Estimating Risk Parameters

Aswath Damodaran

Stern School of Business

44 West Fourth Street

New York, NY 10012

adamodar@stern.nyu.edu

Estimating Risk Parameters

Over the last three decades, the capital asset pricing model has occupied a central and often controversial place in most corporate finance analysts' tool chests. The model requires three inputs to compute expected returns – a riskfree rate, a beta for an asset and an expected risk premium for the market portfolio (over and above the riskfree rate). Betas are estimated, by most practitioners, by regressing returns on an asset against a stock index, with the slope of the regression being the beta of the asset. In this paper, we attempt to show the flaws in regression betas, especially for companies in emerging markets. We argue for an alternate approach that allows us to estimate a beta that reflect the current business mix and financial leverage of a firm.

Risk Parameter Estimation

Most assets that we choose to invest in, financial as well as real, have some exposure to risk. Financial theory and common sense tell us that investments that are riskier need to make higher returns to compensate for risk. Models of risk and return in finance take the view that the risk in an investment should be the risk perceived by a well diversified investor, and that the expected return should be a function of this risk measure. Differences exist, however, between different models in how to measure this market risk. At one end, the capital asset pricing model measures the market risk with a beta measured relative to a market portfolio, and at the other are multi-factor models that measure market risk using multiple betas estimated relative to different factors.

Risk and Return Models

While there are several accepted risk and return models in finance, they all share some common views about risk. First, they all define risk in terms of variance in actual returns around an expected return; thus, an investment is riskless when actual returns are always equal to the expected return. Second, they all argue that risk has to be measured from the perspective of the marginal investor in an asset, and that this marginal investor is well diversified. Therefore, the argument goes, it is only the risk that an investment adds on to a diversified portfolio that should be measured and compensated.

In fact, it is this view of risk that leads risk models to break the risk in any investment into two components. There is a firm-specific component that measures risk that relates only to that investment or to a few investments like it, and a market component that contains risk that affects a large subset or all investments. It is the latter risk that is not diversifiable and should be rewarded.

While all risk and return models agree on this fairly crucial distinction, they part ways when it comes to how measure this market risk. The capital asset pricing model, with assumptions about no transactions cost or private information, concludes that the marginal investor hold a portfolio that includes every traded asset in the market, and that the risk of any investment is the risk added on to this "market portfolio". The expected return from the model is

$$\text{Expected Return} = \text{Riskfree Rate} + \beta_{jM} \text{ (Risk Premium on Market Portfolio)}$$

The arbitrage pricing model, which is built on the assumption that assets should be priced to prevent arbitrage, conludes that there can be multiple sources of market risk, and that the betas relative to each of these sources measures the expected return. Thus, the expected return is:

$$\text{Expected Return} = \text{Riskfree Rate} + \sum_{j=1}^{j=k} \beta_j (\text{Risk Premium}_j)$$

where β_j = Beta of investment relative to factor j

Risk Premium_j = Risk Premium for factor j

Multi-factor models, which specify macro economic variables as these factors take the same form.

Assuming that the riskfree rate is known, these models all require two inputs. The first is the beta or betas of the investment being analyzed, and the second is the appropriate risk premium for the factor or factors in the model. While we examine the issue of risk premium estimation[1] in a companion piece, we will concentrate on the measurement of the risk premium in this paper.

What we would like to measure in the beta

The beta or betas that measure risk in models of risk in finance have two basic characteristics that we need to keep in mind during estimation. The first is that they measure the risk added on to a diversified portfolio, rather than total risk. Thus, it is entirely possible for an investment to be high risk, in terms of individual risk, but to be low risk, in terms of market risk. The second characteristic that all betas share is that they measure the relative risk of an asset, and thus are standardized around one. The market-capitalization weighted average beta across all investments, in the capital asset pricing model, should be equal to one. In any multi-factor model, each beta should have the same property.

Keeping in mind these characteristics, we would like the beta we estimate for an asset to measure the risk added on by that asset to a diversified portfolio. This, of course, raises interesting follow-up questions. When we talk about diversified portfolios, are we referring to a portfolio diversified into just equity or should we include other asset classes? Should we look at diversifying only domestically or should we look globally? In

---

[1] "Estimating Risk Premiums", Aswath Damodaran, Stern School of Business

the CAPM, for instance, with no transactions costs, the diversified portfolio includes all asset classes and is globally diversified. If there are transactions costs and barriers to global investment, the market portfolio may not include all asset classes or be as globally diversified. We would suggest an alternate route to answering these questions. In coming up with a diversified portfolio, we should take the perspective of the marginal investor in the market. The extent to which that marginal investor is diversified should determine the composition of our diversified portfolio.

What we do in practice…

The textbook description of beta estimation is simple. The beta for an asset can be estimated by regressing the returns on any asset against returns on an index representing the market portfolio, over a reasonable time period.

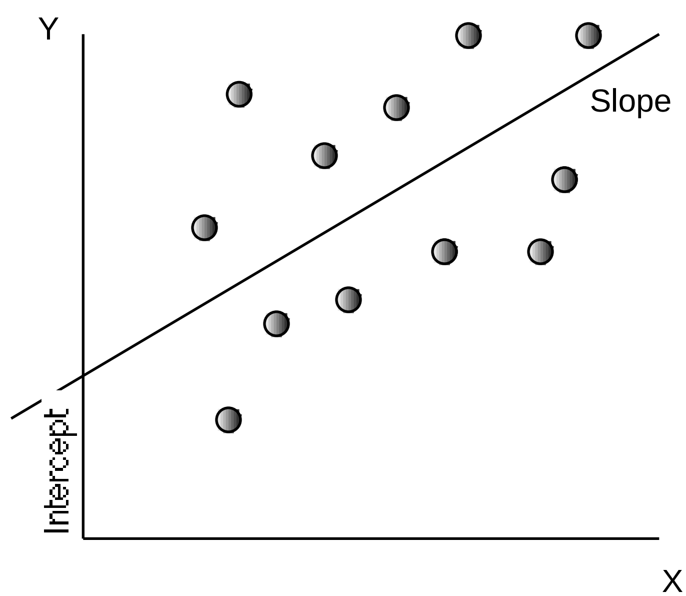

where the returns on the asset represent the Y variable, and the returns on the market index represent the X variable. Note that the regression equation that we obtain is as follows:

$$R_j = a + b R_M$$

Where R_j is the return on investment j, and R_M is the return on the market index. The slope of the regression 'b" is the beta, because it measures the risk added on by that investment to the index used to capture the market portfolio. In addition, it also fulfils the requirement that it be standardized, since the weighted average of the slope coefficients estimated for all of the securities in the index will be one.

In practice, however, there are a number of measurement issues that can color the beta estimate.

1. Choice of a Market Index: In practice, there are no indices that measure or even come close to the market portfolio. Instead, we have equity market indices and fixed income market indices, that measure the returns on subsets of securities in each market. In addition, even these indices are not comprehensive and include only a subset of the securities in each market. Thus, the S&P 500, which is the most widely used index for beta estimation for US companies, includes only 500 of the thousands of equities that are traded in the US market. In many emerging markets, the indices used tend to be even narrower and include only a few dozen large companies. These choices more complex when we consider the possibility of using global equity indices, such as the Morgan Stanley Capital Index, which is a market-weighted composite index that includes most major equity markets. Can the choice of a market index make a difference? The following table, for instance, summarizes betas estimated for Disney, using monthly data from January 1, 1993 to December 31, 1997, using a number of different indices:

Index Used	Beta Calculated
Dow 30	0.99
S&P 500	1.13
NYSE Composite	1.14
Wilshire 5000	1.05
MS Capital Index	1.06

Note that none of these indices include other asset classes, such as fixed income or real assets. This is because indices that include these asset classes are generally not reported on a weekly or a monthly basis.

In terms of making a judgment as to which of these indices gives us the best beta estimate, we would suggest passing it through the "market portfolio" test. In other words, indices that include more securities should provide better estimates than indices that include less, and indices that are market-weighted should yield better estimates that indices that are not. Finally, the index should reflect the extent to which the marginal investor in that market is diversified. Thus, the rationale for the use of the S&P 500 becomes clearer. It includes fewer stocks than the NYSE composite or the Wilshire 5000, but it has an advantage over those indices because it is market weighted, and it includes the 500 largest firms.

2. <u>Choice of a Time Period</u>: Risk and return models are silent on how long a time period one needs to use to estimate betas. Services use periods ranging from two years to five years for beta estimates, with varying results. In fact, using Disney as an example again, we estimated betas for periods ranging from 3 years to 10 years:

Time Period Used	Beta Estimated
3 years	1.04
5 years	1.13
7 years	1.09
10 years	1.18

In choosing a time period for beta estimation, it is worth noting the trade off involved. By going back further in time, we get the advantage of having more observations in the regression, but this could be offset by the fact that the firm itself might have changed its characteristics, in terms of business mix and leverage, over that period. Our objective is not to estimate the best beta we can over the last period but to obtain the best beta we can for the future.

At the most general level, we should be able to go back further in time for firms which have remained fairly stable in terms of business mix and leverage. We should use shorter estimation periods for firms that have restructured, acquired or divested business, or changed their financial leverage over the last few years.

3. Choice of a Return Interval: The final choice that can affects beta estimates is the return interval, used to measure returns historically. Returns can be measured daily, weekly, monthly, quarterly or annually. In fact, with data on intra-day transactions, returns can even be measured using intervals as short as fifteen minutes.

Return Interval Used	*Beta Estimated*
Daily	1.33
Weekly	1.38
Monthly	1.13
Quarterly	0.44
Annual	0.77

Using shorter return intervals increases the number of observations in the regression, for any given time period, but it does come with a cost. Assets do not trade on a continuous basis, and when there is non-trading on the asset, the beta estimated can be affected. In particular, non-trading on an asset during a return period[2] can reduce the measured correlation with the market index, and consequently the beta estimate.

This non-trading problem can be reduced in one of two ways. One way is to use longer return intervals; quarterly and annual returns result in to few observations in

[2] Even if an asset trades during a return period, there can still be a estimation problem associated with the fact that trading on the asset and trading on the index can be non-synchronous. Thus, if the last trade on an

the regression, but monthly returns should provide sufficient observations for firms listed for more than three years. Betas estimated using daily or even weekly returns are likely to have a significant bias due to the non-trading problem, with illiquid firms reporting lower betas than they really should have and liquid firms reporting higher betas than is justified. The other way is to estimate betas using short return interval returns, and then adjusting[3] these betas for the extent of the non-trading.

Post-Regression Beta Adjustments in Practice

As a consequence of different choices in time period used, return intervaling and market index, different services often end up with different estimates of beta for the same firm. In addition, most services adjust their regression betas towards one, using fairly simple techniques. To illustrate, Bloomberg estimates an adjusted beta by doing the following:

Adjusted Beta = Regression Beta (0.67) + 1.00 (0.33)

This effectively pushes all regression beta estimates closer to one. The weights remain the same for all companies, and are not a function of the precision of the beta estimate.

---

asset occurred at 2.30 pm and the index is measured at the close of trading at 4.00 pm, the measured correlation between the asset and the index will be biased downwards.

[3] There are adjustment techniques that can be used. The Scholes-Williams betas are defined as

$$\beta_j = \sum_{k=-1}^{k=+1} \frac{\beta_{jk}}{(1+2\rho)}$$

where ρ is the autocorrelation of the CRSP value-weighted daily market return and β_{jk} are the slope coefficients from three separate OLS regressions,

$$R_{jt} = a_{jk} + \beta_{jk} R_{m,t+k} \text{ for } k = -1, 0, +1$$

This allows the beta estimate to reflect the spill over of returns that often occurs around non-trading. The Dimson beta estimates (β_j) for each firm are obtained by summing the slope coefficients on the five lagged, five leading and contemporaneous returns on a market index in the following OLS regression:

Most other services use similar techniques, with small differences in the weights attached to each number.

Why adjust betas towards one? The rationale can be traced to studies that indicate that, over time, there is a tendency on the part of betas of all companies to move towards one. Intuitively, this should not be surprising. Firms that survive in the market tend to increase in size over time, become more diversified and have more assets in place, producing cash flows. All of these factors should push betas towards one.

Using constant weights to estimate these betas, however, does not make sense. The speed with betas converge on one should vary across companies. Firms that tend to diversify more should see their betas converge on one far faster than firms which stay focused in one business. While conceding the fact that betas for most firms will move towards one over time, we would argue that there is no need to adjust regression betas towards one right now to reflect this tendency. Instead, the betas can change over time, in a valuation or a project analysis, to approach one over time.

There are some services that do consider other fundamentals in estimating betas. They adjust regression betas to reflect other factors that they discover over time to be correlated with risk. This approach combines industry and company-fundamental factors to predict betas. We will consider some of these adjustments later in this paper.

$$R_j = a_j + \sum_{k=-5}^{k=+5} \beta_{jk} R_{m,t+k}$$

This approach also allows for the spill over effect created by non-trading.

The Limitations of Regression Betas

1. The Index Problem

The regression betas reported are clearly affected by estimation choices - the time period used, the return interval and the index. The beta will vary widely depending upon how the regression is set up and run. This problem is exacerbated when we estimate betas for companies in emerging markets. Using the index of the local market, which is still what most services do, provides a regression that looks very good. For instance, the following is a beta estimate for Telebras against the Brazilian index, which is the Bovespa:

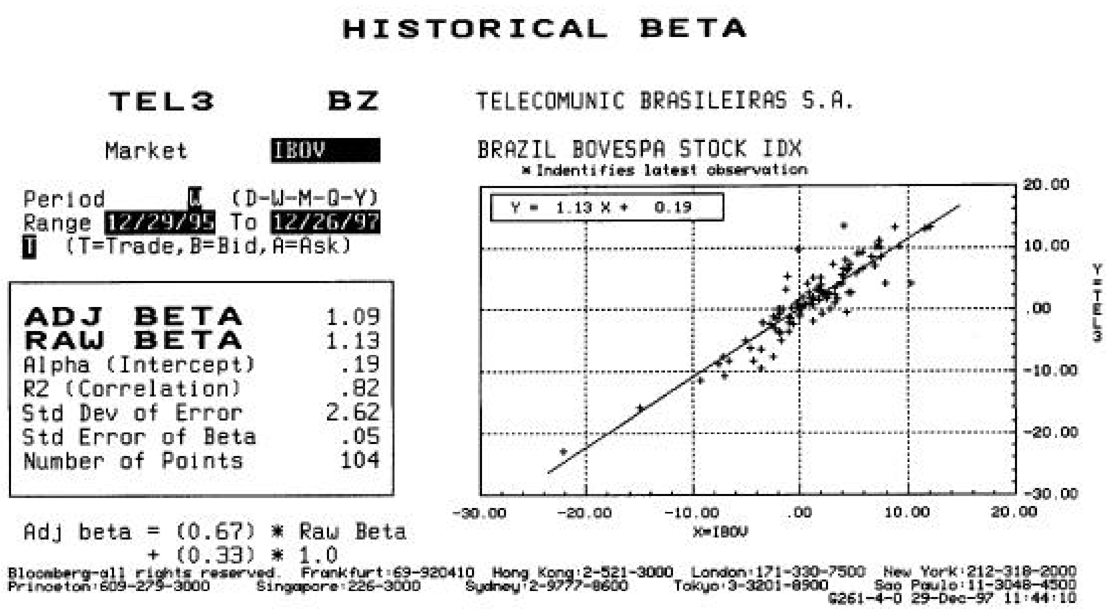

The good fit is deceptive, however. Telebras is 40% or more of the Bovespa, and this has some strange consequences. The first is that the beta estimates for all other Brazilian stocks essentially become regressions of those stocks against Telebras, rather than a diversified stock index. The second is that more than 90% of all stocks on the Brazilian index were reporting betas less than one at the time of this regression. Since it is the weighted average beta that is one, and Telebras has a beta greater than one, this asymmetry in beta estimates becomes possible. The third and most troubling consequence is that it is the smallest, riskiest companies in the Brazilian market that have the lowest betas, while the largest and most estabilished firms have the highest betas.

There are still many who argue that this is, in fact, the best measure of risk in these firms, and that the marginal investor's portfolio in Brazil is likely to be weighted heavily with Telebras. This argument may have resonance in markets where investors invest only in domestic stocks, but is tough to sustain in a global marketplace. The marginal investor in Telebras could very well be, and often is, a portfolio manager in New York or London. This portfolio manager's portfolio is certainly not going to resemble the Bovespa, and the beta estimate obtained relative to it is meaningless. Even if the marginal investor is a Brazilian, it is extremely unlikely that he or she would want to hold a portfolio that resembles the Bovespa.

So what are the options? The right index to use in analysis should be determined by who the marginal investor in the firm is - a good indicator is to look at the largest holders of stock in the company and the markets where the trading volume is heaviest. If the marginal investor is, in fact, a Brazilian investor, it is reasonable to use a well-constructed Brazilian index. If the marginal investor is a US investor, the right index might be the S&P 500[4]. If the marginal investor is a global investor, a more relevant measure of risk may emerge by using the global index. Over time, you would expect global investors to displace local investors as the marginal investors, because they will perceive far less of the risk as market risk and thus pay a higher price for the same security. Thus, one of the ironies of our notion of risk is that Aracruz will be less risky to an overseas investor who has a global portfolio than to a Brazilian investor with all of his or her wealth in Brazilian assets.

2. The Noise Problem

Using a global index or even the S&P 500 may solve the problem of a few stocks dominating the index, but it aggravates the second problem with regression beta estimates. The beta estimate from the regression is noisy, and the range that emerges for the beta is large. The following table summarizes the beta estimates for Disney, using different indices, with the standard errors on each.

Index Used	Beta Calculated	Standard Error of Estimate
Dow 30	0.99	0.20
S&P 500	1.13	0.22
NYSE Composite	1.14	0.24
Wilshire 5000	1.05	0.23
MS Capital Index	1.06	0.18

[4] When running regressions of a non-US stock against dollar denominated indices like the S&P 500 and the Morgan Stanley Capital Index, the returns have to be converted into dollar returns, either by using an ADR listed in the US, or by using exchange rates in each period.

Note, for instance, the beta estimate relative to the S&P 500. The true beta for Disney, with 67% confidence, could be between 0.91 and 1.35. With 95% confidence, this range becomes 0.69 to 1.57. Lest this be viewed as a problem just with Disney, we have graphed the standard errors on beta estimates for all US stocks in the following chart.

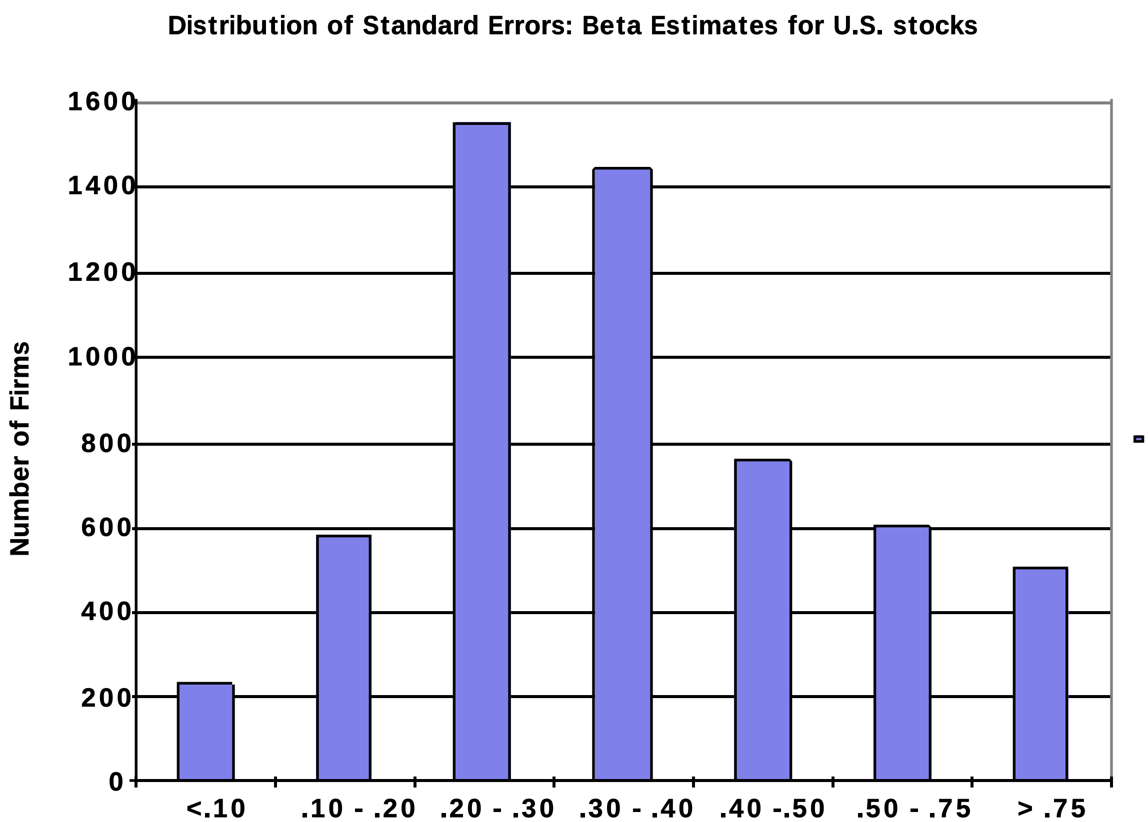

Note the number of beta estimates with standard errors in excess of 0.50. These estimates can be viewed as close to useless, by themselves, since the range that emerges for the beta is so large.

Returning to the Telebras example from the previous section, we could estimate the beta for Telebras using the ADR (which is dollar denominated) and the S&P 500 or the Morgan Stanley Capital Index. The betas that emerge from these regressions, however, have very large standard errors and may not be particularly useful in analysis.

3. The Problem of Firms Changing over Time

Even if a stock does not dominate the index, and the regression beta has a low standard error, there is a final problem with regression beta estimates. They are based upon historical data, and firms change over time. Thus, the regression reflects the firm's characteristics, on average, over the period of the estimation rather than the firm as it exists today. Firms do change over time for three reasons:

- They divest existing businesses, invest in new businesses and acquire firms. In the process, they change their business mix, which will change their beta.
- They also change financial leverage, by adding to or paying off debt. In addition, actions such as the payment of dividends and buying back stock can also affect financial leverage. Finally, changing market values for both debt and equity can cause leverage to change significantly over short periods.
- Even if firms do not change their business mix or financial leverage, they tend to grow over time. As they grow, their operating cost structures will tend to change leading to changes in the betas.

Thus, there are very few firms where regression betas will be good measures of the beta for the current period, let alone for future periods.

Alternatives to Regression Betas

There are three basic alternatives to simple regression betas. One is to use modify the regression betas to reflect the firm's current operating and financial characteristics. The second is to come up with a measure of relative risk (which is what beta is) without using historical prices on the stock and the index. The third is to estimate bottom-up betas, which reflect the businesses a firm is operating in and its current financial leverage.

Modified Regression Betas

Once a regression beta has been estimated it can be modified to reflect our need to estimate a beta for future time periods. We talked about one such adjustment that many services make, where they shrink all betas towards one after the estimation is done. There are, however, some estimation services that adopt far more sophisticated adjustments to make the regression betas reflect the current fundamentals of the firm. Income statement and balance sheet variables are important predictors of beta - high payout is predictive of low beta; high variability of earnings and covariability with economy-wide earnings are predictive of high beta. A series of researchers[5] have looked at the relationship between betas and fundamental variables. Rosenberg and Marathe suggest that fundamental information about a firm can be used in conjunction with historical beta estimates to provide superior predictors of future betas. This approach can be generalized and updated to estimate modified betas for all firms. The following is a regression relating the betas of

[5] Beaver, Kettler and Scholes (1970) examined the relationship between betas and seven variables - dividend payout, asset growth, leverage, liquidity, asset size, earnings variability and the accounting beta. Rosenberg and Guy did a similar analysis of the relationship between betas and financial fundamentals.

NYSE and AMEX stocks in 1996 to four variables - dividend yield, standard deviation in operating income, market capitalization and debt/equity ratio yielded the following.

$$BETA = 0.7997 + 2.28\ \sigma_{OI} - 3.23\ Yield + 0.21\ D/E - .000005\ Market\ Capitalization$$

where,

σ_{OI} = Standard deviation in percentage change in operating income over previous 5 years

Yield = Dividend yield on stock

D/E = Book value of debt/ Market Value of Equity

Market Capitalization: measured as market value of equity (in millions)

This approach has two basic limitations. The first is that the initial beta estimates need to reasonably good, though they can be noisy, for the estimation to work. The second is that the regression of betas against financial fundamentals has a low R-squared. Thus, while modified betas may reflect more of the firm's current characteristics, it does not really solve the noise problem of large standard errors.

Relative Risk Measures

Our objective when we began this process was to come up with a measure that reflected a firm's exposure to market risk, relative to other firms in the market. The regression of stock prices on the market index was the tool we used to arrive at this estimate, and we have argued that it is flawed because of the noise in the estimate. There are relative risk measures that we can compute that do not require historical prices, but do require that we make other assumptions about the nature of risk:

a. Relative Volatility

The relative volatility measures the volatility of an asset's price relative to the average volatility across all assets in that market. Thus the relative volatility of any asset can be defined as follows:

$$\text{Relative Volatility}_j = \frac{\text{Standard Deviation}_j}{\text{Average Standard Deviation across All Assets}}$$

Note that the denominator is not the standard deviation in an index, but an average of the asset-specific standard deviations in the market. While we could compute a market-capitalization weighted average, this would then open us up to the same problems of index domination that we noted earlier in the regressions.

To illustrate, consider the example of Telebras. The annualized standard deviation in Telebras stock prices in 1997 was 48%, while the average annualized standard deviation of a stock listed on the Brazilian index was 60%. The relative volatility for Telebras can be estimated as follows:

$$\text{Relative Volatility}_{\text{Telebras}} = \frac{\text{Standard Deviation}_{\text{Telebras}}}{\text{Average Standard Deviation}_{\text{Brazilian Stock}}} = \frac{48\%}{60\%} = 0.80$$

Note that the relative volatility is standardized around one; relative volatilities greater than one indicate above-average risk, while relative volatilites less than one indicate below-average risk. In addition, the average relative volatility across all stocks will average to one. Finally, the relative volatility is used in much the same way as the traditional beta estimate to compute expected returns:

Expected Return = Riskfree Rate + Relative Volatility (Expected Risk Premium)

What are the advantage of using relative volatility? The term that creates the most noise in the traditional beta estimate is the estimate of correlation between the asset and the market index. The relative volatility measure does not require a correlation measure and hence is less noisy. This comes with a cost, however. The relative volatility measure is based upon the assumption that total risk and market risk exposures are perfectly correlated. In other words, firms with high total risk will also be exposed to high market risk. Finally, decisions made on how to compute the average standard deviation will affect relative volatility. Thus, the relative volatility of Telebras computed relative to the average Brazilian stock will be very different from the relative volatility computed for Telebras using the ADR and the average standard deviation across stocks listed on the S&P 500.

b. Accounting Betas

Another approach to estimate the relative risk parameters from accounting earnings rather than from traded prices. Thus, changes in earnings at a division or a firm, on a quarterly or annual basis, can be regressed against changes in earnings for the market, in the same periods, to arrive at an estimate of a “market beta” to use in the CAPM. When market prices are viewed as too noisy or are unavailable, this approach can still be used to estimate to estimate an "accounting" beta.

While the approach has some intuitive appeal, it suffers from three potential pitfalls. First, accounting earnings tend to be smoothed out relative to the underlying value of the company, resulting in betas that are "biased down", especially for risky firms, or "biased up", for safer firms. In other words, betas are likely to be closer to one for all firms using accounting data. Second, accounting earnings can be influenced by non-operating factors, such as changes in depreciation or inventory methods, and by allocations of corporate expenses at the divisional level. Finally, accounting earnings are measured, at most, once every quarter, and often only once every year, resulting in regressions with few observations and not much power.

Bottom-Up Betas

The beta of a firm might be estimated from a regression but it is determined by fundamental decisions that a firm takes on where to invest, what type of cost structure it plans to maintain and how much debt it takes on. The final approach to beta estimation considers these fundamentals and we call it the bottom-up approach to beta estimation. To understand this approach, we will begin be considering the fundamentals that determine betas and then provide a framework for estimating bottom-up betas.

Determinants of Betas

The beta of a firm is determined by three variables -(1) the type of business(es) the firm is in, (2) the degree of operating leverage in the firm and (3) the firm's financial leverage. While much of the discussion in this section will be couched in terms of CAPM betas, the same analysis can be applied to the betas estimated in the APM and the multi-factor model as well.

Type of Business Since betas measure the risk of a firm relative to a market index, the more sensitive a business is to market conditions, the higher is its beta. Thus, other things remaining equal, cyclical firms can be expected to have higher betas than non-cyclical firms. Other things remaining equal, then, companies involved in housing and automobiles, two sectors of the economy which are very sensitive to economic conditions, will have higher betas than companies which are in food processing and tobacco, which are relatively insensitive to business cycles.

Building on this point, we would also argue that the degree to which a product's purchase is discretionary will affect the beta of the firm manufacturing the product. Thus, the betas of food processing firms, such as General Foods and Kellogg's, should be lower than the betas of specialty retailers, since consumers can defer the purchase of the latter's products during bad economic times.

It is true that firms have only limited control over how discretionary the product or service that they provide is to their customers. There are firms, however, that have used this limited control to maximum effect to make their products less discretionary to buyers, and by extension, lowered their business risk. One approach is to make the product or service a much more intergral and necessary part of everyday life, thus making its purchase more of a requirement. (The push by online services, like America Online, for people to use e-mail and shopping services on the web is designed to make it less of a discretionary purchase.) A second approach is to effectively use advertising and marketing to build brand loyalty. The objective in good advertising, as I see it, is to make discretionary products or services seem like necessities to the target audience. Thus, corporate strategy, advertising and marketing acumen can, at the margin, alter the business risk and betas over time.

Degree of Operating Leverage The degree of operating leverage is a function of the cost structure of a firm, and is usually defined in terms of the relationship between fixed costs and total costs. A firm that has high operating leverage (i.e., high fixed costs relative to total costs), will also have higher variability in earnings before interest and taxes (EBIT) than would a firm producing a similar product with low operating leverage. Other things remaining equal, the higher variance in operating income will lead to a higher beta for the firm with high operating leverage.

This has consequences for major strategic decisions that firms make on future direction. While there is much good that comes from updating plants and getting the latest technology, there might also be a hidden cost. By reducing the flexibility of the firm to respond to economic downturns, it may make the firm riskier.

While operating leverage affects betas, it is difficult to measure the operating leverage of a firm, at least from the outside, since fixed and variable costs are often aggregated in income statements. It is possible to get an approximate measure of the operating leverage of a firm by looking at changes in operating income as a function of changes in sales.

Degree of Operating leverage = % Change in Operating Profit / % Change in Sales

For firms with high operating leverage, operating income should change more than proportionately, when sales change.

Can firms change their operating leverage? While some of a firm's cost structure is determined by the business it is in (an energy utility has to build expensive power plants, and airlines have to lease expensive planes), firms in the United States have become increasingly inventive in lowering the fixed cost component in their total costs. Labor contracts which emphasize flexibility and allow the firm to make its labor costs more sensitive to its financial success, joint venture agreements, where the fixed costs are borne by someone else, and sub-contracting of manufacturing, which reduce the need for expensive plant and equipment, are only some of the manifestations of this phenomenon. While the arguments for such actions may be couched in terms of competitive advantage and flexibility, they do reduce the operating leverage of the firm and its exposure to "market" risk.

Degree of Financial Leverage Other things remaining equal, an increase in financial leverage will increase the equity beta of a firm. Intuitively, the obligated payments on debt increase the variance in net income, with higher leverage increasing income during good times and decreasing income during economic downturns. If all of the firm's risk are borne by the stockholders (i.e., the beta of debt is zero)[6], and debt has a tax benefit to the firm, then,

$$\beta_L = \beta_u \; (1 + (1-t)\,(D/E))$$

where

β_L = Levered Beta for equity in the firm

β_u = Unlevered beta of the firm (i.e., the beta of the firm without any debt)

t = Corporate tax rate

D/E = Debt/Equity Ratio

The unlevered beta of a firm is determined by the types of the businesses in which it operates and its operating leverage. Thus, the equity beta of a company is determined both by the riskiness of the business it operates in, as well as the amount of financial leverage risk it has taken on.

The equity betas estimated for highly levered firms in practice tend to be much lower than the betas estimated from the levered beta equation developed in the preceding section. This difference can be attributed to one or more of the following factors:

a. The beta estimated from a regression of past returns on a stock against market returns will lag the true beta, when the leverage change is recent. This is because the returns are estimated over a long time period (2 to 5 years) and reflect the average leverage over the period rather than the most recent leverage. This can be fixed fairly easily by

---

[6] If debt has market risk (i.e., its beta is greater than zero), this formula can be modified to take it into account. If the beta of debt is β_D , the beta of equity can be written as:

unlevering the beta using the average debt/equity ratio over the regression period, and then relevering it back using the current debt/equity ratio.

a. The equity betas estimated on the assumption that debt has no market risk will overstate the true betas because debt does bear market risk especially at higher debt ratios. This problem can be resolved by estimating a debt beta and calculating the equity beta using the expanded equation

$$\beta_L = \beta_u \ (1+(1-t)(D/E)) - \beta_D \ (1-t)D/E$$

This spreadsheet allows you to estimate the unlevered beta for a firm and compute the betas as a function of the leverage of the firm.

Bottom Up Betas

Breaking down betas into their business, operating leverage and financial leverage components provides us with an alternative way of estimating betas, where we do not need past prices on an individual firm or asset to estimate its beta.

To develop this alternative approach, we need to introduce an additional feature that betas possess that proves invaluable. The beta of two assets put together is a weighted average of the individual asset betas, with the weights based upon market value. Consequently, the beta for a firm is a weighted average of the betas of all of different businesses it is in. Thus, the bottom-up beta for a firm, asset or project can be estimated as follows.

$\beta_L = \beta_u \ (1+(1-t)(D/E)) - \beta_D \ (D/E)$

1. Identify the business or businesses that make up the firm, asset or project.

2. Estimate the unlevered beta(s) for the business or businesses that the firm is involved in. Table 1 summarizes average betas, market debt ratios and unlevered betas by industry in the United States as of March 1997. The simplest approach uses these unlevered betas directly, without adjusting for any differences between the firm being analyzed and the average firm in the sector. When we do this we implicitly assume that all firms in a sector have the same operating leverage. Given that smaller firms tend to have a greater proportion of fixed costs than larger firm, a more discriminating approach would require that we do one of the following:

   - Assume that market capitalization and operating leverage are correlated, and use the unlevered beta of firms with similar market capitalization in estimating the unlevered beta.

   - Calculate the operating leverage of the division or firm being analyzed and compare it to the operating leverage of comparable firms. If the firm being analyzed has a higher proportion of fixed costs than the comparable firms, the unlevered beta should be adjusted upwards (downwards).

3. To calculate the unlevered beta for the firm, take a weighted average of the unlevered betas, using the market values of the different businesses that the firm is involved in. If the market value is not available, use a reasonable proxy such as operating income or revenues.

4. <u>Calculate the leverage for the firm</u>, using market values if available. If not, use the target leverage specified by the management of the firm or industry-typical debt ratios.

5. <u>Estimate the levered beta</u> for the firm (and each of its businesses) using the unlevered beta from step 3 and the leverage from step 4.

This approach provides much better beta estimate for firms for three reasons. The first is that we estimate the unlevered betas, by sector, by averaging across regression betas. While regression betas are noisy and have large standard errors, averaging across regression betas reduces the noise in the estimate. In fact, the standard error of the average beta can be approximated as follows:

$$\text{Standard Error}_{\text{Average Beta}} = \frac{\text{Average Standard Error}_{\text{Beta Estimate}}}{\sqrt{n}}$$

where n is the number of firms in the sector. To illustrate, consider the software sector. Assume that the average standard error for betas estimates in this sector is 0.50, and that there are 225 firms in the sector. The standard error of the average beta estimate can then be estimated as follows:

$$\text{Standard Error}_{\text{Average Software Beta}} = \frac{\text{Average Standard Error}_{\text{Beta Estimate}}}{\sqrt{n}} = \frac{0.50}{\sqrt{225}} = 0.03$$

The second advantage is that the beta estimates reflect the firm as it exists today, since it is computed based upon current weightings for different businesses. In fact, expected changes in business mix can be reflected in beta estimates quite easily with bottom-up betas. The final advantage is that the levered beta is computed using the current financial leverage of the firm, rather than the average leverage over the period of the regression. Thus, the beta can be estimated more accurately for firms which have changed their debt/equity ratio in recent periods.

This approach is generic and can be applied to firms in any market. In fact, bottom-up betas can be estimated for emerging market companies. When making the bottom-up estimates for these firms, however, we have two choices:

- Use bottom-up betas for firms in the same business in the United States, and assume that the relative risk of firms is similar across markets. Thus, if oil companies have an unlevered beta of 0.60 in the United States, the unlevered beta of YPF, the oil company from Argentina, is assumed to be similar.

- Use bottom-up betas for firms in the same business globally, using betas estimated against a common index. While this approach may be more appealing than the first, it is also much more data intensive since service beta estimates for these firms tend to use local indices. This approach becomes far more practical if we use betas estimated for firms against local indices[7], and average out across these betas.

This data set on the web has updated betas and unlevered betas by business sector in the United States.

[7] As mentioned earlier in this paper, betas estimated against local indices can be skewed by the fact that these indices are dominated by one or a few stocks. Again, we draw on the fact that averaging across a large number of beta estimates can reduce or even eliminate some of these estimation errors.

Summing Up

Models of risk and return in finance require us to estimate the exposure of a firm to market risk, relative to other firms in the market. The conventional estimate of this relative risk, measured by regressing stock returns against a market index, is flawed for three reasons - the market index can be dominated by a few stocks, the beta estimate can be noisy and the firm itself might have changed during the course of the regression. While these regression betas can be modified to reflect financial fundamentals and there exist measures of relative risk that do not require a regression, the bottom-up approach has the most promise when it comes to delivering updated betas for most firms. In the bottom-up approach the beta for a firm is estimated as the weighted average of the unlevered betas of the different businesses that the firm operates in, adjusted to reflect both the current operating and financial leverage of the firm.