Inputs Methodology

Prepared for

Portfolio Strategist

September 2007

IbbotsonAssociates
225 North Michigan Avenue Suite 700 Chicago, IL 60601-7676 • (312) 616-1620

Table of Contents

Forecasting Expected Returns

Ibbotson uses a “building block” approach to generate expected return estimates. Historical data are used to calculate risk premia. These premia are added to the risk-free rate to arrive at the expected return unique to each asset class. This approach separates the expected return of each asset class into three components.

Building-Block Component	Description
Real Risk-Free Rate	A return that can be earned without incurring any default or inflation risk.
Expected Inflation	The additional reward demanded to compensate investors for future price increases.
Risk Premia	The additional reward demanded for accepting the return uncertainty associated with investing in a given asset class.

When choosing a risk-free rate, Ibbotson uses treasury yield curve rates with a maturity to match the investment period. The following table outlines the risk-free rates that are applied to various time horizons:

Risk Free Rate

When choosing a risk-free rate, Ibbotson uses treasury yield curve rates with a maturity to match the investment period. The following table outlines the risk-free rates that are applied to various time horizons:

Time Horizon	Years to Maturity	Yield*
Long-term	20	4.91%

*- All data are from the Treasury Department website as reported for December 31, 2006

The risk premia are derived from the historical relationship between the returns of the asset class and the risk-free rate, in this way, past data are incorporated into the assumption of the future returns. Various premia are added to the current risk-free rate in order to forecast the expected return unique to each asset class.

The risk premia are derived from a historical analysis of asset class returns relative to a risk-free investment. In this way, past data (which includes a number of economic scenarios) are incorporated into the assumption of the future returns. Premia are estimated as the difference between returns of certain benchmarks over time. Various premia are added to the current risk-free rate in order to forecast the expected return unique to each asset class.

Ibbotson calculates the risk premia by comparing the historical average returns of the appropriate benchmarks. The average returns are calculated over annual periods and may be income returns or total returns depending on the nature of the benchmark. In general, total returns are used for the equity benchmarks, whereas income returns are used for fixed income benchmarks. Total return is composed of the capital appreciation component and the income return. Income returns are the cash flows actually received by the investor via coupon payments, dividend streams, or other direct payments. Income returns are used for fixed income asset classes because they represent an unbiased estimate of what investors expect to receive for holding these investments. The realization of capital gains and losses is assumed to sum to zero over the time horizon of the investment.

In developing premia, the arithmetic average, as opposed to geometric average, is more appropriate for forecasting. The arithmetic average is the simple average of a return series. This measure incorporates the volatility of the returns (the risk) into expectations for the future and represents the center of the probability distribution. The geometric return is a backward looking statistic and appropriate when measuring actual historical performance. Although a geometric average may be more appropriate for historical numbers, the premia used by Ibbotson are forecasts. Therefore, arithmetic averages are used.[1]

Domestic Equity

Ibbotson estimates the additional return investors holding stocks of various sizes and styles expect to receive over the domestic risk-free rate when forecasting the expected return of the domestic equity asset classes.

The *domestic equity risk premium*[2] is the foundation for the expected returns of domestic equity asset classes. This premium is the excess return gained from holding domestic large-cap equities rather than risk-free securities. Ibbotson estimates the domestic equity risk premium by subtracting the historical arithmetic mean income return of long-term U.S. government bonds from the arithmetic mean total return of large-cap equity represented by the CRSP Deciles 1-2.[3] While the historical risk-free rate is used in the calculation of premia, the current risk-free rate is used to generate expectations of the future. The equity risk premium is applied to all domestic equity asset classes. All domestic equity expected returns are calculated by using the risk-free rate as a base, adding the equity risk premium, and then adding additional premia based on the asset class market capitalization or size, and style.

The calculation of the equity risk premium is outlined below:

Domestic Equity Risk Premium	= Large-Cap Stocks Total Return	- Long-Term Government Bond Income Return

[1] See *Appendix A-Statistical Methods* for arithmetic and geometric average calculations and further description.
[2] All domestic equity premia are estimated using historical data since 1926, unless otherwise noted.
[3] The CRSP Deciles 1-2 is a market-weighted portfolio of domestic equity securities that represent the first and second deciles of all securities by capitalization. See *Appendix B-Benchmark Clarifications* for more information on the construction of CRSP deciles.

Ibbotson uses a *size premium* to quantify the reward associated with investing in domestic equity asset classes with a market capitalization different than that of the CRSP deciles 1-2. Since smaller capitalization equities generally exhibit greater volatility, investors need to be compensated for the additional risk assumed. Investors expect compensation for additional risk in the form of higher expected returns. Empirically, these additional returns are not captured in the domestic equity risk premium estimate. Ibbotson estimates the size premia by taking the difference between the CRSP Deciles 1-2 and the smaller capitalization domestic equity benchmark. Following is the general formula for the size premium:

Size Premium	=Smaller Capitalization Domestic Equity Benchmark Total Return	-CRSP Deciles 1-2 Total Return

A size premium is included when estimating expected returns for the large-cap asset classes due to the nature of their benchmarks. The Russell 1000, contains stock with capitalizations that are different than the large-cap measure (CRSP 1-2) used in the domestic equity risk premium.

The expected return of large-cap growth and value equity represented by the Russell series is estimated as the sum of the risk-free rate, the domestic equity risk premium, a size premium and style premia.[4] Ibbotson estimates the size premium as the difference between the arithmetic mean total return of the composite asset-class benchmark (Russell 1000) and the CRSP Deciles 1-2 total return from 1926-present.

The expected return of the small and mid-cap equity asset classes are estimated as the sum of the risk-free rate, the domestic equity risk premium, and a size premium. The indices that compose the designated benchmarks for these equity asset classes do not have total return history that exists since 1926. These series must be extended to cover the full relevant period. Ibbotson uses returns-based style analysis to determine a proxy portfolio of CRSP Deciles that will represent the short-lived benchmarks for the entire relevant period. The CRSP proxy portfolios resulting from the returns based style analysis of the Russell MidCap and Russell 2000 Indices are averaged to produce a synthetic index composed of the appropriate proportion of the CRSP Indices to proxy the blended asset class.[5] The size premium is computed as the difference between corresponding estimated CRSP proxy and the CRSP Deciles 1-2.

Each of the specific size premia calculations are outlined below:

Large-Cap Size Premium	= Russell 1000 Total Return	- CRSP Deciles 1-2 Total Return
Small & Mid-Cap Size Premium	= Corresponding Estimated CRSP Proxy Total Return	- CRSP Deciles 1-2 Total Return

[4] Details on the development of the style premia can be found later in this section.

[5] When computing risk premia for a benchmark with a limited data history a synthetic series must be constructed. The series is composed of the CRSP Deciles 1-2, 3-5, 6-8, 9-10. The weightings assigned to each decile are determined by a regression analysis of the benchmark against the deciles, and then rounded to the nearest integer. See *Appendix C-Input Clarifications* for further explanation.

Ibbotson uses a *style premium* to quantify the reward for investing in value and growth securities. Relative to growth, value securities are generally defined as having low price/earnings or price/book ratios. Over the period 1975-present, value stocks have outperformed growth stocks in the United States. The additional return attributable to value stocks is not represented in the equity risk premium. Therefore, a style premium must be added.

The table below summarizes the computations and benchmarks used in estimating the style premia.

Ibbotson Style Premium	=		Ibbotson Growth or Value Benchmark Total Return	–	(50% IA Growth Total Return + 50% IA Value Total Return)	
			Benchmark Cap Range Weight		Ibbotson Style Premium	
Style Premium	= Sum	(	CRSP 1-2 (L) CRSP 3-5 (M) CRSP 6-8 (S)	X	L M S	)

Style Premiums are calculated in a series of steps:

1. Ibbotson Style Premiums are calculated for large, mid- and small cap using Ibbotson Style Indices[6].
2. Capitalization range exposure is estimated for each core benchmark (S&P 500, Russell 1000 etc.) using returns-based style analysis against the CRSP size groupings that correspond to the Ibbotson Style Series over the life of the benchmark.[7]
3. Capitalization range weights are applied to the corresponding Ibbotson Style Premiums to calculate the style premium for each benchmark.

Style Premium calculations comprise the complete history of the Ibbotson Style Series from 1969-Present. This method is used in lieu of the straight premium approach (using actual style benchmarks) in order to provide a common period of measurement and more consistent style premium measures across benchmarks.

The real estate equity premium measures the return gained or lost from holding real estate equity rather than large-cap stocks. Ibbotson estimates this premium as the difference in historical returns of the REIT asset-class benchmark and large-cap stocks. The real estate equity premium is estimated since 1972 due to the lack of REIT data. The following table summarizes the real estate equity premium.

Real Estate Equity Premium	=	NAREIT-Equity Index Total Return	- CRSP Deciles 1-2 Total Return

The table below summarizes the expected returns of the domestic equity asset classes, illustrating how they are constructed using the building block approach.

[6] Ibbotson growth and value indices are constructed using CRSP decile groupings: large-cap (deciles 1-2), mid-cap (deciles 3-5), and small-cap (deciles 6-8). Each size grouping is evenly divided into growth and value series based on book-to-price ratio. Portfolios are rebalanced annually in June. Please see Appendix A for more information on CRSP series.

[7] See Appendix B for an explanation on return-based style analysis using CRSP series.

Large-Cap Growth Equity	= Risk-Free Rate	+ Domestic Equity Risk Premium	+ Large-Cap Size Premium	+ Large-Cap Growth Style Premium
Large-Cap Value Equity	= Risk-Free Rate	+ Domestic Equity Risk Premium	+ Large-Cap Size Premium	+ Large-Cap Value Style Premium
Small and Mid-Cap Equity	= Risk-Free Rate	+ Domestic Equity Risk Premium	+ Small or Mid-Cap Size Premium	
Real Estate Equity	= Risk-Free Rate	+ Domestic Equity Risk Premium	+ Real Estate Premium	

International Equity

The foundation for the expected returns of the international equity asset classes is the *international equity risk premium*. The international equity risk premium measures the reward investors receive for holding assets that are domiciled outside the United States. Due to the lack of performance history for international benchmarks, the premium cannot be calculated by using the difference between historical returns with sufficient confidence. Therefore, Ibbotson uses the International Capital Asset Pricing Model (CAPM) to calculate the international equity risk premium.

To use CAPM it is first necessary to calculate the *world equity risk premium.* This premium is the incremental reward demanded by investors for holding a complete basket of risky assets worldwide over the U.S. risk-free rate. This relationship is computed indirectly. Dividing the domestic equity risk premium by the sensitivity of domestic markets to the world equity market derives the world equity risk premium. This method provides a stable "anchor" for the world equity risk premium because it uses the domestic equity risk premium, which implies a data history extending back to 1926.

Using the world equity risk premium and the relationship between international markets and worldwide markets, Ibbotson calculates the international equity risk premium. This removes the market effects of the United States from the premium. The international equity risk premium is calculated by multiplying the world equity risk premium by the sensitivity of international markets to the world equity market.

The relationships previously mentioned are calculated using regression analysis. Ibbotson identifies the relationship between domestic and world markets by regressing the monthly returns of domestic large-cap equity is the CRSP Deciles 1-2, and the benchmark for world equity is the MSCI World Index. The domestic equity risk premium is divided by this resulting beta in order to approximate the world equity risk premium.

The relationship between international markets and worldwide markets is found by regressing the monthly returns of international equity against the monthly return of world equity. The benchmark for international equity is the MSCI EAFE Index. This beta is then multiplied by the world equity risk premium in order to approximate the international equity risk premium.

The following table displays how the expected returns of the international equity asset class is constructed using the building block approach.

International Equity Risk Premium $= \dfrac{\text{Domestic Equity Risk Premium}}{\text{Beta}_{\text{CRSP 1-2 vs. World}}} \times \text{Beta}_{\text{EAFE vs. World}}$

Both regressions use the period 1970-present to compute the international and world betas. The regression equations are as follow:

CRSP Deciles 1-2	$= \alpha + \beta$ (MSCI World) $+ u_I$
MSCI EAFE	$= \alpha + \beta$ (MSCI World) $+ u_I$

International Large-Cap Equity	= Risk-Free Rate	+ International Equity Risk Premium

The emerging markets asset class uses the international risk premium in the calculation of expected return. To incorporate the added return emerging markets investors have received over investors focused on developed countries, an *emerging markets premium* is added.

The preferred method for estimating the emerging markets premium is to take the difference in the return of emerging markets and developed international markets over the longest period for which there is reliable data. However, because the emerging markets series does not begin until 1989, Ibbotson does not have a high degree of confidence in the straight premium approach.

As an alternative method, Ibbotson extrapolates the emerging markets premium from the risk/return trade-off of developed international equities. Ibbotson assumes that emerging markets equity is simply a riskier form of developed international equities. Ibbotson assumes that emerging markets equity is simply a riskier form of developed international equities. The emerging markets risk premium is approximated as a linear function of expected return and standard deviation in relation to that of cash and developed international equity. Graphic presentation of this concept demonstrates that the points for cash, developed international equities, and emerging markets equity fall on a straight line as depicted in the figure below.

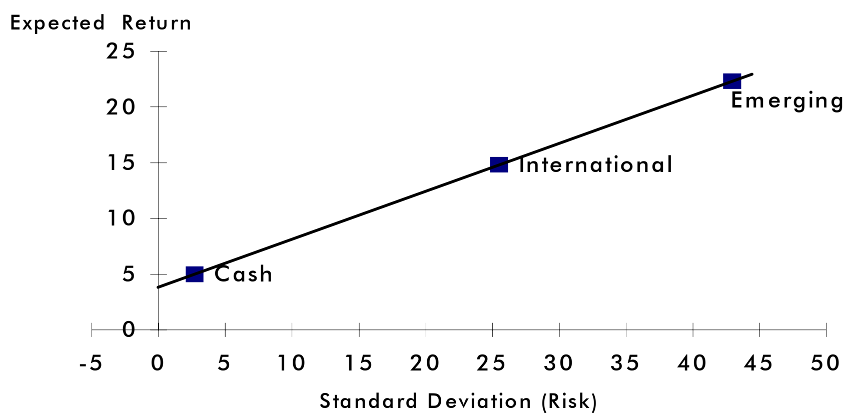

In this model, the formula for the emerging markets premium is as follows:

Emerging Markets Premium	= Int. Eq. Premium – Cash Horizon Prem. S.D. of Int. Eq. – S.D. of Cash	X (S.D. of Emer. Mkts. – S.D. of Int. Eq.)

In applying the above formula, Ibbotson uses the estimates of standard deviation for cash, international equity, and emerging markets equity as developed later in this report.

The following table displays how the expected returns of the international equity asset classes are constructed using the building block approach.

Developed International Equity	=Risk-Free Rate	+ International Equity Risk Premium	
Emerging Markets Equity	=Risk-Free Rate	+ International Equity Risk Premium	+ Emerging Markets Premium

Fixed Income

Ibbotson calculates the premia fixed income investors demand for holding bonds of varying maturity and credit quality when estimating the expected return of fixed income asset classes. All fixed income premia are measured using data since 1970 (unless otherwise noted). Ibbotson feels that data collected before this date do not present a good estimate of current and future fixed income market conditions.

The premia for holding bonds of varying maturity are the *horizon premia.* They measure the excess yield fixed income investors expect to receive in exchange for the additional uncertainty associated with longer time periods and the potential loss of liquidity. The premia for holding bonds of different credit quality are referred as *default premia.* They measure the historical premia demanded for holding corporate bonds that have a risk of default, rather than default-free government bonds of the same maturity and other similar characteristics.

In calculating the fixed income premia, Ibbotson assigns each fixed income asset class a *government bond proxy*[8] that has a maturity equivalent to that of the asset-class benchmark. The government bond proxy attempts to show how a government bond with the same maturity as the asset-class benchmark would have behaved over the full, relevant time period.

Ibbotson attempts to estimate asset class maturity using its benchmark's average maturity over time, rather than its current maturity. Average maturity is an estimate of the benchmark maturity throughout history and is used because the maturity of many benchmarks changes over time. Average maturity considers the dynamic nature of many fixed income benchmarks, whereas current maturity does not.

The government bond proxy is a portfolio of government bond benchmarks that best explains the historical return behavior of the asset-class benchmark. The government bonds used in this analysis are Coleman, Fisher, Ibbotson (CFI) estimated bonds.[9]

In most cases, Ibbotson determines the government bond proxy using returns-based style analysis, a method that compares the returns of a dependent variable (the fixed income benchmark) to the returns of several independent variables (the CFI bonds). Ibbotson uses the CFI bonds for three reasons:

1. Data history extends over the full, relevant period (1970 to present).
2. Estimated nature of the CFI bonds allows Ibbotson to control a number of key variables that affect bond return, including maturity.
3. Both total return and income return data are readily available.

The CFI bond methodology uses historical maturity of 90 days. Ibbotson uses a single CFI bond based on the current maturity of the cash benchmark (90 days) as the asset-class government bond proxy.

The *horizon premia* are the foundations for the fixed income expected returns. As previously mentioned these premiums measure the excess yield long-term fixed income investors expect to receive in exchange for additional uncertainty and the potential loss of liquidity. Ibbotson estimates these premia as the difference in income returns between two government bonds. The first government bond has the same maturity as the benchmark being modeled (the government bond proxy). The second is a CFI bond with a maturity equivalent to the investment horizon. In this case, a 20-year CFI bond was used to approximate the investment horizon. The horizon premia measure excess returns based on bond maturity not credit quality or other factors.

[8] See *Appendix C-Input Clarifications* for a detailed explanation of the government bond proxy specific to each asset class benchmark.
[9] See *Appendix B-Benchmark Clarifications* for a complete description of the Coleman, Fisher, and Ibbotson (CFI) methodology.

The specific premia calculations are outlined below.

Horizon Premium	= CFI Government Bond Proxy[A] Income Return	- CFI Long-Term Government Bond[B] Income Return

[A] Same maturity (average or current) as the asset-class benchmark
[B] Same maturity as the time horizon, i.e. 20 years

The expected return for the U.S. high yield fixed income asset class is the sum of the risk-free rate, the horizon premium, and the default premium. The horizon premium is the income return difference between a CFI government bond with an equivalent maturity to high yield and a twenty-year CFI government bond. The default premium is the difference between the total return of the U.S. high yield bond benchmark, LB high yield index, and the income return of a CFI government bond with an equivalent maturity to the high yield index. The corporate exposure is 100%; therefore, only the difference is necessary.

International bonds may require a *sovereign risk premium* that reflects the inherent differences between government bonds of different nations. The return differential between domestic and foreign bonds is attributable to both currency movements and the risk that foreign government may default on the bond (sovereignty risk). These effects, however, cannot be disentangled from return data alone. In order to isolate these effects, additional data and the development of a model would be necessary that is beyond the scope of the current project. As such, Ibbotson determines the expected return for international bonds based on the risk-free rate and a horizon premium; a sovereign risk premium is not included.

The expected return calculations for each taxable fixed income asset class are outlined in the following table.

Cash Equivalents	= Risk-Free Rate	+ Horizon Premium	
Short-Term Bonds	= Risk-Free Rate	+ Horizon Premium	
Intermediate-Term Bonds	= Risk-Free Rate	+ Horizon Premium	
Long-Term Bonds	= Risk-Free Rate	+ Horizon Premium	
High Yield Bonds	= Risk-Free Rate	+ Horizon Premium	+ Corporate Default Premium
International Bonds	= Risk-Free Rate	+ Horizon Premium	

Tax Exempt Bonds

The historical municipal bond indices are inappropriate measures for calculating premia associated with tax-free bonds, due to the shifting of tax rates over time. To determine an expected return for municipal bonds, Ibbotson first examines the current, general relationship (as expressed by yields) between tax-free and taxable bonds. Ratios are calculated between the current yields of the municipal bond series and the current yields of similar maturity government bonds. One minus these ratios represents the marginal tax rate and other factors (i.e. call and/or default premia) currently priced into municipal bonds. The ratios of the two yields are multiplied by the expected returns of government bonds (calculated based on the principles discussed above in the "Taxable Fixed Income" section), having similar maturities as the municipal series, to arrive at expected returns for municipal bonds and are specified below:

Expected Return of Municipal Fixed Income	= Current Municipal Fixed Income Yield / Current Government Fixed Income Yield*	X Expected Return of Taxable Fixed Income
Where: Expected Return of Taxable Fixed Income	= Risk-Free Rate	+ Bond Horizon Premium

* Having similar maturity as the municipal fixed income benchmark

Forecasting Standard Deviation

Mean-variance analysis requires a quantifiable risk measure for each asset class. Ibbotson uses standard deviation to estimate the risk of each asset class. Standard deviation measures dispersion around an average return.

Ibbotson uses historical data to forecast standard deviation because it provides an unbiased estimate of future volatility. Ideally, Ibbotson uses historical standard deviations using all available and relevant data (1926 and 1970 for equity and fixed income, respectively).

Ibbotson uses the *ratio method* to extend the standard deviation estimates of the shorter-lived asset classes so that they incorporate all relevant economic events. The ratio method attempts to extend the standard deviation estimate for certain asset-class benchmarks using two data series:

1. **Short benchmark** – an asset class benchmark used in this project that does not have historical data over the full, relevant period.
2. **Long proxy** – an index that has historical data over the full, relevant period and is economically similar to the short benchmark, i.e. there is a logical reason to believe that the returns of the two series are highly related.

Ibbotson compares the standard deviation of the short benchmark to that of the long proxy over a common time period via a ratio. The common time period is the inception of the short benchmark (unless otherwise noted). Ibbotson assumes that the relationship between the short benchmark and long proxy is representative of what it would have been if both series had existed over the full relevant period.

The ratio is multiplied by the standard deviation of the long proxy measured over the full, relevant period. The product is an estimate of what the standard deviation of the short benchmark would have been had it existed over the full, relevant period.

Extended Standard Deviation	$= \dfrac{\text{Std Dev of Short Benchmark}_{\text{Short}}}{\text{Std Dev of Long Proxy}_{\text{Short}}}$	X Std Dev of Long $\text{Proxy}_{\text{Long}}$

Equity

Ibbotson uses historical data and the ratio method to estimate the standard deviation of some equity asset classes. The equity asset classes utilized in this project do not have historical return data available from 1926. Ibbotson uses the ratio method to extend the standard deviation estimates of these short-lived asset classes.

The large-cap growth and value asset classes represented by the Russell 1000 growth and value indices only have history from 1979. The series composing the Mid and Small Cap Equity asset classes (Russell MidCap and Russell 2000) also only have return history since 1979. The first step is to find a suitable proxy with data for the full, relevant period. Ibbotson uses returns-based style analysis to estimate a portfolio of CRSP Deciles to represent the Large-Cap growth and value, Mid-Cap and Small-Cap proxies. The resulting portfolios of CRSP Deciles (rounded to the nearest integer) serve as the long proxy in the extension process.

Ibbotson also extends the standard deviation of the MSCI EAFE using the ratio method. The standard deviation of the asset-class benchmark is measured over its lifetime (1970 to present) and compared to the standard deviation of large-cap stocks (CRSP Deciles 1-2). This ratio is multiplied by the full period standard deviation of large-cap stocks.

The standard deviation of the IFCI Emerging Composite is also calculated using the ratio method. The standard deviation is compared to the .

Ibbotson estimates the standard deviation of REITs using the ratio method. However, Ibbotson maintains that fundamental changes in the real estate market have changed the risk characteristics of the asset class. Ibbotson has determined that the full, relevant period for REIT risk estimates begins in 1992. Research conducted by Ibbotson and others indicates that investor perceptions about REITs have changed over time. Ibbotson hold that REIT data prior to 1992 is not relevant when estimating future economic conditions.[10]

The long proxies and common periods used to calculate standard deviation estimates for the equity asset classes are outlined in the following table on the next page.

Asset Class	Short Benchmark	Common Period	Long Proxy
Large-Cap Growth Equity	Russell 1000 Growth Index	1979	Extended Russell 1000 Proxy
Large-Cap Value Equity	Russell 1000 Value Index	1979	Extended Russell 1000 Proxy
Mid-Cap Equity	Russell MidCap	1979	Extended Russell MidCap Proxy
Small-Cap Equity	Russell 2000	1979	Extended Russell 2000 Proxy
Real Estate	NAREIT – Equity	1992	CRSP Deciles 6-10
International Large-Cap Equity	MSCI EAFE	1970	CRSP Deciles 1-2
Emerging Markets	IFCI Emerging Composite	1989	Extended IA Emerging Markets

Fixed Income

The methodologies used to estimate the standard deviation of the fixed income asset classes are the same as those used for equities. When an assigned benchmark does not have a history that covers the full, relevant period (1970 to present). Ibbotson uses alternative processes, including the ratio method, to extend the standard deviation estimates of cash and International/Long-term Municipal Bonds.

Ibbotson creates an adjusted series to estimate the standard deviation of cash equivalents. This adjusted series consists of a CFI 90-day bill (1970-1977) and the Salomon Brother 3-Month T-bill (1978-present). Ibbotson estimates the standard deviation of cash as the historical standard deviation of this series.
Ibbotson uses the ratio methodology for the International and Long-term Municipal Fixed Income asset classes. Ibbotson compares the asset-class benchmark to the asset-class government bond proxy.[11]

The standard deviation for the Short-term, Intermediate-term, Long-term, and High Yield Bonds is found using the benchmark extended from 1970 and forward.

The long proxies and common periods used to calculate standard deviation estimates for the fixed income asset classes are outlined in the following table.

[10] See Appendix E – Real Estate Investment Trust Behavior for more information.

[11] This is the same government bond proxy that was determined in estimating the horizon premium of the intermediate/long-term fixed income asset class, details can be found in *Appendix C*.

Asset Class	Short Benchmark	Common Period	Long Proxy
Cash	Adj. CG 3-Month T-bill	1970	N/A
Short-term Bonds	US 1 Yr Gvt. Bond	1970	N/A
Intermediate-term Bonds	US IT Gvt. Bond	1970	N/A
Long-term Bonds	US LT Gvt. Bond	1970	N/A
High Yield Bonds	LB Domestic High Yield	1970	N/A
International Bonds	CG Non-US 1 + Year Gvt.	1985	IA International Bond Composite
Long-term Municipal Bonds	Muni: LB 20 Yr Muni	1980	CFI Gvt. Bond Proxy (for LB 20 Yr. Gvt. TR)

Forecasting Correlation Coefficients

In the mean variance analysis setting, the standard deviation of a portfolio is based not only on the risk of each asset class, but on the relationship between the returns of asset classes as well. The relationship between the returns of asset classes is measured by the correlation coefficient.

The correlation coefficient measures the degree to which two asset classes' returns change with respect to each other. The statistic can range between positive one (+1.00) and negative one (-1.00) and provides the following information about the relationship between asset classes:

- **Positive One (+1.00):** perfect positive relationship – When the returns of one asset class increase, the returns of the other asset class increase.
- **Negative One (-1.00):** perfect negative relationship – When the returns of one asset class increase, the returns of the other asset class decrease.
- **Zero (0)**: no relationship – Changes in the returns of the asset classes are unrelated.

Ibbotson typically uses correlation coefficients derived from the historical returns of the asset class benchmarks. In this project, annual data is used.

Appendices

Appendix A – Statistical Methods Applied in Forecasting

Several statistical methods were used by Ibbotson to calculate arithmetic mean returns, standard deviations, and correlation coefficients. The standard statistical methods typically used for data over consistent units of time are also included. Whenever possible, Ibbotson uses annual data in the calculation of inputs. For indices or return series that have less than a 20-year history, however, monthly or quarterly data may be used in order to ensure a sufficient number of data points.

All data used in this project includes data current through year-end present. The commentary below describes how annual data has been incorporated into the estimation of optimization inputs.

Arithmetic Mean Return

The arithmetic mean return is the simple average of all the returns in a given period. To calculate the annual arithmetic mean, we need the returns.

$R_{(x),,}\ R_{(x)}(2),\ R_{(x)}(3),\ R_{(x)}(4)\ldots.R_{(x)}(t)$

Where:

$R_{(x)}$ is the return in year t for asset X

The arithmetic mean then, is the sum of all the returns divided by the total number of returns.

$$\text{Mean } \mu_{x*} = \frac{\sum_{t=1}^{n} R_x(t)}{n}$$

Where:

$R_x(t)$	is the return in year t
n	is the number of years in the period being measured
μ_{x*}	is the annual arithmetic mean return

Geometric Mean Return

The geometric mean of a return series over a period is the compound rate of return over the period. The geometric mean return equation is as follows:

Mean $\mu_{x^*} = \left[\prod_{t=1}^{n} (1 + R_x(t)) \right]^{1/n} - 1$

Where:

$R_x(t)$	is the return in year t
n	is the number of years in the period being measured
μ_{x^*}	is the annual geometric mean return

Standard Deviation

The standard deviation of a typical series is the square root of the sum of the squared differences between each return and the mean, divided by the number of data points less 1.

Standard Deviation $\sigma_x = \sqrt{\frac{\sum_{t=1}^{n} [R_x(t) - \mu_{x^*}]^2}{n-1}}$

Where:

$R_x(t)$	is the return in year t
n	is the number of years in the period being measured
μ_{x^*}	is the annual arithmetic mean return

Correlation Coefficients

Typically, the correlation coefficient between two series is determined if the covariance between the two series is calculated and the standard deviations of each are known. To find the covariance of two series, the cross product between series X and series Y must first be determined. The cross product of the differences between all X and Y values and their respective means.

Cross Product $SXY = \sum_{t=1}^{n} [R_x(t) - \mu_{x*}] \times [R_y(t) - \mu_{y*}]$

Where: $R_x(t)$ is the return on asset X in period t

μ_{x*} is the arithmetic mean return of asset X

$R_y(t)$ is the return on asset Y in period t

μ_{y*} is the arithmetic mean return of asset Y

The cross product is used in the calculation of the covariance. The covariance is the cross product of X and Y divided by the number of annual periods minus 1 (n-1).

Covariance Cov (X,Y) $= \frac{SXY}{n-1}$

The covariance is used in the calculation of the correlation coefficient. The correlation coefficient is the linear association between two variables. The coefficient lines between the values of +1 and –1. A correlation of +1 indicates a perfect positive association and –1 a perfect negative association.

Correlation Coefficient $\rho_{x,y} = \frac{\text{Cov (X,Y)}}{\sigma_{\text{x}} \times \sigma_y}$

Where: σ_x is the standard deviation of asset X

σ_y is the standard deviation of asset Y

Cov (X,Y) is the covariance between asset X and asset Y.

Appendix B – Benchmark Clarifications

CRSP Deciles

The CRSP Deciles are taken from the Center for Research in Security Prices at the University of Chicago Graduate School of Business. Like the S&P 500, it is a market value-weighted benchmark of common equity performance.

Decile returns are value weighted and calculated monthly. Security weights are determined using market capitalization based on the shares outstanding and closing price for the last trading day of the previous month. Dividends and split factors are included in the month containing the ex-dividend date. Certain distributions such as spin-offs and rights are reinvested on the ex-dividend date.

The CRSP Cap-Based Indices universe includes common stocks listed on the NYSE, AMEX, and the NASDAQ National Market excluding the following: preferred stocks, unit investment trusts, closed-end funds, real estate investment trusts, americus trusts, international stocks, and American depository receipts.

All eligible companies listed on the NYSE are ranked by market capitalization on the last trading day of each quarter, and then split into ten equally populated groups, or deciles. The capitalization of the largest company in each decile serves as a breakpoint for that decile. When multiple issues of a company trade, the sum of the issue capitalization is used for the company capitalization so that each company is counted only once. The portfolios are reformed every quarter using the price and shares outstanding at the end of the previous quarter. During the quarter, companies move between deciles, since their market capitalization changes while the breakpoints of each decile remain fixed.

Breakpoints are based exclusively on companies with issues traded on the NYSE. For the CRSP series that include securities from the AMEX and NASDAQ over-the-counter market (these are the series used for the Portfolio Strategist project), non-NYSE companies are assigned to appropriate portfolios according to their capitalization in relation to the NYSE decile breakpoints. Thus, the series that include non-NYSE securities are not comprised of true deciles in the sense than an equal number of companies are represented in each of the then portfolios (this is only the case for the NYSE series).

Appendix C – Correlation Extension Process

Ibbotson uses a statistical procedure, described in this appendix, to estimate extended correlations among benchmarks and to develop adjusted historical returns and standard deviations for domestic short-history asset classes. The purpose of this procedure is to use additional available information to model the behavior of "short-history" benchmarks, those benchmarks for which we have less baseline data than for some "long-history benchmarks." The additional information in long-history benchmarks judged to be predictive of a short-history benchmark is incorporated into parameter input estimates through this process.

In addition to the problem of short-history benchmarks, the process of developing a consistent estimate for correlations between asset class benchmarks faces an additional challenge particular to the Ibbotson methodology. Ibbotson uses data going back to 1926 in developing expectations for future equity performance, and data going back to 1970 for fixed income analysis. Thus, correlations between fixed and equity benchmarks can only be based on data since 1970. Correlations among equity benchmarks, however, we believe, should be based on the longest time period available.

The idea behind the extension approach is to model each short-history benchmark as an optimal mix of long-history benchmarks. A statistical model is used to determine this mix. These models are used with the correlations between long-history benchmarks to imply the correlations between short- and long-history benchmarks. Further, the models are used imply one component of the correlations between two short-history benchmarks. The second component of the correlations between short-history benchmarks is determined by the residuals from the modeling process.

This appendix describes the extension process and considers its relationship to other suggested methods of adjusting historical data. Subsection 1 describes the steps of the extension process in more detail. Subsection 2 compares the Ibbotson extension process to some other methods that are used to adjust historical estimates. Subsection 3 provides a technical description of the extension procedure.

1. Conceptual Description of the Extension Process

Each short-history benchmark is modeled by at least one long history benchmark. The model is a simple linear model where the returns of the short-history benchmark are represented as a weighted linear combination of the returns of the long-history benchmarks plus a constant. The constant and the weights are determined by regression analysis. A technical description of the method is provided in Subsection 3(c).

The statistical model can be used to adjust a short-history benchmark's expected return. This is done by using the model to forecast returns for periods when the short-history benchmark did not exist. This is done by applying the weights and constant determined by the model to the long-history benchmarks. The actual and forecast benchmarks are then averaged together. A technical description of the method is provided in Subsection 3(d).

Extended correlations are derived from an extended covariance matrix. The extended covariance matrix is constructed using the covariance matrix of the long-history benchmarks and the covariance matrix of the residuals from the model building process. The calculation of the covariance matrix of long-history benchmarks is a standard procedure and is described in Subsection 3(e). The residuals from the modeling process of Subsection 3(c) are important information regarding the nature of the departure of historical returns from the statistical model of those returns. The covariance of these residuals identifies correlations between short-history benchmarks that cannot be captured through the statistical models. Subsection 3(f) introduces notation relating to computing this matrix.

The extended covariance matrix is computed by combining the two covariance matrices with the statistical models. This process is explained in Subsection 3(g). The statistical models represent short-history benchmarks as linear combinations of long-history benchmarks. The covariances between long-history benchmarks then imply the covariances between short- and long-history benchmarks. The statistical models also imply a component of the covariances between the short-history benchmarks themselves. Finally, the extended correlation matrix is easily determined from the extended covariance matrix.

If the statistical models of the short-history benchmarks were perfect, then the covariances between them could be perfectly determined from the covariances between the long-history benchmarks that are used to model them. Similarly, covariances between short- and long-history benchmarks could also be perfectly inferred from the models. Since the models, however, are only approximations, the residuals from the modeling process must be incorporated into the estimation process as well.

2. Relationship to Other Methods

There are other methods of adjusting historical inputs. Shrinkage estimators, pioneered by Stein (1955), are a large class of such methods. The basic idea of shrinkage estimation is to adjust the mean of a subpopulation toward the mean of a larger, encompassing, population. A relevant example is that the expected return of a small-cap benchmark could be estimated by adjusting the observed historical return in the direction of the mean return for the market as a whole. This type of adjustment of optimization inputs is advocated by Jorion (1986), DiBartolomeo (1991), and Michaud (1998). The Ibbotson extension process has some similarities to and some differences with shrinkage estimation. This subsection provides a brief explanation of why we believe the extension process described here is a more appropriate way to generate optimization inputs.

Shrinkage estimation has some relation to the phenomenon of regression to the mean. Sports statistics provide a common example of regression to the mean. The difference between the best and worst performers can ordinarily be expected to decrease over the course of a season. If this is true, then adjusting early-season performance statistics toward the then-observed mean may provide a better prediction of end-of-season performance than simple averages. Shrinkage estimators provide methods for making this type of adjustment.

When applied to the generation of optimization inputs, shrinkage estimators will tend to adjust performance of benchmarks in the direction of overall market performance. This will reduce the extremes in expected returns, standard deviations, and correlations. The degree of adjustment will be related to the uncertainty associated with a benchmark relative to the market. For example, using the Stein (1955) shrinkage estimator, the degree of adjustment of a benchmark toward the expected market return will be proportional to the variance of the benchmark relative to the variance among benchmarks.

The Ibbotson method can be described in this framework as follows. First, instead of adjusting toward overall market behavior, adjustment is made toward a composite benchmark represented by a model or set of models described by equation (1). Then, instead of making the degree of adjustment proportional to a measure such as the variance of the short-history benchmark, the degree of adjustment essentially becomes proportional to the explanatory power (i.e., R^2) of the model. This is a different type of mechanism since the explanatory power of the model is not necessarily related to the variance of the short-history benchmark. Indeed, to the degree that they are related (e.g., higher variance benchmarks might be associated with lower model R^2 values), the Ibbotson adjustment operates on an opposing principle. The Stein estimator will tend to adjust more toward the mean under a high variance/low R^2 scenario, while the Ibbotson method will adjust less toward the composite benchmark implied by the model. Similarly, under a low variance/high R^2 scenario the Ibbotson method will adjust more toward the composite benchmark, while the Stein estimator will adjust less toward the market return.

The most important advantage of the Ibbotson approach is that it directly addresses the basic problem of the incorporation of information contained in long-history benchmarks into expectations for short-history benchmarks. This stands in distinction to the basic problem addressed by shrinkage estimation as described in the beginning of this subsection. In our judgment, the adjustment to represent an appropriate historical record is more fundamental and of larger absolute magnitude than adjustments to account for phenomena such as regression to the mean.

A further advantage of the Ibbotson method is that professional judgment enters into model development in a more understandable way, principally through the construction of the models of short-history benchmarks. Shrinkage estimators have varying and complex estimation methods. Choosing an appropriate estimator is a nontrivial task. Often, these estimators are Bayesian. This type of estimator requires the specification of appropriate prior beliefs. Making such probability assessments can be a difficult undertaking.

Appendix D – Forecasting the Inflation Rate

Since 1976, Roger Ibbotson and Rex Sinquefield, and later Ibbotson Associates have provided estimates of the market's long-term forecasts of asset class returns and inflation. The market's forecast of inflation is not directly observable. However, it can be inferred from current yields on Treasury bonds and the statistical time series properties of historical data using techniques first developed by Ibbotson and Sinquefield in "Stocks, Bonds, Bills, and Inflation: Simulations of the Future (1976-present)" *(Journal of Business*, July 1976). The methodology described below is Ibbotson Associates' most recent refinement of the Ibbotson-Sinquefield methodology as it applies to expected inflation.

A) Theory

The key insight in the analysis presented here is that investors' long-term inflation forecasts are embedded into long-term risk-free yields. Investors expect to be compensated for the lost purchasing power of the dollar over time.

Compensation for expected lost purchasing power is not the only component of Treasury bond yields. Bondholders also expect to be rewarded for foregoing real consumption for a period of time. This reward can be expressed as the expected real risk-free rate.

Bond yields do not remain constant through time. Since holders of long-term bonds typically do not hold their bonds until maturity, the variability of yields is a source of risk for bondholders. Consider an investor with a one-month investment horizon. The investor can purchase either a one-month Treasury bill (and lock-in a return with perfect certainty) or purchase a Treasury bond and face the risk that at the end of the month, the bond's yield will rise (causing a fall in its value). In order for the bond to be an attractive alternative to the bill, the expected return on the bond must be high enough to compensate the investor for the market risk of holding a longer-term instrument. The spread between the expected returns on bonds and bills leads to the embedding of horizon premia into bond yields.

Thus, observed yields on Treasury bonds are composed of three components: expected inflation, expected real risk-free rates, and horizon premia. None of the three components are observable. In order to estimate expected inflation, we estimate expected real risk-free rates and horizon premiums from statistical relationships evident in historical data and remove them from observed market yields. Expected inflation rates are the residuals from this process.

B) Forward Rates

Consider an investor with a 20-year time horizon. The investor can purchase a 20-year zero-coupon bond and lock-in a 20-year return. Alternatively, the investor can purchase a one-year bill and plan on rolling-over the proceeds into another one-year bill, repeating the process for twenty years. If the investor were not concerned about risk, he would be indifferent between these two strategies if they both had the same expected return. Under these conditions, the yield on the 20-year bond must be comprised of the investor's forecast of one-year yields for the next twenty years. This concept can be formalized with forward rates.

The yield on a zero-coupon bond with T years to maturity can be decomposed into T one-year forward rates as follows:

$$Y(T) = \sqrt[T]{[1+F(1)][1+F(2)]\cdots[1+F(T)]} - 1$$

Where: Y(T) = the yield on the bond; and
F(t) = the one-year forward rate which predicts what the yield on a one-year bond will be at the beginning of year t.

If investors were not concerned about risk, each forward rate would consist of a forecast of inflation and the real risk-free rate for its year. Since investors are concerned about risk, each forward rate also includes a horizon premium. To obtain inflation forecasts for each year in the future, we subtract an estimate of the expected real risk-free and horizon premium from each of the forward rates.

C) Continuously Compounded Rates

To simplify calculations, Ibbotson uses continuously compounded rates in all computations. This requires conversion of all variables from discrete rates to continuously compounded rates. Once the continuously compounded rate of inflation is forecast, Ibbotson converts it back to the more familiar discrete form.

The formula for converting a rate of return from discrete rate to continuous rate is:

$$r = \ln(1+R)$$

Where:
r = the continuously compounded rate; and,
R= the discrete rate.

Bond yields are typically expressed in semiannual form. The formula for converting a semi-annual yield to a continuous yield is:

$$y = 2\ln\left(1+\frac{Y}{2}\right)$$

Where:
y = the continuous yield; and
Y = the semiannual yield.

Appendix E – Real Estate Investment Trust Behavior

Real estate investment trusts are closed-end investments that invest primarily in real properties. Unlike direct investments in real estate, REITs are relatively liquid in that they can be easily bought and sold and do not require substantial initial investments. Furthermore, real estate investment trusts are traded on major stock exchanges. In the past, however, these qualities have dominated the behavior of real estate investment trusts. The performance of this asset class has been better explained by the equity markets these securities are traded in, rather than the underlying investments in real property.

In the past, the risk and return characteristics of real estate investment trusts have been very similar to small-cap stocks. The high correlation between REITs and domestic stocks indicates that market movements have affected real estate investment trusts. Early in the 1990's, however, the market's perception of these securities began to shift. In turn, shifts in the market began to alter the behavior patterns of this asset class. Since 1992, the REIT market has more than quadrupled and investors have begun to view these investments more as real estate and less as simply domestic equity.

To incorporate these changes in the behavior of real estate investment trusts, an alternative period is used as the basis for correlation and standard deviation estimates. Ibbotson believes that the NAREIT-Equity Index measured since 1992 is the most appropriate proxy for real estate investment trusts correlation coefficients and standard deviation estimates.