CHAPTER 1

Discount Rates and Returns

The most powerful force in the universe is compound interest.
—Albert Einstein

ESTIMATING RETURNS

The total return on an investment in any security is the percentage change in the value of the asset including dividends over a specific interval of time. Assuming asset value is captured in the market price P and dividends by d, then the one period *total return*, r_1, is equal to $100 * \frac{(P_1 + d_1 - P_0)}{P_0}$ percent, in which the subscripts index time. For simplicity, we will ignore dividends, which gives us the price return, which, in decimal form, is equal to $r_1 = \frac{(P_1 - P_0)}{P_0} = \left(\frac{P_1}{P_0}\right) - 1$. Thus, $1 + r_1 = \left(\frac{P_1}{P_0}\right)$ is the *gross return* (the return *plus* the initial one-dollar outlay in the security) on the investment for one period, and r_1 is the *net return;* it is the return on a \$1 investment. We can geometrically link returns to get the time equivalent of a longer-term investment. For example, suppose that the period under study is one month and that $1 + r_1 = \left(\frac{P_1}{P_0}\right)$ is therefore the one-month return. We can annualize this return by assuming the investment returns this amount in each month. Compounding this for one year is a product yielding the amount:

$$r_A = (1 + r_1)^{12} - 1$$

Here, r_1 is the monthly return, while r_A is the annualized equivalent. On the other hand, we may observe a time series of past monthly returns (called *trailing returns*), which we geometrically link to estimate an annualized figure, that is, the previous 12 monthly returns generate an annual return given by:

$$1 + r_A = (1 + r_1)(1 + r_2)(1 + r_3)\cdots(1 + r_{11})(1 + r_{12})$$

We can similarly link quarterly returns to estimate an annual equivalent, for example, $1 + r_A = (1 + r_1)(1 + r_2)(1 + r_3)(1 + r_4)$, and we can do the same for weekly, daily, or any frequency for that matter, to achieve a lower frequency equivalent return. Focusing once again on monthly gross returns, annualization is a compounded return that is the product of monthly relative prices, each measuring price appreciation from the previous month, that is, by generalizing from the fact that if $1 + r_1 = \left(\frac{P_1}{P_0}\right)$, then the following must also be true:

$$1 + r_A = \left(\frac{P_1}{P_0}\right)\left(\frac{P_2}{P_1}\right)\left(\frac{P_3}{P_2}\right)\cdots\left(\frac{P_{11}}{P_{10}}\right)\left(\frac{P_{12}}{P_{11}}\right)$$

Upon canceling, this reduces to the following, which is consistent with our definition of gross return given earlier.

$$1 + r_A = \frac{P_{12}}{P_0}$$

This suggests that we can calculate the gross return over any period by taking the ratio of the market values and ignoring all intermediate market values. Similarly, we can solve for any intervening periodic average return by using the power rule; in this case, if the annual return is $1 + r_A$, then the *geometric average* monthly return r_M must be

$$(1 + r_A)^{\frac{1}{12}} = (1 + r_M)$$

For example, the average monthly return necessary to compound to a 15 percent annual return must be approximately 1.17 percent per month:

$$(1.15)^{\frac{1}{12}} = (1.011715)$$

EXAMPLE 1.1

Quarterly returns to the Russell 3000 Domestic Equity Index for the years 2005 to 2007 are given in Table 1.1.

TABLE 1.1 Russell 3000 Dom Eq Index

Date	Return (%)
2005Q1	–2.20
2005Q2	2.24
2005Q3	4.01
2005Q4	2.04
2006Q1	5.31
2006Q2	–1.98
2006Q3	4.64
2006Q4	7.12
2007Q1	1.28
2007Q2	5.77
2007Q3	1.55
2007Q4	–3.34

a. Geometrically link these quarterly returns to generate annual returns.

b. Calculate the return for the three-year period.

c. What is the *arithmetic average* annual return for these three years?

d. What is the *arithmetic average* quarterly return over this three-year period?

SOLUTIONS

(Refer to Table 1.2)

a. Annual Return (Column E) $= \prod_{n=0}^{n+3} (Column\ C) - 1$

b. The return for the three-year period $= \prod [Column\ C] - 1 = 0.29 =$ 29%; annualized, this is $\prod [Column\ C]^{\frac{1}{3}} - 1 = 1.088896 - 1 =$ 8.9% geometric average.

c. The arithmetic average annual return $= AVG(Column\ E) = 9.0\%$.

d. The arithmetic average quarterly return $= AVG[Column\ B] = 2.2\%$.

(*continued*)

(continued)

TABLE 1.2 Geometric Returns

Quarter	Return (%)	Return/ 100 + 1	Geometric Mean	Annualized Return (%)
(A)	(B)	(C)	(D)	(E)
2005Q1	–2.20	0.98		
2005Q2	2.24	1.02		
2005Q3	4.01	1.04		
2005Q4	2.04	1.02	1.06	6.10
2006Q1	5.31	1.05		
2006Q2	–1.98	0.98		
2006Q3	4.64	1.05		
2006Q4	7.12	1.07	1.16	15.70
2007Q1	1.28	1.01		
2007Q2	5.77	1.06		
2007Q3	1.55	1.02		
2007Q4	–3.34	0.97	1.05	5.10

Go to the companion website for more details.

GEOMETRIC AND ARITHMETIC AVERAGES

If things weren't already complicated enough, we now see that there are two distinct averages—geometric as well as arithmetic. It is important to understand the difference. If you want to know what an asset actually returned, then geometrically link the N gross returns over the relevant time. And, upon doing that, if you then want to know what the average return (geometric) was for each period in the return series, then take the Nth root and subtract one. Using a trailing series of the past 12 monthly returns as an example, we get:

$$1 + r_A = (1 + r_1)(1 + r_2)(1 + r_3) \cdots (1 + r_{11})(1 + r_{12})$$

The annual return is r_A. The *geometric average* of the monthly returns is

$$(1 + r_A)^{\frac{1}{12}} = [(1 + r_1)(1 + r_2)(1 + r_3) \cdots (1 + r_{11})(1 + r_{12})]^{\frac{1}{12}}$$

Clearly, this is different from the *arithmetic average:*

$$\bar{r} = \frac{r_1 + r_2 + r_3 + \cdots + r_{12}}{12}$$

The difference is not subtle. For example, suppose we observe a sequence of four returns {0.9, 0.1, –0.9, 0.2}. The arithmetic average is 0.075 (7.5 percent), while the geometric average is –29 percent! Why the large discrepancy? If you had a dollar invested over these four periods, the return you would have received would have been affected to a greater degree (in a negative way) by the third period's negative 90 percent return, that is, you would have lost 90 percent of your accumulated investment by the end of the third period and then earned a 20 percent return on whatever was left for the final period. The arithmetic average, however, places equal weight on all returns and, therefore, the impact of the large negative return is diluted by $1/N$. As the sample size increases, the impact of a single bad return declines asymptotically and it does not matter if that single bad return occurred early or late in the sample. In reality, that is not how money is earned and that is why we use geometric averages. In this example, the investment indeed earned an average –29 percent return in each period. Had you invested a dollar at the beginning of the first period, that dollar would have shrunk to about $0.25 in four periods. This is certainly not an amount implied by the arithmetic mean.

We will not prove the following formally, but it is intuitive that, in general, as the variance in the individual periodic returns declines, so does the difference between the arithmetic and geometric means. In the limit, if the four returns in our example were identical, then the arithmetic and geometric means would also be identical. Otherwise, it can be shown that the arithmetic mean is always greater than the geometric mean because the arithmetic mean ignores the correlations across returns over time. The takeaway is that these two measures tend to diverge in value as volatility in returns rises.

CAVEATS TO RETURN EXTRAPOLATION

Practitioners prefer to compare annualized returns and therefore extrapolate higher frequency returns (daily, monthly, and quarterly) to annual frequency. This practice, though common, can be misleading and it is important to know why. Consider the set of monthly returns to the S&P 500 index given in Table 1.3. The returns are for the year 2006.

TABLE 1.3 Monthly S&P 500 Returns

Month	Monthly Return	Gross Monthly Return	Quarterly Return
(A)	(B)	(C = B + 1)	(D)
Jan-06	0.0132	1.0132	
Feb-06	–0.0016	0.9984	
Mar-06	0.0134	1.0134	1.0251
Apr-06	0.0065	1.0065	
May-06	–0.0093	0.9907	
Jun-06	–0.0286	0.9714	0.9686
Jul-06	0.0056	1.0056	
Aug-06	0.0214	1.0214	
Sep-06	0.0238	1.0238	1.0515
Oct-06	0.0346	1.0346	
Nov-06	0.0185	1.0185	
Dec-06	0.0200	1.0200	1.0749

We can geometrically link these monthly returns to get an annualized return, which is computed by taking the product of the gross monthly returns in column C and subtracting one, yielding 12.23 percent. This is the return that you would have received had you held the index for those 12 months. The table also geometrically links monthly returns to estimate quarterly returns, which are given in column D of the table. These quarterly returns are then geometrically linked by taking their product (and subtracting one) to get an annual return, also equal to 12.23 percent.

Now imagine that it is April 1, 2006, and having just observed the March return, we estimate the first-quarter return for 2006 at 2.51 percent. Your supervisor wants to know what this is on an annualized basis. In response, you compute $\left(1 + r_{Q1}\right)^4 - 1$, which upon substitution, computes to $(1.0251)^4 - 1 = 10.42$ percent. What you have done is extrapolate a higher frequency return (quarterly) to a lower frequency estimate (annual). The implicit assumption in extrapolation is that the return observed for the period just realized (first quarter 2006) will hold for the remaining three quarters. (This is what is referred to as a *naïve forecast*). In general, this will not be the case and the extrapolated return will therefore most likely contain errors. In the example given, you can readily see the error embodied in each of these quarterly extrapolations by comparing them to the true observed annual return in column C.

It is also important to realize that extrapolation generates more measurement error the greater the difference in the frequencies we extrapolate between. For example, if we extrapolate the monthly returns, that is,

$(1+r_M)^{12}-1$, then clearly, a single extraordinary monthly return will translate into an even more extraordinary annual return. Practitioners seem to know this and that is why higher frequency returns (weekly or daily) are generally never annualized. Returns are random. Extrapolation is not. The greater the difference between the frequencies we extrapolate between, the less we believe in randomness and the more we believe that the current observation portends all future observations. We are all most certainly aware of this problem but, nevertheless, we continue to extrapolate. It is important to remind ourselves and our colleagues of the weaknesses in these numbers.

DISCOUNTING PRESENT VALUES OF CASH FLOW STREAMS

As individual consumers, we are always trying to maximize our intertemporal utilities by trading off future and present consumption. That is, we will consume a dollar's worth of goods today if we feel that the satisfaction we receive from doing so exceeds the satisfaction we'd get had we saved that dollar and consumed it somewhere in the future. The decision to consume intertemporally therefore depends on our abilities to compare wealth today with future wealth, which is what we mean when we talk about the *time value of money*. A dollar cash amount invested in the future will be worth $C=(1+r)$ after, say, one year. Therefore, the present value P of a cash flow C to be received one period from now is the future C discounted at rate r:

$$P=\frac{C}{(1+r)}$$

Alternatively, investing P for one period at rate r will generate value equal to $P(1+r)=C$. The present value of a cash flow received two time periods from now is therefore:

$$P=\frac{C}{(1+r)^2}$$

If the cash flow is received more than once (say, three periods), then it has present value:

$$P=\frac{C}{(1+r)}+\frac{C}{(1+r)^2}+\frac{C}{(1+r)^3}$$

An example of discrete discounting is *net present value* (NPV), which is present value minus the initial outlay. The NPV function in Excel is: = NPV(rate, cash flow 1, cash flow 2, . . .). Suppose, for instance, that you were to undertake an investment that requires an initial cash outlay of \$100 but will return dividends over the next two years in the amount of \$100 per year with certainty. Suppose the opportunity cost of capital is 10 percent (this is the return you could have received had you invested your \$100 in the market instead). Then the net present value is:

$$NPV = -100 + \frac{100}{1.1} + \frac{100}{(1.1)^2} = -100 + 90.91 + 82.64 = \$73.55$$

What exactly is the *discount rate?* It is the rate at which we are willing to trade present for future consumption. For example, suppose you are waiting to receive $C = \$100$ one year from now. Rather than wait, you agree to receive a smaller amount $P = \$90$ now. The smaller amount is consistent with consumers' preference for present versus future consumption; it suggests we are impatient, that we discount future gains (for a whole host of reasons) or more specifically, that we have our own set of time preferences that determine our individual decisions to consume our wealth intertemporally. We examine derivation of the discount rate in more rigorous detail in Chapter 4. The implication in this example is that the interest rate that is consistent with your time preferences is 11 percent and that your discount rate is ${}^{1}/_{1.11} = 0.9$, which is your willingness to trade the future \$100 for current consumption worth \$90. That is, you discount the future at 10 percent. The converse argument is that you would be willing to give up \$90 today only if you knew you'd receive in exchange an amount of \$100 one year from now.

Discount rates and market returns are obviously linked. Returns are determined by changes in the market prices of assets that more fundamentally reflect market participants' utility preferences that manifest themselves through the interaction of supply and demand. In this sense, returns can be thought of as an aggregate of all of our revealed preferences, that is, our attitudes regarding present over future consumption. We discount cash flows using observed market rates and we use different market rates to discount different types of cash flows, for example, risk-free Treasury rates to discount riskless bond coupons and more risky equity returns to discount private equity cash flows. For now, we will abstract from these details and study only the process of discounting. Generalizing the cash flow discounting problem, then, to t periods, we get a sum of periodic discounted cash flows:

$$P = \sum \left[\frac{C_t}{(1+r)^t} \right]$$

Note how we index C by time. Suppose, now, that the rate r is *compounded discretely* m times per year (that is, you receive a fraction $\frac{r}{m}$ percent, m times each year). Instead of t periods, we now have $m * t$ periods to discount, each at $\frac{r}{m}$ percent. Thus,

$$P = \sum \left[\frac{C_t}{\left(1 + \frac{r}{m}\right)^{mt}} \right]$$

The quantity $\left(1 + \frac{r}{m}\right)^m$ has a limit as m goes to infinity, that is, as interest is paid continuously. This limit is very important. It is

$$e^r = \lim_{m \to \infty} \left(1 + \frac{r}{m}\right)^m$$

More fundamentally, recall that

$$e = \lim_{m \to \infty} \left(1 + \frac{1}{m}\right)^m$$

Therefore,

$$e^{-\mathrm{rt}} = \lim_{m \to \infty} \frac{1}{\left(1 + \frac{r}{m}\right)^{mt}}$$

Continuous compounding is therefore the limit of discrete compounding. For example, semiannual compounding (that is, interest paid twice each year) is

$$\left(1 + \frac{r}{2}\right)^2 = \left(1 + \frac{r}{2}\right)\left(1 + \frac{r}{2}\right)$$

Compounding interest quarterly over the year,

$$\left(1 + \frac{r}{4}\right)^4 = \left(1 + \frac{r}{4}\right)\left(1 + \frac{r}{4}\right)\left(1 + \frac{r}{4}\right)\left(1 + \frac{r}{4}\right)$$

Finally, paying interest monthly over the year is equal to the already familiar geometric return, or annual equivalent, equal to $\left(1 + \frac{r}{12}\right)^{12}$.

Taking the compounding frequency to the limit results in continuous compounding, e^r. This means that an amount C received at the end of time t with continuous compounding has a present value:

$$P = Ce^{-rt}$$

The discount rate in this case is e^{-rt}. Let me digress a bit on this concept. Assume I have capital to be invested right now in the amount P_0, and at the end of one period, it grows to P_1. Thus, $P_1 = P_0 e^{rt}$. Taking natural logs and noting that $t = 1$ for this example gives us the following:

$$\ln(P_1) = \ln(P_0) + r$$

Equivalently,

$$\ln\left(\frac{P_1}{P_0}\right) = r$$

where r in this case is the rate of return on the investment. It also determines the discount rate in the sense that it represents the opportunity cost of investment, that is, r is what I give up if I choose to consume P_0 today. Had I invested (saved) it, it would have grown to P_1 in one period.

In general then, an equivalent amount C earns the following over time t with continuous compounding:

$$C = Pe^{rt}$$

Thus, the process of evaluating future obligations as a present value problem is referred to as *discounting*. The present value of the future monetary amount (C) to be received is less than the face value of that amount because the future is discounted, reflecting, among other things, time preferences (a dollar today is worth more to me than a dollar to be received sometime in the future).

I define the k-period discount rate in discrete time as

$$d_k = \left(1 + \frac{r}{m}\right)^{-mk}$$

In continuous time, we have

$$d_k = e^{-rk}$$

It should be obvious that future value is the inverse function of present value. For example, let $(C_1, C_2, \ldots C_n)$ refer to a cash flow stream. Assume each cash flow is received at the beginning of the period and that the interest rate is constant at r. Then the *future value* (FV) is the sum of the compounded cash flow values:

$$FV = C_1(1 + r)^n + C_2(1 + r)^{n-1} + \cdots + C_n$$

Likewise, it should be clear that FV has present value P equal to $\frac{FV}{(1+r)^n}$, that is,

$$P = C_1 + \frac{C_2}{1+r} + \cdots + \frac{C_n}{(1+r)^n}$$

And, therefore, with compounding m times per period:

$$P = \sum_{k=0}^{n} \left[\frac{C_k}{\left(1+\frac{r}{m}\right)^{mk}} \right]$$

This relationship can be written more compactly with continuous compounding:

$$P = \sum C_k e^{-rk}$$

Although we develop this concept more fully in Chapter 2, this is our first pricing model. It is a simple discounted cash flow model with certain (riskless) cash flows.

INTERNAL RATE OF RETURN AND YIELD TO MATURITY

We now assume that the discount rate is endogenous, in which case we solve for the rate that equates two sets of cash flows. Suppose you make an investment in a business equal to P_0 dollars. This investment is expected to yield a stream of cash flows for n periods equal to $C_1, C_2, \ldots C_n$. The internal rate of return (IRR) is the discount rate, which makes the two streams, the outflow P_0 and the present value of the inflows $C_1, C_2, \ldots C_n$, equivalent. That is, the IRR is the value of r that discounts the following set of cash flows to the initial outlay P_0:

$$P_0 = \frac{C_1}{(1+r)} + \frac{C_2}{(1+r)^2} + \cdots + \frac{C_n}{(1+r)^n}$$

Since P_0 is an outlay, hence, a negative cash flow, then the preceding equation is the same as:

$$0 = -P_0 + \frac{C_1}{(1+r)} + \frac{C_2}{(1+r)^2} + \cdots + \frac{C_n}{(1+r)^n}$$

Notice that this is a net present value, too, but with the important exception that the discount rate is exogenous in NPV problems. The internal rate of return, on the other hand, is the solution objective. The internal rate of return gets its name from understanding that it is the interest rate implied by the internal cash flow stream to the firm. In a sense, it is the firm's required rate of return necessary to achieve a breakeven level on its investments. It is therefore not a market rate. Firms use the IRR most often to compare alternative investments.

In general, the IRR is difficult to solve because it doesn't have an analytic solution; rather, one must resort to iterative techniques to arrive at a solution. Most software packages' solvers use some form of Newton's method to solve for the IRR (see the IRR and XIRR functions in Microsoft Excel).

EXAMPLE 1.2

Let's put numbers in Example 1.1. Assume the initial outlay is \$100 and we expect to receive cash flows in years 1 to 4 equal to (\$50, \$0, \$100, \$100). Then the IRR is the rate that solves:

$$0 = -100 + \frac{50}{(1+r)} + \frac{0}{(1+r)^2} + \frac{100}{(1+r)^3} + \frac{100}{(1+r)^4}$$

This is equal to:

$$0 = -100 + 50C + 100C^3 + 100C^4$$

where $C = \frac{1}{(1+r)}$. This is a polynomial of order four. The solution is r = IRR = 39 percent, whose details can be found on the chapter spreadsheet.

What does this mean? Well, suppose again that this is your firm. Then this rate discounts your cash flows to a present value equal to your outlay of \$100. This is a pretty good rate of return if all other investments generate cash flows with IRRs less than 39 percent. Thinking differently, if your firm requires an annual rate of return over four years on their cash flows equal to 39 percent, then a \$100 investment with the stated cash flows will meet that requirement.

EXAMPLE 1.3

Suppose the required return is 39 percent on a \$100, four-year investment with expected cash flows in each of the four years given by (\$0, \$0, \$100, \$150). What is the estimated IRR (r in the denominator of the following equation) and will this investment be undertaken?

$$0 = -100 + \frac{0}{(1+r)} + \frac{0}{(1+r)^2} + \frac{100}{(1+r)^3} + \frac{150}{(1+r)^4}$$

Example 1.3 Table

Initial Outlay	–100
Year 1 Cash Flow	0
Year 2 Cash Flow	0
Year 3 Cash Flow	100
Year 4 Cash Flow	150
Required Rate of Return (RRR)	**39%**
Internal Rate of Return	**29%**
RRR = IRR	**FALSE**
Check the IRR Calculation	
Initial Outlay	–100
C =	0
C^2 =	0
C^3 =	46.289
C^4 =	53.711
SUM =	**0**

In this example, the IRR = 29 percent, which is below the required return of 39 percent. Therefore, the investment should not be undertaken.

Now let's briefly jump ahead and look at the similarity between the IRR and what bond analysts call the *yield to maturity*. Suppose you lend \$100 for a period of five years. The borrower promises to pay you \$25 in each of those five years to expunge his debt. What is the rate of return that equates the present value of the creditor's payments to the loan amount? We set this up as:

$$100 = \frac{25}{(1+r)} + \frac{25}{(1+r)^2} + \frac{25}{(1+r)^3} + \frac{25}{(1+r)^4} + \frac{25}{(1+r)^5}$$

(*continued*)

(*continued*)

The IRR that solves this problem is 8 percent. As a lender, you therefore receive an 8 percent annual return on your investment—in this case, a loan. This is essentially a bond, and in the world of bonds, the 8 percent is the yield to maturity. The yield to maturity is a return that is equal to an IRR. Thus, bond yields are IRRs.

Go to the companion website for more details.

REAL AND NOMINAL RETURNS

Inflation erodes the real returns on cash flows. For example, if the inflation rate is 4 percent, then a nominal cash flow of \$1 has a real inflation-adjusted value of $\frac{\$1.00}{\$1.04} = \$0.96$. Using our discounting rules already developed, we can generalize that a nominal *gross* return $(1 + r)$ has an equivalent real *gross* return of:

$$1 + r_0 = \frac{1 + r}{1 + f}$$

where f is the inflation rate (in our example, 4 percent) and r_0 is the real *net* return.

Note that if the inflation rate is $f = 0$, then the nominal return is identical to the real return. Also, note that we can simplify this equation for the real gross return to get at the net return as follows:

$$r_0 = \frac{(r - f)}{(1 + f)}$$

The important point is that inflation affects the relevant discount rate that one uses to value a cash flow stream. We return to this topic when we look at Treasury Inflation Protected Securities, or TIPS, in the next chapter.

SUMMARY

Returns measure growth rates in asset value over time. They are the observed reward to postponing present for future consumption. Returns may be measured discretely over any frequency such as daily, monthly, or quarterly and higher frequency returns can be converted into lower frequencies—annualizing monthly returns is one such case. The notion of compounding is linked to how often interest is paid; thus, annual returns can represent a single yearly payment, or higher frequency returns can be geometrically linked, or compounded, to form annualized equivalents. In the limit, continuously compounded returns represent the continuous payment of interest. In sum, we can work with returns over any interval and extrapolate those returns to either any longer interval of time or average them to any subinterval of time. An example of extrapolation is annualizing discrete monthly returns, and an example of averaging is finding the geometric average monthly return from an annual return. Caveats relate to the implicit assumption that observed returns will hold into the future.

Discounting links cash flows over time. The discounted present value of a cash flow to be received in the future is the result of finding the amount of cash in present dollars that, when invested at the discount rate, will grow to an amount stipulated by the future cash flow. Discount rates are intimately linked to returns; in the simplest case, the discount rate is the reciprocal of the gross return and the discount rate may be applied discretely or continuously. The role that discount rates play in the trade-off of present over future consumption is a topic that I develop more fully in Chapter 4.

The internal rate of return is an application of discounting in cash management and the yield-to-maturity on a coupon-paying bond is itself an internal rate of return. We can thus link the subject of Chapter 2 on bond pricing to the discount function developed in this chapter. In fact, almost all asset-pricing models will rely on some form of discounting since they all involve the valuation of cash flows that occur over time.

Finally, we recognize the impact that inflation has on the value of cash flows, which requires us to distinguish between inflation adjusted (real) returns and nominal returns and model them accordingly.