Chapter I. Preliminaries

1. Introduction

Knowing the factors that determine the prices of financial and non-financial assets is important in making decisions at the firm level; it is also crucial in understanding financial market signals as they apply to business activity, financial stability, and fluctuations in aggregate wealth and the distribution of wealth. When we consider "asset" pricing we often have in mind stock prices. However, asset pricing in general also applies to other financial assets, for instance, bonds and derivatives, to non-financial assets such as gold, real estate, and oil, and to collectibles like art, coins, baseball cards, etc. We will mostly think of the assets as being stocks, but we will also see that many of the pricing principles that apply to stocks apply to other assets as well.

We may think of any asset as generating (usually) risky, future payoffs distributed over time. The value of the asset can be viewed as the present value of the payoffs or cash flows, properly discounted for risk and time lags. At a partial equilibrium level of analysis, both the risk-adjustment method and its quantitative value, as well as the time-lag adjustment method and value are taken as given. For our purposes we will be interested mostly in general equilibrium models in which the proper risk and time adjustments are obtained endogenously.

Obtaining values for asset prices is quite similar to obtaining asset returns. In particular, for a given stream of direct payoffs and derived asset prices, the current one-period asset return can found as the sum of current price plus current payoff divided by previous-period price. Conversely, for given asset returns and given stream of direct payoffs, the asset prices can be backed out. In most cases it will be more convenient for us to focus on asset returns than asset prices.

Two prominent asset pricing theories, the Arbitrage Pricing Theory and Merton's Intertemporal Capital Asset Pricing Model, provide a very general solution for the factors that affect expected asset returns. It is a common misperception to view more general models as superior to less general models. The problem with the aforementioned APT and ICAPM is that they are so general that they have virtually no empirical content. I would characterize this shortcoming as the key problem in the asset pricing field. In these notes we will not fail to cover the general asset pricing models, but we will emphasize other asset pricing models that are less general but are based on specific general equilibrium models, with specific, economically interesting, factors determining the asset returns.

Aside from the "general" asset pricing models and the "specific" asset pricing models discussed above, a third type of asset pricing model may be distinguished: the "pricing by arbitrage" type. This type of model prices assets by comparison. One creates a "portfolio" that mimics the payoffs of a particular primary asset. The mimicking portfolio is then priced by arbitrage. The portfolio may be constructed such that one may infer the price of a derivative asset (an option, for instance). This derivative asset is priced only in comparison to the primary asset. We will pay little attention to this type of asset pricing in the first financial economics course. It will be covered extensively in our third financial economics course.

In the remainder of the chapter we will discuss issues related to adjustment of cash flows to account for time preference in section 2 and issues related to adjustment of cash flows to account for risk in section 3.

2. THE TIME VALUE OF MONEY

(a) The Subjective Rate of Time Preference

In a standard consumer choice model, one may think of life-time utility U as being determined by consumption levels c_i in the different T periods that the consumer has left to live $U = U(c_1, c_2, \ldots, c_T)$. Time preference (alternatively impatience, or discounting the future) is said to exist if future consumption, in a sense to be made precise in the following, is less valuable to the individual than is current consumption. Traditionally, explicit dynamic models have assumed time separable preferences, which imply a constant rate of time preference. However, other dynamic preference specifications, in which the rate of time preference changes over time are becoming increasingly popular. It is therefore important to consider a general definition for time preference that allows variation over time. Define the *discount rate* (or *rate of time preference* or *degree of impatience*) at time t as ρ_t. Then we can define the (one-period ahead) *discount factor* as $\beta_t = 1/(1+\rho_t)$. Thus, if the discount rate at time t is positive $\rho_t > 0$, then the discount factor is less than one, $\beta_t < 1$; if the discount rate increases (falls), the discount factor falls (increases).

The one-period discount factor at time t is defined exactly in a discrete-time formulation as:

$$\text{(1)} \qquad \beta_t(c_1, \ldots, c_{t-1}, c, c, \ldots, c_T) \equiv -\frac{dc_t}{dc_{t+1}}\Big|_{dU=0,\; c_t = c_{t+1} = c}$$

$$= \frac{\partial U(c_1, \ldots, c_{t-1}, c, c, \ldots, c_T)/\partial c_{t+1}}{\partial U(c_1, \ldots, c_{t-1}, c, c, \ldots, c_T)/\partial c_t} .$$

The second equality follows since $dU = 0 = (\partial U/\partial c_t)\, dc_t + (\partial U/\partial c_{t+1})\, dc_{t+1}$. The discount factor, in words, is thus given as the *Marginal Rate of Intertemporal Substitution* between two points in time at the location where consumption at the two points in time is equal. In intuitive terms, it indicates, given initially that $c_t = c_{t+1}$, how much of c_t the individual is willing to sacrifice for one extra unit of c_{t+1}. Note that, in general, the discount factor may depend on consumption levels at all points in time.[1]

For a *time-separable* intertemporal utility specification, utility can be written as:

$$\text{(2)} \qquad U(c_1, \ldots, c_T) = \sum_{t=1}^{T} \beta^{t-1} u(c_t).$$

The discount factor, then, can be obtained from the definition in equation (1) as:

[1] A typical misconception is that time variation in the discount rate implies time-inconsistent behavior. This is not the case as long as the form of the discount function in equation (1) does not change over time. That is, if the discount factor only depends on the stream of consumption and not explicitly on time [see Strotz (1956)]. For more, relatively non-technical, information about discounting see Price (1993).

(3) $\qquad \beta_t(c_1, \ldots, c_{t-1}, c, c, \ldots, c_T) \;=\; \beta.$

Thus, discount rate and factor are constant for the time-separable utility formulation.

Figure 1 shows that the discount factor can be obtained as the absolute value of the slope of the indifference curve along the 45^o degree line emanating from the origin. The slope dc_t/dc_{t+1} should be smaller than one for positive time preference to exist. Time separability implies that the slope dc_t/dc_{t+1} does not depend on the level of consumption at times other than *t* or *t+1*. For the commonly used time separable form in equation (2), the slope of dc_t/dc_{t+1} along the 45^o degree line also does not depend on the level of $c_t = c_{t+1}$, which implies that the slope is the same for each indifference curve, as shown in Figure 1.

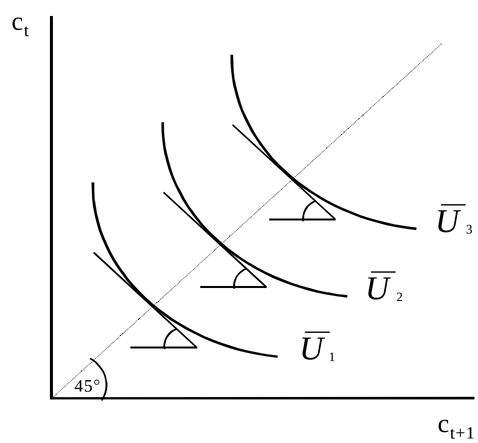

Figure 1
Time Preference
Utility is time-separable, as indicated by equal slopes along the 45° line.

For positive time preference, it makes sense to discount future cash flows more than current cash flows. A natural discount factor for an individual valuing the payoffs of a particular asset would appear to be the subjective rate of time preference. However, we will see in the following that (1) individual preferences should have no impact on the discount rate used in valuing an asset; and (2) even if preferences of all individuals are time-separable and identical, the proper discount rate in asset valuation is not generally equal to β.

(b) The Objective Rate of Time Preference.

For simplicity consider here only the two-period world originally studied by the Classical economist Irving Fisher. For a perfectly competitive world, Figure 2 displays the standard general equilibrium allocation, but for the composite consumption good at two points in time (c_1 and c_2) rather than two different consumption goods at one point in time.

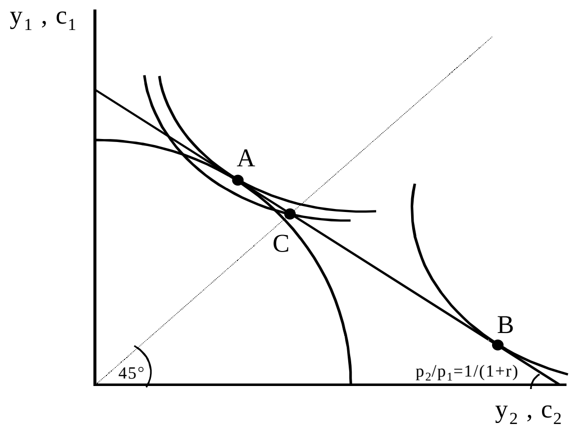

Figure 2
Objective Discount Factor and Fisher Separation
The objective rate of time preference will be the same for any two consumers. The objective and subjective rates of time preference will differ for the representative consumer.

At point *A* we have the standard condition for a Pareto Optimal allocation sustained in competitive equilibrium by an equilibrium price ratio:

$$(4) \qquad MRIT(c_1, c_2) = p_2/p_1 = MRIS(c_1, c_2),$$

the Marginal Rate of Intertemporal Transformation is equal to the Marginal Rate of Intertemporal Substitution and this equality is obtained in general equilibrium through the equilibrium price ratio p_2/p_1. This equilibrium price ratio represents the *objective* (or market) *discount factor,* generating the objective rate of time preference. The analysis here presumes of course that an aggregate production possibilities frontier and an aggregate preference ordering exist, or, similarly, that a representative firm and representative consumer exist. Note that we can interpret

$$(5) \qquad p_2/p_1 = 1/(1+r),$$

where r is the real interest rate: saving one unit of consumption currently can buy you one more unit in the next period.[2] If we assume existence of a representative firm and consumer, then the two equalities in (4) can be derived from standard optimization problems:

$\underset{y_1, y_2}{Max} \; y \equiv p_1 y_1 + p_2 y_2, \quad subject \; to \quad T(y_1, y_2, \bar{T}) = 0$: for given inputs $\bar{T}$ maximize profit, subject to

the transformation technology, yielding, in market equilibrium ($y_1 = c_1$, $y_2 = c_2$), the first expression in (4), and:

[2] One may verify that the Fisher analysis determines the market interest rate by setting savings (postponing consumption to the next period) equal to physical investment (reserving inputs for production in the next period).

$\underset{c_1,\, c_2}{Max}\ \ U(c_1, c_2)\,, \quad subject\ to\ \ p_1 y_1 + p_2 y_2 = y:$ maximize the, possibly non-time-separable, utility subject to a budget constraint, which yields the second equality in (4).

Consider the preferences of an individual consumer, taking market prices as given. In Figure 2, view the allocation *A* as that for a representative consumer *a*. For a consumer *b* that differs from the representative consumer *a* but has an identical wealth endowment, the optimal allocation may be found at point *B*. Note that $MRIS_A = MRIS_B$; this fact illustrates our first result of interest, namely that individual preferences are irrelevant for the discount factor that should guide the individual's decisions, including the decision how to value a particular asset. This is known as the *Fisher Separation* result. It also applies in a corporate finance situation where the way to maximize the value of the firm to shareholders should be independent of individual preferences but instead should be based on market prices, such as the market discount rate. The argument is that any feasible trade for the individual must lie along the intertemporal budget line. Any allocation along the budget line that differs from point *B*, of course will change the realized marginal rate of intertemporal substitution but will also move the individual to a lower indifference curve.

A second result may be seen by considering point *C* in Figure 2 which is the intersection between the market discount rate line (and the budget line for the representative consumer) and the 45^o line. It indicates also the subjective discount factor of the representative consumer, given as the slope of the indifference curve at point *C*. If the preferences are intertemporally separable then this slope is equal to the constant β. This is clearly different from the slope of the budget line.

Thus, even for the representative consumer, the objective and subjective rates of time preference will differ. When preferences are intertemporally separable both are different constants. The objective discount factor is given by $MRIS(c_1, c_2) = \beta u^{/}(c_2)/u^{/}(c_1)$ which equals $1/(1 + r)$ and differs generally from β which equals by definition $1/(1 + \rho)$, since the *MRIT* may easily imply an allocation where c_1 and c_2 are not equal. Hence, the representative consumer will discount cash flows differently (namely by the constant objective discount factor) from the way she will discount utility (which, for intertemporally separable preferences, is discounted by the constant subjective discount factor). In simple (but imprecise) terms one may say that wealth is discounted differently than consumption.

(c) Present Value and Value Additivity

As we have seen, in the certainty context of the Fisher model we can use the real interest rate to discount future consumption streams. Take now a real cash flow to be received (with certainty) at a time *s* periods in the future. What is its present value? We can obtain the correct answer in two different ways. First, using the idea of opportunity cost. By postponing the cash flow by *s* periods you forego interest. The Future Value at time *s* of an initial investment X_0 is given as $FV_s = (1+r)^s X_0$, where we account for periodic compounding and we assume a constant real interest rate. If we turn this around and set the Future Value at time *s* equal to a certain cash flow CF_s , fixed in advance, then we can find the Present Value at time *0*, PV_0 , as

$$(6) \qquad PV_0 = \frac{CF_s}{(1 + r)^s}\ .$$

Note that, in a perfect market environment, the opportunity you sacrifice–your next best alternative used for calculating the opportunity cost– must be economically identical to the project you evaluate. Given that the return on this next best alternative is equal to r , we can view the present value calculation as simply representing arbitrage: the cash flow of the project should be valued such that both this project and the alternative generate the same return.

As a second, related, approach to calculating present value, one could extend the logic of the Fisher analysis to a multi-period context. Assuming that the objective discount factor remains constant we then would have the objective discount factor as $p_s/p_0 = (p_s/p_{s-1})(p_{s-1}/p_{s-2})...(p_1/p_0) = [1/(1+r)]^s$. This is the appropriate factor to discount cash flows, yielding again equation (6).

In the perfectly competitive Fisher world it also follows that each component of a stream of cash flows could be bought or sold separately. Thus, the present value of a stream of cash flows is just the sum of the present values of the separate cash flows. This property is called *Value Additivity*. Accordingly, we can write the present value of a stream of future payments as:

$$(7) \qquad PV_0 = \sum_{s=1}^{\infty} \frac{CF_s}{(1+r)^s} .$$

In practice, Treasury Bonds are broken up into their separate components: coupon payments and the zero-coupon remainder (the latter called a "strip"). As dealers find it worth their while to incur the transaction costs to do this, it appears that Value Additivity holds only approximately in actuality.

(d) The Gordon Growth Model

The cash flows in equation (7) could follow any pattern. Usually they will be zero beyond some point *T*. In many cases cash flows can only be approximated and one has but vague ideas of its pattern over time. These vague ideas are often best summarized by an initial value, that is quite precisely known, and an average growth rate, maybe less precisely known.

The "Gordon Growth Model" makes the practical assumptions of a known initial cash flow X_1 and a known constant growth rate g. It is a useful model not only because it provides a reasonable approximation to the entire future of cash flows, it also generates a simple expression for the Present Value. With these assumptions, equation (7) becomes:

$$(8) \qquad PV_0 = \sum_{s=1}^{\infty} \frac{(1+g)^{s-1} X_1}{(1+r)^s} = \frac{X_1}{r-g} .$$

That the second equality holds is easy to see by realizing that we have an infinite geometric series with a constant geometric factor of decline of $(1+g)/(1+r)$. Clearly, if cash flows grow at a higher or equal rate compared to the real rate of interest, $g\$ r$, the present value is infinite as the individual present value components do not decline over time.

Equation (8) can be quite useful in a variety of cases. One application is obtained if we apply the Gordon Growth Model to the overall stock market. Strictly taken the model only applies for a fairly certain growth rate.

However, as we will find later in this chapter, one may just discount expected cash flows for risk by discounting by the appropriate expected real stock return (the proper opportunity cost for the riskiness of this particular project) rather than the real interest rate.

The only cash flows generated by stocks in the aggregate are the dividend payments since we can ignore the pure distribution effects rendered by capital gains or losses from stock trading. Thus, the value of a stock should be the present value of the stream of dividends generated. In equation (8) we may thus interpret X_1 as a typical annual aggregate dividend payment and PV_0 as the current price of a stock index. The *price-dividend ratio* is then equal to $PV_0/X_0 \approx PV_0/X_1$. If average price-dividend ratios for S&P 500 firms equal around 33. This implies from equation (8) that $r - g = 0.03$. Historically, real stock returns have been around 0.08 (8%). Hence, stock valuations imply that $g = 0.05$, that dividends growth is expected on a permanent basis to equal 5%. This number substantially exceeds the around 2% historical growth rate of real per capita GDP and the similar projections for its growth. Explanations may be that the growth of profitability on a permanent basis is expected to exceed real growth of GDP; that the risk premium inherent in stock returns has fallen permanently; that we have a permanently altered economy; or that stocks are simply overvalued.

(e) Inflation and Valuation

A stylized fact is the observation that stocks are not a perfect hedge for anticipated inflation. Inflation, anticipated as well as unanticipated, appears to have a clear negative impact on real stock returns. Here we will see that, at first sight, inflation should have no impact on stock prices and returns.

Assume that future dividends in real terms are unaffected by inflation π. This is a reasonable assumption since firm profits in real terms should not change–both costs and revenues should change with inflation, causing nominal profits to rise with inflation and real profits to be unchanged. Using the information of the present value calculation in equation (8) for simplicity, it is assumed that real cash flow X_1 is not affected by inflation and thus equals $(1+\pi)X_1$ in nominal terms. Due to inflation, nominal dividends will be growing at a higher rate g_n and we now discount the nominal cash flows by the nominal return i. Equation (9) then gives the valuation for the Gordon Growth Model when inflation is introduced:

$$(9) \qquad PV_0 = \sum_{s=1}^{\infty} \frac{(1+g_n)^{s-1}(1+\pi)X_1}{(1+i)^s} = \frac{(1+\pi)X_1}{(1+i)-(1+g_n)} = \frac{X_1}{r-g} .$$

The second equality follows by the definitions of real interest rate r: $(1+i)/(1+\pi) = 1+r$ and real growth rate g: $(1+g_n)/(1+\pi) = 1+g$. It follows in the context of our assumptions that (a) current stock prices are not affected by anticipated inflation and that (b) one may alternatively discount nominal payoffs with the nominal interest rate or real payoffs with the real interest rate.

Empirically, it appears that both anticipated and unanticipated increases in inflation lower stock prices, so that stocks appear to be poor inflation hedges. This is clearly counter to our discussion which implies that inflation's effect on stock prices should be neutral. The most prominent explanation for this apparent non-neutrality is by Fama (1981):

since stocks react quickly to new information, they decrease on news of a business cycle downturn, before this downturn happens. But a downturn in activity, given constant money growth, implies higher inflation (in a classical macro model). Thus, higher inflation is associated with (but does not cause!) lower stock prices; which is what we find empirically.

(f) Compounding

Consider the calculation of two different types of "portfolio" returns: (1) n different assets held for one period; (2) one asset held over n periods.

In the first case, no compounding occurs since the return on asset i cannot be reinvested in asset j during the one period. Thus the portfolio return should be calculated as the *arithmetic* average:

$$(10) \qquad r_t(n \text{ assets}) = \sum_{i=1}^{n} s_{it}\, r_{it} \ ,$$

where the s_{it} indicate the fraction of total wealth invested in asset i during period t . Thus the portfolio return is found as a weighted arithmetic average.

In the second case, the return on the asset for a period is assumed to be fully reinvested in the same asset in the subsequent period. If no reinvestment takes place (and the returns are kept under the mattress) then it is appropriate to calculate the per-period return as an arithmetic average. But, in the more typical case of reinvestment in the same asset, compounding should be taken into consideration:

$$(11) \qquad 1 + r_i(n \text{ periods}) = \left(\prod_{t=1}^{n} (1 + r_{it}) \right)^{1/n} .$$

Here the per-period return on asset i is found as the geometric average of the returns in each of the n periods. The rationale for this definition is, of course, that receiving this geometric average n periods in a row yields exactly the same final wealth as receiving the actual returns in each period.

If the returns are fairly small then the two approaches yield very similar results. Intuitively, this makes sense since compounding is a second-order effect (interest upon interest) which becomes negligible if the first-order effects are small. Mathematically, we often make logarithmic approximations, taking advantage of the fact that $\log\,(1 + x) \approx x$ for small x. Applying this approximation to equation (11), and taking advantage of several of the properties of logs, yields:

$$(12) \qquad r_i(n \text{ periods}) \approx (1/n) \sum_{t=1}^{n} r_{it} \ .$$

Equation (12) gives of course an equal-weighted version of the result in equation (10).

The compounding in equation (11) is discrete: the return is paid only at the end of each period. In many cases returns trickle in on a fairly continuous basis. And banks, for instance, should pay and receive interest on a daily basis (pretty much continuous compared to the semi-annual coupon payments made by most bonds) since one of their opportunities is buying or selling federal funds, which provide daily interest.

Consider an interest payment r to be made once at the end of a period. It will provide a total (gross) return of

$1 + r$ at the end of the period. However, if this payment is paid semi-periodically, i.e., in two equal parts, and the early payment is reinvested, then the gross payoff at the end of the period becomes $[1 + (r/2)]^2$. More generally, if interest is paid n times during the period, the gross payoff becomes $[1 + (r/z)]^z$. If the same z-equal- payments scenario holds in each of n periods, we get $[1 + (r/z)]^{zn}$.

Continuous compounding means that we have to take the limit of z going to infinity. Employing the definition of e, this yields the identity:

$$\text{(13)} \qquad \lim_{z \to \infty} [1 + (r/z)]^{zn} = e^{rn} .$$

Both sides of equation (13) indicate of course the total return. If one wants to know the average return per period, it is necessary to take the n-th root on both sides, yielding e^r for the right-hand term.

If we compare the one-period discrete compounding gross return $1 + r_d$ to the one-period continuous compounding gross return e^r, it should be clear that in order for both end-of-period gross returns to be equal, $1 + r_d = e^r$, we need $r_d > r$. Additionally, if we take the log of both sides of the previous equation, we get

$$\text{(14)} \qquad r = \log(1 + r_d) \approx r_d .$$

Taking logs provides an approximation for the discrete-compounding return but an exact expression for the continuous compounding return.

A further advantage of continuous compounding is that, under some circumstances, it allows one to work in log terms, which in turn allow linear expressions. Consider for instance a case where all dividends earned on a portfolio are automatically reinvested in the same portfolio. The value at time t of the portfolio , with dividends reinvested, is given as p_t^{div}. The gross return then is equal to p_{t+1}^{div}/p_t^{div}. Given that the return is continuously compounded, we have:

$$\text{(15)} \qquad e^{r_{t+1}^c} = p_{t+1}^{div}/p_t^{div} \quad \leftrightarrow \quad r_{t+1}^c = \ln p_{t+1}^{div} - \ln p_t^{div},$$

where the latter inequality follows by taking logs on both sides of the first inequality. Many data sets express stock prices or stock-index prices as prices "with dividends reinvested"; for such data (continuously compounded) returns are calculated most easily by subtracting the log of the previous price observations from the log of the current price observation.

Lastly, note that a drawback of using continuously compounded returns is that calculating weighted average portfolio returns is not straightforward. While for discrete compounding one may use equation (10) above, for continuous compounding there is no easy expression akin to equation (10) for cross-sectional portfolio returns. Using a simple equal-weighted two asset example:

$$\text{(16)} \qquad e^{r_t^c} = (1/2)\, e^{r_{1t}^c} + (1/2)\, e^{r_{2t}^c}; \quad \text{but} \quad r_t^c \neq (1/2)\, r_{1t}^c + (1/2)\, r_{2t}^c .$$

Thus, continuous compounding may be more convenient in calculating time series returns, while discrete compounding may be easier in calculating cross-sectional returns.

As an example, we can apply continuous compounding to the Gordon Growth Model where we now assume that a fixed payment X is paid continuously through time. The value of this asset then is given as:

$$\int_0^\infty e^{-rn}(e^{gn}X)\,dn = X\int_0^\infty e^{(g-r)n}\,dn = \frac{X}{g-r}[e^{(g-r)n}]|_0^\infty = \frac{X}{r-g}\,. \tag{17}$$

Here we obtain essentially the same result as in the discrete-compounding case except that all constants need to be interpreted as continuous-compounding-based.

(g) Remaining Issues in Time Discounting

Even for deterministic cash flows, the previous discussion has not exhausted the issues that may arise in valuing payments that are delayed.

One issue is liquidity. There should be a substantial difference in valuing a future cash flow of an asset that is highly marketable as compared to an asset that could not be sold until the payment is due. For the latter type of asset (possibly your own human capital before obtaining your degree, or software in development), the liquidation value is important. In case of a personal emergency, when the project has to be abandoned for whatever reason, its liquidation value would matter.

As a partial answer to this liquidity issue, consider first the opportunity cost of holding a perfectly marketable asset. As the asset is riskless and could be sold for true value on short notice, the proper opportunity cost is the rate of interest of a short-term riskless bond, such as a T-Bill. On the other hand, the opportunity cost of holding an asset whose intermediate value may vary is more like the rate of return on a long-term riskless bond, such as a Government Bond. Based on the term structure of interest rates literature, the rate on a long-term bond usually exceeds that on a short-term bond, so that the less liquid asset will be worth less.

A further issue is related to variability of the opportunity cost, the discount rate. If the variability is known or predictable, it is not a big issue, the discount factor will just differ by period. An uncertain opportunity cost, however, is a different matter, especially if the opportunity cost is correlated with the cash flow. It seems then that, our simple time-discounting method requires not only known cash flows but also a known pattern of the opportunity cost in order to be strictly valid. The impact of uncertainty on discounting cash flows is introduced next and is further examined in Chapter V.

(h) Applications and Exercises

1. Consider the Gordon Growth Model with a finite stream of cash flows. Specifically, assume that cash flows grow at a constant rate until time T and are zero after time T.

   (a) Derive the Present Value of the stream of payments under discrete compounding.
   (b) Derive the Present Value of the stream of payments under continuous compounding.

2. Derive the Gordon Growth Model. Explain how inflation is taken into consideration in this model. Assume that projected inflation is 3%, projected stock returns are 10% and the growth rate of earnings is projected to be 5%. Calculate what the price/earnings ratio should be for an average stock according to the Gordon Growth Model.

3. Prove that the *n*-period discount factor at time t, β_t^n [defined as in equation (1) but with $c_t = c_{t+i}$ for all $0 \le i \le n-1$], is equal to the product of one-period discount factors: $\prod_{i=0}^{n-1} \beta_{t+i}$.

3. ACCOUNTING FOR RISK

(a) One-Sided and Two-Sided Risk

The standard use of the concept of risk as used by finance academicians relates to a *two-sided* risk: it considers the impact of losses as well as gains, relative to expectation. This concept of risk was first formalized precisely by Rothschild and Stiglitz (1970). For a risky cash flow, they define an *increase in risk* as adding independent noise to the cash flow for a *given expected value* of the cash flow, which, they show, is equivalent to putting more weight in the tails of the cash flow distribution. They prove that any risk averse investor (that is, one who has concave utility over consumption) will dislike such an increase in risk. (In fact they show that, *only* for such a definition of risk, any risk averse investor dislikes more risk)[3].

The Rothschild-Stiglitz definition of risk is an intellectually satisfying one. However, it is of little practical use in pricing assets for the following reasons. First, it provides only a partial ordering of risky prospects. That is, many risks cannot be compared in this way since one can not separate out an independent noise by which these risks differ. Thus, some investors would prefer the one risk, some would prefer the other risk. Second, even if two risks can be ranked in the Rothschild-Stiglitz way, this is only an ordinal ranking and it is difficult to put a cardinal number on the difference for pricing purposes. Third, this general definition of risk is individual based and ignores the Fisher Separation result that implies for competitive markets. Thus, in the following we resort to more operational definitions of risk. Conceptually, however, we continue to think of risk in these general terms.

In the business community, the implicit concept of risk is often a *one-sided* one. "Putting your money at risk" means that you may lose it. It doesn't really speak to what happens when you gain. This concept of risk is associated with *default risk* (as opposed to the two-sided *market risk*) and occurs in its purest form in a risky bond or junk bond. It differs from the academic definition of risk in that it matters even for a risk-neutral individual. It further is difficult to work with since introducing one-sided risk implies a change in the mean of the payoff as well as in the variance of the

[3] This concept of risk considers an increase in payoff variability for a given mean payoff and is often referred to as a "mean preserving spread." A slight generalization is the concept of "second-order stochastic dominance": A "less risky" payoff distribution second-order stochastically dominates a "more risky" payoff distribution if all risk averse individuals prefer it. This generalizes the Rothschild-Stiglitz concept of risk since this may occur not only if the second payoff distribution is a mean preserving spread but also if the mean of the second payoff is not preserved and is lower.

payoff. In the following we always employ the two-sided concept of risk. We show next how to take one-sided “risk” into account.

Assume that the market is risk neutral so that two-sided risk is irrelevant (in exercise 1 below we drop the assumption of a risk neutral market). Imagine a firm issuing a perpetual bond with constant annual coupon payment X and a probability of bankruptcy of F. When the firm goes bankrupt no further coupon payments will be paid. What is the value of this perpetual bond? Given a risk-neutral market, all we need to do is to calculate the expected cash-flow for each period and discount by the risk free interest rate. The expected cash flow for, say, period n is easily calculated as $E(CF_n) = (1-F)^{n-1}X$. Thus, the (present) value V of the perpetual bond equals:

$$(1) \qquad V = \frac{X}{1+r} + \frac{(1-F)X}{(1+r)^2} + \frac{(1-F)^2 X}{(1+r)^3} .... = \frac{X}{r+F} .$$

Where the second equality holds pretty much along the lines of the Gordon Growth Model (but with F replacing $-g$). Equation (1) is as if the probability of default is added as part of the discount rate: you discount payments further in the future by more because of impatience and because you are less likely to receive them.

In summary, one-sided “risk” may be taken into account by simply adjusting the *expected value* of the cash flow. Clearly, a higher probability of default lowers expected cash flows by more the further in the future they are expected to be received. This affects therefore investors even in a risk neutral market. We will see that, in a realistic risk averse market, one-sided “risk” may still be riskless according to certain asset pricing models if it is uncorrelated with market returns and therefore disappears in a well-diversified portfolio. Alternatively, if we do not want to adjust the expected value of the cash flow to account for default risk, equation (1) shows that we may be able to augment the risk free discount rate with the default probability to account for one-sided risk. Here the cash flow should be interpreted of course, not as an expected cash flow, but as the cash flow assuming that nothing goes wrong.

(b) Risk Adjustment in Discounting

Suppose that we want to value a risky payment to be received one period from now. First we address the issue from a purely theoretical perspective: What compensation C should an individual receive to be indifferent between the risky cash flow X and a certain cash flow with the same expected value $E(X)$. Using the standard expected utility criterion we find C from:

$$(2) \qquad E[u(X)] = u[E(X) - C].$$

Figure 3 shows graphically how to obtain C given a two-point distribution for X. Note that the risk introduced here is a specific example of a Rothschild-Stiglitz increase in risk. Once we have the value for C we could obtain the appropriate discount rate μ_X for X from the following:

$$(3) \qquad \frac{E(X)}{1+\mu_X} = \frac{E(X) - C}{1+r} .$$

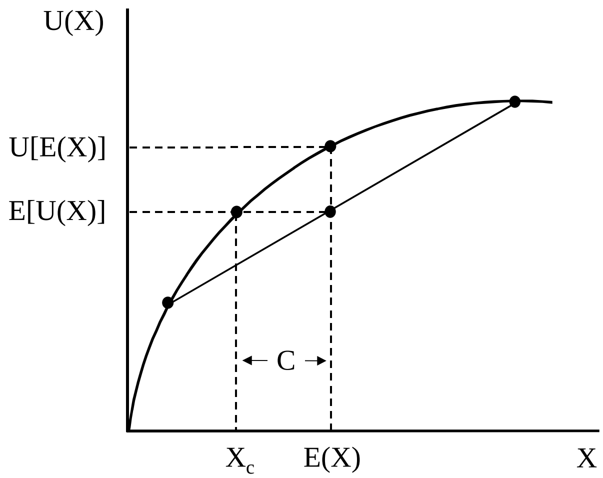

Figure 3
Risk Adjustment
For a two-point distribution, the amount, X_c, is the certainty equivalent amount. The risk adjustment or insurance premium needed to obtain X_c is given by the interval C.

We then find μ_X directly from equation (3) as:

$$\text{(4)} \qquad \mu_X = r + \frac{C(1+r)}{E(X) - C},$$

where the term after the "+" sign is the risk premium in the "required return" for *X*.

Adjusting cash flows for risk in this manner is subject to similar problems as the use of the Rothschild-Stiglitz definition of risk for this purpose: It is difficult to apply and it ignores the Fisher Separation result. Notice for instance that *C* will differ by individual, and, even for identical individuals, depends generally on initial wealth. Next we consider the standard market-based approach for discounting risky cash flows.

A market-based approach is based on the celebrated CAPM (Capital Asset Pricing Model) which we discuss extensively in the next chapter. According to the CAPM the expected return on any asset *i* is determined as

$$\text{(5)} \qquad \mu_i = r + \lambda \operatorname{Cov}(r_i, r_m),$$

where the covariance term indicates the covariance between the return on asset *i* and the return on the overall market. It reflects the fact that only "systematic" risk that is correlated with the market matters since other "non-systematic" risk may be diversified away. Further, λ indicates the "price of risk" which is determined at the market level.

One may now use the "arbitrage" principle and the opportunity cost idea to price the cash flow X: if one were to invest amount *I* in asset *X'*, with known expected return μ_X, the expected forward value would be:

$E(FV_{X'}) = I\,(1 + \mu_X)$. Assuming that X' has the same risk characteristics as asset X we can then set the required return for X equal to the expected return for X' : $\mu_{X'} = \mu_X$. Then equating the expected forward value of X' to the expected cash flow of X we obtain the (present) value of asset X as:

$$(6) \qquad V_X = \frac{E(X)}{1 + \mu_X} \ .$$

Thus, discounting should be accomplished by using the expected return on a "similar" asset as the discount rate. Of course, what is similar depends on the particular asset pricing model. For the CAPM, "similar" means for the returns of both assets to have the same covariance with the market. If we apply the CAPM equation in (5) to equation (6) we can obtain the Certainty Equivalence expression for the value of asset X, as follows.

Bear in mind first that the return on asset X is defined as $r_X = X/V_X$ given that the asset is priced correctly. Then we can write $\mathrm{Cov}(r_X, r_m) = \mathrm{Cov}(X, r_m)\,/\,V_X$. Applying equation (5) to equation (6) and solving for V_X based on this covariance expression we obtain:

$$(7) \qquad V_X = \frac{E(X) - \lambda\,\mathrm{Cov}(X, r_m)}{1 + r} \ .$$

Thus, one may account for risk by using the appropriate "opportunity cost" discount rate. Or, under the assumptions of the CAPM, one may adjust the expected cash flow for the systematic risk inherent in the cash flow to yield the certainty equivalent value of the expected cash flow and then discount using the risk free rate.

There are aspects of accounting for risk in multiple-period present value calculations that we have not yet addressed and which are quite complicated, even in the context of the CAPM. These aspects are related to cross-period correlations in cash flows and to uncertainty about discount rates. We deal with these issues in Chapter V (addressing results of Fama (1976) and Constantinides (1980)) and further where we treat the issue of intertemporal asset pricing.

(c) Applications and Exercises

1. Consider the stock with value V of a company that each year faces the same earnings distribution X. The earnings are all paid to the stock holders. There is a probability F in each year that the company fails, in which case of course its stock becomes valueless.

   (a) Explain intuitively that the value of the stock can be written as:

   $$V = \frac{E(X) + (1 - F)V}{1 + \mu_X}.$$

   (b) Compare the expression of V here to equation (1).

   (c) Assuming that the CAPM applies, solve for V and obtain the Certainty Equivalence form of adjusting for risk.