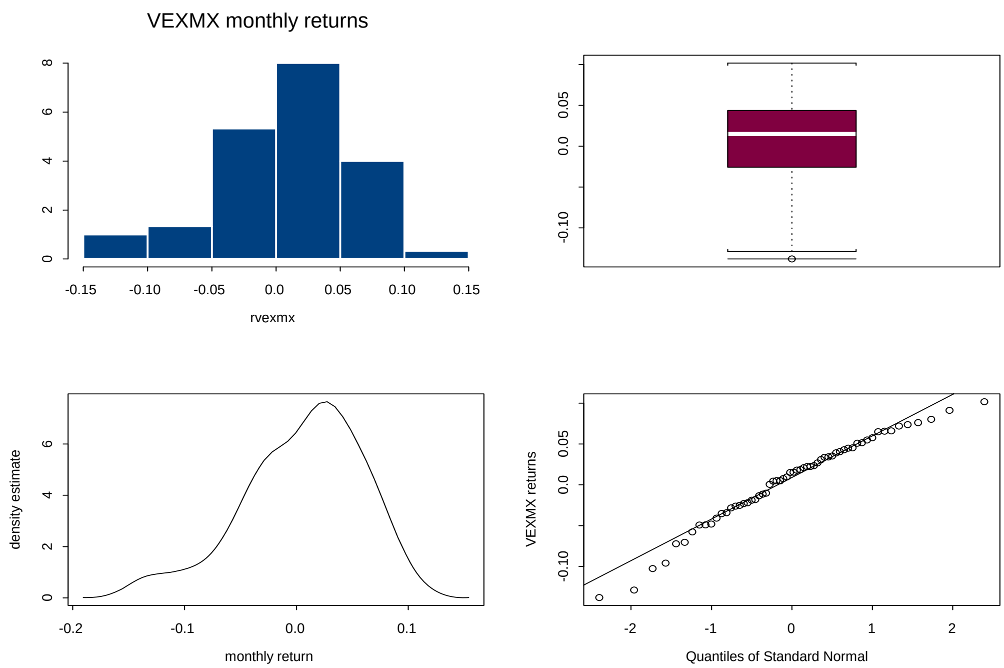

Based on this information, do you think the monthly cc returns on vexmx are normally distributed? Briefly justify your answer.

The histogram is slightly left skewed (long-left tail) and the estimated skewness is moderately negative. Also, the boxplot shows a negative outliers and the qq-plot departs from linearity in the left tail. The excess kurtosis is only slightly greater than 3. Hence, There is some evidence against the normal distribution for vexmx returns. A formal test may be computed using the JB statistic

```
> JB = nobs * (x.skew^2 + 0.25 * x.ekurt^2) / 6
> JB
[1] 4.39
```

Since the JB statistic is less than 6, we cannot reject the null hypothesis that the returns on vexmx are normally distributed at the 5% level.

i) Below are the sample autocorrelation functions for vexmx and vbltx. Using the information in these graphs, would you say that the CER model assumption that returns are uncorrelated over time is appropriate? Briefly justify your answer.

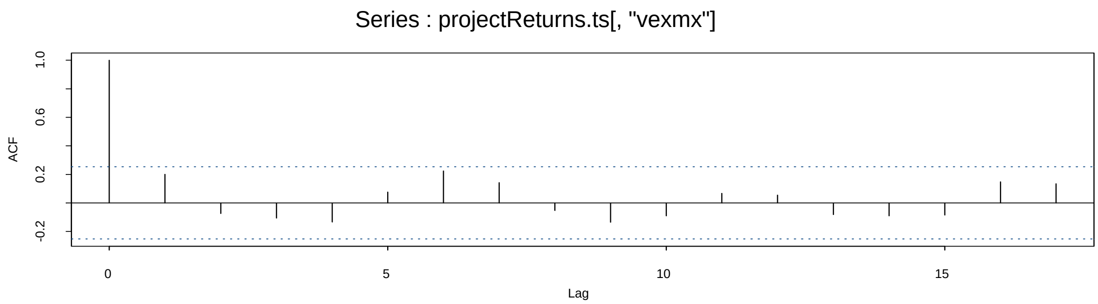

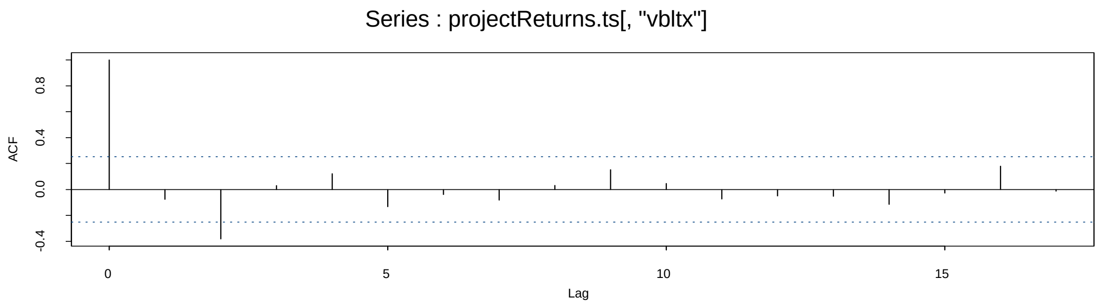

For vexmx, none of the sample autocorrelations are outside of the standard error bands so we cannot reject the null hypothesis that the returns on vexmx are uncorrelated over time. For vbltx, we only see that the 2^{nd} lag autocorrelation is outside the standard error bands. Therefore, we can reject the null hypothesis that the returns on vbltx are uncorrelated over time at the 5% level.

j) Below are graphs showing 24-month rolling estimates of μ and σ along with the monthly returns for vexmx and vbltx (in each graph, the higher dotted line represents the rolling standard deviations and the lower solid line represents the rolling means). Based on these graphs, would you say that the CER model assumption that μ and σ are constant over time is appropriate? Briefly justify your answer.

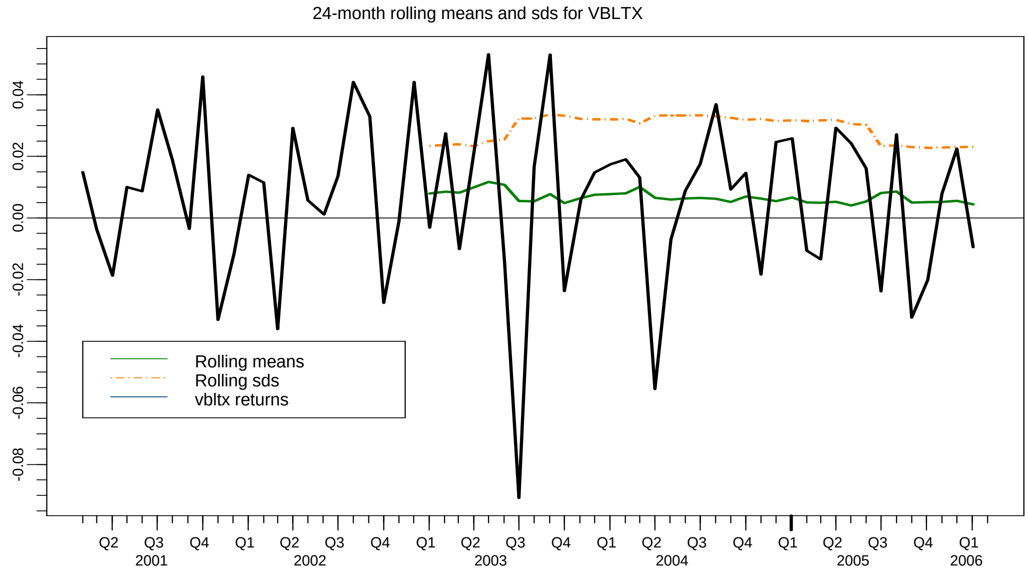

The 24-month rolling estimates of the mean and sd for vbltx are fairly constant over time, supporting the CER model assumptions.

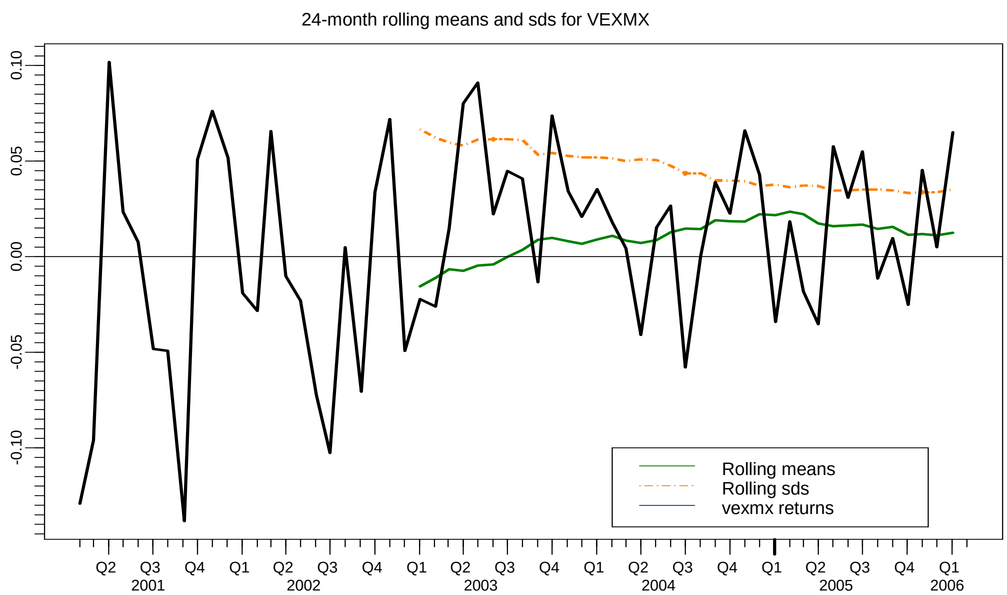

For vexmx, the 24-month rolling mean and sd estimates are clearly changing over time. The mean starts out negative and trend upward to a positive value at the end of the sample. The sd shows a downward trend over the sample. Therefore, the CER model assumption of constant parameters is not reasonable for the returns on vexmx.