

Chapter 1

Return Calculations

Updated: June 24, 2014

In this Chapter we cover asset return calculations with an emphasis on equity returns. Section 1.1 covers basic time value of money calculations. Section 1.2 covers asset return calculations, including both simple and continuously compounded returns. Section 1.3 illustrates asset return calculations using R .

1.1 The Time Value of Money

This section reviews basic time value of money calculations. The concepts of future value, present value and the compounding of interest are defined and discussed.

1.1.1 Future value, present value and simple interest.

Consider an amount $\$V$ invested for n years at a *simple interest rate* of R per annum (where R is expressed as a decimal). If compounding takes place only at the end of the year the *future value* after n years is:

$$FV_n = \$V(1 + R) \times \cdots \times (1 + R) = \$V \cdot (1 + R)^n. \quad (1.1)$$

Over the first year, $\$V$ grows to $\$V(1+R) = \$V + \$V \times R$ which represents the initial principle $\$V$ plus the payment of simple interest $\$V \times R$ for the year. Over the second year, the new principle $\$V(1+R)$ grows to $\$V(1+R)(1+R) = \$V(1 + R)^2$, and so on.

Example 1 *Future value with simple interest.*

Consider putting \$1000 in an interest checking account that pays a simple annual percentage rate of 3%. The future value after $n = 1, 5$ and 10 years is, respectively,

$$\begin{aligned} FV_1 &= \$1000 \cdot (1.03)^1 = \$1030, \\ FV_5 &= \$1000 \cdot (1.03)^5 = \$1159.27, \\ FV_{10} &= \$1000 \cdot (1.03)^{10} = \$1343.92. \end{aligned}$$

Over the first year, \$30 in interest is paid; over three years, \$159.27 in interest is accrued; over five years, \$343.92 in interest is accrued ■

The future value formula (1.1) defines a relationship between four variables: FV_n , V , R and n . Given three variables, the fourth variable can be determined. Given FV_n , R and n and solving for V gives the *present value* formula:

$$V = \frac{FV_n}{(1 + R)^n}. \quad (1.2)$$

Given FV_n , n and V , the annual interest rate on the investment is defined as:

$$R = \left(\frac{FV_n}{V} \right)^{1/n} - 1. \quad (1.3)$$

Finally, given FV_n , V and R we can solve for n :

$$n = \frac{\ln(FV_n/V)}{\ln(1 + R)}. \quad (1.4)$$

The expression (1.4) can be used to determine the number years it takes for an investment of \$ V to double. Setting $FV_n = 2V$ in (1.4) gives:

$$n = \frac{\ln(2)}{\ln(1 + R)} \approx \frac{0.7}{R},$$

which uses the approximations $\ln(2) = 0.6931 \approx 0.7$ and $\ln(1 + R) \approx R$ for R close to zero (see the Appendix). The approximation $n \approx 0.7/R$ is called the *rule of 70*.

Example 2 *Using the rule of 70.*

The table below summarizes the number of years it takes for an initial investment to double at different simple interest rates.

R	$\ln(2)/\ln(1 + R)$	$0.7/R$
0.01	69.66	70.00
0.02	35.00	35.00
0.03	23.45	23.33
0.04	17.67	17.50
0.05	14.21	14.00
0.06	11.90	11.67
0.07	10.24	10.00
0.08	9.01	8.75
0.09	8.04	7.77
0.10	7.28	7.00

■

1.1.2 Multiple compounding periods.

If interest is paid m times per year then the future value after n years is:

$$FV_n^m = \$V \cdot \left(1 + \frac{R}{m}\right)^{m \cdot n}.$$

$\frac{R}{m}$ is often referred to as the *periodic interest rate*. As m , the frequency of compounding, increases the rate becomes continuously compounded and it can be shown that future value becomes

$$FV_n^c = \lim_{m \rightarrow \infty} \$V \cdot \left(1 + \frac{R}{m}\right)^{m \cdot n} = \$V \cdot e^{R \cdot n},$$

where $e^{(\cdot)}$ is the exponential function and $e^1 = 2.71828$.

Example 3 *Future value with different compounding frequencies.*

If the simple annual percentage rate is 10% then the value of \$1000 at the end of one year ($n = 1$) for different values of m is given in the table below.

Compounding Frequency	Value of \$1000 at end of 1 year ($R = 10\%$)
Annually ($m = 1$)	1100
Quarterly ($m = 4$)	1103.8
Weekly ($m = 52$)	1105.1
Daily ($m = 365$)	1105.515
Continuously ($m = \infty$)	1105.517

■

The continuously compounded analogues to the present value, annual return and horizon period formulas (1.2), (1.3) and (1.4) are:

$$\begin{aligned}
 V &= e^{-Rn} FV_n, \\
 R &= \frac{1}{n} \ln \left(\frac{FV_n}{V} \right), \\
 n &= \frac{1}{R} \ln \left(\frac{FV_n}{V} \right).
 \end{aligned}$$

1.1.3 Effective annual rate

We now consider the relationship between simple interest rates, periodic rates, effective annual rates and continuously compounded rates. Suppose an investment pays a periodic interest rate of 2% each quarter. This gives rise to a simple annual rate of 8% ($2\% \times 4$ quarters). At the end of the year, \$1000 invested accrues to

$$\$1000 \cdot \left(1 + \frac{0.08}{4} \right)^{4 \cdot 1} = \$1082.40.$$

The *effective annual rate*, R_A , on the investment is determined by the relationship

$$\$1000 \cdot (1 + R_A) = \$1082.40.$$

Solving for R_A gives

$$R_A = \frac{\$1082.40}{\$1000} - 1 = 0.0824,$$

or $R_A = 8.24\%$. Here, the effective annual rate is the simple interest rate with annual compounding that gives the same future value that occurs with simple interest compounded four times per year. The effective annual rate is greater than the simple annual rate due to the payment of interest on interest.

The general relationship between the simple annual rate R with payments m time per year and the effective annual rate, R_A , is

$$(1 + R_A) = \left(1 + \frac{R}{m}\right)^m.$$

Given the simple rate R , we can solve for the effective annual rate using

$$R_A = \left(1 + \frac{R}{m}\right)^m - 1. \quad (1.5)$$

Given the effective annual rate R_A , we can solve for the simple rate using

$$R = m \left[(1 + R_A)^{1/m} - 1 \right].$$

The relationship between the effective annual rate and the simple rate that is compounded continuously is

$$(1 + R_A) = e^R.$$

Hence,

$$\begin{aligned} R_A &= e^R - 1, \\ R &= \ln(1 + R_A). \end{aligned}$$

Example 4 *Determine effective annual rates.*

The effective annual rates associated with the investments in Example 2 are given in the table below:

Compounding Frequency	Value of \$1000 at end of 1 year ($R = 10\%$)	R_A
Annually ($m = 1$)	1100	10%
Quarterly ($m = 4$)	1103.8	10.38%
Weekly ($m = 52$)	1105.1	10.51%
Daily ($m = 365$)	1105.515	10.55%
Continuously ($m = \infty$)	1105.517	10.55%

Example 5 *Determine continuously compounded rate from effective annual rate*

Suppose an investment pays a periodic interest rate of 5% every six months ($m = 2, R/2 = 0.05$). In the market this would be quoted as having an annual percentage rate, R_A , of 10%. An investment of \$100 yields $\$100 \cdot (1.05)^2 = \110.25 after one year. The effective annual rate, R_A , is then 10.25%. To find the continuously compounded simple rate that gives the same future value as investing at the effective annual rate we solve

$$R = \ln(1.1025) = 0.09758.$$

That is, if interest is compounded continuously at a simple annual rate of 9.758% then \$100 invested today would grow to $\$100 \cdot e^{0.09758} = \110.25 ■

1.2 Asset Return Calculations

In this section, we review asset return calculations given initial and future prices associated with an investment. We first cover simple return calculations, which are typically reported in practice but are often not convenient for statistical modeling purposes. We then describe continuously compounded return calculations, which are more convenient for statistical modeling purposes.

1.2.1 Simple Returns

Consider purchasing an asset (e.g., stock, bond, ETF, mutual fund, option, etc.) at time t_0 for the price P_{t_0} , and then selling the asset at time t_1 for the price P_{t_1} . If there are no intermediate cash flows (e.g., dividends) between t_0 and t_1 , the rate of return over the period t_0 to t_1 is the percentage change in price:

$$R(t_0, t_1) = \frac{P_{t_1} - P_{t_0}}{P_{t_0}}. \quad (1.6)$$

The time between t_0 and t_1 is called the *holding period* and (1.6) is called the *holding period return*. In principle, the holding period can be any amount of time: one second; five minutes; eight hours; two days, six minutes, and two seconds; fifteen years. To simply matters, in this chapter we will assume that the holding period is some increment of calendar time; e.g., one day, one month or one year. In particular, we will assume a default holding period of one month in what follows.

Let P_t denote the price at the end of month t of an asset that pays no dividends and let P_{t-1} denote the price at the end of month $t - 1$. Then the one-month *simple net return* on an investment in the asset between months $t - 1$ and t is defined as

$$R_t = \frac{P_t - P_{t-1}}{P_{t-1}} = \% \Delta P_t. \quad (1.7)$$

Writing $\frac{P_t - P_{t-1}}{P_{t-1}} = \frac{P_t}{P_{t-1}} - 1$, we can define the *simple gross return* as

$$1 + R_t = \frac{P_t}{P_{t-1}}. \quad (1.8)$$

The one-month gross return has the interpretation of the future value of \$1 invested in the asset for one-month. Unless otherwise stated, when we refer to returns we mean net returns. Since asset prices must always be non-negative (a long position in an asset is a limited liability investment), the smallest value for R_t is -1 or -100% .

Example 6 *Simple return calculation.*

Consider a one-month investment in Microsoft stock. Suppose you buy the stock in month $t - 1$ at $P_{t-1} = \$85$ and sell the stock the next month for

$P_t = \$90$. Further assume that Microsoft does not pay a dividend between months $t - 1$ and t . The one-month simple net and gross returns are then

$$R_t = \frac{\$90 - \$85}{\$85} = \frac{\$90}{\$85} - 1 = 1.0588 - 1 = 0.0588,$$

$$1 + R_t = 1.0588.$$

The one-month investment in Microsoft yielded a 5.88% per month return. Alternatively, \$1 invested in Microsoft stock in month $t - 1$ grew to \$1.0588 in month t ■

Multi-period returns

The simple two-month return on an investment in an asset between months $t - 2$ and t is defined as

$$R_t(2) = \frac{P_t - P_{t-2}}{P_{t-2}} = \frac{P_t}{P_{t-2}} - 1.$$

Writing $\frac{P_t}{P_{t-2}} = \frac{P_t}{P_{t-1}} \cdot \frac{P_{t-1}}{P_{t-2}}$ the two-month return can be expressed as:

$$\begin{aligned} R_t(2) &= \frac{P_t}{P_{t-1}} \cdot \frac{P_{t-1}}{P_{t-2}} - 1 \\ &= (1 + R_t)(1 + R_{t-1}) - 1. \end{aligned}$$

Then the simple two-month gross return becomes:

$$1 + R_t(2) = (1 + R_t)(1 + R_{t-1}) = 1 + R_{t-1} + R_t + R_{t-1}R_t,$$

which is a product of the two simple one-month gross returns and not one plus the sum of the two one-month returns. Hence,

$$R_t(2) = R_{t-1} + R_t + R_{t-1}R_t.$$

If, however, R_{t-1} and R_t are small then $R_{t-1}R_t \approx 0$ and $1 + R_t(2) \approx 1 + R_{t-1} + R_t$ so that $R_t(2) \approx R_{t-1} + R_t$.

Adding two simple one-period returns to arrive at a two-period return when one-period returns are large can lead to very misleading results. For example, suppose that $R_{t-1} = 0.5$ and $R_t = -0.5$. Adding the two one-period returns gives a two-period return of zero. However, the actual two-period

return is $R_t(2) - 1 = (1.5)(0.5) - 1 = -0.25$. This example highlights the fact that equal positive and negative percentage changes do not affect wealth symmetrically.

In general, the k -month gross return is defined as the product of k one-month gross returns:

$$\begin{aligned} 1 + R_t(k) &= (1 + R_t)(1 + R_{t-1}) \cdots (1 + R_{t-k+1}) \\ &= \prod_{j=0}^{k-1} (1 + R_{t-j}). \end{aligned} \quad (1.9)$$

Example 7 *Computing two-period returns*

Continuing with the previous example, suppose that the price of Microsoft stock in month $t - 2$ is \$80 and no dividend is paid between months $t - 2$ and t . The two-month net return is

$$R_t(2) = \frac{\$90 - \$80}{\$80} = \frac{\$90}{\$80} - 1 = 1.1250 - 1 = 0.1250,$$

or 12.50% per two months. The two one-month returns are

$$\begin{aligned} R_{t-1} &= \frac{\$85 - \$80}{\$80} = 1.0625 - 1 = 0.0625, \\ R_t &= \frac{\$90 - 85}{\$85} = 1.0588 - 1 = 0.0588, \end{aligned}$$

and the geometric average of the two one-month gross returns is

$$1 + R_t(2) = 1.0625 \times 1.0588 = 1.1250.$$

■

Portfolio Returns

Consider an investment of $\$V$ in two assets, named asset A and asset B. Let x_A denote the fraction or share of wealth invested in asset A, and let x_B denote the remaining fraction invested in asset B. The dollar amounts invested in assets A and B are $\$V \times x_A$ and $\$V \times x_B$, respectively. We assume that the investment shares add up to 1, so that $x_A + x_B = 1$. The collection of investment shares (x_A, x_B) defines a portfolio. For example, one

portfolio may be $(x_A = 0.5, x_B = 0.5)$ and another may be $(x_A = 0, x_B = 1)$. Negative values for x_A or x_B represent short sales. Let $R_{A,t}$ and $R_{B,t}$ denote the simple one-period returns on assets A and B. We wish to determine the simple one-period return on the portfolio defined by (x_A, x_B) . To do this, note that at the end of period t , the investments in assets A and B are worth $\$V \times x_A(1 + R_{A,t})$ and $\$V \times x_B(1 + R_{B,t})$, respectively. Hence, at the end of period t the portfolio is worth

$$\$V \times [x_A(1 + R_{A,t}) + x_B(1 + R_{B,t})].$$

Hence, $x_A(1 + R_{A,t}) + x_B(1 + R_{B,t})$ defines the gross return on the portfolio. The portfolio gross return is equal to a weighted average of the gross returns on assets A and B, where the weights are the portfolio shares x_A and x_B .

To determine the portfolio rate of return, re-write the portfolio gross return as

$$1 + R_{p,t} = x_A + x_B + x_A R_{A,t} + x_B R_{B,t} = 1 + x_A R_{A,t} + x_B R_{B,t},$$

since $x_A + x_B = 1$ by construction. Then the portfolio rate of return is

$$R_{p,t} = x_A R_{A,t} + x_B R_{B,t},$$

which is equal to a weighted average of the simple returns on assets A and B, where the weights are the portfolio shares x_A and x_B .

Example 8 *Compute portfolio return*

Consider a portfolio of Microsoft and Starbucks stock in which you initially purchase ten shares of each stock at the end of month $t - 1$ at the prices $P_{msft,t-1} = \$85$ and $P_{sbux,t-1} = \$30$, respectively. The initial value of the portfolio is $V_{t-1} = 10 \times \$85 + 10 \times 30 = \$1,150$. The portfolio shares are $x_{msft} = 850/1150 = 0.7391$ and $x_{sbux} = 30/1150 = 0.2609$. Suppose at the end of month t , $P_{msft,t} = \$90$ and $P_{sbux,t} = \$28$. Assuming that Microsoft and Starbucks do not pay a dividend between periods $t - 1$ and t , the one-period returns on the two stocks are

$$R_{msft,t} = \frac{\$90 - \$85}{\$85} = 0.0588,$$

$$R_{sbux,t} = \frac{\$28 - \$30}{\$30} = -0.0667.$$

The one-month rate of return on the portfolio is then

$$R_{p,t} = (0.7391)(0.0588) + (0.2609)(-0.0667) = 0.02609,$$

and the portfolio value at the end of month t is

$$V_t = V_{t-1}(1 + R_{p,t}) = \$1,100 \times (1.02609) = \$1,180$$

■

In general, for a portfolio of n assets with investment shares x_i such that $x_1 + \dots + x_n = 1$, the one-period portfolio gross and simple returns are defined as

$$1 + R_{p,t} = \sum_{i=1}^n x_i(1 + R_{i,t}) \quad (1.10)$$

$$R_{p,t} = \sum_{i=1}^n x_i R_{i,t} \quad (1.11)$$

Adjusting for dividends

If an asset pays a dividend, D_t , sometime between months $t - 1$ and t , the *total net return* calculation becomes

$$R_t^{total} = \frac{P_t + D_t - P_{t-1}}{P_{t-1}} = \frac{P_t - P_{t-1}}{P_{t-1}} + \frac{D_t}{P_{t-1}}, \quad (1.12)$$

where $\frac{P_t - P_{t-1}}{P_{t-1}}$ is referred as the *capital gain* and $\frac{D_t}{P_{t-1}}$ is referred to as the *dividend yield*. The *total gross return* is

$$1 + R_t^{total} = \frac{P_t + D_t}{P_{t-1}}. \quad (1.13)$$

The formula (1.9) for computing multiperiod return remains the same except that one-period gross returns are computed using (1.13).

Example 9 *Compute total return when dividends are paid*

Consider a one-month investment in Microsoft stock. Suppose you buy the stock in month $t - 1$ at $P_{t-1} = \$85$ and sell the stock the next month for

$P_t = \$90$. Further assume that Microsoft pays a \$1 dividend between months $t - 1$ and t . The capital gain, dividend yield and total return are then

$$\begin{aligned} R_t &= \frac{\$90 + \$1 - \$85}{\$85} = \frac{\$90 - \$85}{\$85} + \frac{\$1}{\$85} \\ &= 0.0588 + 0.0118 \\ &= 0.0707 \end{aligned}$$

The one-month investment in Microsoft yields a 7.07% per month total return. The capital gain component is 5.88% and the dividend yield component is 1.18% ■

Adjusting for Inflation

The return calculations considered so far are based on the *nominal* or current prices of assets. Returns computed from nominal prices are nominal returns. The *real return* on an asset over a particular horizon takes into account the growth rate of the general price level over the horizon. If the nominal price of the asset grows faster than the general price level then the nominal return will be greater than the inflation rate and the real return will be positive. Conversely, if the nominal price of the asset increases less than the general price level then the nominal return will be less than the inflation rate and the real return will be negative.

The computation of real returns on an asset is a two step process:

- Deflate the nominal price of the asset by the general price level
- Compute returns in the usual way using the deflated prices

To illustrate, consider computing the real simple one-period return on an asset. Let P_t denote the nominal price of the asset at time t and let CPI_t denote an index of the general price level (e.g. consumer price index) at time t ¹. The deflated or real price at time t is

$$P_t^{\text{Real}} = \frac{P_t}{CPI_t},$$

¹The CPI is usually normalized to 1 or 100 in some base year. We assume that the CPI is normalized to 1 in the base year for simplicity.

and the real one-period return is

$$\begin{aligned} R_t^{\text{Real}} &= \frac{P_t^{\text{Real}} - P_{t-1}^{\text{Real}}}{P_{t-1}^{\text{Real}}} = \frac{\frac{P_t}{CPI_t} - \frac{P_{t-1}}{CPI_{t-1}}}{\frac{P_{t-1}}{CPI_{t-1}}} \\ &= \frac{P_t}{P_{t-1}} \cdot \frac{CPI_{t-1}}{CPI_t} - 1. \end{aligned} \quad (1.14)$$

The one-period gross real return is

$$1 + R_t^{\text{Real}} = \frac{P_t}{P_{t-1}} \cdot \frac{CPI_{t-1}}{CPI_t}. \quad (1.15)$$

If we define inflation between periods $t - 1$ and t as

$$\pi_t = \frac{CPI_t - CPI_{t-1}}{CPI_{t-1}} = \% \Delta CPI_t, \quad (1.16)$$

then $1 + \pi_t = \frac{CPI_t}{CPI_{t-1}}$ and (1.14) may be re-expressed as

$$R_t^{\text{Real}} = \frac{1 + R_t}{1 + \pi_t} - 1. \quad (1.17)$$

Example 10 *Compute real return*

Consider, again, a one-month investment in Microsoft stock. Suppose the CPI in months $t - 1$ and t is 1 and 1.01, respectively, representing a 1% monthly growth rate in the overall price level. The real prices of Microsoft stock are

$$P_{t-1}^{\text{Real}} = \frac{\$85}{1} = \$85, \quad P_t^{\text{Real}} = \frac{\$90}{1.01} = \$89.1089,$$

and the real monthly return is

$$R_t^{\text{Real}} = \frac{\$89.10891 - \$85}{\$85} = 0.0483.$$

The nominal return and inflation over the month are

$$R_t = \frac{\$90 - \$85}{\$85} = 0.0588, \quad \pi_t = \frac{1.01 - 1}{1} = 0.01.$$

Then the real return computed using (1.17) is

$$R_t^{\text{Real}} = \frac{1.0588}{1.01} - 1 = 0.0483.$$

Notice that simple real return is almost, but not quite, equal to the simple nominal return minus the inflation rate:

$$R_t^{\text{Real}} \approx R_t - \pi_t = 0.0588 - 0.01 = 0.0488.$$

■

Annualizing returns

Very often returns over different horizons are annualized, i.e., converted to an annual return, to facilitate comparisons with other investments. The annualization process depends on the holding period of the investment and an implicit assumption about compounding. We illustrate with several examples.

To start, if our investment horizon is one year then the annual gross and net returns are just

$$1 + R_A = 1 + R_t(12) = \frac{P_t}{P_{t-12}} = (1 + R_t)(1 + R_{t-1}) \cdots (1 + R_{t-11}),$$

$$R_A = R_t(12).$$

In this case, no compounding is required to create an annual return.

Next, consider a one-month investment in an asset with return R_t . What is the annualized return on this investment? If we assume that we receive the same return $R = R_t$ every month for the year, then the gross annual return is

$$1 + R_A = 1 + R_t(12) = (1 + R)^{12}.$$

That is, the annual gross return is defined as the monthly return compounded for 12 months. The net annual return is then

$$R_A = (1 + R)^{12} - 1.$$

Example 11 *Compute annualized return from one-month return*

In the first example, the one-month return, R_t , on Microsoft stock was 5.88%. If we assume that we can get this return for 12 months then the annualized return is

$$R_A = (1.0588)^{12} - 1 = 1.9850 - 1 = 0.9850,$$

or 98.50% per year. Pretty good! ■

Now, consider a two-month investment with return $R_t(2)$. If we assume that we receive the same two-month return $R(2) = R_t(2)$ for the next six two-month periods, then the gross and net annual returns are

$$\begin{aligned} 1 + R_A &= (1 + R(2))^6, \\ R_A &= (1 + R(2))^6 - 1. \end{aligned}$$

Here the annual gross return is defined as the two-month return compounded for 6 months.

Example 12 *Compute annualized return from two-month return*

Suppose the two-month return, $R_t(2)$, on Microsoft stock is 12.5%. If we assume that we can get this two-month return for the next 6 two-month periods then the annualized return is

$$R_A = (1.1250)^6 - 1 = 2.0273 - 1 = 1.0273$$

or 102.73% per year ■

Now suppose that our investment horizon is two years. That is, we start our investment at time $t - 24$ and cash out at time t . The two-year gross return is then $1 + R_t(24) = \frac{P_t}{P_{t-24}}$. What is the annual return on this two-year investment? The process is the same as computing the effective annual rate. To determine the annual return we solve the following relationship for R_A :

$$\begin{aligned} (1 + R_A)^2 &= 1 + R_t(24) \implies \\ R_A &= (1 + R_t(24))^{1/2} - 1. \end{aligned}$$

In this case, the annual return is compounded twice to get the two-year return and the relationship is then solved for the annual return.

Example 13 *Compute annualized return from two-year return*

Suppose that the price of Microsoft stock 24 months ago is $P_{t-24} = \$50$ and the price today is $P_t = \$90$. The two-year gross return is $1 + R_t(24) = \frac{\$90}{\$50} = 1.8000$ which yields a two-year net return of $R_t(24) = 0.80 = 80\%$. The annual return for this investment is defined as

$$R_A = (1.800)^{1/2} - 1 = 1.3416 - 1 = 0.3416,$$

or 34.16% per year ■

1.2.2 Continuously Compounded Returns

In this section we define continuously compounded returns from simple returns, and describe their properties.

One-period Returns

Let R_t denote the simple monthly return on an investment. The *continuously compounded monthly return*, r_t , is defined as:

$$r_t = \ln(1 + R_t) = \ln\left(\frac{P_t}{P_{t-1}}\right), \quad (1.18)$$

where $\ln(\cdot)$ is the natural log function². To see why r_t is called the continuously compounded return, take the exponential of both sides of (1.18) to give:

$$e^{r_t} = 1 + R_t = \frac{P_t}{P_{t-1}}.$$

Rearranging we get

$$P_t = P_{t-1}e^{r_t},$$

so that r_t is the continuously compounded growth rate in prices between months $t - 1$ and t . This is to be contrasted with R_t , which is the simple growth rate in prices between months $t - 1$ and t without any compounding. Furthermore, since $\ln\left(\frac{x}{y}\right) = \ln(x) - \ln(y)$ it follows that

$$\begin{aligned} r_t &= \ln\left(\frac{P_t}{P_{t-1}}\right) \\ &= \ln(P_t) - \ln(P_{t-1}) \\ &= p_t - p_{t-1}, \end{aligned}$$

where $p_t = \ln(P_t)$. Hence, the continuously compounded monthly return, r_t , can be computed simply by taking the first difference of the natural logarithms of monthly prices.

Example 14 Compute continuously compounded returns

²The continuously compounded return is always defined since asset prices, P_t , are always non-negative. Properties of logarithms and exponentials are discussed in the appendix to this chapter.

Using the price and return data from Example 1, the continuously compounded monthly return on Microsoft stock can be computed in two ways:

$$\begin{aligned}r_t &= \ln(1.0588) = 0.0571, \\r_t &= \ln(90) - \ln(85) = 4.4998 - 4.4427 = 0.0571.\end{aligned}$$

Notice that r_t is slightly smaller than R_t . Why?

Given a monthly continuously compounded return r_t , is straightforward to solve back for the corresponding simple net return R_t :

$$R_t = e^{r_t} - 1. \tag{1.19}$$

Hence, nothing is lost by considering continuously compounded returns instead of simple returns. Continuously compounded returns are very similar to simple returns as long as the return is relatively small, which it generally will be for monthly or daily returns. Since R_t is bounded from below by -1 , the smallest value for r_t is $-\infty$. This, however, does not mean that you could lose an infinite amount of money on an investment. The actual amount of money lost is determined by the simple return (1.19).

For modeling and statistical purposes it is often much more convenient to use continuously compounded returns due to the additivity property of multiperiod continuously compounded returns discussed in the next subsection.

Example 15 *Determine simple return from continuously compounded return*

In the previous example, the continuously compounded monthly return on Microsoft stock is $r_t = 5.71\%$. The simple net return is then

$$R_t = e^{.0571} - 1 = 0.0588.$$

■

Multi-Period Returns

The relationship between multi-period continuously compounded returns and one-period continuously compounded returns is more simple than the relationship between multi-period simple returns and one-period simple returns.

To illustrate, consider the two-month continuously compounded return defined as:

$$r_t(2) = \ln(1 + R_t(2)) = \ln\left(\frac{P_t}{P_{t-2}}\right) = p_t - p_{t-2}.$$

Taking exponentials of both sides shows that

$$P_t = P_{t-2}e^{r_t(2)}$$

so that $r_t(2)$ is the continuously compounded growth rate of prices between months $t - 2$ and t . Using $\frac{P_t}{P_{t-2}} = \frac{P_t}{P_{t-1}} \cdot \frac{P_{t-1}}{P_{t-2}}$ and the fact that $\ln(x \cdot y) = \ln(x) + \ln(y)$ it follows that

$$\begin{aligned} r_t(2) &= \ln\left(\frac{P_t}{P_{t-1}} \cdot \frac{P_{t-1}}{P_{t-2}}\right) \\ &= \ln\left(\frac{P_t}{P_{t-1}}\right) + \ln\left(\frac{P_{t-1}}{P_{t-2}}\right) \\ &= r_t + r_{t-1}. \end{aligned}$$

Hence the continuously compounded two-month return is just the sum of the two continuously compounded one-month returns. Recall, with simple returns the two-month return is a multiplicative (geometric) sum of two one-month returns.

Example 16 *Compute two-month continuously compounded returns.*

Using the data from Example 2, the continuously compounded two-month return on Microsoft stock can be computed in two equivalent ways. The first way uses the difference in the logs of P_t and P_{t-2} :

$$r_t(2) = \ln(90) - \ln(80) = 4.4998 - 4.3820 = 0.1178.$$

The second way uses the sum of the two continuously compounded one-month returns. Here $r_t = \ln(90) - \ln(85) = 0.0571$ and $r_{t-1} = \ln(85) - \ln(80) = 0.0607$ so that

$$r_t(2) = 0.0571 + 0.0607 = 0.1178.$$

Notice that $r_t(2) = 0.1178 < R_t(2) = 0.1250$ ■

The continuously compounded k -month return is defined by

$$r_t(k) = \ln(1 + R_t(k)) = \ln\left(\frac{P_t}{P_{t-k}}\right) = p_t - p_{t-k}.$$

Using similar manipulations to the ones used for the continuously compounded two-month return, we can express the continuously compounded k -month return as the sum of k continuously compounded monthly returns:

$$r_t(k) = \sum_{j=0}^{k-1} r_{t-j}. \quad (1.20)$$

The additivity of continuously compounded returns to form multiperiod returns is an important property for statistical modeling purposes.

Portfolio Returns

The continuously compounded portfolio return is defined by (1.18), where R_t is computed using the portfolio return (1.11). However, notice that

$$r_{p,t} = \ln(1 + R_{p,t}) = \ln\left(1 + \sum_{i=1}^n x_i R_{i,t}\right) \neq \sum_{i=1}^n x_i r_{i,t}, \quad (1.21)$$

where $r_{i,t}$ denotes the continuously compounded one-period return on asset i . If the portfolio return $R_{p,t} = \sum_{i=1}^n x_i R_{i,t}$ is not too large then $r_{p,t} \approx R_{p,t}$ otherwise, $R_{p,t} > r_{p,t}$.

Example 17 Compute continuously compounded portfolio returns

Consider a portfolio of Microsoft and Starbucks stock with $x_{msft} = 0.25$, $x_{sbux} = 0.75$, $R_{msft,t} = 0.0588$, $R_{sbux,t} = -0.0503$ and $R_{p,t} = -0.02302$. Using (1.21), the continuous compounded portfolio return is

$$r_{p,t} = \ln(1 - 0.02302) = \ln(0.977) = -0.02329.$$

Using $r_{msft,t} = \ln(1 + 0.0588) = 0.0572$ and $r_{sbux,t} = \ln(1 - 0.0503) = -0.05161$, notice that

$$x_{msft}r_{msft} + x_{sbux}r_{sbux} = -0.02442 \neq r_{p,t}$$

■

Adjusting for Dividends

The continuously compounded one-period return adjusted for dividends is defined by (1.18), where R_t is computed using (1.12)³.

Example 18 *Compute continuously compounded total return*

From example 9, the total simple return using (1.12) is $R_t = 0.0707$. The continuously compounded total return is then

$$r_t = \ln(1 + R_t) = \ln(1.0707) = 0.0683.$$

■

Adjusting for Inflation

Adjusting continuously compounded nominal returns for inflation is particularly simple. The continuously compounded one-period real return is defined as:

$$r_t^{\text{Real}} = \ln(1 + R_t^{\text{Real}}). \quad (1.22)$$

Using (1.15), it follows that

$$\begin{aligned} r_t^{\text{Real}} &= \ln\left(\frac{P_t}{P_{t-1}} \cdot \frac{CPI_{t-1}}{CPI_t}\right) \\ &= \ln\left(\frac{P_t}{P_{t-1}}\right) + \ln\left(\frac{CPI_{t-1}}{CPI_t}\right) \\ &= \ln(P_t) - \ln(P_{t-1}) + \ln(CPI_{t-1}) - \ln(CPI_t) \\ &= \ln(P_t) - \ln(P_{t-1}) - (\ln(CPI_t) - \ln(CPI_{t-1})) \\ &= r_t - \pi_t^c, \end{aligned} \quad (1.23)$$

where $r_t = \ln(P_t) - \ln(P_{t-1}) = \ln(1 + R_t)$ is the nominal continuously compounded one-period return and $\pi_t^c = \ln(CPI_t) - \ln(CPI_{t-1}) = \ln(1 + \pi_t)$ is the one-period continuously compounded growth rate in the general price level (continuously compounded one-period inflation rate). Hence, the real continuously compounded return is simply the nominal continuously compounded return minus the the continuously compounded inflation rate.

³Show formula from Campbell, Lo and MacKinlay.

Example 19 *Compute continuously compounded real return*

From example 10, the nominal simple return is $R_t = 0.0588$, the monthly inflation rate is $\pi_t = 0.01$, and the real simple return is $R_t^{\text{Real}} = 0.0483$. Using (1.22), the real continuously compounded return is

$$r_t^{\text{Real}} = \ln(1 + R_t^{\text{Real}}) = \ln(1.0483) = 0.047.$$

Equivalently, using (1.23) the real return may also be computed as

$$r_t^{\text{Real}} = r_t - \pi_t^c = \ln(1.0588) - \ln(1.01) = 0.047.$$

■

Annualizing Continuously Compounded Returns

Just as we annualized simple monthly returns, we can also annualize continuously compounded monthly returns. For example, if our investment horizon is one year then the annual continuously compounded return is just the sum of the twelve monthly continuously compounded returns:

$$\begin{aligned} r_A &= r_t(12) = r_t + r_{t-1} + \cdots + r_{t-11} \\ &= \sum_{j=0}^{11} r_{t-j}. \end{aligned}$$

The average continuously compounded monthly return is defined as:

$$\bar{r}_m = \frac{1}{12} \sum_{j=0}^{11} r_{t-j}.$$

Notice that

$$12 \cdot \bar{r}_m = \sum_{j=0}^{11} r_{t-j}$$

so that we may alternatively express r_A as

$$r_A = 12 \cdot \bar{r}_m.$$

That is, the continuously compounded annual return is twelve times the average of the continuously compounded monthly returns.




Figure 1.1: Monthly adjusted closing prices on Microsoft and Starbucks stock over the period March, 1993 through March, 2008.

As another example, consider a one-month investment in an asset with continuously compounded return r_t . What is the continuously compounded annual return on this investment? If we assume that we receive the same return $r = r_t$ every month for the year, then the annual continuously compounded return is just 12 times the monthly continuously compounded return:

$$r_A = r_t(12) = 12 \cdot r.$$

1.3 Return Calculations in R

This section discusses the calculation and manipulation of returns in R. We first discuss the use of core R functions to compute returns. We then we discuss some R packages that contain functions for return calculations.

The examples in this section are based on the monthly adjusted closing price data for Microsoft and Starbucks stock over the period March, 1993 through March, 2008, downloaded from finance.yahoo.com, in the files `msftPrices.csv` and `sbuxPrices.csv`, respectively. These prices can be imported into `data.frame` objects in R using the function `read.csv()`:

```
> loadPath = "C:\\Users\\ezivot\\Documents\\FinBook\\Data\\"
```

```
> sbux.df = read.csv(file=paste(loadPath, "sbuxPrices.csv", sep=""),
+                   header=TRUE, stringsAsFactors=FALSE)
> msft.df = read.csv(file=paste(loadPath, "msftPrices.csv", sep=""),
+                   header=TRUE, stringsAsFactors=FALSE)
```

1.3.1 Return Calculations Using Core R Functions

In the `data.frame` objects `sbux.df` and `msft.df`, the first column contains the date information and the second column contains the adjusted closing price data (adjusted for dividend and stock splits):

```
> colnames(sbux.df)
[1] "Date"      "Adj.Close"
> head(sbux.df)
      Date Adj.Close
1 3/31/1993    1.19
2  4/1/1993    1.21
3  5/3/1993    1.50
4  6/1/1993    1.53
5  7/1/1993    1.48
6  8/2/1993    1.52
```

Notice that the dates do not always match up with the last trading day of the month. Figure 1.1 shows the monthly prices of the two stocks created using

```
> plot(msft.df$Adj.Close, type = "l", lty = "solid", lwd = 2,
+      col = "blue", ylab = "return")
> lines(sbux.df$Adj.Close, lty = "dotted", lwd = 2, col = "black")
> legend(x = "topleft", legend=c("MSFT", "SBUX"), lwd = 2,
+      lty = c("solid", "dotted"), col = c("blue", "black"))
```

In Figure 1.1 the x-axis, unfortunately, does not show the dates. We will see later how better time series plots can be created.

To compute simple monthly returns use

```
> n = nrow(sbux.df)
> sbux.ret = sbux.df$Adj.Close[2:n]/sbux.df$Adj.Close[1:n-1] - 1
> msft.ret = msft.df$Adj.Close[2:n]/msft.df$Adj.Close[1:n-1] - 1
> head(cbind(sbux.ret, msft.ret))
```

```

      sbux.ret msft.ret
[1,]  0.01681 -0.07407
[2,]  0.23967  0.08444
[3,]  0.02000 -0.04918
[4,] -0.03268 -0.15948
[5,]  0.02703  0.01538
[6,]  0.12500  0.09596

```

To compute continuously compounded returns use

```

> sbux.ccret = log(1 + sbux.ret)
> msft.ccret = log(1 + msft.ret)

```

or, equivalently, use

```

> sbux.ccret = log(sbx.df$Adj.Close[2:n]/sbx.df$Adj.Close[1:n-1])
> sbux.ccret = log(sbx.df$Adj.Close[2:n]/sbx.df$Adj.Close[1:n-1])
> head(cbind(sbx.ccret, msft.ccret))

```

```

      sbux.ccret msft.ccret
[1,]    0.01667   -0.07696
[2,]    0.21484    0.08107
[3,]    0.01980   -0.05043
[4,]   -0.03323   -0.17374
[5,]    0.02667    0.01527
[6,]    0.11778    0.09163

```

1.3.2 R Packages for Return Calculations

There are several R packages that contain functions for creating, manipulating and plotting returns. An up-to-date list of packages relevant for finance is given in the finance task view on the comprehensive R archive network (CRAN). This section briefly describes some of the functions in the **PerformanceAnalytics**, **quantmod**, **zoo**, and **xts** packages.

The **PerformanceAnalytics** package, written by Brian Peterson and Peter Carl, contains functions for performance and risk analysis of financial portfolios. Table 1.1 summarizes the functions in the package for performing return calculations and for plotting financial data.

The functions in **PerformanceAnalytics** work best with the financial data being represented as zoo objects (see the zoo package for details). A **zoo**

Function	Description
<code>CalculateReturns</code>	Calculate returns from prices
<code>Drawdowns</code>	Find the drawdowns and drawdown levels
<code>maxDrawdown</code>	Calculate the maximum drawdown from peak
<code>Return.annualized</code>	Calculate an annualized return
<code>Return.cummulative</code>	Calculate a compounded cumulative return
<code>Return.excess</code>	Calculate returns in excess of a risk free rate
<code>Return.Geltner</code>	Calculate Geltner liquidity-adjusted return
<code>Return.read</code>	Read returns data with different date formats
<code>chart.CumReturns</code>	Plot cumulative returns over time
<code>chart.Drawdown</code>	Plot drawdowns over time
<code>chartRelativePerformance</code>	Plot relative performance among multiple assets
<code>chart.TimeSeries</code>	Plot time series
<code>charts.PerformanceSummary</code>	Combination of performance charts.

Table 1.1: PerformanceAnalytics return calculation and plotting functions

object is a rectangular data object with an associated time index. The time index can be any ordered object, but it is usually some object that represents dates (e.g., `Date`, `yearmon`, `yearqtr`, `POSIXct`, `timeDate`) To convert the `data.frame` objects `msft.df` and `sbux.df` to `zoo` objects use:

```
> library(zoo)
> dates.sbux = as.yearmon(sbux.df$Date, format="%m/%d/%Y")
> dates.msft = as.yearmon(msft.df$Date, format="%m/%d/%Y")
> sbux.z = zoo(x=sbux.df$Adj.Close, order.by=dates.sbux)
> msft.z = zoo(x=msft.df$Adj.Close, order.by=dates.msft)
> class(sbux.z)
[1] "zoo"
> head(sbux.z)
Mar 1993 Apr 1993 May 1993 Jun 1993 Jul 1993 Aug 1993
      1.19      1.21      1.50      1.53      1.48      1.52
```

For monthly time series the `zoo` "yearmon" class, which represents dates in terms of the month and year, is the most convenient time index. For general daily time series, use the core R "Date" class as the time index.

There are several advantages of using "zoo" objects to represent a time series data. For extracting observations, "zoo" objects can be subsetted using dates and the `window()` function:

```
> sbux.z[as.yearmon(c("Mar 1993", "Mar 1994"))]
Mar 1993 Mar 1994
      1.19      1.52
> window(sbux.z, start=as.yearmon("Mar 1993"), end=as.yearmon("Mar 1994"))
Mar 1993 Apr 1993 May 1993 Jun 1993 Jul 1993 Aug 1993 Sep 1993 Oct 1993
      1.19      1.21      1.50      1.53      1.48      1.52      1.71      1.67
Nov 1993 Dec 1993 Jan 1994 Feb 1994 Mar 1994
      1.39      1.39      1.50      1.45      1.52
```

Two or more "zoo" objects can be merged together and aligned to a common date index:

```
> sbuxMsft.z = merge(sbux.z, msft.z)
> head(sbuxMsft.z)
      sbux.z msft.z
Mar 1993  1.19  2.43
Apr 1993  1.21  2.25
```




Figure 1.2: Time series plot for zoo object.

May 1993	1.50	2.44
Jun 1993	1.53	2.32
Jul 1993	1.48	1.95
Aug 1993	1.52	1.98

As illustrated in Figures 1.2 and 1.3, time series graphs with dates on the axes can be easily created:

```
# two series on same graph
> plot(msft.z, lwd=2, col="blue", ylab="Prices", xlab="Months")
> lines(sbx.z, col="black", lwd=2, lty="dotted")
> legend(x="topleft", legend=c("MSFT", "SBUX"), col=c("blue", "black"),
+       lwd=2, lty=c("solid", "dotted"))
# two series in two separate panels
> plot(sbxMsft.z, lwd=c(2,2), plot.type="multiple",
+     col=c("black", "blue"), lty=c("solid", "dotted"),
+     ylab=c("SBUX", "MSFT"), main="")
```

With prices in a zoo object, simple and continuously compounded returns can be easily computed using the `diff()` and `lag()` functions:




Figure 1.3: Multi-panel time series plot for zoo objects.

```
> sbuxRet.z = diff(sbux.z)/lag(sbux.z, k=-1)
> sbuxRetcc.z = diff(log(sbux.z))
> head(merge(sbuxRet.z, sbuxRetcc.z))
```

		sbuxRet.z	sbuxRetcc.z
Apr	1993	0.01681	0.01667
May	1993	0.23967	0.21484
Jun	1993	0.02000	0.01980
Jul	1993	-0.03268	-0.03323
Aug	1993	0.02703	0.02667
Sep	1993	0.12500	0.11778

Simple and continuously compounded returns can also be computed using the PerformanceAnalytics function `CalculateReturns()`:

```
> sbuxMsftRet.z = CalculateReturns(sbuxMsft.z, method="simple")
> head(sbuxMsftRet.z)
```

		sbux.z	msft.z
Apr	1993	0.01681	-0.07407
May	1993	0.23967	0.08444
Jun	1993	0.02000	-0.04918

```

Jul 1993 -0.03268 -0.15948
Aug 1993  0.02703  0.01538
Sep 1993  0.12500  0.09596
> sbuxMsftRetcc.z = CalculateReturns(sbuxMsft.z, method="compound")
> head(sbuxMsftRetcc.z)
      sbux.z  msft.z
Apr 1993  0.01667 -0.07696
May 1993  0.21484  0.08107
Jun 1993  0.01980 -0.05043
Jul 1993 -0.03323 -0.17374
Aug 1993  0.02667  0.01527
Sep 1993  0.11778  0.09163

```

The **PerformanceAnalytics** package contains several useful charting functions specially designed for financial time series. Fancy time series plots with event labels and shading can be created with `chart.TimeSeries()`:

```

# dates are formatted the same way they appear on the x-axis
> shading.dates = list(c("Jan 98", "Oct 00"))
> label.dates = c("Jan 98", "Oct 00")
> label.values = c("Start of Boom", "End of Boom")
> chart.TimeSeries(msft.z, lwd=2, col="blue", ylab="Price",
+                 main="The rise and fall of Microsoft stock",
+                 period.areas=shading.dates, period.color="yellow",
+                 event.lines=label.dates, event.labels=label.values,
+                 event.color="black")

```

Two or more investments can be compared by showing the growth of \$1 invested over time using the function `chart.CumReturns()`:

```

> chart.CumReturns(sbuxMsftRet.z, lwd=2, main="Growth of $1",
+                 legend.loc="topleft")

```

quantmod

The `quantmod` package for R, written by Jeff Ryan, is designed to assist the quantitative trader in the development, testing, and deployment of statistically based trading models. See www.quantmod.com for more information about the `quantmod` package. Table 1.1 summarizes the functions in the




Figure 1.4: Fancy time series chart created with `chart.TimeSeries()`.

package for retrieving data from the internet, performing return calculations, and plotting financial data.

Several functions in **quantmod** (i.e., those starting with `get`) automatically download specified data from the internet and import this data into R objects (typically "xts" objects). For example, to download from `finance.yahoo.com` all of the available daily data on Yahoo! stock (ticker symbol YHOO) and create the "xts" object YHOO use the `getSymbols()` function:

```
> library(quantmod)
> getSymbols("YHOO")
[1] "YHOO"
> class(YHOO)
[1] "xts" "zoo"
> colnames(YHOO)
[1] "YHOO.Open"      "YHOO.High"      "YHOO.Low"       "YHOO.Close"
[5] "YHOO.Volume"   "YHOO.Adjusted"
> start(YHOO)
[1] "2007-01-03"
> end(YHOO)
```




Figure 1.5: Growth of \$1 invested in Microsoft and Starbucks stock.

```
[1] "2009-10-02"
```

```
> head(YH00)
```

	YH00.Open	YH00.High	YH00.Low	YH00.Close	YH00.Volume
2007-01-03	25.85	26.26	25.26	25.61	26352700
2007-01-04	25.64	26.92	25.52	26.85	32512200
2007-01-05	26.70	27.87	26.66	27.74	64264600
2007-01-08	27.70	28.04	27.43	27.92	25713700
2007-01-09	28.00	28.05	27.41	27.58	25621500
2007-01-10	27.48	28.92	27.44	28.70	40240000

```
YH00.Adjusted
```

2007-01-03	25.61
2007-01-04	26.85
2007-01-05	27.74
2007-01-08	27.92
2007-01-09	27.58
2007-01-10	28.70

To create the daily price-volume chart of the YH00 data shown in Figure 1.6, use the `chartSeries()` function:

```
> chartSeries(YH00,theme=chartTheme('white'))
```

Function	Description
<code>chartSeries</code>	Create financial charts
<code>getDividends</code>	Download dividend data from Yahoo!
<code>getFX</code>	Download exchange rates from Oanda
<code>getFinancials</code>	Download financial statements from google
<code>getMetals</code>	Download metals prices from Oanda
<code>getQuote</code>	Download quotes from various sources
<code>getSymbols</code>	Download data from various sources
<code>periodReturn</code>	Calculate returns from prices

Table 1.2: quantmod return calculation and plotting functions

1.4 Notes and Complements

This chapter describes basic asset return calculations with an emphasis on equity calculations. Campbell, Lo and MacKinlay (1997) provide a nice treatment of continuously compounded returns. A useful summary of a broad range of return calculations is given in Watsham and Parramore (1998). A comprehensive treatment of fixed income return calculations is given in Stigum (1981), and the official source of fixed income calculations is the so-called “The Pink Book”.



Figure 1.6: Daily price-volume plot for Yahoo! stock created with the `quantmod` function `chartSeries()`.

1.5 Appendix: Properties of exponentials and logarithms

The computation of continuously compounded returns requires the use of natural logarithms. The natural logarithm function, $\ln(\cdot)$, is the inverse of the exponential function, $e^{(\cdot)} = \exp(\cdot)$, where $e^1 = 2.718$. That is, $\ln(x)$ is defined such that $x = \ln(e^x)$. Figure 1.7 plots e^x and $\ln(x)$. Notice that e^x is always positive and increasing in x . $\ln(x)$ is monotonically increasing in x and is only defined for $x > 0$. Also note that $\ln(1) = 0$ and $\ln(-\infty) = 0$.

The exponential and natural logarithm functions have the following properties

1. $\ln(x \cdot y) = \ln(x) + \ln(y)$, $x, y > 0$
2. $\ln(x/y) = \ln(x) - \ln(y)$, $x, y > 0$
3. $\ln(x^y) = y \ln(x)$, $x > 0$
4. $\frac{d \ln(x)}{dx} = \frac{1}{x}$, $x > 0$
5. $\frac{d}{dx} \ln(f(x)) = \frac{1}{f(x)} \frac{d}{dx} f(x)$ (chain-rule)




Figure 1.7: Exponential and natural logarithm functions.

6. $e^x e^y = e^{x+y}$

7. $e^x e^{-y} = e^{x-y}$

8. $(e^x)^y = e^{xy}$

9. $e^{\ln(x)} = x$

10. $\frac{d}{dx} e^x = e^x$

11. $\frac{d}{dx} e^{f(x)} = e^{f(x)} \frac{d}{dx} f(x)$ (chain-rule)

Bibliography

- [1] Campbell, J., A. Lo, and C. MacKinlay (1997), *The Econometrics of Financial Markets*, Princeton University Press.
- [2] *Handbook of U.W. Government and Federal Agency Securities and Related Money Market Instruments*, “The Pink Book”, 34th ed. (1990), The First Boston Corporation, Boston, MA.
- [3] Stigum, M. (1981), *Money Market Calculations: Yields, Break Evens and Arbitrage*, Dow Jones Irwin.
- [4] Watsham, T.J. and Parramore, K. (1998), *Quantitative Methods in Finance*, International Thomson Business Press, London, UK.
- [5] Wüertz, D., Chalabi, Y., Chen, W., and Ellis, A. (2009). *Portfolio Optimization with R/Rmetrics*, eBook, Rmetrics Association & Finance Online.