CHAPTER 5

RISK ADJUSTED VALUE

Risk-averse investors will assign lower values to assets that have more risk associated with them than to otherwise similar assets that are less risky. The most common way of adjusting for risk to compute a value that is risk adjusted. In this chapter, we will consider four ways in which we this risk adjustment can be made. The first two approaches are based upon discounted cash flow valuation, where we value an asset by discounting the expected cash flows on it at a discount rate. The risk adjustment here can take the form of a higher discount rate or as a reduction in expected cash flows for risky assets, with the adjustment based upon some measure of asset risk. The third approach is to do a post-valuation adjustment to the value obtained for an asset, with no consideration given for risk, with the adjustment taking the form of a discount for potential downside risk or a premium for upside risk. In the final approach, we adjust for risk by observing how much the market discounts the value of assets of similar risk.

While we will present these approaches as separate and potentially self-standing, we will also argue that analysts often employ combinations of approaches. For instance, it is not uncommon for an analyst to estimate value using a risk-adjusted discount rate and then attach an additional discount for liquidity to that value. In the process, they often double count or miscount risk.

Discounted Cash Flow Approaches

In discounted cash flow valuation, the value of any asset can be written as the present value of the expected cash flows on that asset. Thus, the value of a default free government bond is the present value of the coupons on the bond, discounted back at a riskless rate. As we introduce risk into the cash flows, we face a choice of how best to reflect this risk. We can continue to use the same expected cash flows that a risk-neutral investor would have used and add a risk premium to the riskfree rate to arrive at a risk-adjusted discount rate to use in discounting the cash flows. Alternatively, we can continue to use the risk free rate as the discount rate and adjust the expected cash flows for risk; in effect, we replace the uncertain expected cash flows with certainty equivalent cash flows.

The DCF Value of an Asset

We buy most assets because we expect them to generate cash flows for us in the future. In discounted cash flow valuation, we begin with a simple proposition. The value of an asset is not what someone perceives it to be worth but is a function of the expected cash flows on that asset. Put simply, assets with predictable cash flows should have higher values than assets with volatile cash flows. There are two ways in which we can value assets with risk:

- The value of a risky asset can be estimated by discounting the expected cash flows on the asset over its life at a risk-adjusted discount rate:

$$\text{Value of asset} = \frac{E(CF_1)}{(1+r)} + \frac{E(CF_2)}{(1+r)^2} + \frac{E(CF_3)}{(1+r)^3} ..... + \frac{E(CF_n)}{(1+r)^n}$$

  where the asset has a n-year life, $E(CF_t)$ is the expected cash flow in period t and r is a discount rate that reflects the risk of the cash flows.

- Alternatively, we can replace the expected cash flows with the guaranteed cash flows we would have accepted as an alternative (certainty equivalents) and discount these certain cash flows at the riskfree rate:

$$\text{Value of asset} = \frac{CE(CF_1)}{(1+r_f)} + \frac{CE(CF_2)}{(1+r_f)^2} + \frac{CE(CF_3)}{(1+r_f)^3} ..... + \frac{CE(CF_n)}{(1+r_f)^n}$$

  where $CE(CF_t)$ is the certainty equivalent of $E(CF_t)$ and r_f is the riskfree rate.

The cashflows will vary from asset to asset -- dividends for stocks, coupons (interest) and the face value for bonds and after-tax cashflows for a investment made by a business. The principles of valuation do not.

Using discounted cash flow models is in some sense an act of faith. We believe that every asset has an intrinsic value and we try to estimate that intrinsic value by looking at an asset's fundamentals. What is intrinsic value? Consider it the value that would be attached to an asset by an all-knowing analyst with access to all information available right now and a perfect valuation model. No such analyst exists, of course, but we all aspire to be as close as we can to this perfect analyst. The problem lies in the fact that none of us ever gets to see what the true intrinsic value of an asset is and we therefore have no way of knowing whether our discounted cash flow valuations are close to the mark or not.

Risk Adjusted Discount Rates

Of the two approaches for adjusting for risk in discounted cash flow valuation, the more common one is the risk adjusted discount rate approach, where we use higher discount rates to discount expected cash flows when valuing riskier assets, and lower discount rates when valuing safer assets.

Risk and Return Models

In the last chapter, we examined the development of risk and return models in economics and finance. From the capital asset pricing model in 1964 to the multi-factor models of today, a key output from these models is the expected rate of return for an investment, given its risk. This expected rate of return is the risk-adjusted discount rate for the asset's cash flows. In this section, we will revisit the capital asset pricing model, the arbitrage-pricing model and the multi-factor model and examine the inputs we need to compute the required rate of return with each one.

In the capital asset pricing model, the expected return on an asset is a function of its beta, relative to the market portfolio.

Expected Return = Riskfree Rate + Market Beta * Equity Risk Premium

There are two inputs that all assets have in common in risk and return models. The first is the riskfree rate, which is the rate of return that you can expect to make with certainty on an investment. This is usually measured as the current market interest rate on a default-free (usually Government) security; the U.S. Treasury bond rate or bill rate is used as the long term or short-term riskfree rate in U.S. dollars. It is worth noting that the riskfree rate will vary across currencies, since the expected inflation rate is different with each currency. The second is the equity risk premium, which can be estimated in one of two ways. The first is a historical risk premium, obtained by looking at returns you would have earned on stocks, relative to a riskless investment, and the other is to compute a forward-looking or implied premium by looking at the pricing of stocks, relative to the cash flows you expect to get from investing in them. In chapter 3, we estimated both for the U.S. market and came up with 4.80% for the former and 4.09% for the latter in early 2006, relative to the treasury bond rate. The only risk parameter that is investment-specific is the beta, which measures the covariance of the investment with the market portfolio. In practice, it is estimated by other regressing returns on the investment (if it is publicly traded) against returns on a market index, or by looking at the betas of other publicly traded firms in the same business. The latter is called a bottom-up beta and generally yields more reliable estimates than a historical regression beta, which, in addition to being backward looking, also yields betas with large error terms. Appendix 5.1 provides a more detailed description of the steps involved in computing bottom-up betas.

Consider a simple example. In January 2006, the ten-year treasury bond rate in the United States was 4.25%. At that time, the regression beta for Google was 1.83, with a standard error of 0.35, and the bottom-up beta for Google, looking at other internet firms was 2.25. If we accept the latter as the best estimate of the beta, the expected return on Google stock, using the implied risk premium of 4.09%, would have been:

Expected return on Google = 4.25% + 2.25 (4.09%) = 13.45%

If you were valuing Google's equity cash flows, this would have been the risk adjusted discount rate that you would have used.[1]

The arbitrage pricing and multi-factor models are natural extensions of the capital asset pricing model. The riskfree rate remains unchanged, but risk premiums now have to be estimated for each factor; the premiums are for the unspecified market risk factors in the arbitrage pricing model and for the specified macro economic risk factors in the multi-factor models. For individual investments, the betas have to be estimated, relative to each factor, and as with the CAPM betas, they can come from examining historical returns data on each investment or by looking at betas that are typical for the business that the investment is in.

---

[1] When firms are funded with a mix of equity and debt, we can compute a consolidated cost of capital that is weighted average of the cost of equity (computed using a risk and return model) and a cost of debt (based upon the default risk of the firm). To value the entire business (rather than just the equity), we would discount the collective cashflows generated by the business for its equity investors and lenders at the cost of capital.

As we noted in chapter 4, the risk and return models in use share the common assumption of a marginal investor who is well diversified and measure risk as the risk added on to a diversified portfolio. They also share a common weakness insofar as they make simplifying assumptions about investor behavior – that investors have quadratic utility functions, for instance- or return distributions – that returns are log-normally distributed. They do represent a convenient way of adjusting for risk and it is no surprise that they are in the toolboxes of most analysts who deal with risky investments.

Proxy Models

In chapter 4, we examined some of the variables that have historically characterized stocks that have earned high returns: small market capitalization and low price to book ratios are two that come to mind. We also highlighted the findings of Fama and French, who regressed returns on stocks against these variables, using data from 1963 to 1990, to arrive at the following result for monthly returns:

$$\text{Return}_j = 1.77\% - 0.11 \ln\left(\text{MV}_j\right) + 0.35 \ln\left(\frac{\text{BV}_j}{\text{MV}_j}\right)$$

where

Return_j = Monthly Return on company j

$\ln(\text{MV}_j)$ = Natural log of the Market Value of Equity of company j

ln(BV/MV) = Natural log of ratio of Book Value to Market Value of Equity

Plugging in a company's market value and book to price ratio into this equation will generate an expected return for that investment, which, in turn, is an estimate of the risk-adjusted discount rate that you could use to value it. Thus, the expected monthly return for a company with a market value of equity of \$ 500 million and a book value of equity of \$ 300 million can be written as:

Expected Monthly Return = 1.77% -0.11 ln(500) + 0.35 ln (300/500) = 0.9076%

Annualized, this would translate into an expected annual return of 11.45%:

Expected Annual Return = $(1.009076)^{12}-1$ = .1145 or 11.45%

This would be the risk-adjusted discount rate that you would use the value the company's cash flows (to equity investors).

In recent years, there have been other variables that have been added to proxy models. Adding price momentum, price level and trading volume have been shown to improve the predictive power of the regression; strong stock price performance in the last six months, low stock price levels and low trading volume are leading indicators of high returns in the future.

Proxy models have developed a following among analysts, especially those whose primary focus is valuing companies. Many of these analysts use an amalgam of risk and return models and proxy models to generate risk-adjusted discount rates to use in valuing stocks; for instance, the CAPM will be used to estimate an expected return for a small company and a small-stock premium (usually based upon historical return premium earned by small stocks relative to the market index) is added on to arrive at the “”right” discount rate for a small company. The approach has been less useful for those who are called upon to analyze either real or non-traded investments, since the inputs to the model (market capitalization and price to book ratio) require a market price.

Implied Discount Rates

For assets that are traded in the market, there is a third approach that can be used to estimate discount rates. If we are willing to make estimates of the expected cash flows on the asset, the risk-adjusted discount rate can be backed out of the market price. Thus, if an asset has a market value of \$ 1000, expected cash flow next year of \$100 and a predicted growth rate of 3% in perpetuity, the risk-adjusted discount rate implied in the price can be computed as follows:

Market Value = Expected cash flow next year/ (Risk adjusted Discount Rate – Growth)

1000 = 100/(r - .03)

Solving for r, we obtain a risk-adjusted discount rate of 13%.

While the implied discount rate does away with the requirements of making assumptions about investor utility and return distributions of the risk and return models, and the dependence on historical patterns underlying the proxy models, it has two critical flaws that have prevented its general usage:

1. It requires that the investment be traded and have a market price. Thus, it cannot be used without substantial modification for a non-traded asset.
2. Even if the asset has a market price, this approach assumes that the market price is correct. Hence, it becomes useless to an analyst who is called upon to make a

judgment on whether the market price is correct; put another way, using the implied discount rate to value any risky asset will yield the not surprising conclusion that everything is always fairly priced.

There are interesting ways in which practitioners have got around these problems. One is to compute implied risk adjusted discount rates for every asset in a class of risky assets – all cement companies, for example – and to average the rate across the assets. Implicitly, we are assuming that the assets all have equivalent risk and that they should therefore all share the same average risk-adjusted rate of return. The other is to compute risk-adjusted discount rates for the same asset for each year for a long period and to average the rate obtained over the period. Here, the assumption is that the risk adjusted discount rate does not change over time and that the average across time is the best estimate of the risk adjusted rate today.

General Issues

While the use of risk adjusted discount rates in computing value is widespread in both business valuation and capital budgeting, there are a surprising number of unresolved or overlooked issues in their usage.

a. Single period models and Multi period projects: The risk and return models that we usually draw upon for estimating discount rates such as the CAPM or the APM are single period models, insofar as they help you forecast expected returns for the next period. Most assets have cash flows over multiple periods and we discount these cash flows at the single period discount rate, compounded over time. In other words, when we estimate the risk-adjusted return at Google to be 13.45%, it is an expected return for the next year. When valuing Google, we discount cash flows in years 2, 3 and beyond using the same discount rate. Myers and Turnbull (1977) note that this is appropriate only if we assume that the systematic risk of the project (its beta in the CAPM) and the market risk premium do not change over time.[2] They also go on to argue that this assumption will be violated when a business or asset has growth potential, since the systematic risk (beta) of growth is likely to be higher than the systematic risk of investments already made and that this

---

[2] Myers, S.C. and S.M. Turnbull, 1977, Capital Budgeting and the Capital Asset Pricing Model: Good News and Bad New, Journal of Finance, v32, 321-333.

will cause the systematic risk of an asset to change over time. One approximation worth considering in this scenario is to change the risk-adjusted discount rate each period to reflect changes in the systematic risk.

b. Composite Discount Rate versus Item-specific discount rate: In most discounted cash flow valuations, we estimate the expected cash flows of the asset by netting all outflows against inflows and then discount these cash flows using one risk adjusted cost of capital. Implicitly, we are assuming that all cash flow items have equivalent exposure to systematic risk, but what if this assumption is not true? We could use different risk-adjusted discount rates for each set of cash flows; for instance, revenues and variable operating expenses can be discounted at the cost of capital whereas fixed operating expenses, where the firm may have pre-committed to making the payments, can be discounted back at a lower rate (such as the cost of debt). The question, though, is whether the risk differences are large enough to make a difference. At the minimum, the one or two cash flow items that diverge most from the average risk assumption (underlying the risk adjusted cost of capital) can be separately valued.

c. Negative versus Positive Cash flows: Generally, we penalize riskier assets by increasing the discount rate that we use to discount the cash flows. This pre-supposes that the cash flows are positive. When cash flows are negative, using a higher discount rate will have the perverse impact of reducing their present value and perhaps increasing the aggregate value of the asset. While some analysts get around this by discounting negative cash flows at the riskfree rate (or a low rate variant) and positive cash flows at the risk adjusted discount rate, they are being internally inconsistent in the way they deal with risk. In our view, any value benefits that accrue from discounting negative cash flows at the risk adjusted rate will be more than lost when the eventual positive cash flows are discounted back at the same risk adjusted rate, compounded over time. Consider, for instance, a growth business with negative cash flows of \$ 10 million each year for the first 3 years and a terminal value of \$ 100 million at the end of the third year. Assume that the riskfree rate is 4% and the risk-adjusted discount rate is 10%. The value of the firm using the riskfree rate for the first 3 years and the risk-adjusted rate only on the terminal value is as follows:

$$\text{Value of firm} = \frac{-10}{(1.04)^1} + \frac{-10}{(1.04)^2} + \frac{-10}{(1.04)^3} + \frac{100}{(1.04)^3} = 61.15$$

Note that the terminal value is being discounted back at the riskfree rate for 3 years.[3] In contrast, the value of the same firm using the risk-adjusted discount rate on all of the cash flows is as follows:

$$\text{Value of firm} = \frac{-10}{(1.10)^1} + \frac{-10}{(1.10)^2} + \frac{-10}{(1.10)^3} + \frac{100}{(1.10)^3} = 50.26$$

Put another way, it is reasonable to discount back negative cash flows at a lower rate, if they are more predictable and stable, but not just because they are negative.

Certainty Equivalent Cashflows

While most analysts adjust the discount rate for risk in DCF valuation, there are some who prefer to adjust the expected cash flows for risk. In the process, they are replacing the uncertain expected cash flows with the certainty equivalent cashflows, using a risk adjustment process akin to the one used to adjust discount rates.

Misunderstanding Risk Adjustment

At the outset of this section, it should be emphasized that many analysts misunderstand what risk adjusting the cash flows requires them to do. There are analysts who consider the cash flows of an asset under a variety of scenarios, ranging from best case to catastrophic, assign probabilities to each one, take an expected value of the cash flows and consider it risk adjusted. While it is true that bad outcomes have been weighted in to arrive at this cash flow, it is still an expected cash flow and is not risk adjusted. To see why, assume that you were given a choice between two alternatives. In the first one, you are offered $ 95 with certainty and in the second, you will receive $ 100 with probability 90% and only $50 the rest of the time. The expected values of both alternatives is $95 but risk averse investors would pick the first investment with guaranteed cash flows over the second one.

If this argument sounds familiar, it is because it is a throwback to the very beginnings of utility theory and the St. Petersburg paradox that we examined in chapter 2.

In that chapter, we unveiled the notion of a certainty equivalent, a guaranteed cash flow that we would accept instead of an uncertain cash flow and argued that more risk averse investors would settle for lower certainty equivalents for a given set of uncertain cash flows than less risk averse investors. In the example given in the last paragraph, a risk averse investor would have settled for a guaranteed cash flow of well below $95 for the second alternative with an expected cash flow of $95.

The practical question that we will address in this section is how best to convert uncertain expected cash flows into guaranteed certainty equivalents. While we do not disagree with the notion that it should be a function of risk aversion, the estimation challenges remain daunting.

Utility Models: Bernoulli revisited

In chapter 2, we introduced the first (and oldest) approach to computing certainty equivalents, rooted in the utility functions for individuals. If we can specify the utility function of wealth for an individual, we are well set to convert risky cash flows to certainty equivalents for that individual. For instance, an individual with a log utility function would have demanded a certainty equivalent of $79.43 for the risky gamble presented in the last section (90% chance of $ 100 and 10% chance of $ 50):

Utility from gamble = .90 ln(100) + .10 ln(50) = 4.5359

Certainty Equivalent = $\exp^{4.5359}$ = $93.30

The certainty equivalent of $93.30 delivers the same utility as the uncertain gamble with an expected value of $95. This process can be repeated for more complicated assets, and each expected cash flow can be converted into a certainty equivalent.[4]

One quirk of using utility models to estimate certainty equivalents is that the certainty equivalent of a positive expected cash flow can be negative. Consider, for instance, an investment where you can make $ 2000 with probability 50% and lose $ 1500 with probability 50%. The expected value of this investment is $ 250 but the

---

[3] There are some who use the risk adjusted rate only on the terminal value but that would be patently unfair since you would be using two different discount rates for the same time periods. The only exception would be if the negative cash flows were guaranteed and the terminal value was uncertain.

[4] Gregory, D.D., 1978, Multiplicative Risk Premiums, Journal of Financial and Quantitative Analysis, v13, 947-963. This paper derives certainty equivalent functions for quadratic, exponential and gamma distributed utility functions and examines their behavior.

certainty equivalent may very well be negative, with the effect depending upon the utility function assumed.

There are two problems with using this approach in practice. The first is that specifying a utility function for an individual or analyst is very difficult, if not impossible, to do with any degree of precision. In fact, as we noted in chapter 3, most utility functions that are well behaved (mathematically) do not seem to explain actual behavior very well. The second is that, even if we were able to specify a utility function, this approach requires us to lay out all of the scenarios that can unfold for an asset (with corresponding probabilities) for every time period. Not surprisingly, certainty equivalents from utility functions have been largely restricted to analyzing simple gambles in classrooms.

Risk and Return Models

A more practical approach to converting uncertain cash flows into certainty equivalents is offered by risk and return models. In fact, we would use the same approach to estimating risk premiums that we employed while computing risk adjusted discount rates but we would use the premiums to estimate certainty equivalents instead.

Certainty Equivalent Cash flow = Expected Cash flow/ (1 + Risk Premium in Risk-adjusted Discount Rate)

Consider the risk-adjusted discount rate of 13.45% that we estimated for Google in early 2006:

Expected return on Google = 4.25% + 2.25 (4.09%) = 13.45%

Instead of discounting the expected cash flows on the stock at 13.45%, we would decompose the expected return into a risk free rate of 4.25% and a compounded risk premium of 8.825%.[5]

$$\text{Compounded Risk Premium} = \frac{(1+\text{ Risk adjusted Discount Rate})}{(1+\text{ Riskfree Rate})} - 1 = \frac{(1.1345)}{(1.0425)} - 1 = .08825$$

If the expected cash flow in years 1 and 2 are $ 100 million and $ 120 million respectively, we can compute the certainty equivalent cash flows in those years:

---

[5] A more common approximation used by many analysts is the difference between the risk adjusted discount rate and the risk free rate. In this case, that would have yielded a risk premium of 9.2% (13.45% - 4.25% = 9.20%)

Certainty Equivalent Cash flow in year 1 = \$ 100 million/1.08825 = \$ 91.89 million

Certainty Equivalent Cash flow in year 2 = \$120 million/ 1.08825^2 = \$ 101.33 million

This process would be repeated for all of the expected cash flows and it has two effects. Formally, the adjustment process for certainty equivalents can be then written more formally as follows (where the risk adjusted return is r and the riskfree rate is r_f:[6]

$$CE\ (CF_t) = \alpha_t\ E(CF_t) = \frac{(1+r_f)^t}{(1+r\ )^t} E(CF_t)$$

This adjustment has two effects. The first is that expected cash flows with higher uncertainty associated with them have lower certainty equivalents than more predictable cash flows at the same point in time. The second is that the effect of uncertainty compounds over time, making the certainty equivalents of uncertain cash flows further into the future lower than uncertain cash flows that will occur sooner.

Cashflow Haircuts

A far more common approach to adjusting cash flows for uncertainty is to "haircut" the uncertain cash flows subjectively. Thus, an analyst, faced with uncertainty, will replace uncertain cash flows with conservative or lowball estimates. This is a weapon commonly employed by analysts, who are forced to use the same discount rate for projects of different risk levels, and want to even the playing field. They will haircut the cash flows of riskier projects to make them lower, thus hoping to compensate for the failure to adjust the discount rate for the additional risk.

In a variant of this approach, there are some investors who will consider only those cashflows on an asset that are predictable and ignore risky or speculative cash flows when valuing the asset. When Warren Buffet expresses his disdain for the CAPM and other risk and return models and claims to use the riskfree rate as the discount rate, we suspect that he can get away with doing so because of a combination of the types of companies he chooses to invest in and his inherent conservatism when it comes to estimating the cash flows.

---

[6] This equation was first derived in a paper in 1966: Robichek, A.A. and S. C. Myers, 1966, Conceptual Problems in the Use of Risk Adjusted Discount Rates, Journal of Finance, v21, 727-730.

While cash flow haircuts retain their intuitive appeal, we should be wary of their usage. After all, gut feelings about risk can vary widely across analysts looking at the same asset; more risk averse analysts will tend to haircut the cashflows on the same asset more than less risk averse analysts. Furthermore, the distinction we drew between diversifiable and market risk that we drew in the last chapter can be completely lost when analysts are making intuitive adjustments for risk. In other words, the cash flows may be adjusted downwards for risk that will be eliminated in a portfolio. The absence of transparency about the risk adjustment can also lead to the double counting of risk, especially when the analysis passes through multiple layers of analysis. To provide an illustration, after the first analyst looking at a risky investment decides to use conservative estimates of the cash flows, the analysis may pass to a second stage, where his superior may decide to make an additional risk adjustment to the cash flows.

Risk Adjusted Discount Rate or Certainty Equivalent Cash Flow

Adjusting the discount rate for risk or replacing uncertain expected cash flows with certainty equivalents are alternative approaches to adjusting for risk, but do they yield different values, and if so, which one is more precise? The answer lies in how we compute certainty equivalents. If we use the risk premiums from risk and return models to compute certainty equivalents, the values obtained from the two approaches will be the same. After all, adjusting the cash flow, using the certainty equivalent, and then discounting the cash flow at the riskfree rate is equivalent to discounting the cash flow at a risk adjusted discount rate. To see this, consider an asset with a single cash flow in one year and assume that r is the risk-adjusted cash flow, r_f is the riskfree rate and RP is the compounded risk premium computed as described earlier in this section.

$$\text{Certainty Equivalent Value} = \frac{CE}{(1+r_f)} = \frac{E(CF)}{(1+RP)(1+r_f)} = \frac{E(CF)}{\frac{(1+r)}{(1+r_f)}(1+r_f)} = \frac{E(CF)}{(1+r)}$$

This analysis can be extended to multiple time periods and will still hold.[7] Note, though, that if the approximation for the risk premium, computed as the difference between the risk-adjusted return and the risk free rate, had been used, this equivalence will no longer hold. In that case, the certainty equivalent approach will give lower values for any risky asset and the difference will increase with the size of the risk premium.

Are there other scenarios where the two approaches will yield different values for the same risky asset? The first is when the risk free rates and risk premiums change from time period to time period; the risk-adjusted discount rate will also then change from period to period. Robichek and Myers, in the paper we referenced earlier, argue that the certainty equivalent approach yields more precise estimates of value in this case. The other is when the certainty equivalents are computed from utility functions or subjectively, whereas the risk adjusted discount rate comes from a risk and return model. The two approaches can yield different estimates of value for a risky asset. Finally, the two approaches deal with negative cash flows differently. The risk adjusted discount rate discounts negative cash flows at a higher rate and the present value becomes less negative as the risk increases. If certainty equivalents are computed from utility functions, they can yield certainty equivalents that are negative and become more negative as you increase risk, a finding that is more consistent with intuition.[8]

Hybrid Models

Risk-adjusted discount rates and certainty equivalents come with pluses and minuses. For some market-wide risks, such as exposure to interest rates, economic growth and inflation, it is often easier to estimate the parameters for a risk and return model and the risk adjusted discount rate. For other risks, especially those occur infrequently but can have a large impact on value, it may be easier to adjust the expected cash flows. Consider, for instance, the risk that a company is exposed to from an investment in India, China or any other large emerging market. In most periods, the investment will like an investment in a developed market but in some periods, there is the potential for major political and economic disruptions and consequent changes in value.

---

[7] The proposition that risk adjusted discount rates and certainty equivalents yield identical net present values is shown in the following paper: Stapleton, R.C., 1971, Portfolio Analysis, Stock Valuation and Capital Budgeting Decision Rules for Risky Projects, Journal of Finance, v26, 95-117.

[8] Beedles, W.L., 1978, Evaluating Negative Benefits, Journal of Financial and Quantitative Analysis, v13, 173-176.

While we can attempt to incorporate this risk into the discount rate,[9] it may be easier to adjust the cash flows for this risk, especially if the possibility of insuring against this risk exists. If so, the cost of buying insurance can be incorporated into the expenses, and the resulting cash flow is adjusted for the insured risk (but not against other risks). An alternate approach to adjusting cash flows can be used if a risk is triggered by a specific contingency. For instance, a gold mining company that will default on its debt if the gold price drops below $250 an ounce can either obtain or estimate the cost of a put option on gold, with a strike price of $250, and include the cost when computing cash flows.

The biggest dangers arise when analysts use an amalgam of approaches, where the cash flows are adjusted partially for risk, usually subjectively and the discount rate is also adjusted for risk. It is easy to double count risk in these cases and the risk adjustment to value often becomes difficult to decipher. To prevent this from happening, it is best to first categorize the risks that a project faces and to then be explicit about how the risk will be adjusted for in the analysis. In the most general terms, risks can then be categorized as follows in table 5.1.

Table 5.1: Risks: Types and Adjustment

Type of Risk	*Examples*	*Risk adjustment in valuation*
Continuous market risk where buying protection against consequences is difficult or impossible to do	Interest rate risk, inflation risk, exposure to economic cyclicality	Adjust discount rate for risk
Discontinuous market risk, with small likelihood of occurrence but large economic consequences	Political risk, Risk of expropriation, Terrorism risk	If insurance markets exist, include cost of insurance as operating expense and adjust cash flows. If not, adjust the discount rate.
Market risk that is contingent on a specific occurrence	Commodity price risk	Estimate cost of option required to hedge against risk, include as operating expense and adjust cash flows.
Firm specific risks	Estimation risk, Competitive risk,	If investors in the firm are diversified, no risk

[9] Damodaran, A., 2002, Investment Valuation, John Wiley and Sons. Several approaches for adjusting discount rates for country risk are presented in this book.

	Technology risk	adjustment needed. If investors not diversified, follow the same rules used for market risk.

We will use a simple example to illustrate the risk-adjusted discount rate, the certainty equivalent and the hybrid approaches. Assume that Disney is considering investing in a new theme park in Thailand and that table 5.2 contains the estimates of the cash flows that they believe that they can generate over the next 10 years on this investment.

Table 5.2: Expected Cash Flows form Bangkok Disney (in U.S. dollars)

Year	*Annual Cashflow*	*Terminal Value*
0	-$2,000	
1	-$1,000	
2	-$880	
3	-$289	
4	$324	
5	$443	
6	$486	
7	$517	
8	$571	
9	$631	
10	$663	$7,810

Note that the cash flows are estimated in dollars, purely for convenience and that the entire analysis could have been done in the local currency. The negative cash flows in the first 3 years represent the initial investment and the terminal value is an estimate of the value of the theme park investments at the end of the tenth year.

We will first estimate a risk-adjusted discount rate for this investment, based upon both the riskiness of the theme park business and the fact that the theme parks will be located in Thailand, thus exposing Disney to some additional political and economic risk.

Cost of capital = Risk free Rate + Business risk premium + Country Risk premium

=4% + 3.90% + 2.76% = 10.66%

The business risk premium is reflective of the non-diversifiable or market risk of being in the theme park business,[10] whereas the country risk premium reflects the risk involved in the location.[11] Appendix 1 includes a fuller description of these adjustments. The risk-adjusted value of the project can be estimated by discounting the expected cash flows at the risk-adjusted cost of capital (in table 5.3).

Table 5.3: Risk-Adjusted Value: Risk-adjusted Discount Rate approach

Year	*Annual Cashflow*	*Salvage Value*	*Present Value @10.66%*
0	-$2,000		-$2,000
1	-$1,000		-$904
2	-$880		-$719
3	-$289		-$213
4	$324		$216
5	$443		$267
6	$486		$265
7	$517		$254
8	$571		$254
9	$631		$254
10	$663	$7,810	$3,077
		Risk adjusted Value =	$751

As an alternative, lets try the certainty equivalent approach. For purposes of simplicity, we will strip the total risk premium in the cost of capital and use this number to convert the expected cash flows into certainty equivalents in table 5.4:

$$\text{Risk premium in cost of capital} = \frac{1 + \text{Risk} - \text{adjusted Cost of capital}}{1 + \text{Riskfree Rate}} - 1$$

$$= 1.1066/1.04\text{-}1 = 6.4038\%$$

Table 5.4: Certainty Equivalent Cash Flows and Risk Adjusted Value

Year	*Annual Cashflow*	*Salvage Value*	*Certainty Equivalent*	*Present value @ 4%*
0	-$2,000		-$2,000	-$2,000
1	-$1,000		-$940	-$904
2	-$880		-$777	-$719

[10] For a more detailed discussion of the computation, check Damodaran, A., 2005, Applied Corporate Finance, Second Edition, John Wiley and Sons.

[11] The additional risk premium was based upon Thailand's country rating and default spread as a country, augmented for the additional risk of equity. The details of this calculation are also in Damodaran, A., 2005, Applied Corporate Finance, Second Edition, John Wiley and Sons.

3	-$289		-$240	-$213
4	$324		$252	$216
5	$443		$324	$267
6	$486		$335	$265
7	$517		$335	$254
8	$571		$348	$254
9	$631		$361	$254
10	$663	$7,810	$4,555	$3,077
			Risk-adjusted Value=	$751

Note that the certainty equivalent cash flows are discounted back at the riskfree rate to yield the same risk-adjusted value as in the first approach. Not surprisingly, the risk-adjusted value is identical with this approach.[12]

Finally, let us assume that we could insure at least against country risk and that the after-tax cost of buying this insurance will be $150 million a year, each year for the next 10 years. Reducing the expected cash flows by the after-tax cost of insurance yields the after-tax cash flows in table 5.5.

Table 5.5: Expected Cash Flows after Insurance Payments

Year	*Annual Cashflow*	*Salvage Value*	*Insurance Payment*	*Adjusted Cashflow*	*PV @ 7.90%*
0	-$2,000		$150	-$2,150	-$2,150
1	-$1,000		$150	-$1,150	-$1,066
2	-$880		$150	-$1,030	-$885
3	-$289		$150	-$439	-$350
4	$324		$150	$174	$128
5	$443		$150	$293	$200
6	$486		$150	$336	$213
7	$517		$150	$367	$216
8	$571		$150	$421	$229
9	$631		$150	$481	$243
10	$663	$7,810	$150	$8,324	$3,891
					$670

These cash flows are discounted back at a risk-adjusted discount rate of 7.90% (i.e. without the country risk adjustment) to arrive at the present value in the last column. The risk-adjusted value in this approach of $670 million is different from the estimates in the first two approaches because the insurance market's perceptions of risk are different from those that gave rise to the country risk premium of 2.76% in the first two analyses.

DCF Risk Adjustment: Pluses and Minuses

There are good reasons why risk adjustment is most often done in a discounted cash flow framework. When the risk adjustment is made through a risk and return model, whether it is the CAPM, the arbitrage pricing model or a multi-factor model, the effect is transparent and clearly visible to others looking at the valuation. If they disagree with the computation, they can change it. In addition, the models are explicit about the risks that are adjusted for and the risks that do not affect the discount rate. In the CAPM, for instance, it is only the risks that cannot be diversified away by a well-diversified investor that are reflected in the beta.

There are, however, costs associated with letting risk and return models carry the burden of capturing the consequences of risk. Analysts take the easy way out when it comes to assessing risk, using the beta or betas of assets to measure risk and them moving on to estimate cash flows and value, secure in the comfort that they have already considered the effects of risk and its consequences for value. In reality, risk and return models make assumptions about how both markets and investors behave that are at odds with actual behavior. Given the complicated relationship between investors and risk, there is no way that we can capture the effects of risk fully into a discount rate or a cashflow adjustment.

Post-valuation Risk Adjustment

A second approach to assessing risk is to value a risky investment or asset as if it had no risk and to then adjust the value for risk after the valuation. These post-valuation adjustments usually take the form of discounts to assessed value, but there are cases where the potential for upside from risk is reflected in premiums.

---

[12] Using the approximate risk premium of 6.66% (Risk-adjusted cost of capital minus the riskfree rate) would have yielded a value of $661 million.

It is possible to adjust for all risk in the post-valuation phase – discount expected cash flows at a riskfree rate and then apply a discount to that value - but the tools that are necessary for making this adjustment are the same ones we use to compute risk-adjusted discount rates and certainty equivalents. As a consequence, it is uncommon, and most analysts who want to adjust for risk prefer to use the conventional approach of adjusting the discount rates or cash flows. The more common practice with post-valuation adjustments is for analysts to capture some of the risks that they perceive in a risk-adjusted discount rate and deal with other risks in the post-valuation phase as discounts or premiums. Thus, an analyst valuing a private company will first value it using a high discount rate to reflect its business risk, but they apply an illiquidity discount to the computed value to arrive at the final value estimate.

In this section, we will begin by looking at why analysts are drawn to the practice of post-valuation discounts and premiums and follow up by taking a closer look at some of the common risk adjustments. We will end the section by noting the dangers of what we call value garnishing.

Rationale for post-valuation adjustments

Post-valuation risk discounts reflect the belief on the part of analysts that conventional risk and return models short change or even ignore what they see as significant risks. Consider again the illiquidity discount. The CAPM and multi-factor models do not explicitly adjust expected returns for illiquidity. In fact, the expected return on two stocks with the same beta will be equal, even though one might be widely traded and liquid and the other is not. Analysts valuing illiquid assets or businesses therefore feel that they are over valuing these investments, using conventional risk and return models; the illiquidity discount is their way of bringing the estimated value down to a more "reasonable" number.

The rationale for applying post-valuation premiums is different. Premiums are usually motivated by the concern that the expected cash flows do not fully capture the potential for large payoffs in some investments. An analyst who believes that there is synergy in a merger and does not feel that the cash flows reflect this synergy will add a premium for it to the estimated value.

Downside Risks

It is not uncommon to see valuations where the initial assessments of value of a risky asset are discounted by 30% or 40% for one potential downside risk or another. In this section, we will examine perhaps the most common of these discounts – for illiquidity or lack of marketability – in detail and the dangers associated with the practice.

1. Illiquidity Discount

When you take invest in an asset, you generally would like to preserve the option to liquidate that investment if you need to. The need for liquidity arises not only because your views on the asset value change over time – you may perceive is as a bargain today but it may become over priced in the future - but also because you may need the cash from the liquidation to meet other contingencies. Some assets can be liquidated with almost no cost – Treasury bills are a good example – whereas others involve larger costs – stock in a lightly traded over-the-counter stock or real estate. With investments in a private business, liquidation cost as a percent of firm value can be substantial. Consequently, the value of equity in a private business may need to be discounted for this potential illiquidity. In this section, we will consider measures of illiquidity, how much investors value illiquidity and how analysts try to incorporate illiquidity into value.

Measuring illiquidity

You can sell any asset, no matter how illiquid it is perceived to be, if you are willing to accept a lower price for it. Consequently, we should not categorize assets into liquid and illiquid assets but allow for a continuum on liquidity, where all assets are illiquid but the degree of illiquidity varies across them. One way of capturing the cost of illiquidity is through transactions costs, with less liquid assets bearing higher transactions costs (as a percent of asset value) than more liquid assets.

With publicly traded stock, there are some investors who undoubtedly operate under the misconception that the only cost of trading is the brokerage commission that they pay when they buy or sell assets. While this might be the only cost that they pay explicitly, there are other costs that they incur in the course of trading that generally dwarf the commission cost. When trading any asset, they are three other ingredients that go into the trading costs.

- The first is the spread between the price at which you can buy an asset (the dealer's ask price) and the price at which you can sell the same asset at the same point in time (the dealer's bid price). For heavily traded stocks on the New York Stock Exchange, this cost will be small (10 cents on a $ 50 stock, for instance) but the costs will increase as we move to smaller, less traded companies. A lightly traded stock may have an ask price of $2.50 and a bid price of $ 2.00 and the resulting bid-ask spread of 50 cents will be 20% of the ask price.
- The second is the price impact that an investor can create by trading on an asset, pushing the price up when buying the asset and pushing it down while selling. As with the bid-ask spread, this cost will be highest for the least liquid stocks, where even relatively small orders can cause the price to move. It will also vary across investors, with the costs being higher for large institutional investors like Fidelity who have to buy and sell large blocks of shares and lower for individual investors.
- The third cost, which was first proposed by Jack Treynor in his article[13] on transactions costs, is the opportunity cost associated with waiting to trade. While being a patient trader may reduce the first two components of trading cost, the waiting can cost profits both on trades that are made and in terms of trades that would have been profitable if made instantaneously but which became unprofitable as a result of the waiting.

It is the sum of these costs, in conjunction with the commission costs that makes up the trading cost on an asset.

If the cost of trading stocks can be substantial, it should be even larger for assets that are not traded regularly such as real assets or equity positions in private companies.

- Real assets can range from gold to real estate to fine art and the transactions costs associated with trading these assets can also vary substantially. The smallest transactions costs are associated with commodities – gold, silver or oil – since they tend to come in standardized units and are widely traded. With residential real estate, the commission that you have to pay a real estate broker or salesperson can be 5-6% of the value of the asset. With commercial real estate, commissions may be smaller

---

[13] This was proposed in his article titled What does it take to win the trading game? published in the *Financial Analysts Journal*, January-February 1981.

for larger transactions, but they will be well in excess of commissions on financial assets. With fine art or collectibles, the commissions become even higher. If you sell a Picasso through one of the auction houses, you may have to pay15-20% of the value of the painting as a commission. Why are the costs so high? The first reason is that there are far fewer intermediaries in real asset businesses than there are in the stock or bond markets. The second is that real estate and fine art are not standardized products. In other words, one Picasso can be very different from another, and you often need the help of experts to judge value. This adds to the cost in the process.

- The trading costs associated with buying and selling a private business can range from substantial to prohibitive, depending upon the size of the business, the composition of its assets and its profitability. There are relatively few potential buyers and the search costs (associated with finding these buyers) will be high. Later in this chapter, we will put the conventional practice of applying 20-30% illiquidity discounts to the values of private businesses under the microscope.
- The difficulties associated with selling private businesses can spill over into smaller equity stakes in these businesses. Thus, private equity investors and venture capitalists have to consider the potential illiquidity of their private company investments when considering how much they should pay for them (and what stake they should demand in private businesses in return).

In summary, the costs of trading assets that are usually not traded are likely to be substantial.

Theoretical Backing for an Illiquidity Discount

Assume that you are an investor trying to determine how much you should pay for an asset. In making this determination, you have to consider the cashflows that the asset will generate for you and how risky these cashflows are to arrive at an estimate of intrinsic value. You will also have to consider how much it will cost you to sell this asset when you decide to divest it in the future. In fact, if the investor buying the asset from you builds in a similar estimate of transactions cost she will face when she sells it, the value of the asset today should reflect the expected value of all future transactions cost to all future holders of the asset. This is the argument that Amihud and Mendelson used in

1986, when they suggested that the price of an asset would embed the present value of the costs associated with expected transactions costs in the future.[14] In their model, the bid-ask spread is used as the measure of transactions costs and even small spreads can translate into big illiquidity discounts on value, if trading is frequent. The magnitude of the discount will be a function of investor holding periods and turnover ratios, with shorter holding periods and higher turnover associated with bigger discounts. In more intuitive terms, if you face a 1% bid-ask spread and you expect to trade once a year, the value of the asset today should be reduced by the present value of the costs your will pay in perpetuity. With a 8% discount rate, this will work out to roughly an illiquidity discount of 12.5% (.01/.08).

What is the value of liquidity? Put differently, when does an investor feel the loss of liquidity most strongly when holding an asset? There are some who would argue that the value of liquidity lies in being able to sell an asset, when it is most overpriced; the cost of illiquidity is not being able to do this. In the special case, where the owner of an asset has the information to know when this overpricing occurs, the value of illiquidity can be considered an option, Longstaff presents an upper bound for the option by considering an investor with perfect market timing abilities who owns an asset on which she is not allowed to trade for a period (t). In the absence of trading restrictions, this investor would sell at the maximum price that an asset reaches during the time period and the value of the look-back option estimated using this maximum price should be the outer bound for the value of illiquidity.[15] Using this approach, Longstaff estimates how much marketability would be worth as a percent of the value of an asset for different illiquidity periods and asset volatilities. The results are graphed in figure 5.1:

---

[14] Amihud, Y. and H. Mendelson, 1986, Asset Pricing and the Bid-ask Spread, *Journal of Financial Economics,* v 17, 223-250.

[15] Longstaff, F.A., 1995, How much can marketability affect security values? *Journal of Finance*, v 50, 1767-1774.

Figure 5.1: Upper bounds on Marketability Discount - Option Pricing Model

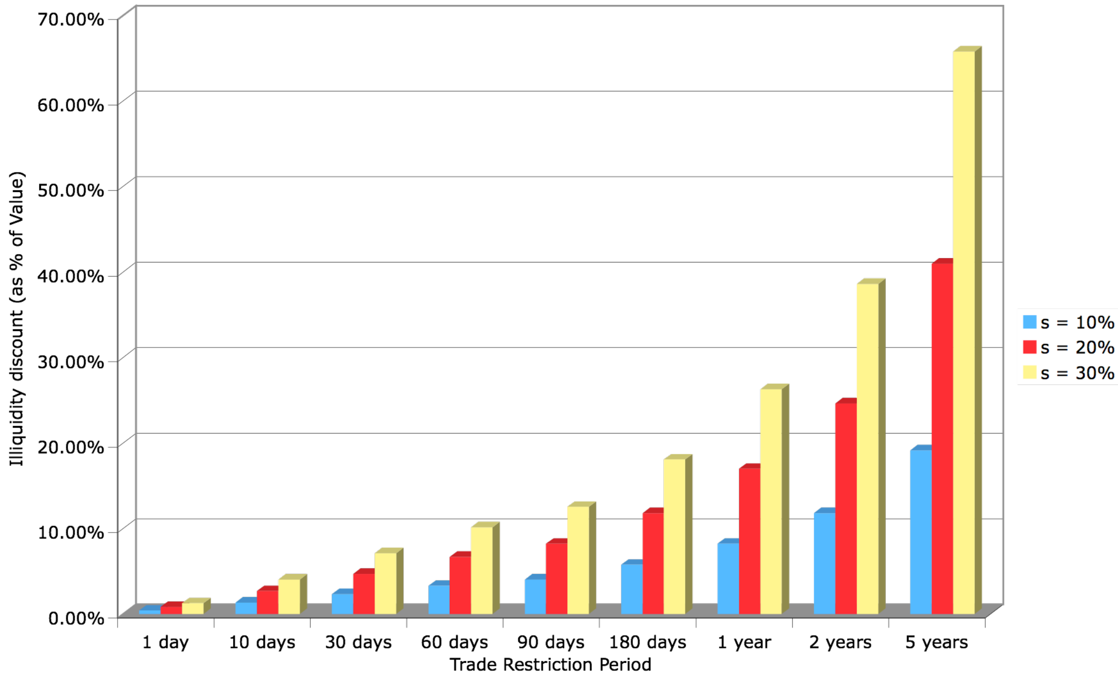

It is worth emphasizing that these are upper bounds on the value of illiquidity since it is based upon the assumption of a perfect market timer. To the extent that investors are unsure about when an asset has reached its maximum price, the value of illiquidity will be lower than these estimates. The more general lessons will still apply. The cost of illiquidity, stated as a percent of firm value, will be greater for more volatile assets and will increase with the length of the period for which trading is restricted.

Empirical Evidence that Illiquidity Matters

If we accept the proposition that illiquidity has a cost, the next question becomes an empirical one. How big is this cost and what causes it to vary across time and across assets? The evidence on the prevalence and the cost of illiquidity is spread over a number of asset classes.

a. Bond Market: There are wide differences in liquidity across bonds issued by different entities, and across maturities, for bonds issued by the same entity. These differences in liquidity offer us an opportunity to examine whether investors price liquidity and if so, how much, by comparing the yields of liquid bonds with otherwise similar illiquid bonds.

Amihud and Mendelson compared the yields on treasury bonds with less than six months left to maturity with treasury bills that have the same maturity.[16] They concluded that the yield on the less liquid treasury bond was 0.43% higher on an annualized basis than the yield on the more liquid treasury bill, a difference that they attributed to illiquidity. A study of over 4000 corporate bonds in both investment grade and speculative categories concluded that illiquid bonds had much higher yield spreads than liquid bonds. Comparing yields on these corporate bonds, the study concluded that the yield increases 0.21% for every 1% increase in transactions costs for investment grade bonds, whereas the yield increases 0.82% for every 1% increase in transactions costs for speculative bonds.[17] Looking across the studies, the consensus finding is that liquidity matters for all bonds, but that it matters more with risky bonds than with safer bonds.

b. Publicly Traded Stocks: It can be reasonably argued that the costs associated with trading equities are larger than the costs associated with trading treasury bonds or bills. It follows therefore that some of the equity risk premium, that we discussed in chapter 4, has to reflect these additional transactions costs. Jones, for instance, examines bid-ask spreads and transactions costs for the Dow Jones stocks from 1900 to 2000 and concludes that the transactions costs are about 1% lower today than they were in the early 1900s and that this may account for the lower equity risk premium in recent years.[18] Within the stock market, some stocks are more liquid than others and studies have looked at the consequences of these differences in liquidity for returns. The consensus conclusion is that investors demand higher returns when investing in more illiquid stocks. Put another way, investors are willing to pay higher prices for more liquid investments relative to less liquid investments.

c. Restricted Stocks: Much of the evidence on illiquidity discounts comes from examining "restricted stock" issued by publicly traded firms. Restricted securities are

---

[16] Amihud, Y., and H. Mendelson, 1991, Liquidity, Maturity and the Yield on U.S. Treasury Securities, *Journal of Finance*, 46, 1411-1425.

[17] Chen, L., D.A. Lesmond and J. Wei, 2005, Corporate Yield Spreads and Bond Liquidity, Working Paper, SSRN.

[18] This becomes clear when we look at forward-looking or implied equity risk premiums rather than historical risk premiums. The premiums during the 1990s averaged about 3%, whereas there were more than 5% prior to 1960. Jones, C.M., 2002, A Century of Stock Market Liquidity and Trading Costs, Working Paper, Columbia University.

securities issued by a publicly traded company, not registered with the SEC, and sold through private placements to investors under SEC Rule 144. They cannot be resold in the open market for a one-year holding period[19], and limited amounts can be sold after that. When this stock is issued, the issue price is set much lower than the prevailing market price, which is observable, and the difference can be viewed as a discount for illiquidity. The results of two of the earliest and most quoted studies that have looked at the magnitude of this discount are summarized below:

- Maher examined restricted stock purchases made by four mutual funds in the period 1969-73 and concluded that they traded an average discount of 35.43% on publicly traded stock in the same companies.[20]
- Silber examined restricted stock issues from 1981 to 1988 and found that the median discount for restricted stock is 33.75%.[21] He also noted that the discount was larger for smaller and less healthy firm, and for bigger blocks of shares.

Other studies confirm these findings of a substantial discount, with discounts ranging from 30-35%, though one recent study by Johnson did find a smaller discount of 20%.[22] These studies have been used by practitioners to justify large marketability discounts, but there are reasons to be skeptical. First, these studies are based upon small sample sizes, spread out over long time periods, and the standard errors in the estimates are substantial. Second, most firms do not make restricted stock issues and the firms that do make these issues tend to be smaller, riskier and less healthy than the typical firm. This selection bias may be skewing the observed discount. Third, the investors with whom equity is privately placed may be providing other services to the firm, for which the discount is compensation.

d. Private Equity: Private equity and venture capital investors often provide capital to private businesses in exchange for a share of the ownership in these businesses. Implicit in these transactions must be the recognition that these investments are not liquid. If private equity investors value liquidity, they will discount the value of the private

---

[19] The holding period was two years prior to 1997 and has been reduced to one year since.
[20] Maher, J.M., 1976, Discounts for Lack of Marketability for Closely Held Business Interests, *Taxes*, 54, 562-571.
[21] Silber, W.L., 1991, Discounts on Restricted Stock: The Impact of Illiquidity on Stock Prices, *Financial Analysts Journal*, v47, 60-64.

business for this illiquidity and demand a larger share of the ownership of illiquid businesses for the same investment. Looking at the returns earned by private equity investors, relative to the returns earned by those investing in publicly traded companies, should provide a measure of how much value they attach to illiquidity. Ljungquist and Richardson estimate that private equity investors earn excess returns of 5 to 8%, relative to the public equity market, and that this generates about 24% in risk-adjusted additional value to a private equity investor over 10 years. They interpret it to represent compensation for holding an illiquid investment for 10 years.[23] Das, Jagannathan and Sarin take a more direct approach to estimating private company discounts by looking at how venture capitalists value businesses (and the returns they earn) at different stages of the life cycle. They conclude that the private company discount is only 11% for late stage investments but can be as high as 80% for early stage businesses. [24]

Illiquidity Discounts in Practice

The standard practice in many private company valuations is to either use a fixed illiquidity discount for all firms or, at best, to have a range for the discount, with the analyst's subjective judgment determining where in the range a particular company's discount should fall. The evidence for this practice can be seen in both the handbooks most widely used in private company valuation and in the court cases where these valuations are often cited. The genesis for these fixed discounts seems to be in the early studies of restricted stock that we noted in the last section. These studies found that restricted (and therefore illiquid) stocks traded at discounts of 25-35%, relative to their unrestricted counterparts, and private company appraisers have used discounts of the same magnitude in their valuations.[25] Since many of these valuations are for tax court, we

---

[22] B. A. Johnson,1999, Quantitative Support for Discounts for Lack of Marketability, *Business Valuation Review*, v16, 152-55 .
[23] Ljungquist, A. and M. Richardson, 2003, The Cashflow, Return and Risk Characteristics of Private Equity, Working Paper, Stern School of Business.
[24] Das, S., M. Jagannathan and A. Sarin, 2002, The Private Equity Discount: An Empirical Examination of the Exit of Venture Capital Companies, Working Paper, SSRN.
[25] In recent years, some appraisers have shifted to using the discounts on stocks in IPOs in the years prior to the offering. The discount is similar in magnitude to the restricted stock discount.

can see the trail of "restricted stock" based discounts littering the footnotes of dozens of cases in the last three decades.[26]

In recent years, analysts have become more creative in their measurement of the illiquidity discount. They have used option pricing models and studies of transactions just prior to initial public offerings to motivate their estimates and been more willing to estimate firm-specific illiquidity discounts.[27] Appendix 2 describes some of the approaches used to compute liquidity discounts.

2. Other Discounts

While illiquidity discounts are the most common example of post-valuation discounts, there are other risks that also show up as post-valuation adjustments. For instance, analysts valuing companies that are subject to regulation will sometimes discount the value for uncertainty about future regulatory changes and companies that have exposure to lawsuits for adverse judgments on these cases. In each of these cases, analysts concluded that the risk was significant but difficult to incorporate into a discount rate. In practice, the discounts tend to be subjective and reflect the analyst's overall risk aversion and perception of the magnitude of the risk.

Upside Risks

Just as analysts try to capture downside risk that is missed by the discount rates in a post-valuation discount, they try to bring in upside potential that is not fully incorporated into the cashflows into valuations as premiums. In this section, we will examine two examples of such premiums – control and synergy premiums – that show up widely in acquisition valuations.

---

[26] As an example, in one widely cited tax court case (McCord versus Commissioner, 2003), the expert for the taxpayer used a discount of 35% that he backed up with four restricted stock studies.

[27] One common device used to compute illiquidity discounts is to value an at-the-money put option with the illiquidity period used as the life of the option and the variance in publicly traded stocks in the same business as the option volatility. The IPO studies compare prices at which individuals sell their shares in companies just prior to an IPO to the IPO price; the discounts range from 40-60% and are attributed to illiquidity.

1. Control Premium

It is not uncommon in private company and acquisition valuations to see premiums of 20% to 30% attached to estimated value to reflect the "value of control'. But what exactly is this premium for? The value of controlling a firm derives from the fact that you believe that you or someone else would operate the firm differently from the way it is operated currently. When we value a business, we make implicit or explicit assumptions about both who will run that business and how they will run it. In other words, the value of a business will be much lower if we assume that it is run by incompetent managers rather than by competent ones. When valuing an existing company, private or public, where there is already a management in place, we are faced with a choice. We can value the company run by the incumbent managers and derive what we can call a status quo value. We can also revalue the company with a hypothetical "optimal" management team and estimate an optimal value. The difference between the optimal and the status quo values can be considered the value of controlling the business.

If we apply this logic, the value of control should be much greater at badly managed and run firms and much smaller at well-managed firms. In addition, the expected value of control will reflect the difficulty you will face in replacing incumbent management. Consequently, the expected value of control should be smaller in markets where corporate governance is weak and larger in markets where hostile acquisitions and management changes are common.

Analysts who apply control premiums to value are therefore rejecting the path of explicitly valuing control, by estimating an optimal value and computing a probability of management change, in favor of a simpler but less precise approximation. To prevent double counting, they have to be careful to make sure that they are applying the premium to a status quo value (and not to an optimal value). Implicitly, they are also assuming that the firm is badly run and that its value can be increased by a new management team.

2. Synergy Premium

Synergy is the additional value that is generated by combining two firms, creating opportunities that would not been available to these firms operating independently. Operating synergies affect the operations of the combined firm and include economies of scale, increased pricing power and higher growth potential. They generally show up as higher expected cash flows. Financial synergies, on the other hand, are more focused and include tax benefits, diversification, a higher debt capacity and uses for excess cash. They sometimes show up as higher cash flows and sometimes take the form of lower discount rates.

Since we can quantify the impact of synergy on cash flows and discount rates, we can explicitly value it. Many analysts, though, are either unwilling or unable to go through this exercise, arguing that synergy is too subjective and qualitative for the estimates to be reliable. Instead, they add significant premiums to estimated value to reflect potential synergies.

The Dangers of Post-valuation Adjustments

Though the temptation to adjust value for downside and upside risk that has been overlooked is strong, there are clearly significant dangers. The first is that these risks can be easily double counted, if analysts bring their concerns about the risk into the estimation of discount rates and cash flows. In other words, an analyst valuing an illiquid asset may decide to use a higher discount rate for that asset because of its lack of marketability, thus pushing down value, and then proceed to apply a discount to that value. Similarly, an analyst evaluating an acquisition may increase the growth rate in cash flows to reflect the control and synergy benefits from the acquisition and thus increase value; attaching control and synergy premiums to this value will risk double counting the benefits.

The second problem is that the magnitude of the discounts and premiums are, if not arbitrary, based upon questionable evidence. For instance, the 20% control premium used so often in practice comes from looking at the premiums ((over the market price) paid in acquisitions, but these premiums reflect not just control and synergy and also any overpayment on acquisitions. Once these premiums become accepted in practice, they are seldom questioned or analyzed.

The third problem is that adjusting an estimated value with premiums and discounts opens the door for analysts to bring their biases into the number. Thus, an analyst who arrives at an estimate of $100 million for the value of a company and feels it is too low can always add a 20% control premium to get to $ 120 million, even though it may not be merited in this case.

Relative Valuation Approaches

The risk adjustment approaches we have talked about in this chapter have been built around the premise that assets are valuing using discounted cash flow models. Thus, we can increase the discount rate, replace uncertain cash flows with certainty equivalent numbers or apply discounts to estimated value to bring risk into the value. Most valuations, in practice, are based upon relative valuation, i.e., the values of most assets are estimated by looking at how the market prices similar or comparable assets. In this section, we will examine how analysts adjust for risk when doing relative valuation.

Basis for Approach

In relative valuation, the value of an asset is derived from the pricing of 'comparable' assets, standardized using a common variable. Included in this description are two key components of relative valuation. The first is the notion of comparable or similar assets. From a valuation standpoint, this would imply assets with similar cash flows, risk and growth potential. In practice, it is usually taken to mean other companies that are in the same business as the company being valued. The other is a standardized price. After all, the price per share of a company is in some sense arbitrary since it is a function of the number of shares outstanding; a two for one stock split would halve the price. Dividing the price or market value by some measure that is related to that value will yield a standardized price. When valuing stocks, this essentially translates into using multiples where we divide the market value by earnings, book value or revenues to arrive at an estimate of standardized value. We can then compare these numbers across companies.

The simplest and most direct applications of relative valuations are with real assets where it is easy to find similar assets or even identical ones. The asking price for a Mickey Mantle rookie baseball card or a 1965 Ford Mustang is relatively easy to estimate given that there are other Mickey Mantle cards and 1965 Ford Mustangs out there and that the prices at which they have been bought and sold can be obtained. With equity valuation, relative valuation becomes more complicated by two realities. The first is the absence of similar assets, requiring us to stretch the definition of comparable to include companies that are different from the one that we are valuing. After all, what company in the world is similar to Microsoft or GE? The other is that different ways of standardizing prices (different multiples) can yield different values for the same company.

Risk Adjustment

The adjustments for risk in relative valuations are surprisingly rudimentary and require strong assumptions to be justified. To make matters worse, the adjustments are often implicit, rather than explicit, and completely subjective.

a. Sector comparisons: In practice, analysts called upon to value a software company will compare it to other software companies and make no risk adjustments. Implicit is the assumption that all software firms are of equivalent risk and that their price earnings ratios can therefore be compared safely. As the risk characteristics of firms within sectors diverge, this approach will lead to misleading estimates of value for firms that have more or less risk than the average firm in the sector; the former will be over valued and the latter will be under valued.

b. Market Capitalization or Size: In some cases, especially in sectors with lots of firms, analysts will compare a firm only to firms of roughly the same size (in terms of revenues or market capitalization). The implicit assumption is that smaller firms are riskier than larger firms and should trade at lower multiples of earnings, revenues and book value.

c. Ratio based Comparisons: An approach that adds a veneer or sophistication to relative valuation is to compute a ratio of value or returns to a measure of risk. For instance, portfolio managers will often compute the ratio of the expected return on an investment to its standard deviation; the resulting “Sharpe ratio” and can be considered a measure of the returns you can expect to earn for a given unit of risk. Assets that have higher Sharpe ratios are considered better investments.

d. Statistical Controls: We can control for risk in a relative valuation statistically. Reverting to the software sector example, we can regress the PE ratios of software companies against their expected growth rates and some measure of risk (standard deviation in stock price or earnings, market capitalization or beta) to see if riskier firms are priced differently from safer firms. The resulting output can be used to estimate predicted PE ratios for individual companies that control for the growth potential and risk of these companies.

DCF versus Relative Valuation

It should come as no surprise that the risk adjustments in relative valuation do not match up to the risk adjustments in discounted cash flow valuation. The fact that risk is usually considered explicitly in discounted cash flow models gives them an advantage over relative valuations, with its ad-hoc treatment of risk. This advantage can be quickly dissipated, though, if we are sloppy about how we risk adjust the cash flows or discount rates or if we use arbitrary premiums and discounts on estimated value.

The nature of the risk adjustment in discounted cash flow valuation makes it more time and information intensive; we need more data and it takes longer to adjust discount rates than to compare a software company's PE to the average for the software sector. If time and/or data is scarce, it should come as no surprise that individuals choose the less precise risk adjustment procedure embedded in relative valuation.

There is one final difference. In relative valuation, we are far more dependent on markets being right, at least on average, for the risk adjustment to work. In other words, even if we are correct in our assessment that all software companies have similar risk exposures, the market still has to price software companies correctly for the average price earnings ratio to be a good measure of an individual company's equity value. We may be dependent upon markets for some inputs in a DCF model – betas and risk premiums, for instance – but the assumption of market efficiency is less consequential.

The Practice of Risk Adjustment

In this chapter, we have described four ways of adjusting for risk: use a higher discount rate for risky assets, reduce uncertain expected cash flows, apply a discount to estimated value and look at how the market is pricing assets of similar risk. Though each of these approaches can be viewed as self-standing and sufficient, analysts often use more than one approach to adjust for risk in the same valuation. In many discounted cash flow valuations, the discount rate is risk-adjusted (using the CAPM or multi-factor model), the cash flow projections are conservative (reflecting a cash flow risk adjustment), the terminal value is estimated using a multiple obtained by looking at comparable companies (relative valuation risk adjustment) and there is a post-valuation discount for illiquidity.

At the risk of repeating what we said in an earlier section, using multiple risk adjustment procedures in the same valuation not only makes it difficult to decipher the effect of the risk adjustment but also creates the risk of double counting or even triple counting the same risk in value.

Conclusion

With risk-adjusted values, we try to incorporate the effect of risk into our estimates of asset value. In this chapter, we began by looking at ways in which we can do this in a valuation. First, we can estimate a risk-adjusted discount rate, relying if need be on a risk and return model which measures risk and converts it into a risk premium. Second, we can discount uncertain expected cash flows to reflect the uncertainty; if the risk premium computed in a risk and return model is used to accomplish this, the value obtained in this approach will be identical to the one estimated with risk adjusted discount rates. Third, we can discount the estimated value of an asset for those risks that we believe have not been incorporated into the discount rate or the cash flows. Finally, we can use the market pricing of assets of similar risk to estimate the value for a risky asset. The difficulty of finding assets that have similar risk exposure leads to approximate solutions such as using other companies in the same business as the company being valued.

Appendix 5.1: Adjusting Discount Rates for Country Risk

In many emerging markets, there is very little historical data and the data that exists is too volatile to yield a meaningful estimate of the risk premium. To estimate the risk premium in these countries, let us start with the basic proposition that the risk premium in any equity market can be written as:

Equity Risk Premium = Base Premium for Mature Equity Market + Country Premium

The country premium could reflect the extra risk in a specific market. This boils down our estimation to answering two questions:

1. What should the base premium for a mature equity market be?
2. How do we estimate the additional risk premium for individual countries?

To answer the first question, we will make the argument that the US equity market is a mature market and that there is sufficient historical data in the United States to make a reasonable estimate of the risk premium. In fact, reverting back to our discussion of historical premiums in the US market, we will use the geometric average premium earned by stocks over treasury bonds of 4.82% between 1928 and 2003. We chose the long time period to reduce standard error, the treasury bond to be consistent with our choice of a riskfree rate and geometric averages to reflect our desire for a risk premium that we can use for longer term expected returns. There are three approaches that we can use to estimate the country risk premium.

1. *Country bond default spreads*: While there are several measures of country risk, one of the simplest and most easily accessible is the rating assigned to a country's debt by a ratings agency (S&P, Moody's and IBCA all rate countries). These ratings measure default risk (rather than equity risk), but they are affected by many of the factors that drive equity risk – the stability of a country's currency, its budget and trade balances and its political stability, for instance[28]. The other advantage of ratings is that they come with default spreads over the US treasury bond. For instance, Brazil was rated B2 in early 2004 by Moody's and the 10-year Brazilian C-Bond, which is a dollar denominated bond was priced to yield 10.01%, 6.01% more than the interest rate

[28] The process by which country ratings are obtained is explained on the S&P web site at http://www.ratings.standardpoor.com/criteria/index.htm.

(4%) on a 10-year treasury bond at the same time.[29] Analysts who use default spreads as measures of country risk typically add them on to both the cost of equity and debt of every company traded in that country. For instance, the cost of equity for a Brazilian company, estimated in U.S. dollars, will be 6.01% higher than the cost of equity of an otherwise similar U.S. company. If we assume that the risk premium for the United States and other mature equity markets is 4.82%, the cost of equity for a Brazilian company can be estimated as follows (with a U.S. Treasury bond rate of 4% and a beta of 1.2).

Cost of equity = Riskfree rate + Beta *(U.S. Risk premium) + Country Bond Default Spread

= 4% + 1.2 (4.82%) + 6.01% = 15.79%

In some cases, analysts add the default spread to the U.S. risk premium and multiply it by the beta. This increases the cost of equity for high beta companies and lowers them for low beta firms.

2. *Relative Standard Deviation*: There are some analysts who believe that the equity risk premiums of markets should reflect the differences in equity risk, as measured by the volatilities of these markets. A conventional measure of equity risk is the standard deviation in stock prices; higher standard deviations are generally associated with more risk. If you scale the standard deviation of one market against another, you obtain a measure of relative risk.

$$\text{Relative Standard Deviation}_{\text{Country X}} = \frac{\text{Standard Deviation}_{\text{Country X}}}{\text{Standard Deviation}_{\text{US}}}$$

This relative standard deviation when multiplied by the premium used for U.S. stocks should yield a measure of the total risk premium for any market.

$$\text{Equity risk premium}_{\text{Country X}} = \text{Risk Premum}_{\text{US}} * \text{Relative Standard Deviation}_{\text{Country X}}$$

Assume, for the moment, that you are using a mature market premium for the United States of 4.82% and that the annual standard deviation of U.S. stocks is 20%. The

[29] These yields were as of January 1, 2004. While this is a market rate and reflects current expectations, country bond spreads are extremely volatile and can shift significantly from day to day. To counter this volatility, the default spread can be normalized by averaging the spread over time or by using the average default spread for all countries with the same rating as Brazil in early 2003.

annualized standard deviation[30] in the Brazilian equity index was 36%, yielding a total risk premium for Brazil:

$$\text{Equity Risk Premium}_{\text{Brazil}} = 4.82\% * \frac{36\%}{20\%} = 8.67\%$$

The country risk premium can be isolated as follows:

$$\text{Country Risk Premium}_{\text{Brazil}} = 8.67\% - 4.82\% = 3.85\%$$

While this approach has intuitive appeal, there are problems with using standard deviations computed in markets with widely different market structures and liquidity. There are very risky emerging markets that have low standard deviations for their equity markets because the markets are illiquid. This approach will understate the equity risk premiums in those markets.

3. *Default Spreads + Relative Standard Deviations*: The country default spreads that come with country ratings provide an important first step, but still only measure the premium for default risk. Intuitively, we would expect the country equity risk premium to be larger than the country default risk spread. To address the issue of how much higher, we look at the volatility of the equity market in a country relative to the volatility of the bond market used to estimate the spread. This yields the following estimate for the country equity risk premium.

$$\text{Country Risk Premium} = \text{Country Default Spread} * \left(\frac{\sigma_{\text{Equity}}}{\sigma_{\text{Country Bond}}}\right)$$

To illustrate, consider the case of Brazil. As noted earlier, the dollar denominated bonds issued by the Brazilian government trade with a default spread of 6.01% over the US treasury bond rate. The annualized standard deviation in the Brazilian equity index over the previous year was 36%, while the annualized standard deviation in the Brazilian dollar denominated C-bond was 27%[31]. The resulting additional country equity risk premium for Brazil is as follows:

---

[30] Both the US and Brazilian standard deviations were computed using weekly returns for two years from the beginning of 2002 to the end of 2003. While you could use daily standard deviations to make the same judgments, they tend to have much more noise in them.

[31] The standard deviation in C-Bond returns was computed using weekly returns over 2 years as well. Since there returns are in dollars and the returns on the Brazilian equity index are in real, there is an inconsistency

$$\text{Brazil's Country Risk Premium} = 6.01\%\left(\frac{36\%}{27\%}\right) = 7.67\%$$

Note that this country risk premium will increase if the country rating drops or if the relative volatility of the equity market increases. It is also in addition to the equity risk premium for a mature market. Thus, the total equity risk premium for a Brazilian company using the approach and a 4.82% premium for the United States would be 12.49%.

Why should equity risk premiums have any relationship to country bond spreads? A simple explanation is that an investor who can make 11% on a dollar-denominated Brazilian government bond would not settle for an expected return of 10.5% (in dollar terms) on Brazilian equity. Both this approach and the previous one use the standard deviation in equity of a market to make a judgment about country risk premium, but they measure it relative to different bases. This approach uses the country bond as a base, whereas the previous one uses the standard deviation in the U.S. market. This approach assumes that investors are more likely to choose between Brazilian government bonds and Brazilian equity, whereas the previous one approach assumes that the choice is across equity markets.

The three approaches to estimating country risk premiums will generally give you different estimates, with the bond default spread and relative equity standard deviation approaches yielding lower country risk premiums than the melded approach that uses both the country bond default spread and the equity and bond standard deviations. In the case of Brazil, for instance, the country risk premiums range from 3.85% using the relative equity standard deviation approach to 6.01% for the country bond approach to We believe that the larger country risk premiums that emerge from the last approach are the most realistic for the immediate future, but that country risk premiums may decline over time. Just as companies mature and become less risky over time, countries can mature and become less risky as well.

---

here. We did estimate the standard deviation on the Brazilian equity index in dollars but it made little difference to the overall calculation since the dollar standard deviation was close to 36%.

Appendix 5.2: Estimating the Illiquidity Discount

In conventional valuation, there is little scope for showing the effect of illiquidity. The cashflows are expected cashflows, the discount rate is usually reflective of the risk in the cashflows and the present value we obtain is the value for a liquid business. With publicly traded firms, we then use this value, making the implicit assumption that illiquidity is not a large enough problem to factor into valuation. In private company valuations, analysts have been less willing (with good reason) to make this assumption. The standard practice in many private company valuations is to apply an illiquidity discount to this value. But how large should this discount be and how can we best estimate in? This is a very difficult question to answer empirically because the discount in private company valuations itself cannot be observed. Even if we were able to obtain the terms of all private firm transactions, note that what is reported is the price at which private firms are bought and sold. The value of these firms is not reported and the illiquidity discount is the difference between the value and the price. In this section, we will consider four approaches that are in use – a fixed discount (with marginal and subjective adjustments for individual firm differences), a firm-specific discount based upon a firm's characteristics, a discount obtained by estimating a synthetic bid-ask spread for an asset and an option-based illiquidity discount.

a. Fixed Discount

The standard practice in many private company valuations is to either use a fixed illiquidity discount for all firms or, at best, to have a range for the discount, with the analyst's subjective judgment determining where in the range a particular company's discount should fall. The evidence for this practice can be seen in both the handbooks most widely used in private company valuation and in the court cases where these valuations are often cited. The genesis for these fixed discounts seems to be in the early studies of restricted stock that we noted in the last section. These studies found that restricted (and therefore illiquid) stocks traded at discounts of 25-35%, relative to their unrestricted counterparts, and private company appraisers have used discounts of the same magnitude in their valuations.[32] Since many of these valuations are for tax court, we can see the trail of "restricted stock" based discounts littering the footnotes of dozens of cases in the last three decades.[33]

As we noted in the last section, some researchers have argued that these discounts are too large because of the sampling bias inherent in using restricted stock and that they should be replaced with smaller discounts. In recent years, the courts have begun to look favorably at these arguments. In a 2003 case[34], the Internal Revenue Service, often at the short end of the illiquidity discount argument, was able to convince the judge that the conventional restricted stock discount was too large and to accept a smaller discount.

b. Firm-specific Discount

Much of the theoretical and empirical discussion in this chapter supports the view that illiquidity discounts should vary across assets and business. In particular, with a private company, you would expect the illiquidity discount to be a function of the size and the type of assets that the company owns. In this section, we will consider the determinants of the illiquidity discount and practical ways of estimating it.

Determinants of Illiquidity Discounts

With any asset, the illiquidity discount should be a function of the number of potential buyers for the asset and the ease with which that asset can be sold. Thus, the illiquidity discount should be relatively small for an asset with a large number of potential buyers (such as real estate) than for an asset with a relatively small number of buyers (an expensive collectible). With private businesses, the illiquidity discount is likely to vary across both firms and buyers, which renders rules of thumb useless. Let us consider first some of the factors that may cause the discount to vary across firms.

1. *Liquidity of assets owned by the firm*: The fact that a private firm is difficult to sell may be rendered moot if its assets are liquid and can be sold with no significant loss

---

[32] In recent years, some appraisers have shifted to using the discounts on stocks in IPOs in the years prior to the offering. The discount is similar in magnitude to the restricted stock discount.

[33] As an example, in one widely cited tax court case (McCord versus Commissioner, 2003), the expert for the taxpayer used a discount of 35% that he backed up with four restricted stock studies.

[34] The court case was McCord versus Commissioner. In the case, the taxpayer's expert argued for a discount of 35% based upon the restricted stock studies. The IRS argued for a discount of 7%, on the basis

in value. A private firm with significant holdings of cash and marketable securities should have a lower illiquidity discount than one with factories or other assets for which there are relatively few buyers.

2. *Financial Health and Cash flows of the firm*: A private firm that is financially healthy should be easier to sell than one that is not healthy. In particular, a firm with strong earnings and positive cash flows should be subject to a smaller illiquidity discount than one with losses and negative cash flows.

3. Possibility of going public in the future: The greater the likelihood that a private firm can go public in the future, the lower should be the illiquidity discount attached to its value. In effect, the probability of going public is built into the valuation of the private firm. To illustrate, the owner of a private e-commerce firm in 1998 or 1999 would not have had to apply much of a illiquidity discount to his firm's value, if at all, because of the ease with which it could have been taken public in those years.

4. Size of the Firm: If we state the illiquidity discount as a percent of the value of the firm, it should become smaller as the size of the firm increases. In other words, the illiquidity discount should be smaller as a percent of firm value for private firms like Cargill and Koch Industries, which are worth billions of dollars, than it should be for a small firm worth $5 million.

5. Control Component: Investing in a private firm is decidedly more attractive when you acquire a controlling stake with your investment. A reasonable argument can be made that a 51% stake in a private business should be more liquid than a 49% stake in the same business.[35]

The illiquidity discount is also likely to vary across potential buyers because the desire for liquidity varies among investors. It is likely that those buyers who have deep pockets, longer time horizons and see little or no need to cash out their equity positions will attach much lower illiquidity discounts to value, for similar firms, than buyers that do not possess these characteristics. The illiquidity discount is also likely to vary across time, as the market-wide desire for liquidity ebbs and flows. In other words, the illiquidity discount attached to the same business will change over time even for the same buyer.

---

that a big portion of the observed discount in restricted stock and IPO studies reflects factors other than liquidity. The court ultimately decided on an illiquidity discount of 20%.

Estimating Firm-Specific Illiquidity Discount

While it is easy to convince skeptics that the illiquidity discount should vary across companies, it is much more difficult to get consensus on how to estimate the illiquidity discount for an individual company. In this section, we revert back to the basis for the fixed discount studies and look and look for clues on why discounts vary across companies and how to incorporate these differences into illiquidity discounts.

i. Restricted Stock Studies

Earlier in the chapter, we looked at studies of the discount in restricted stock. One of the papers that we referenced by Silber (1991) examined factors that explained differences in discounts across different restricted stock by relating the size of the discount to observable firm characteristics including revenues and the size of the restricted stock offering. He reported the following regression.

$$LN(RPRS) = 4.33 + 0.036 \ln(REV) - 0.142 LN(RBRT) + 0.174 DERN + 0.332 DCUST$$

where,

RPRS = Restricted Stock Price/ Unrestricted stock price = 1 – illiquidity discount

REV = Revenues of the private firm (in millions of dollars)

RBRT = Restricted Block relative to Total Common Stock (in %)

DERN = 1 if earnings are positive; 0 if earnings are negative;

DCUST = 1 if there is a customer relationship with the investor; 0 otherwise;

The illiquidity discount tends to be smaller for firms with higher revenues, decreases as the block offering decreases and is lower when earnings are positive and when the investor has a customer relationship with the firm. These findings are consistent with some of the determinants that we identified in the previous section for the illiquidity premium. In particular, the discounts tend to be smaller for larger firms (at least as measured by revenues) and for healthy firms (with positive earnings being the measure of financial health). This would suggest that the conventional practice of using constant discounts across private firms is wrong and that we should be adjusting for differences across firms.

Consider again the regression that Silber presents on restricted stock. Not only does it yield a result specific to restricted stock, but it also provides a measure of how

---

[35] For more on the value of control, see the companion paper on the value of control.

much lower the discount should be as a function of revenues. A firm with revenue of $20 million should have an illiquidity discount that is 1.19% lower than a firm with revenues of $10 million. Thus, we could establish a benchmark discount for a profitable firm with specified revenues (say $10 million) and adjust this benchmark discount for individual firms that have revenues much higher or lower than this number. The regression can also be used to differentiate between profitable and unprofitable firms. Figure 14.6 presents the difference in illiquidity discounts across both profitable and unprofitable firms with different revenues, using a benchmark discount of 25% for a firm with positive earnings and $10 million revenues.

Figure 14.6: Illiquidity Discounts: Base Discount of 25% for profitable firm with $ 10 million in revenues

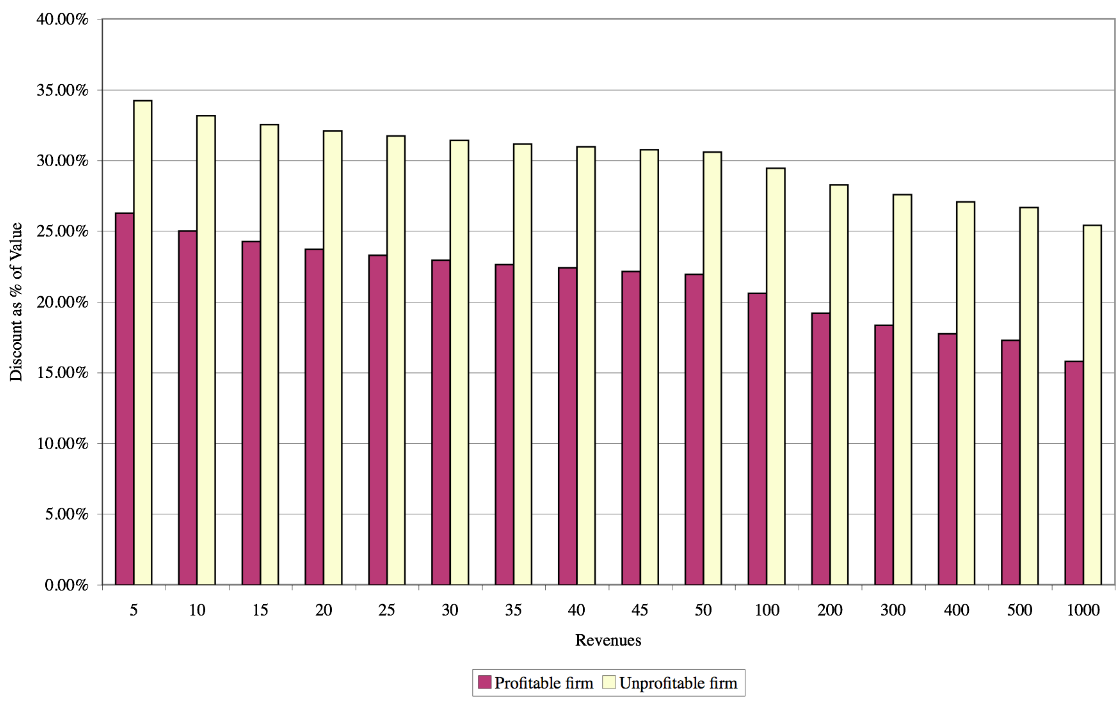

There are clearly dangers associated with extending a regression run on a small number of restricted stocks to estimating discounts for private firms, but it does provide at least a road map for adjusting discount factors.

(http;//www.damodaran.com: Look under research/papers) *ii. Private Placements*

Just as Silber considered fundamental factors that cause restricted stock discounts to vary across firms, Bajaj et al. (referenced earlier) considered various fundamental factors that may cause illiquidity discounts to vary across firms in private placements. Their regression, run across 88 private placements between 1990 and 1995 is summarized below:

DISC = 4.91% + 0.40 SHISS -0.08 Z -7.23 DREG + 3.13 SDEV R^2 = 35.38%

(0.89) (1.99) (2.51) (2.21) (3.92)

Where

DISC = Discount on the Market Price

SHISS = Private Placement as percent of outstanding shares

Z = Altman Z-Score (for distress)

DREG = 1 if registered; 0 if unregistered (restricted stock)

SDEV = Standard deviation of returns

Other things remaining equal, the discount is larger for larger private placements (as a percent of outstanding stocks) by risky and distressed firms and smaller for safer firms. As noted before, the discount is larger for restricted stock than for registered stock. Hertzel and Smith (also referenced earlier) ran a similar regression with 106 private placements between 1980 and 1987 and also found larger private placement discounts at more distressed, riskier and smaller firms.

These regressions are a little more difficult to adapt for use with private company valuations since they are composite regressions that include registered private placements (where there is no illiquidity). However, the results reinforce the Silber regression findings that troubled or distressed firms should have larger illiquidity discounts than healthy firms.

There are legitimate criticisms that can be mounted against the regression approach. The first is that the R squared of these regressions is moderate (30-40%) and that the estimates will have large standard errors associated with them. The second is that the regression coefficients are unstable and likely to change over time. While both criticisms are valid, they really can be mounted against any cross sectional regression and cannot be used to justify a constant discount for all firms. After all, these regressions clearly reject the hypothesis that the discount is the same across all firms.

c. Synthetic Bid-ask Spread

The biggest limitation of using studies based upon restricted stock or private placements is that the samples are small. We would be able to make far more precise estimates if we could obtain a large sample of firms with illiquidity discounts. We would argue that such a sample exists, if we consider the fact that an asset that is publicly traded is not completely liquid. In fact, liquidity varies widely across publicly traded stock. A small company listed over-the-counter is much less liquid that a company listed on the New York Stock Exchange which in turn is much less liquid that a large capitalization company that is widely held. If, as we argued earlier, the bid-ask spread is a measure of the illiquidity of a stock, we can compute the spread as a percent of the market price and relate it to a company's fundamentals. While the bid-ask spread might only be a quarter or half a dollar, it looms as a much larger cost when it is stated as a percent of the price per unit. For a stock that is trading at $2, with a bid-ask spread of 1/4, this cost is 12.5%. For higher price and very liquid stocks, the illiquidity discount may be less than 0.5% of the price, but it is not zero. What relevance does this have for illiquidity discounts on private companies? Think of equity in a private company as a stock that never trades. On the continuum described above, you would expect the bid-ask spread to be high for such a stock and this would essentially measure the illiquidity discount.

To make estimates of the illiquidity discounts using the bid-ask spread as the measure, you would need to relate the bid-ask spreads of publicly traded stocks to variables that can be measured for a private business. For instance, you could regress the bid-ask spread against the revenues of the firm and a dummy variable, reflecting whether the firm is profitable or not, and extend the regression done on restricted stocks to a much larger sample. You could even consider the trading volume for publicly traded stocks as an independent variable and set it to zero for a private firm. Using data from the end of 2000, for instance, we regressed the bid-ask spread against annual revenues, a dummy variable for positive earnings (DERN: 0 if negative and 1 if positive), cash as a percent of firm value and trading volume.

Spread = 0.145 – 0.0022 ln (Annual Revenues) -0.015 (DERN) – 0.016 (Cash/Firm Value) – 0.11 ($ Monthly trading volume/ Firm Value)

Plugging in the corresponding values – with a trading volume of zero – for a private firm should yield an estimate of the synthetic bid-ask spread for the firm. This synthetic spread can be used as a measure of the illiquidity discount on the firm.

d. Option-Based Discount

In an earlier section, we examined an option-pricing based approach, which allowed you to estimate an upper bound for the illiquidity discount, by assuming an investor with perfect market timing skills. There have been attempts to extend option pricing models to valuing illiquidity, with mixed results. In one widely used variation, liquidity is modeled as a put option for the period when an investor is restricted from trading. Thus, the illiquidity discount on value for an asset where the owner is restricted from trading for 2 years will be modeled as a 2-year at-the-money put option.[36] There are several flaws, both intuitive and conceptual, with this approach. The first is that liquidity does not give you the right to sell a stock at today's market price anytime over the next 2 years. What it does give you is the right to sell at the prevailing market price anytime over the next 2 years.[37] The second (and smaller) problem is that option pricing models are based upon continuous price movements and arbitrage and it is difficult to see how these assumptions will hold up for an illiquid asset.

The value of liquidity ultimately has to derive from the investor being able to sell at some pre-determined price during the non-trading period rather than being forced to hold until the end of the period. The look-back option approach that assumes a perfect market timer, explained earlier in the chapter, assumes that the sale would have occurred at the high price and allows us to estimate an upper bound on the value. Can we use option pricing models to value illiquidity without assuming perfect market timing. Consider one alternative. Assume that you have a disciplined investor who always sells investments, when the price rises 25% above the original buying price. Not being able to

---

[36] In a 1993 study, David Chaffe used this approach to estimate illiquidity discounts rangings from 28-49% for an asset, using the Black Scholes option pricing model and volatilities ranging from 60 to 90% for the underlying asset.

[37] There is a simple way to illustrate that this put option has nothing to do with liquidity. Assume that you own stock in a liquid, publicly traded company and that the current stock price is $ 50. A 2-year put option on this stock with a strike price of $ 50 will have substantial value, even though the underlying stock is completely liquid. The value has nothing to do with liquidity but is a price you are willing to pay for insurance.

trade on this investment for a period (say, 2 years) undercuts this discipline and it can be argued that the value of illiquidity is the produce of the value of the put option (estimated using a strike price set 25% above the purchase price and a 2 year life) and the probability that the stock price will rise 25% or more over the next 2 years.

If you decide to apply option pricing models to value illiquidity in private businesses, the value of the underlying asset (which is the private business) and the standard deviation in that value will be required inputs. While estimating them for a private business is more difficult to do than for a publicly traded firm, we can always use industry averages.