Notes on Present Value, Discounting, and Financial Transactions

Professor John Yinger
The Maxwell School
Syracuse University
Version 2.0

Introduction

These notes introduce the concepts of present value and discounting and explain how these concepts can be applied to mortgages, annuities and pensions, and bonds. The concepts are illustrated with numerous examples. The concepts of inflation, uncertainty, and risk are briefly considered at the end.

Compound Interest

Suppose I put \$100 in a savings account on January 1 and it earns interest at a 3% rate per year. For now, suppose that this interest is paid all at once at the end of the year. Then if I do not withdraw any money, at the end of the year my account will contain (\$100 + .03×\$100) = \$100×(1.03) = \$103. (The symbol "×" indicates multiplication.)

Now consider what happens if another year is added to the analysis. At a 3% interest rate, \$100 in my account turns into \$103 at the end of the first year. This \$103 turn into \$100×(1.03) ×(1.03) = \$100×$(1.03)^2$ = \$100×(1.0609) = \$106.09 at the end of the second year. This bank account involves a concept known as ***compound interest***. In the second year, interest is paid not only on the original \$100, but also on the \$3 of interest received in the second year—a concept known as compounding.

These concepts can be expressed in a general way. Let V_0 be the initial value of an investment (such as my savings account) and let i be the rate of return on this investment (such as the 3% interest on my savings account). Then the value of the account at time y is:

$$V_y = V_0 \times (1+i)^y \qquad (1)$$

Note that y and i both refer to the interval that determines compounding. If the return is compounded once a year, then y indicates years and i indicates the annual interest rate. If the return is compounded every month, then y stands for month and i is the monthly interest rate, which is the annual interest rate divided by 12.

The frequency of compounding makes a small difference, which grows over time. Consider \$100 invested in two different savings accounts, both of which have a 6% annual interest rate but one of which compounds once a year and the other compounds monthly. At the end of the first year, the first account contains \$100×(1.06) = \$106, whereas the second contains \$100×$(1+.06/12)^{12}$ = \$106.17. (The "/" indicates "divided by.") A third account with a 6 percent rate and daily compounding would contain $\$100\times(1+.06/365)^{365} = \$106.18$. If these three accounts are held for 5 years, their values are \$133.82, \$134.89, and \$134.98, respectively.

Present Value and Discounting

Now suppose someone asks me what I would be willing to pay today to receive \$103 at the end of the year. This question defines the concept of ***present value***, which is what an amount received in the future is worth to me today. The answer is contained in the above expression; if \$100 received today turns into \$103 in a year, then the receipt of \$103 a year from now is worth \$103/1.03 = \$100 today. In this example, in other words, the present value of \$103 received a year from now is \$100.

We can also ask about the present value of funds received at the end of two years. The present value of \$106.09 received in two years is $\$106.09/[(1.03)^2] = \$100$. With an interest rate of 3%, in other words, I would be equally happy with \$100 received now, \$103 received a year from now, or \$106.09 received two years from now.

To write down a general form for these concepts, let *FV* stand for future value, *PV* stand for present value. Then we can write down an alternative version of equation (1) and use it to define present value:

$$FV_y = PV \times (1+i)^y \tag{2}$$

and

$$PV = \frac{FV_y}{(1+i)^y} \tag{3}$$

With equation (3), we can translate the value of our investment in time period *y* into a value today, that is, into a present value. As in the case of equation (1), these results apply to any time period for compounding, so long as *y* and *i* are defined by the same unit of time.

Wait a minute, you might say, what if I do not want to wait for two years to get my money? In the above example, how can \$106.09 received in two years possibly be as valuable to me as \$100 received today? The answer lies in the interest rate, which actually has two meanings, not just one. The first meaning is the return on a particular account. In this example, I receive 3% interest on my savings account. The second meaning is the ***opportunity cost*** of tieing up my money in one activity or investment instead of another.

Suppose I could also invest my money in a stock market account that I expect will yield a return of 6 percent per year. (The notion of risk is ignored here, but considered in a later section.) Then the opportunity cost of placing my money in my savings account is 6 percent—the return I give up—not 3 percent, the return I receive in that account. With this alternative investment possibility, \$106.09 received in two years is not worth \$100 to me today. Instead, it is only worth $\$106.09/[(1.06)^2] = \$94.42$.

Thus, it is appropriate to use equation (1) or (2) to calculate the future values of an investment that yields return *i*, but to use equation (3) to calculate present value based on the return on one's best investment, say *r*, which may not be the same as *i*. In symbols,

$$PV = \frac{FV_y}{(1+r)^y} \tag{4}$$

For example, with an opportunity cost rate of *r*, the present value of V_0 invested at *i* percent for *y* time periods is:

$$PV = \frac{V_0 \times (1+i)^y}{(1+r)^y} \tag{5}$$

Equation (4) is enormously useful because it allow us to observe returns received in different years in the future and translate them into a common metric, namely present value. For example, we can compare on investment with returns at annual rate of 6.25% for 3 years with another investment that has returns at an annual rate of 6% for 4 years. If the opportunity cost rate, *r*, equals 5%, then the PV per dollar invested equals $(1.0625)^3/[(1.05)^3]$ =1.0361 for the first investment and $(1.06)^4/[(1.05)^4]$ = 1.0386 for the second. Even though the second investment has a smaller annual return, the fact that it produces an above-opportunity-cost return for one more year makes it the better investment of the two.

Finally, it is possible to define an individual's opportunity cost rate, also called her ***discount rate***, in an even more general way as the value an individual places on future consumption relative to current consumption. A person with a high discount rate values current consumption far more than future consumption, whereas a person with a low discount rate receives roughly equal satisfaction from consumption today and consumption in the future. A high discount rate, that is, a high value for *r* in equation (5), leads people to reject investments unless they have a high rate of return, but many investments would be worthwhile for a person whose discount rate is low.

Mortgages

For the purposes of these notes, a mortgage is a contract in which a lender writes a large check to a homeowner in exchange for a promise from the homeowner to make monthly payments for a certain period of time. These payments include interest at a certain rate on the outstanding balance of the loan. In practice, of course, mortgages come in many forms and are used for many purposes. Some mortgages have variable interest rates, for example, homeowners often take out mortgages to fund home improvements, and businesses often take out mortgages to purchase equipment or real estate. Moreover, when mortgage interest rates decline, many homeowners take out a new mortgage (at the new low rates) to pay off their old mortgage (which has a high rate). This is called refinancing. I am not going to consider these possibilities here. Instead, I will analyze a standard, fixed-rate, home-purchase mortgage.

It is also worth pointing out that once a mortgage has been issued, it is an asset that can be bought and sold. A lender who issues a mortgage may sell it to another lender or to an institution that specializes in buying and repackaging mortgages for investors. Selling a mortgage is equivalent to selling the right to receive the monthly payments that are associated with it. These issues are beyond the scope of these notes, which simply explain the logic of a standard mortgage.

To analyze a mortgage, we need to define some terms.

M_0 is the amount of the mortgage, that is, the size of the check the lender writes to the homeowner;

m is the contractual monthly interest rate, which is the annual rate divided by 12;

P is the size of the monthly payment; and

Y is the term of the loan, that is, the number of months it takes to pay it off.

Now the key to analyzing a mortgage is to recognize that the amount that is still owed on the mortgage gradually declines over time, until it reaches zero at the end of year Y. The amount of the loan can be thought of as the amount that is owed at the initiation of the mortgage, which is why we call this amount M_0. Now let M_y be the amount that is still owed at the end of month y, where y goes from 1 to Y. For a 30-year mortgage, for example, y goes from 1 to 12×30 = 360. To find a formula for M_y, we have to recognize (a) that the borrower owes interest at the end of each period based on the outstanding balance at the beginning of the period and (b) that the borrower makes a payment equal to P.

With these points in mind, we can write:

$$M_1 = M_0(1+m) - P \tag{6}$$

The first term on the right side indicates the initial balance and the interest owed on it and the second term indicates the borrower's payment. (To keep the equation simple, the × sign indicating that M_0 is multiplied by (1+m) is implicit in (6) and in the equations that follow.) Part of the payment covers the interest that is due at the end of the first period (mM_0) and the rest goes to reducing the loan balance. After the first payment, therefore, the balance declines by (P - mM_0).

Following the same logic, we can write down an expression for the outstanding balance at the end of month 2:

$$M_2 = M_1(1+m) - P \tag{7}$$

In the second period, therefore, the mortgage balance declines by (P - mM_1). Because M_1 is smaller than M_0, the interest portion is smaller and the decline in the balance is larger in the second period than in the first. Over the lifetime of a mortgage, the share of the payment that goes to interest gradually declines and the share that goes to reducing principal gradually rises.

Now we can substitute equation (6) into equation (7) to obtain:

$$M_2 = M_1(1+m) - P = [M_0(1+m) - P](1+m) - P$$
$$= M_0(1+m)^2 - P(1+m) - P \tag{8}$$

One more time:

$$M_3 = M_2(1+m) - P = [M_0(1+m)^2 - P(1+m) - P](1+m) - P$$
$$= M_0(1+m)^3 - P(1+m)^2 - P(1+m) - P \tag{9}$$

This process can easily be generalized to obtain a formula for the balance remaining at the end of any period, *y*. To be specific

$$M_y = M_0(1+m)^y - P(1+m)^{y-1} - P(1+m)^{y-2} - P(1+m)^{y-3} - \ldots - P \tag{10}$$

The series of dots indicates the set of terms that we have not written out—all the terms with exponents between *y*-3 and 1. The number of these terms obviously depends on the value of *y* we are looking at. With *y* = 10, for example, there are 6 terms in this set.

Two more steps lead to the formula for a mortgage. The first step is to use a little algebraic magic. We can re-write equation (10) as follows:

$$M_y = M_0(1+m)^y - P \times S \tag{11}$$

where

$$S = (1+m)^{y-1} + (1+m)^{y-2} + (1+m)^{y-3} + \ldots + 1 \tag{12}$$

A sum of this type can be re-arranged so that only a few terms remain. To see, how, first multiply *S* by (1+*m*) to obtain

$$S(1+m) = (1+m)^y + (1+m)^{y-1} + (1+m)^{y-2} + \ldots + (1+m) \tag{13}$$

Now if we subtract equation (12) from equation (13) we find that all the middle terms in the string all cancel out and what remains is

$$S(1+m) - S = (1+m)^y - 1 \tag{14}$$

With just a little bit of algebra, this expression can be solved for *S*:

$$S = \frac{(1+m)^y - 1}{m} \tag{15}$$

Substituting equation (15) back into (11) results in:

$$M_y = M_0(1+m)^y - P\left(\frac{(1+m)^y - 1}{m}\right) \tag{16}$$

The second step is to recognize that the balance due at the end of month *Y* has to be zero. In other words, the mortgage has to be fully paid off at the end of its term! So when we evaluate equation (16) at the value of $y = Y$, the left side has to equal zero. In symbols:

$$\begin{aligned} &M_Y = 0 = M_0(1+m)^Y - P\left(\frac{(1+m)^Y - 1}{m}\right) \\ &\text{or} \\ &P\left(\frac{(1+m)^Y - 1}{m}\right) = M_0(1+m)^Y \end{aligned} \tag{17}$$

The final steps just involve re-arranging this result to figure out how the payment, *P*, depends on the mortgage amount, M_0; the interest rate, *m*, and the term, *Y*. We want to multiply both sides of equation (17) by *m* and divide both sides by $(1+m)^Y$. Remember that $1/(1+m)^Y=(1+m)^{-Y}$. These steps lead to:

$$P = \left(\frac{m}{1-(1+m)^{-Y}}\right)M_0 \tag{18}$$

In short, in a fixed-rate mortgage contract, the payment is proportional to the mortgage amount, and the proportion, which is shown in equation (18), depends on the mortgage's interest rate and term.

This is a useful equation to know. If you ever borrow money with a fixed-rate loan, you can use equation (18) to see if the lender accurately calculated the monthly payment! Suppose you are taking out a mortgage of \$200,000 to buy your dream house. If you have a 30-year mortgage at a 5% annual interest rate, then your monthly payment will be

$$\left(\frac{.05/12}{1-[1+(.05/12)]^{-(30\times 12)}}\right)\times \$200{,}000 = \$1{,}073.64$$

Finally, note that you can also use this formula to figure out the mortgage you can afford (M_0) for a given monthly payment, *P*, and given mortgage characteristics, *m* and *Y*.

Lenders sometimes offer borrowers a choice between option A, a mortgage with a relatively high interest rate, and option B, a mortgage with a relatively low interest rate combined with “points” paid up front. In this context, a point is 1 percent of the mortgage amount, M_0. Because the points in option B are paid at the beginning of the mortgage, their impact is felt no matter how long one lives in the house. In contrast, the lower interest rates in option B do not yield any benefit unless a person lives in the house for a long time. Let us say that people who expect to move in the near future have a short time horizon. People with a short time horizon experience the costs of option B (points up front) but not the benefits (a long stream of lower interest rates). The same is true for people with a high discount rate, because the long-term benefits of a lower rate are heavily discounted. Both of these groups should pick option A. For people with a long time horizon and a low discount rate, that is, people who expect to be in the house for the long haul, the benefits from a long stream of lower interest rates outweighs the cost of the up-front points. These people should pick option B.

Annuities and Pensions

An annuity is a payment stream supported by an initial fund. The most important type of annuity is a pension. For many pensions, a person contributes money into a fund during her working life and then draws a pension payment—an annuity—from the fund after she retires. Annuities do not have to be pensions, however.

It is not correct to calculate the annuity by dividing the fund by the number of pay periods because the fund earns interest along the way. As a result, the payments are actually larger than the fund divided by the number of periods.

An annuity is designed so that the fund will be exhausted at the end of the expected number of pay periods. We have to use the qualifier “expected” here because an annuity generally continues even if the expected number of periods is exceeded. In the case of a pension, for example, the pension payment will be calculated on the basis of an expected lifetime, but will continue even if a person lives longer than expected. The people who set up annuities must account for this type of uncertainty, but we are not going to do that here. Instead, we will just refer to an expected number of pay periods.

As it turns out, except for the uncertainty about timing, a pension is just the inverse of a mortgage. You will be happy to know that we do not have to do any new algebra! An annuity is designed so that the fund is exhausted at the end of the expected number of pay periods—just like a mortgage is paid off at the end of its term. Moreover, an annuity must account for interest received from the amount that remains in the fund—just like a mortgage must account for interest payments on the amount that is still owed.

So let F be the amount of money in a fund supporting an annuity. This money is given to a fund administrator by a retiree (or other investor) to invest in exchange for a promise to pay the retiree an annuity of $\$A$ per month for a period of L months. Suppose g is the interest rate that the fund administrator expects to earn on this fund. Then plugging these terms into equation (18) we find that

$$A = \left(\frac{g}{1-(1+g)^{-L}} \right) F \qquad (19)$$

In some cases, an investor or potential retiree might want to know how large the fund has to be to support a given annuity. The answer to this question (which is analogous to the question of how large a mortgage one can afford at a given payment) can be found by rearranging equation (19):

$$F = \left(\frac{1-(1+g)^{-L}}{g} \right) A \qquad (20)$$

If I want a pension of $2,000 per month for 25 years and I expect the investments from the fund to earn a 6% return per year, then the required fund is

$$\left(\frac{1-[1+(.06/12)]^{-(25\times 12)}}{.06/12} \right) \times \$2,000 = \$2,185,987.92$$

I'll have to save up!

Pension planners have to conduct calculations such as this one to ensure that a pension fund will have enough money in it when person retires to cover promised pensions. The pension fund may receive contributions both from an employee and an employer, and the amount in the fund obviously depends on the size of these contributions, along with the growth in the employee's wages during her working life. These notes do not pursue the issue of pension planning in this sense, but interested students may want to consult the "Spreadsheet on the Algebra of Pensions" that is posted on the "Computer Programs" tab on my web site.

Bonds

A bond is an example of a "certificate of indebtedness." A person who buys a bond (from a corporation or a municipal government, for example) receives the legal assurance that she will receive a stream of payments from the institution that issued the bond. Thus, a bond is a form of investment with a particular pattern of returns for the person who buys it. A bond is, of course, quite different from a stock. A person who buys a stock becomes one of the owners of a company, with the right to vote on certain matters and the right to receive dividend payments, which are much less predictable than bond interest payments. The holder of a bond has no ownership rights.

A bond is similar to a mortgage; the person who buys the bond is like the lender and the institution issuing the bond is like the borrower. Local governments issue bonds, for example, to raise funds for infrastructure projects, such as bridges or schools. Moreover bonds, like mortgages, come in many different forms, including forms with variable interest rates and bonds that can be redeemed before the maturity date. These notes just examine basic, fixed-rate bonds.

However, bonds have one characteristic that makes them quite different from a mortgage, namely that the payments do not alter the principal, that is, they do not alter the amount upon which interest is calculated. As a result, bonds retain a principal amount at the end of the contract, the maturity date; this principal amount must be repaid to the investor.

More specifically, three characteristics are associated with each bond:

(1) the face value, *F*, (also called the par value or the redemption value or the value at maturity), which is the amount upon which the interest payments are calculated;

(2) the coupon rate, *c*, which indicates the interest to be paid as a percentage of *F*; and

(3) the years to maturity, *N*, which is the number of years during which the investor is entitled to receive interest and also the number of years until the bond can be redeemed.

As an aside, bonds are generally used to cover the cost of a project and they are generally issued in sets with different maturities. More specifically, they are usually issued in serial form, which means that some of the bonds in a set have a maturity of 1 year (N=1), others have N=2, and so on all the way up to the highest selected maturity, say N^*. This approach eases the repayment burden on the issuer. If each maturity up to N^* has the same number of bonds, then the issuer will only have to pay back a share equal to $1/N^*$ of its bonds each year.

The key to understanding bonds is to think about what an investor would pay for a bond that has a certain face value, *F*, coupon rate, *c*, and maturity, *N*. The answer is that an investor will pay the present value of the stream of benefits from owning the bond. Any good investor knows about present value! Suppose an investor has an alternative, similar investment, perhaps a U.S. Treasury Bill, that offers an interest rate *r*. Then *r* is the opportunity cost of investing in bonds, and the investor's willingness to pay is the present value of the benefits from holding the bond discounted at rate *r*.

The present value of a bond is not the same thing as its face value. In fact, these two "values" may be quite different. To see why, we have to recognize that the present value of benefits from owning a bond consists of an interest payment of *cF* per year and a redemption payment of *F* in year *N*. (In practice, interest payments are often made twice a year; if you want to account for this you just need to express everything in "half years" instead of years, with an interest rate of $c/2$, where *c* is the annual interest rate.) Thus, the present value of a bond, and hence its market price, is

$$P = \frac{cF}{(1+r)} + \frac{cF}{(1+r)^2} + ... + \frac{cF}{(1+r)^N} + \frac{F}{(1+r)^N} \qquad (21)$$

Note that the last term in this expression does not have *c* in it; this term is the present value received by the investor when the issuer of the bond returns the face value to the investor, which is called redemption. Note also that this expression, like an earlier one, has a series of dots, which indicate the terms for interest received between years 2 and *N*.

This expression can be simplified using the same algebraic trick that led to equation (15). Define the sum of all the terms on the right side with cF in the numerator as C, for coupon payments. Then subtract C from $C\times(1+r)$ and solve the result for C. Replacing all the cF terms with this expression for C leads to:

$$
\begin{aligned}
P &= \left(\frac{1-(1+r)^{-N}}{r}\right)cF + \frac{F}{(1+r)^N} \\
&= F\left(c\left(\frac{1-(1+r)^{-N}}{r}\right) + \left(\frac{1}{(1+r)^N}\right)\right)
\end{aligned}
\tag{22}
$$

So there is the echo of the mortgage formula in the price of a bond! Again, the difference is that the principal in the bond does not decline in absolute terms (although it does decline in present value!) and is repaid in year N. The price of a bond thus reflects the present value of the coupon payments (the first term in equation (22)) plus the present value of the redemption amount (the second term).

In the second line of equation (22) we can see that P is proportional to F. If this proportion is greater than 1, the bond is said to sell at a premium. If it is less than one, the bond is said to sell at a discount. Consider a bond with F = \$5,000; c = 4%, and N = 20 for investors with an opportunity cost rate of 6%. Then the amount this investor will bid on the bond, P, is

$$\$5,000\times\left(.04\times\left(\frac{1-1.06^{-20}}{.06}\right)+\left(\frac{1}{1.06^{20}}\right)\right)=\$3,853.01$$

This bond is priced at a discount (\$3,853.01 < \$5,000) because the rate of return it offers is below the opportunity cost rate. Investors will not buy the bond unless it yields a capital gain, that is, unless the redemption value they receive in year 20 is considerably greater than the price they have to pay.

Sometimes an investor who owns a bond wants to sell it at a given price. In this case, P is known but the rate of return on the bond is not. (This contrasts with the preceding discussion, in which the opportunity cost to investing a bond, r, was known, but the price was not.) This changes the way we use equation (22), because r can now be interpreted as the rate of return on a bond with price P, face value F, coupon rate C, and maturity N. This rate is called the ***internal rate of return***. When comparing bonds (or other assets) an investor wants to select the bonds with the highest internal rate of return.

The problem is that equation (22) is highly nonlinear and cannot easily be solved for r. Fortunately, however, finding the internal rate of return is such a common problem that most calculators and spreadsheet programs include a function that calculates r once P, F, C, and N have been entered.

Some municipal bonds push the issuer's payments into the future. This approach lowers the tax burden of repayments in the short run, but raises this burden in the long run. These bonds therefore undermine to some degree the main purpose of issuing bonds—to smooth out the tax burden of infrastructure projects—but they may also enhance the appeal of the bonds to certain classes of investors. This higher appeal may translate into a lower required return, which corresponds to a lower the cost to the municipality. A zero-coupon bond, for example, sets the coupon rate, c, equal to zero. In this case equation (22) indicates that the price of the bond is:

$$\text{Zero-coupon bond:} \quad P = \frac{F}{(1+r)^N} \tag{23}$$

This type of bond is designed for investors who want to receive their returns in the form of capital gains, not interest, probably for income tax reasons. Capital gains are indicated by the difference between F and P, so the price of a zero coupon bond obviously has to be heavily discounted, that is, P must be far below F, or else no investor would buy it.

Another type of bond that has been used by some school districts to postpone repayments into the future is a capital appreciation bond (CAB). With this type of bond, the interest payments are placed in a fund and reinvested at the coupon rate. The investor claims the money in this fund, along with the face value, at the maturity date. Thus, all returns to this investment are received in year N and discounted by $(1+r)^N$. The payment into the fund each year is cF. The first payment grows to $cF(1+c)^{N-1}$ by the time the bond matures; the second payment grows to is $cF(1+c)^{N-2}$; the last payment of cF is paid in year N and does not grow. Let Z be the sum of these accumulated amounts in year N. Then the same algebraic steps that led to equation (15), imply that $Z = cF[(1+c)^N - 1]/c = F[(1+c)^N - 1]$. The price of the bond equals the present value of $(F + Z)$ or

$$\text{Capital Appreciation Bond:} \quad P = \frac{F + F\left((1+c)^N - 1\right)}{(1+r)^N} = F\left(\frac{(1+c)^N}{(1+r)^N}\right) \tag{24}$$

When $c = r$, the price of a capital appreciation bond equals its face value. Moreover, a CAB is similar to a zero-coupon bond in that the CAB interest rate, c, is algebraically equivalent to the implicit annual capital gains rate in a zero coupon bond. (This equivalence may not be recognized by the Internal Revenue Service, which means that these bonds may appeal to different investors than zero-coupon bonds). To put it another way, a zero-coupon bond will sell for the same amount as a CAB if the capital gains reflected in its face value equal the accumulated interest for the CAB; that is, if its face value equals the CAB face value multiplied by $(1+c)^N$.

A rarer (and more dubious) type of bond is a perpetuity. The holder of this type of bond is entitled to receive interest payments indefinitely, which is equivalent to having an infinite number of years until maturity. Because $[1/(1+r)^{\infty}]=[(1+r)^{-\infty}] = 0$, equation (22) indicates that the price of this type of bond is:

$$\text{Perpetuity:} \quad P = \frac{cF}{r} \tag{25}$$

The ratio of price to face value in this case obviously depends directly on the coupon rate, because the bond can never be redeemed. The bond sells at a premium if $c > r$, that is, if it yields a return greater than the investor's discount rate, and it sells at a discount if $c < r$. In either case, the outcome is that the annual interest from the bond, cF, equals the annual interest that could be earned with the same amount of money invested in an alternative asset, rP.

Inflation

Inflation is a general increase in prices, holding product quality constant. Prices increase by at least 2% almost every year, and sometimes increase much more. The U.S. Bureau of Labor Statistics collects data on prices and calculates various price indexes relative to some base year, such as the well-known Consumer Price Index. This index is a relative price level multiplied by 100 to make it more readable. A price index of 110, for example, indicates that prices are 10 percent higher than they were in the base year. Actual dollars received are called nominal dollars. Dividing nominal dollars by the price index (without the 100 scalar) yields real dollars for a given base year. (Students who want to learn more about inflation and how it is measured may want to consult my "Notes on Price Indexes and Inflation," which are available on my web site at the "Teaching" tab.) When people think about benefits to be received in the future, therefore, they recognize that a given amount of money received in the future will not go as far as the same amount of money received today. To put it another way, dollars will have lower purchasing power in the future.

Because inflation affects the purchasing power of money received in the future, it becomes part of the discount rate. Suppose my discount rate is 5% without considering inflation. Then suppose that I expect the rate of inflation to be 3%. Under these circumstances, I no longer consider $105 received at the end of the year to be equivalent to $100 received now, because I know that the purchasing power of a dollar has declined. In fact, I need $105×1.03 = $108.15 to feel that I have an equivalent amount.

The no-inflation discount rate is also known as a ***real*** discount rate, whereas the rate that incorporates inflation is called a ***nominal*** discount rate. Interest rates that are observed in the market place are all nominal interest rates because they incorporate the inflation expectations of market participants. In the above example, the real discount rate is 5% and the nominal discount rate is $[(1+.05)\times(1+.03) - 1] = [(1.05)\times(1.03) - 1] = .0815 = 8.15\%$. A simpler way to calculate the nominal interest rate, which is an approximation to this exact formula, is simply to add the expected rate of inflation to the no-inflation discount rate. To continue the example, the approximate nominal discount rate is 5% + 3% = 8%.

When thinking about inflation it is important to be consistent, that is, to make certain that the dollar amounts and the discount rates are both in either nominal or real terms. It makes no sense to use a nominal discount rate to find the present value of a real return. The examples given above express everything in nominal terms. A nominal return of $108.15 has a present value of \$100 when the nominal discount rate is 8.15%. It is also possible to obtain the same answer using real values. We can assume that the beginning of the year is the base year, so a nominal value of \$108.15 equals a real value of \$108.15/1.03 = \$105. Now using a real discount rate of 5%, the present value of \$105 is \$100. Because we observe nominal rates, not real rates, in the market place, we would have to calculate the real discount rate to follow this approach. Using the approximation in the previous paragraph, the real discount rate is 8% - 3% = 5%.

To write down a general form for this real/nominal relationship, let i_R be the real interest rate, i_N be the nominal interest rate, and p be the anticipated rate of inflation. Then

$$i_N = (1+i_R)(1+p)-1 = i_R + i_R p + p \tag{26}$$

and

$$i_R = \left(\frac{1+i_N}{1+p}\right) - 1 = \frac{i_N - p}{1+p} \tag{27}$$

Because i_R and p are both small numbers, their product is very small, and these equations are sometimes approximated as $i_N = (i_R + p)$ and $i_R = (i_N - p)$.

Risk and Uncertainty

The financial transactions examined in these notes are all designed with a view toward the future, which cannot, of course, be predicted with perfect accuracy. A homebuyer may or may not default on her mortgage loan. A retiree may live much longer than expected when her pension was designed. A company or a public agency may go bankrupt and be unable to pay off its bondholders. In common parlance, the future is uncertain and financial transactions involve risks. (Following the work of economist Frank Knight, some scholars define these terms in a different way: risk arises when the probabilities of various outcomes are known, even if the outcomes themselves are not, whereas uncertainty refers to cases in which even the underlying probabilities are not known. In these notes I stick to the more conventional definitions.)

The presence of risk and uncertainty lead to two new central concepts in the analysis of financial transactions. The first concept is ***expected value***. The expected value of a financial transaction is the sum across possible outcomes of the probability of an outcome multiplied by its present value. In the case of a mortgage, for example, a default results in a loss of interest payments and to a foreclosure process that might be expensive but that eventually transfers the ownership of the house from the borrower to the lender. The net benefit to the lender could be positive or negative, depending on the magnitude of their costs and the value of the house on which they foreclose. The expected value of mortgage returns in year y is the probability of a default multiplied by the present value of the net benefits from default plus the probability of no default multiplied by the present value of the mortgage payments. The concept of expected value therefore makes it possible to introduce a range of outcomes with known probabilities into the analysis of financial transactions, but it can lead to some pretty complicated algebra! You may be relieved to learn that this type of algebra is beyond the scope of these notes.

The second key concept is ***risk aversion***, which is said to exist when a person prefers a (relatively) certain set of outcomes with a given expected value to a (relatively) uncertain set of outcomes with the same expected value. Consider an investor who is trying to decide between buying one set if mortgages with a relatively low mortgage rate and a relatively low probability of default and buying another set of mortgages with a relatively high mortgage rate and a relatively high probability of default. Suppose the expected values of the two investments are the same because the high interest rate on the second investment compensates for the high risk of default. An investor who is risk averse will always prefer the first investment to the second.

Scholars and financial analysts have developed elaborate mathematical tools to account for risk in financial transactions. These tools go way beyond simple algebra and are way beyond the scope of these notes.

The concept of risk aversion also leads to three important ideas about financial markets. First, the widespread presence of risk aversion in a market implies that risky assets will not attract investors unless they offer a higher return than offered by low-risk assets with returns that have the same expected value. In other words, investors have to be compensated for risk.

Second, risk-averse investors often seek to minimize their risk by diversifying their portfolio. In this context, a diversified portfolio is one that contains a series of investments with returns that are not highly correlated with each other, so that bad outcomes for one investment are unlikely to be accompanied by bad returns for another. If you are risk averse, you should think about diversifying your portfolio!

Third, because many investors are risk-averse, they are willing to pay something for a financial agreement that insures them against certain kinds of financial risks. As a result, modern financial markets contain institutions that specialize in accepting risk or, to put it another way, that specialize in providing this type of insurance. The expansion of this type of institution in the mortgage market, along with the lack of regulation of such institutions, contributed to the financial crisis of 2008. That is a subject for a different set of notes altogether.