FAQs
Tips ‘n Techniques
Trouble-Shooting

For use with Expert Allocator and Expert Performance Analyst

by Investment Technologies

Telephone 212/724-7535 Fax 212/208-4384
Support Telephone 203/364-9915 Fax 203/547-6164
e-mail: support@InvestmentTechnologies.com Website: www.InvestmentTechnologies.com

Contents

Section 1: FAQ's on The Expert Allocator

Does correlation have the same effect on downside deviation as on standard deviation?

Yes, combining assets with less than perfect positive correlation reduces a portfolio's downside deviation.

The following statistics illustrate this beneficial effect using a two-asset model. Return and risk statistics were calculated under three scenarios: perfect positive correlation (+1.00), zero correlation, and perfect negative correlation (-1.00). Notice how negative correlation between the two assets shows the greatest reduction in portfolio risk, regardless of risk measure.

			Risk for 50/50 Mix where Correlation is:		
	Asset 1	Asset 2	+1.00	0.00	-1.00
Expected return	5.0	10.0	7.5	7.5	7.5
Standard deviation	5.0	10.0	7.5	5.6	2.6
Downside deviation (10% goal)	6.9	7.0	6.8	5.5	3.5

Does The Expert Allocator test the correlation matrix for positive definiteness or semi positive definiteness?

The Expert Allocator tests correlation assumptions for positive definiteness, which is a more stringent test than semi-definitiveness. If a correlation coefficient is inconsistent with one or more other values, the correlation matrix is not positive-definite and the Status indicator in the top left corner of the Correlations worksheet changes from "Consistent" to "Inconsistent."

While optimization is not prohibited so that you can determine the potential impact of the conflict, you may get invalid results or a system crash. As a practical matter, you should proceed very carefully if this constraint has been violated.

Note: All correlation matrices generated by the Forecasting Tools (Historical Returns and Economic Scenarios) are guaranteed to be positive definite.

In The Expert Allocator, how is a portfolio's standard deviation adjusted for different holding periods?

A portfolio's holding period standard deviation is calculated as follows:

Holding period standard deviation $= \sigma_p / \sqrt{hp}$

where:

σ_p = annual standard deviation,

hp = holding period, in years

For example, if a portfolio's annual standard deviation is 15.0%, then the standard deviation for a 5-year holding period is 6.71% (15.0/2.236) and for 10-year holding period, 4.74% (15.0/3.162).

How do I incorporate tax effects in my optimization?

Incorporate the effect of taxes on portfolio optimization in the Allocator as follows:

1. *Specify tax information in the Tax Details worksheet.* Click the Tax Rates tab to enter federal, state, and local taxes for capital gains and ordinary income. Click the Asset Proxy tab to indicate, for up to 14 asset proxies, the proportion of total (pre-tax) return from capital gains and/or income and annual turnover.

2. *Assign an asset proxy to each asset in the Risk and Return* worksheet for purposes of calculating after-tax values. To change the asset proxy assigned to an asset, double-click on the list box in the Asset Proxy column or click the down arrowhead to display the choices, and click on the desired setting. To view after-tax values, click After Tax Values.

3. *Select After-Tax as the Tax Environment option in the Optimization Parameters worksheet* to optimize using after-tax return and risk assumptions. Compare the effects of taxes by optimizing two frontiers, one with Pre-Tax and the other with After-Tax.

Does The Expert Allocator assume the portfolio rebalances at the end of each year or at the end of the holding period?

It is assumed that portfolios are rebalanced annually.

Are asset assumptions for expected return and standard deviation expressed in annual terms or in terms of the holding period? For example, for a 5-year holding period, would you input an annual standard deviation or a 5-year annualized standard deviation?

Asset assumptions should be expressed in annual (one year) format regardless of the optimization holding period.

Can I use downside deviation with normal return distributions?

Yes. Downside deviation measures a portfolio's risk for a given goal and holding period, regardless of the shape (normal or skewed) of the portfolio's return distribution. Standard deviation as an optimization risk measure, on the other hand, is typically used only in conjunction with the normal distribution.

Our client's target portfolio sat on the efficient frontier, but its asset mix was not the same as any of the efficient portfolios on the frontier even though it had the same expected return and risk as one of the efficient portfolios. Is this possible?

Theoretically there are an infinite number of portfolios along the efficient frontier. The Expert Allocator identifies up to 40 portfolios along the frontier. Therefore, it is possible to find a portfolio that lies along the frontier but is not included in the 40 portfolios identified by the optimizer.

How can I incorporate The Expert Allocator's charts and tables in my presentations?

Incorporating charts and tables in slide presentations, spreadsheets, or documents is fast and easy and creates a very effective presentation. Charts can be copied to the Clipboard or saved as a bitmap or metafile and then pasted into your document. Tables can be saved as Excel worksheets or text files and inserted into your documents.

Charts:

To copy a chart to the Clipboard or to a file, in the Analysis Window select Export from the File menu option, then select Pane1, Pane2, or Pane3 and one of the following:

- To Bitmap File (`BMP`) to save the chart as a bitmap (`.BMP`) file. The Export to Bitmap dialog box appears. Enter a filename in the File Name list box and click **OK**. The file is saved automatically with a `.BMP` extension.
- To Metafile (`WMF`) to save the chart as a metafile (`.WMF`) file. The Export to Metafile dialog box appears. Enter a filename in the File Name list box and click **OK**. The file is saved automatically with a `.WMF` extension.
- To Clipboard (`BMP`) to copy the chart to the Windows Clipboard as a bitmap.
- To Clipboard (`WMF`) to copy the chart to the Windows Clipboard as a metafile.

Tables:

To export table information shown in Pane 1, 2, or 3, choose Export from the File menu option and one of the following:

- to Text File to export a table to an ASCII text file. The Export to ASCII File dialog box appears. Type the name of the text file in the Filename box (a `.TXT` extension is automatically assigned) and press Enter or click **OK**.

- to Excel (XLS) to export a table to an Excel spreadsheet. The Export to Excel File dialog box appears. Type the name of the Excel spreadsheet in the Filename box (an .XLS extension is automatically assigned) and press Enter or click **OK**.

How do I create a model using my own return, risk and correlation assumptions?

To create a model and enter your own assumptions, click **New** on the Desktop. Then click **Risk and Return** to display the Risk and Return Assumptions worksheet.

Add asset classes to the model by clicking **Add**. Each newly created asset appears at the bottom of the asset class list, and the following default values are provided: 0.0% Expected Return, 0.01% Standard Deviation, -0.01% 90th Percentile Return, 0.01% 10th Percentile Return, Tax Free Asset Proxy, and perfect positive (+100) correlations.

You may change these values by entering new values in the respective worksheet cells. To change standard deviation, make sure Ret Distribution Shape is set to Normal. To change 90th and 10th Percentile Returns, make sure Ret Distribution Shape is set to True. Change default asset names by double-clicking the name and entering the desired name.

How can I change the Analysis Window to suit my preferences?

There are a number of things you can do to customize the look of the Analysis Window including customizing charts and tables and resizing panes.

Customizing Chart Displays:

You can customize a chart's appearance in the Analysis Window using style templates and the Customize Chart dialog box. Many of the chart's attributes, such as fonts, colors, and background, can be changed

Customizing Table Displays:

Use any of the following to change the table display:

- use the vertical and horizontal scroll bars in the Pane to scroll up and down or left and right.
- Click **Zoom** in the Control Panel to display Pane 1 full-screen.
- Enlarge the Pane 1 by dragging the split bar.

- Reduce the font size by selecting **Change Table Fonts** from the **Style** menu.
- Shorten the names of portfolio statistics by selecting **Display Abbreviated Statistics Names** from the **Style** menu.

Resizing Panes:

Each Pane can be resized to your preferences. Using your mouse, point to the split bar between the Panes. Hold down the left mouse button as you drag the split bar between Panes 2 and 3 up or down. Drag the split bar between Pane 1 and Panes 2 and 3 right or left. Resizing Panes does not affect printing.

To save the currently displayed pane sizes, select **Save Pane Sizes** from the **Style** menu. To restore previously saved pane sizes, select **Load Default Pane Sizes** from the **Style** menu.

Both Panes 2 and 3 disappear when you click the **Zoom** button in the Control Panel and Pane 1 is displayed full-screen. To redisplay Panes 2 and 3 when Pane 1 is zoomed, either click the **Unzoom** button, make a selection from the Pane2 or Pane3 menus, or select **Load Default Pane Sizes** from the **Style** menu to restore previously saved settings.

In The Expert Allocator, I want to create a model with the Historical Returns tool using the same series as both asset and benchmark. How can I do this?

Using the Historical Returns tool, a series cannot be both an asset and a benchmark simultaneously. The workaround is to copy the series—i.e., create a new one—directly in the Historical Returns tool. Here's how:

1. In the Historical Returns window, select all the series you want included in the model so that they appear in the Selected Series panel.
2. In the Benchmark drop-down list box, select the desired series for the benchmark.
3. Click **Weight**. The Weight Series window is displayed.
4. In the Selected Series panel, double-click the benchmark series.
5. Click **Equal Weight**.
6. In the New Name text box, name the series. Make sure you give it a different name than the original series.
7. Click **Save**.
8. Assign the appropriate categories in the Specify Categories dialog box and click **OK**.

9. When done, click **Close**. The Historical Returns window is displayed.

10. Select the newly-created series so that it appears in the Selected Series panel.

Now the original series will be the benchmark and the newly-created series will be the asset.

I've created a series using the Weight feature in Historical Returns. How can I update this series?

When you create a series using Weight, only the historical monthly (or quarterly) returns are saved—not the series and percentages used to create the new index. So you'll need to update the series as needed. To do so, use **Weight**—the same way you originally created the series—and overwrite the existing series. (Tip: For easy reference, name your series to reflect its composition, for example, 60% S&P/40% LB Agg.)

Can I override the asset assumptions calculated by The Expert Allocator's Forecasting Tools?

Yes. All return, risk and correlation values can be changed by the user in the Risk and Return worksheet and the Correlation worksheet.

Can I use both monthly and quarterly series in my analysis?

Yes. This means that you can mix manager results with index returns.

What's the difference between efficient and optimal?

All portfolios on the efficient frontier are efficient in that they offer the greatest expected return for each level of risk, but only one is optimal for a particular investment situation. Theoretically, the optimal portfolio is the portfolio that offers the greatest expected utility given the investor's risk profile. In other words, it best reflects the investor's trade-off between return and risk.

The Expert Allocator provides a number of ways to find and select the optimal portfolio.

1. If you know which portfolio is optimal, click **Go To** and enter the portfolio number.

2. Click **Find** to identify an efficient portfolio that meets specific criteria. The

   Find Efficient Portfolio dialog box appears. Select the desired frontier by clicking the appropriate radio button in the Frontier frame. This is the frontier on which you wish to identify an efficient portfolio. Then choose a portfolio statistic from the drop-down list box, and the minimum and maximum values (based on the frontier's optimization parameters) for efficient portfolios on the selected Frontier will be displayed in the Value frame.

   - Find the efficient portfolio with the minimum value for that portfolio statistic by clicking the Minimum radio button.
   - Find the portfolio with the maximum value by clicking the Maximum radio button.
   - Enter any value within the Minimum-Maximum range by clicking the Enter radio button and entering a value in the text box.
   - Find the portfolio with the same value as an Other portfolio, the Compared portfolio, any efficient portfolio or asset by clicking Closest to and selecting a portfolio or asset from the drop-down list.

My forecasts are typically stated in terms of return and standard deviation. How can I incorporate asymmetry into these forecasts?

One of the Expert Allocator's most powerful features is the ability to manage the shape of an asset's return distribution. This is accomplished by adjusting the 10th and 90th percentile returns to capture the desired skewness.

In the Expert Allocator, you can combine your symmetrical (return and standard deviation) forecasts and the asset's historical distribution characteristics. Here are two different ways:

1. Use **Shift Distribution** to change the level of the expected return but maintain the historical shape of the distribution.

2. Set the Return Distribution Shape to Normal and enter the expected return and standard deviation. Make note of the 90th and 10th percentile returns for the normal distribution. Set the Shape back to True, enter the normal 90th and 10th percentile values, and then modify them as desired, insuring that they satisfy the condition for a lognormal distribution.

How is Shift Distribution in the Risk and Return worksheet used?

Shift Distribution is particularly useful when working with asymmetrical distributions since you can change the Expected Return without affecting the shape of the distribution (i.e., skewness and degree of variability) or violating the conditions for consistency in Expected Return, 90th Percentile Return, 10th Percentile Return values.

Follow these steps:

1. Create a model using Historical Returns.
2. In the Risk and Return worksheet, click anywhere in the Asset's row and click **Shift Distribution**.
3. Enter the amount by which you want to increase or decrease the expected return. Select a positive number to increase the expected return and shift the distribution to the right; select a negative number to decrease the expected return and shift the distribution to the left.
4. Click **OK**.

The Expected Return, 90th Percentile Return, and 10th Percentile return are adjusted automatically by the amount entered. Standard deviation and Vol Skew remain unchanged.

Can I do plain vanilla mean-variance optimization?

Yes. Enter your asset assumptions for expected return and standard deviation in the Risk and Return worksheet by setting all return distributions to normal (click the All Normal button) and then entering values in the Expected Return and Standard Deviation fields. To optimize, in the Optimization Parameters worksheet select Standard Deviation as the Optimization Risk Measure and Normal as the Return Distribution Shape.

Can I include my fund's liabilities in an optimization?

Yes. Just enter the expected return and risk assumptions for your fund's liabilities in the Benchmark row in the Risk and Return worksheet (make sure to click the Include check box) and the correlation assumptions in the Correlations worksheet.

Can I optimize against a benchmark?

Yes. Just enter the expected return and risk assumptions for your benchmark in the Benchmark row in the Risk and Return worksheet (make sure to click the Include check box) and the correlation assumptions in the Correlations worksheet.

How is annualized standard deviation, using historical returns, calculated in The Expert Allocator?

In both systems, the unannualized standard deviation is calculated as follows:

$$\sqrt{\frac{\sum (R - r_i)^2}{(n - 1)}}$$

where:

R = average return
r_i = monthly or quarterly return
n = number of returns

To annualize, multiply unannualized values by the square root of the frequency (12 for monthly returns, 4 for quarterly returns) of returns.

Standard deviations for the S&P 500 from January, 1937 through December, 1997 are shown below. Notice that the frequency of returns and the denominator in the standard deviation formula can affect the final result.

Frequency	Denominator	Unannualized Std Dev	Annualized Std Dev
Annual	n - 1	17.6588	17.6588
Annual	n	17.5135	17.5135
Quarterly	n - 1	8.1985	16.3971
Quarterly	n	8.1817	16.3634
Monthly	n - 1	4.5278	15.6847
Monthly	n	4.5247	15.6740

Section 2: : FAQ's on The Expert Performance Analyst

How is annualized standard deviation, using historical returns, calculated in The Expert Performance Analyst?

In both systems, the unannualized standard deviation is calculated as follows:

$$\sqrt{\frac{\sum (R - r_i)^2}{(n - 1)}}$$

where:

R = average return
r_i = monthly or quarterly return
n = number of returns

To annualize, multiply unannualized values by the square root of the frequency (12 for monthly returns, 4 for quarterly returns) of returns.

Standard deviations for the S&P 500 from January, 1937 through December, 1997 are shown below. Notice that the frequency of returns and the denominator in the standard deviation formula can affect the final result.

Frequency	Denominator	Unannualized Std Dev	Annualized Std Dev
Annual	n - 1	17.6588	17.6588
Annual	n	17.5135	17.5135
Quarterly	n - 1	8.1985	16.3971
Quarterly	n	8.1817	16.3634
Monthly	n - 1	4.5278	15.6847
Monthly	n	4.5247	15.6740

I've created a series using the Weight feature in the Series tab in the Analyst. How can I update this series?

When you create a series using Weight, only the historical monthly (or quarterly) returns are saved—not the series and percentages used to create the new index. So you'll need to update the series as needed. To do so, use **Weight**—the same way you originally created the series—and overwrite the existing series. (Tip: For easy reference, name your series to reflect its composition, for example, 60% S&P/40% LB Agg.)

How can I incorporate The Expert Performance Analyst's charts and tables in my presentations?

Incorporating charts and tables in slide presentations, spreadsheets, or documents is fast and easy and creates a very effective presentation. Charts can be copied to the Clipboard or saved as a bitmap or metafile and then pasted into your document. Tables can be saved as Excel worksheets or text files and inserted into your documents.

Charts:

To copy a chart to the Clipboard or to a file, in the Chart Window select Export from the File menu option, then:

- To Bitmap File (`BMP`) to save the chart as a bitmap (`.BMP`) file. The Export to Bitmap dialog box appears. Enter a filename in the File Name list box and click **OK**. The file is saved automatically with a `.BMP` extension.
- To Metafile (`WMF`) to save the chart as a metafile (`.WMF`) file. The Export to Metafile dialog box appears. Enter a filename in the File Name list box and click **OK**. The file is saved automatically with a `.WMF` extension.

- To Clipboard (BMP) to copy the chart to the Windows Clipboard as a bitmap.
- To Clipboard (WMF) to copy the chart to the Windows Clipboard as a metafile.

Tables:

To export table information, choose Export from the File menu option and one of the following:

- to Text File to export a table to an ASCII text file. The Export to ASCII File dialog box appears. Type the name of the text file in the Filename box (a .TXT extension is automatically assigned) and press Enter or click **OK**.

- to Excel (XLS) to export a table to an Excel spreadsheet. The Export to Excel File dialog box appears. Type the name of the Excel spreadsheet in the Filename box (an .XLS extension is automatically assigned) and press Enter or click **OK**.

Section 3: FAQ's on Data Series

Can I include my own data in The Expert Allocator and Expert Performance Analyst?

Yes. You can create your own series by importing returns from an Excel worksheet or adding them directly into the IT Data Editor.

1. Copy-and-paste historical returns into the preformatted Excel worksheet provided for your convenience in your `SSIWIN` directory (`DEDTPLMO.XLS` for monthly returns and `DEDTPLQU.XLS` for quarterly returns). Save your file as an Excel 5.0/95 Workbook. Then import these series using **Update Database** in the Data Editor. For more information, refer to the *User's Guide for IT Data Services,* Chapters 2 and 4.

2. Add series manually in the Data Editor using **Edit Time Series**. In the Time Series Editor, highlight the module to which you're adding the new series and then click **Add**. In the Add New Series dialog box, enter the information including series name, beginning and ending dates, frequency, and categories (optional). Then click **OK** to return to the Time Series Editor. Click **Edit Returns** and enter historical returns. When done, click **OK**. For more information, refer to the *User's Guide for IT Data Services,* Chapter 2.

What's the easiest way to update my own series?

There are two ways to update your own data:

☑ *From an Excel worksheet:* If your series was originally created from an Excel file, simply add the more recent return information to the original file and click **Update Database** in the IT Data Editor. Otherwise, for easy handling, use one of the pre-

formatted Excel templates provided in your `SSIWIN` directory (`DEDTPLMO.XLS` for monthly returns and `DEDTPLQU.XLS` for quarterly returns).

☑ *Manually in the IT Data Editor:* Updating or modifying return data manually in the IT Data Editor is a two-step process:

1. Click **Edit Header** and, in the Edit Series Header dialog box, specify the Ending Date. When done, click **OK**.

2. Click **Edit Returns** and, in the Edit Series Return dialog box, enter the updated monthly or quarterly returns. When done, click **OK**.

Section 4: Post-Modern Portfolio Theory

What is Post-Modern Portfolio Theory (PMPT)?

PMPT is an extension of the basic ideas of Modern Portfolio Theory:

☑ Investors are risk averse

☑ Effective diversification can increase returns without increasing risk.

The roots of PMPT go back to Harry Markowitz and William F. Sharpe, who in 1959 and 1963, respectively, identified the shortcomings of mean-variance. Without the computing power to resolve these difficulties, though, they stayed with mean-variance.

In the 1980's Dr. Frank Sortino at the Pension Research Institute, San Francisco State University, drew on the work of finance theorists Vijay Bawa and Peter Fishburn to extend the Markowitz/Sharpe framework to incorporate the first PMPT-based risk-measurement systems.

In 1988 Sponsor-Software Systems, Inc. (now Investment Technologies) licensed the PMPT technology from The Pension Research Institute. Since then, Investment Technologies has conducted significant additional research and has fully incorporated PMPT into its range of software products.

How does PMPT deal with risk?

In PMPT, risk is defined as a combination of the likelihood and consequences of failing to meet a target return. For example, a portfolio has a 15% chance of not achieving its

long-term goal to beat the S&P 500. And if it does fail to achieve this goal, its expected shortfall is 75 basis points.

These two characteristics--likelihood and consequence of failure--can be represented together by a single statistic, target semi-deviation, commonly known as downside risk.

Formally, downside risk is the squared deviations of returns below the target. The squaring of the shortfalls represents an escalating aversion to risk, i.e., a return twice as far from the target represents four times as much risk. This is consistent with behavior of investors and managers observed by social scientists.

How does PMPT deal with skewed distributions?

PMPT works both with symmetrical, normal return distributions and with significantly skewed distributions.

If a distribution is symmetrical and the target return lies close to the average return, PMPT analysis is the same as traditional mean-variance (MPT) analysis. As the target return becomes higher or lower than the average and/or the distribution becomes less symmetrical, the differences and advantages of PMPT become clear.

How is PMPT incorporated in the Experts?

The Expert Allocator and The Expert Performance Analyst provide all the features and capabilities of PMPT.

In PMPT, returns are represented by the three-parameter lognormal distribution which replaces the traditional two-parameter model. This three-parameter distribution permits both left and right skewness and consequently captures a wide range of asymmetry.

For The Expert Allocator, the parameters required to define an asset's risk and return characteristics are the expected return and the 10th- and 90th-percentile returns. The software automatically fits the appropriate lognormal curve to these three data points. (The user is not required to specify the standard deviation or semi-deviation).

In The Expert Allocator, you select the risk measure and distribution type. When standard deviation and symmetrical distributions are selected, the optimization will produce results indistinguishable from quadratic optimizers.

The table below shows the four optimization alternatives available in The Expert Allocator.

Risk Measure	Asset Return Distribution: Symmetric	Asset Return Distribution: Asymmetric
Standard Deviation	Standard Mean-Variance	Extended Mean-Variance
Downside Risk	Partial PMPT	Full PMPT

Does PMPT give different results than MPT?

Yes, PMPT can give different results than MPT. These differences will be greater the more asymmetrical the asset return distributions and the further the goal return is from the average return.

Consider the following example using stocks, bonds, and cash equivalents. Asset assumptions are based on return, risk, and correlations for the five year period 1991 to 1995. Two optimizations were performed—PMPT and MPT. The PMPT efficient frontier uses the true shape of the asset distributions, downside deviation as the optimization risk measure with a 10% goal, and one year holding period. The MPT frontier uses normal return distributions and standard deviation as the risk measure. Both PMPT and MPT frontiers are plotted in Figure 1.

Notice that these frontiers are very different in terms of the composition of the portfolios on each efficient frontier (shown on the right side of Figure 1). Comparing the left-most ends of the frontiers, the MPT frontier shows a greater concentration in cash (red) and less in stocks (green) than the PMPT frontier.

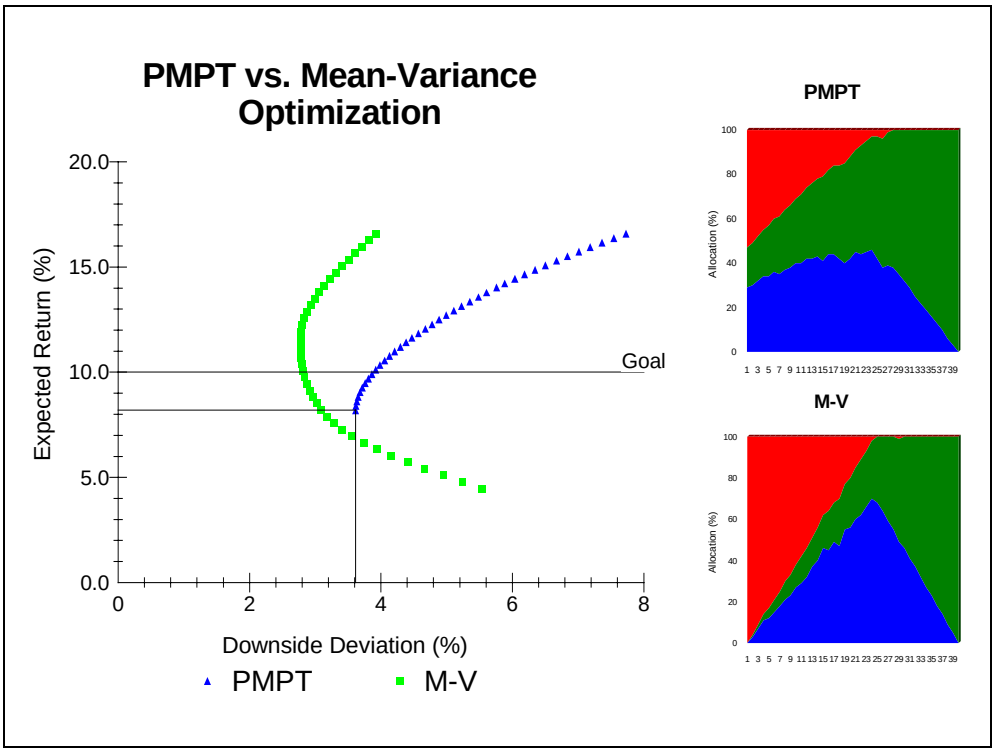

Figure 1

For additional illustrations examining the effect on optimization results of distribution shape, holding period, and goal, download this PowerPoint presentation (`PMPT.PPT`).

Is there really skewness in asset returns?

Yes. Investment Technologies, in conjunction with William H. Seaver, Ph.D., Seaver Statistical Services, conducted an in-depth study of historical returns for five major indices. The results are dramatic: Non-normality was detected in three out of every five periods analyzed. All of the indices showed significant occurrences of non-normality, with Russell 2000 showing the greatest divergence from normal and EAFE, the least.

These findings have profound implications for investment practitioners using standard deviation to measure the risk of their portfolios: *using the normal distribution can lead to false conclusions about a portfolio's true risk.*

In this study, we examined the distributional characteristics of five major market indexes representing broad asset classes: the S&P 500 for large-cap stocks, Russell 2000 for small-cap stocks, MSCI EAFE for foreign stocks, Lehman Aggregate for bonds, and 90-day Treasury Bills for cash equivalents.

Using monthly data from 1937 through 1993 for S&P 500 and T-Bills and inception dates for other indexes (early to mid-1970's), a total of 161 subperiods ranging from 3 to 57 years were reviewed, providing a full representation of market conditions.

Index	Dates	Years
S&P 500	1937-1993	57
Russell 2000	1976-1993	18
MSCI EAFE	1970-1993	24
Lehman Agg	1973-1993	21
T Bills	1937-1993	57

Not Normal 60% of the Time

As the following chart shows, the results are dramatic. Overall, returns for the five indices were not normally distributed more than 60% of the time. That means that standard deviation was *not* an appropriate representation of risk in three out of every five time periods. The frequency of non-normality for individual indices ranged from 45.5% (almost one out of every two periods) for MSCI EAFE to 85.7% (six out of every seven periods) for Russell 2000.

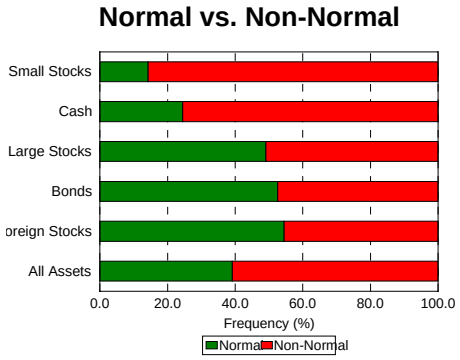

The length of the time period is an important factor in explaining the shape of return distributions. The longer the time period, the more likely return distributions are non-normal. For all indices, returns were not normally-distributed almost all (97%) the time for 11-57 year periods, nearly three-quarters (72.7%) of the time for 6-10 year periods, and 43% of the time for 3-5 year periods.

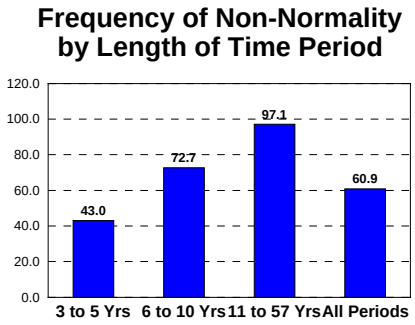

Yes, Skewness Exists

Skewness and kurtosis are the standard statistics used to measure any non-normality in data such as monthly performance returns. Of these, skewness is the more important.

Skewness measures the direction and degree of asymmetry of a return distribution. A normal distribution is bell-shaped, a positively skewed distribution has a longer tail on the right, higher-return side of the curve, and a negatively skewed distribution has a longer tail to the left, lower-return side.

When there is positive skewness, performance was actually achieved with less risk than indicated by standard deviation. Conversely with negative skewness, where performance was actually achieved with more risk than standard deviation suggests. Thus standard deviation overstates risk for positively skewed return distributions and understates risk for negatively skewed return distributions.

Skewness was found to be a contributing factor in almost half (47.2%) of the time periods analyzed. Interestingly, returns for the three stock indices were virtually entirely negatively skewed while returns for the fixed income indices were predominantly positively skewed.

Kurtosis, the other contributor to non-normality, is a measure of how peaked a set of returns is compared to a normal distribution. Positive kurtosis indicates a distribution with a high peak, while negative kurtosis indicates a flattened distribution.

Kurtosis was found to be a factor with about the same frequency as skewness (47.2%).

PMPT--A Better Way

With The Expert Allocator and The Expert Performance Analyst, investment practitioners can automatically take skewness into account in both asset allocation analysis and performance measurement. Both Experts offer Post Modern Portfolio

Theory tools to evaluate the true performance of investment managers, mutual funds, and other portfolios without the restrictions imposed by standard deviation.

With The Expert Allocator for asset allocation, you can incorporate skewness both in forecasts for asset returns (by using the true rather than the normal distribution) and in optimization (by using downside deviation rather than standard deviation as the risk measure).

The tools of PMPT have direct application in performance measurement as well. The Expert Performance Analyst helps you to understand how performance results were achieved and whether there was compensation for the level of risk taken.

What is downside deviation?

Downside deviation is a risk measure, used both in asset allocation and performance measurement, that is stated in relation to an *investor's goal* and defines risk as not achieving that goal. In this context, a portfolio's risk may be perceived differently by different investors with different goals.

Downside deviation distinguishes between "good" and "bad" risk: Good risk comes from returns that are greater than the goal, and bad risk comes from returns that are less than the goal. In addition, the farther returns fall below the goal the greater the risk. Thus, downside deviation distinguishes between "good" and "bad" returns by assigning risk only to those returns below the goal. It measures the portion of total variability that is below the goal (i.e., to the left on the return scale of an asset's return distribution curve). The greater the value of downside deviation, the greater the risk.

Because risk is measured against the same reference point (i.e., the same goal), portfolio risk can be compared directly using downside deviation. Thus a portfolio with 4.0%

downside deviation is twice as risky as a portfolio with 2.0% downside deviation and half as risky as a portfolio with 8.0% downside deviation.

A portfolio's downside deviation will vary depending on the goal (for example, a portfolio's downside deviation is 3.4% at a 6% goal and 6.2% at a 10% goal). Downside deviation for a portfolio will increase as the goal increases.

Downside deviation is a function of the *probability* and *consequences* of not achieving the goal. These are the two components of downside deviation and in The Allocator are termed downside probability and average downside deviation, respectively.

Downside deviation can be used to measure how effectively a portfolio delivers return for each unit of risk taken. The Sortino ratio measures return-per-unit-of-risk and is calculated as the difference between the portfolio's expected return and the goal divided by downside deviation.

How do I interpret downside deviation? What does the downside deviation percentage really mean?

Like standard deviation, the larger the value, the greater the riskiness of the portfolio. But there is opportunity to provide a richer explanation with downside deviation. Take results for two hypothetical portfolios shown in Figure 1 below. We'll assume an investor evaluating these portfolios has a 10% goal and a five-year holding period. Figure 2 shows the return distributions for both portfolios.

	Aggressive Portfolio	Moderate Portfolio
Expected Return	11.50%	7.88%
Standard Deviation	15.68%	7.62%
Sharpe Ratio	0.48	0.48
Downside Risk Statistics for 5 Year Holding Period at 10% Goal:		
Downside Deviation	4.35%	3.83%
Downside Probability	38.36%	72.03%
Average Downside Deviation	5.36%	3.63%
Expected Downside Return	4.64%	6.37%
Downside Magnitude (99th pctl)	16.72%	11.14%
Sortino Ratio	0.34	-0.55

Figure 1

At a 10% goal, the Aggressive Portfolio has a downside deviation of 4.35%. This value can be used to make a statement about the portfolio's riskiness compared to other portfolios at the same goal. For example, this portfolio is 13% more risky than the Moderate Portfolio.

There is about a four-in-ten chance (38.36% downside probability) of falling short of goal with the Aggressive Portfolio. In the event of failure, the average shortfall is 5.36% (average downside deviation), and the average return is expected to be 4.64% (expected downside return). In a worst case scenario (at the 99th percentile), the portfolio will underperform the goal by nearly 17% (16.72% downside magnitude).

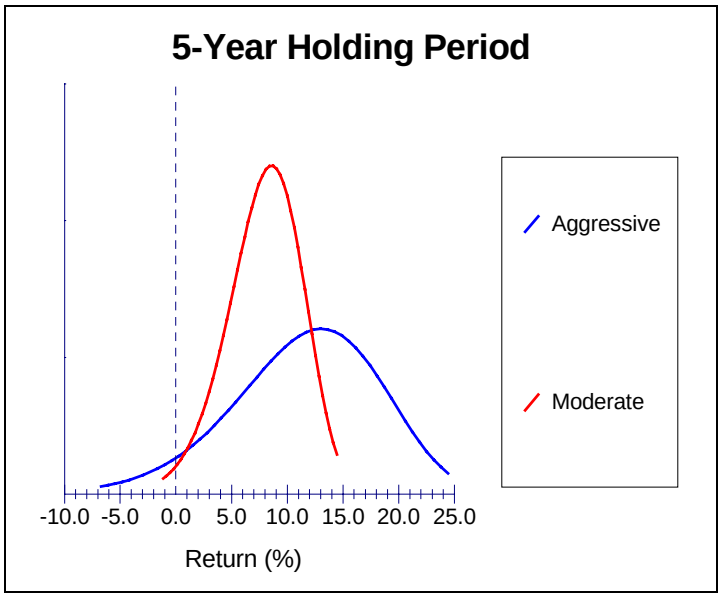

Figure 2

The Moderate Portfolio has almost a three-in-four chance (72.03% downside probability) of not meeting the goal, with an expected downside return of 6.37%. Worst case scenario calls for a return that is 11.14% below goal.

Expressing expected downside return in wealth terms can also be informative. For example, assuming an initial $1 million investment, goal wealth after five years would be $1,610,510. With the Aggressive Portfolio, the expected downside return is $1,254,840, about $355,000 less goal. With the Moderate Portfolio, the expected downside return is $1,361,604, or about $249,000 below goal.

The tradeoff, then, is that the Aggressive Portfolio has a better chance than the Moderate Portfolio of making goal (38% compared to 72% downside probability) but the

consequences, in the event of failure, are greater ($355,000 average shortfall compared to $249,000).

On a risk-adjusted basis, the Aggressive Portfolio provides a better Sortino ratio (0.34 versus –0.55 for the Moderate Portfolio). This is because the Aggressive Portfolio's proportionately greater expected return more than compensates for the slightly greater risk.

Do results differ if I compare portfolios using mean-variance and downside risk analyses?

Yes, the conclusions you draw based on these two techniques can be different in many cases. An example may help to illustrate.

Return and risk characteristics of Value and Growth styles for the 10-year period ending June 30, 1998 are presented in the following table. Using standard deviation and downside deviation, there are striking differences in how styles compare in terms of risk and risk-adjusted performance.

Mean-Variance Analysis:

On the basis of return, Growth outperformed Value during the period, earning 19.15% versus 17.44%, respectively. Growth, however, showed greater variability, with a 13.52% standard deviation compared to Value's 11.31%. According to rule of thumb, Growth ranged between 32.67% and 5.63% two-thirds of the time, a wider range than for Value (28.75% and 6.13%).

	Growth	Value
Annualized Return	19.15%	17.44%
Standard Deviation	13.52%	11.31%
Sharpe Ratio	1.01	1.06

On the basis of return-per-unit-of-risk, Value performed slightly better than Growth. The Sharpe ratio for Value was 1.06 compared to 1.01 for Growth.

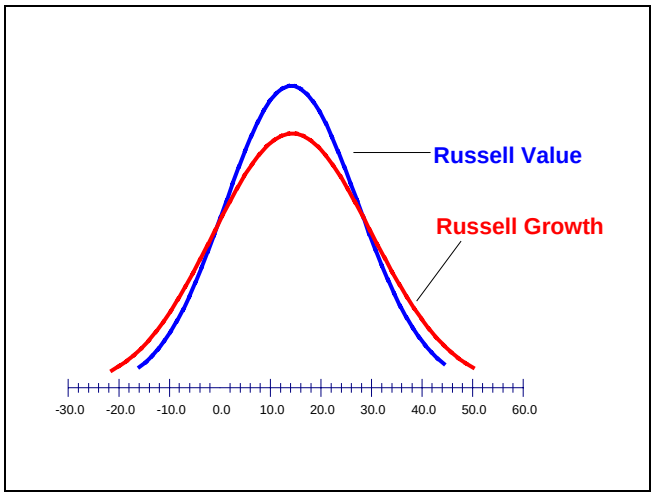

Figure 1 Normal Distributions

Conclusion based on mean-variance analysis: Growth earned 170 basis points more than Value, but this greater return was achieved with proportionately more risk; on a risk-adjusted basis, therefore, they both performed about the same. *Scorecard*: Value wins by a hair.

Downside Risk Analysis:

Recall that downside deviation is stated in relation to a goal. Downside risk statistics for Growth and Value were calculated at three goals to illustrate different return goals for different investors: 0% (an objective not to lose money), 8% (a moderate return objective), and 12% (a more aggressive return objective). These results are shown below.

Downside Deviation	Growth	Value
at 0% goal	2.04%	8.44%
at 8% goal	4.61%	10.49%
at 12% goal	6.37%	11.75%

Growth was actually substantially less risky than Value at all goals. Downside deviation for Growth was one-quarter as risky at a 0% goal (2.04% versus 8.44%) and about half as risky at a 12% goal (6.37% versus 11.75%).

The reason for this is that a larger portion of Growth's variability is on the upside. This is reflected in the shape of its return distribution illustrated in Figure 2.

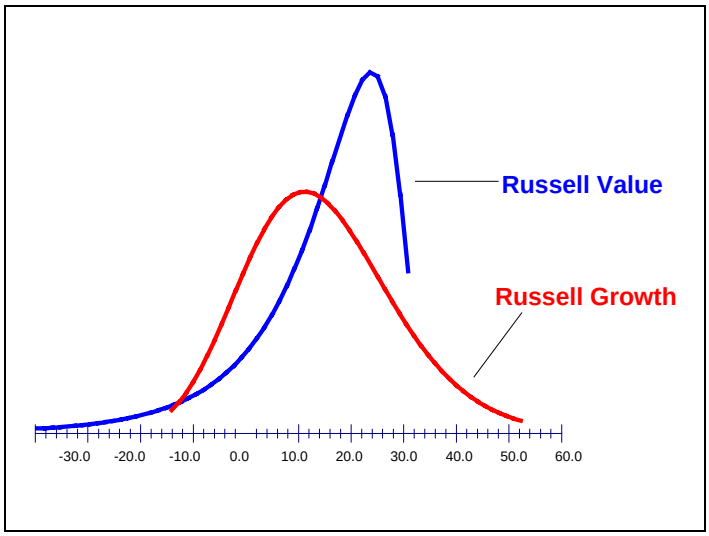

Figure 2 True Distributions

Notice that Growth's distribution is positively-skewed (longer tail at the higher end) whereas Value's is negatively-skewed (longer tail at the lower end). Because its distribution is positively-skewed, Growth was actually less risky and Value was actually more risky than their standard deviation values suggest.

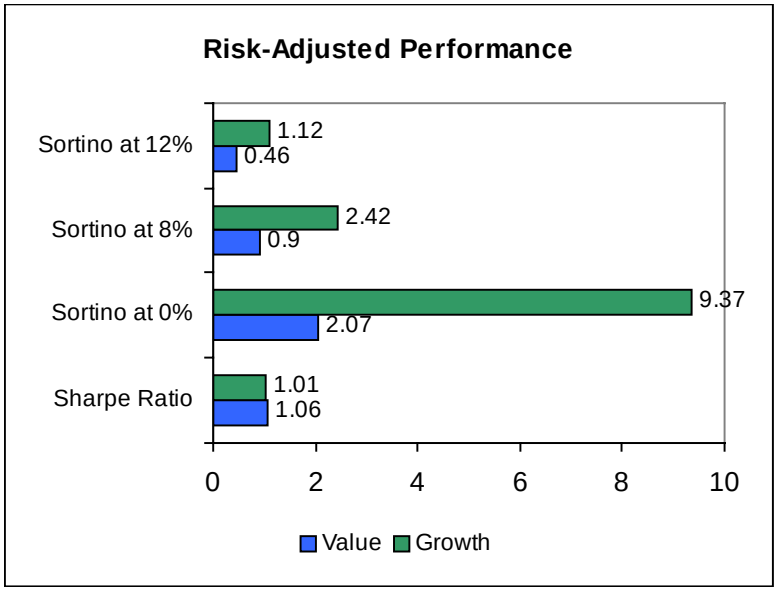

Figure 3 Comparison of Risk Adjusted Performance

On the basis of return-per-unit-of-risk, Growth outperformed Value by a significant margin. Using the Sortino ratio, Growth exceeded Value at all goals (Figure 3).

Conclusion based on downside risk analysis: Growth earned a higher return with significantly less risk than Value at all goals. *Scorecard*: Growth wins hands down.

As you can see, downside risk analysis provides effective and informative yardsticks in measuring portfolio performance for several reasons:

- Downside deviation considers the investor's goal, standard deviation does not.
- Downside deviation defines risk in accordance with an investor's perception of risk, i.e., failure to meet the goal; standard deviation measures the dispersion of returns around the average.
- Downside deviation recognizes a portfolio's true return distribution (bell-shaped or skewed); standard deviation assumes all distributions are bell-shaped.

Does the normal distribution's rule-of-thumb describing the dispersion of returns apply to the skewed distributions as well?

No. The rule of thumb—that two-thirds (67%) of returns fall within plus or minus one standard deviation of the mean, and 95% fall within plus or minus two standard deviations of the mean—does not apply to skewed distributions.

The confidence bands for asymmetric distributions are different from those for the normal distribution. The strongest statement we can make about asymmetric distributions is that three-fourths (75%) of the returns fall within plus or minus two standard deviations of the mean, and eight-ninths (89%) of the returns fall within plus or minus three standard deviations. For distributions relatively close to normal, these confidence bands will actually be higher.