Practical Risk Analysis for Portfolio Managers and Traders

Ananth Madhavan*

Jian Yang†

Current Version: April 2, 2003

Abstract

Once a fairly esoteric subject, risk analysis and measurement has become a critical function for both portfolio managers and traders. Yet, accurate measurement and analysis of risk presents many practical challenges including the choice of risk model, pitfalls in portfolio optimization, horizon mismatches, and out-of-sample testing. This article provides a detailed overview of recent developments in risk analysis and modeling, with a focus on practical applications for both portfolio managers and traders. We demonstrate that these tools can provide invaluable insights regarding portfolio risk, but must be applied with considerable care. Risk analysis, as it stands today, is as much an art as a science.

* ITG Inc., 380 Madison Avenue, New York, NY 10017, Tel: (212)-444-6361; E-mail: amadhava@itginc.com

† ITG Inc., 44 Farnsworth Street, Boston MA 02210, Tel: (617)-692-6860; E-mail: jyang@itginc.com

The authors thank Ken Gosier for expert research assistance.

40103-87702

Once a fairly esoteric subject, risk analysis and measurement has become critical for both portfolio managers and traders. In part, this phenomenon is motivated by increased concern over risk by plan sponsors and investors, the growth of portfolio trading, heightened volatility, and increased appreciation of quantitative tools by investment managers and traders. Historical or subjective estimates of volatility are poor indicators of future volatility, generating strong interest in quantitative, forward-looking models of risk.

Investment managers use formal risk models to both predict the future and understand the past. Specifically, risk models are used for three purposes: (1) **Measurement**, i.e., as *forward-looking* tools to predict *future* volatility, (2) **Analysis**, i.e., as tools to *understand* risk exposures and construct optimal portfolios, and (3) **Evaluation**, i.e., as guides to *past* performance, evaluating *actual* returns in relation to the risks incurred to achieve them. Yet, despite the increased use of complex portfolio analytical tools such as risk models, the accurate measurement and analysis of risk presents many practical challenges. This article provides a detailed overview of recent developments in risk analysis and modeling, with a focus on practical applications.

We begin with a discussion of the basics of risk analysis, focusing on the value of formal models over historical estimates of risk. We also compare and contrast the various types of risk models and their relative strengths and weaknesses from a practical perspective. We then turn to our major focus, i.e., practical issues in risk modeling. We present a series of examples to illustrate some of the practical challenges in risk analysis. We cover the three elements of risk analysis described above (forecasting, analysis, and evaluation) with illustrations drawn from practical experience. In particular, we highlight the importance of matching the horizon of the risk model to the risk under consideration. For example, an active hedge fund following a statistical arbitrage strategy with a large daily turnover is likely to be concerned with risk over short periods of time while a value manager who rebalances quarterly might care about risk on a monthly basis or longer. We also illustrate a methodology to testing risk forecasts on an out-of-sample basis. Of special interest, we highlight the increased use of risk models for a variety of trading applications, including formulating trading strategies and pre-trade analysis.

We conclude that although risk models can provide invaluable insights into risk, they need to be applied with care. Risk analysis, as it stands today, is as much an art as a science.

RISK MODELS

Risk is usually defined as the volatility of returns, measured by standard deviation. For a portfolio of *n* assets with weights *? 1,…,? n*, risk is

$$Std(r) = \sqrt{\sum_{i=1}^{n} w_i^2 Var(r_i) + \sum_{i \neq j} w_i w_j Cov(r_i, r_j)}, \qquad (1)$$

where $Var(r_i)$ denotes the variance of the return of asset *i* and $Cov(r_i, r_j)$ is the covariance between asset returns r_i and r_j. The variance of a portfolio depends not only on the individual variances of the assets, but also on the covariances between the components of the portfolio. As an extreme example, a portfolio consisting of two securities – a long position in a stock and with an appropriately chosen position in a put option on the same stock – will have lower risk than either of its components because the two securities are negatively correlated. For practitioners, the main drawbacks to the focus on volatility is that it ignores other moments of the distribution of returns and treats both "upside" and "downside" risk symmetrically. For this reason, practitioners typically begin with volatility but also look at other risk metrics as well.

Pitfalls in Using Historical Data to Form Risk Estimates

Portfolio risk can be computed straightforwardly using historical data. But using historical volatility to assess risk presents several problems. For large portfolios, we need a long time-series of observations to obtain reliable estimates. As the size of the portfolio grows, the number of observations needed increases dramatically while estimating the variances and covariances needed for a large portfolio becomes asymptotically large. For a portfolio with N assets, we require $N(N+1)/2$ inputs to compute the variance-covariance matrix of asset returns. Even modest numbers of securities create problems. For example, a manager who follows 250 stocks needs a total of 31,375 variances and covariances as inputs to construct the variance matrix. The problem of the estimation and "upkeep" of large numbers parameters at the security level is known as the **curse of dimensionality**. Dimensionality also matters because portfolio optimization requires inverting the variance matrix, which necessitates that a **rank condition** be satisfied, i.e., requires that the columns of the matrix are linearly independent. But if the dimension of the matrix is large, this condition in turn requires an unrealistically long time-series of observations. Even if available, such a time-series implies that the estimates of risk incorporate events in the distant past.

Another problem with using sample variances is that the past does not necessarily provide a good indication of the future. Sample covariance estimates are completely data-driven and hence have no way to capture future changes in return volatility or co-movements arising from changes in the underlying return generating process. For example, although we know that volatility is serially correlated, there is no easy method to capture the time-varying features of volatility within a sample covariance approach. Finally, the sample covariance approach cannot yield risk decompositions essential for meaningful risk analysis and investment decisions such as hedging, benchmarking, optimization, performance attribution, and segmented analysis.

Why Use Factor Risk Models?

Factor risk models provide a solution to the problems above. First, they offer tractability and a solution to the "Curse of Dimensionality" by imposing logical and testable factor structures. Second, factor risk models present the feasibility to reveal the covariations among future stock returns and to capture time-varying features of volatility, and therefore provide predictability in applications ranging from computing risk estimates to forming portfolio benchmarks. Finally, factor risk models allow users to implement meaningful hedging and benchmarking by decomposing risk into alternative risk sources represented by different risk factors and assess performance in terms of portfolio returns in relation to the risks incurred to generate them.

To understand how factor risk models work, consider a simple one-factor model,

$$r_i = m_i + b_i F + u_i, \tag{2}$$

where r_i is the return of asset i, F is the common factor, β_i is a constant, and u_i the idiosyncratic shock. When F is selected as the market return, r_m, the model becomes the classic Capital Asset Pricing Model (CAPM). The single factor model makes assumptions, most importantly that the shock term of asset i is uncorrelated with the factor and the shocks of other assets. Thus, securities are related only through responses (exposures) to the market (index). It is straightforward to show that the variance of returns is $s_i^2 = b_i^2\ s_m^2 + Var(u_i)$. The advantage of specifying a risk model is that it greatly eases the task of computing inputs for portfolio risk calculations. In the special case above, it is straightforward to show that the covariance between any two assets i and j is $Cov(r_i, r_j) = b_i b_j s_m^2$. Thus, to compute the covariance matrix for a portfolio of N assets, we require N betas, N firm-specific volatilities; and 1 estimate of market volatility, a total of $2N$+1 estimates. For a 250 stock portfolio, this implies 501 parameters versus 31,375 parameters if we use actual returns directly. Risk models solve the curse of dimensionality.

Factor models also provide a convenient framework to capture volatility clustering and excess kurtosis prevalent in financial data. One approach is to utilize the class of Generalized Autoregressive Conditional Heteroskedasticity (GARCH) models. In particular, if e_t is the return innovation (surprise) at time t for a particular asset, then under a GARCH(1,1) process:

$$Var(e_t) = c + a \cdot e_{t-1}^2 + b \cdot Var(e_{t-1}). \tag{3}$$

Applying conditional volatility models and using implied volatility make the risk model forward-looking over the appropriate horizon.

Finally, the market factor model provides the insight of risk by decomposing risk into different components, the systematic risk component due to exposure to the market and the idiosyncratic risk component. The relative importance of the idiosyncratic risk component to the systematic risk component informs the traditional fund managers of the level of diversification. Nevertheless, a hedge fund manager is more interested in comparing the exposures of a group of stocks to make long-short hedging decisions, namely market neutral strategy.

The single-factor CAPM has well-known deficiencies (Chan and Lakonishok, 1993). An alternative is the a multi-factor model, which has a theoretical basis in Arbitrage Pricing Theory:

$$r_i = \sum_{k=1}^{K} f_k b_{ik} + e_i, \tag{4}$$

where f_k is the return of factor k, $ß_{ik}$ the loading of stock i on factor k, and e_i is an idiosyncratic shock. The point about risk models reducing dimensionality still holds in the multifactor framework. Given (4), the variance matrix of asset returns at time t has the form:

$$V_t = B\ F_t\ B' + S_t, \tag{5}$$

where F_t is the (K×K) covariance matrix of the risk factor returns, B is the (N×K) exposure matrix, and $å_t$ is the (N×N) diagonal matrix of disturbance term variances. So, we require $K(K+1)/2$+(N×K)+N

estimates to compute the variance matrix. In the case of a 250 stock portfolio, a four-factor model requires 1,260 parameters versus 31,375 using historical estimates.

Fundamental versus Statistical Risk Models

Two approaches exist to characterizing factors as either statistical or pre-specified fundamental factors. Pre-specified factors are readily observable from market data, such as Treasury interest rates, while statistical factors are computed from return data using Principal Components Analysis or a similar algorithm. See, e.g., Connor and Korajczyk (1986) and Jones (2001). Both approaches have pros and cons.

Statistical models can work well in a rapidly changing market or over short horizons. Because the "factors" are statistically derived, they might capture the influence of subtle economic variables that might not be omitted from consideration due to data or modeling considerations. For example, Madhavan (2002) describes a growing literature on commonality in liquidity, which some argue affects asset returns. Others believe momentum is a factor in returns. Statistical models might capture the influence of such variables, which are difficult to measure. By contrast, fundamental models rely on pre-specified factors that might lose relevance because of changes in the economic or regulatory environment. Statistical models may even be the only choice if the fundamental data are not available or not reliable, as in some thinly traded or foreign assets. These models have the additional advantage of relative ease of estimation since they require only time-series return data. Finally, recent advances (Jones, 2001; Connor, Korajczyk, and Linton, 2003) yield more realistic estimates fom statistical models relative to their predecessors. On the con side, the criticism of statistical models centers on the difficulty in relating the derived factors to intuitive economic variables. It is difficult to use statistical risk factors for backtesting, benchmarking, performance attribution, and related analyses.

Selecting Among Fundamental Risk Models

Multiple fundamental factor models fall into three main categories, each with their own advantages and disadvantages: (1) Macroeconomic factor models, (2) Cross sectional fundamental factor models, and (3) time series fundamental factor models. Basic features of these models are summarized in EXHIBIT 1.

EXHIBIT 1: Risk Models Types

Model	Statistical	Macroeconomic	Cross-sectional	Time-Series

			Fundamental	Fundamental
Type	Statistical	Pre-specified	Pre-specified	Pre-specified
Structure	Time series	Time series	Cross-sectional	Time series

Within the class of pre-specified factors, macroeconomic factors generally do poorly because they do not capture significant variations in stock returns. Intuitively, macroeconomic indicators are relatively stable, and stock price behavior is more dynamic than suggested by changes in macroeconomic variables, such as inflation and consumption growth. These factors, however, are useful for evaluating risks in a global, multi-asset context.

Although the differences among risk models in terms of the choice of risk factors are widely recognized by practitioners, the distinction between cross-sectional and time-series estimation has not received much attention despite its obvious importance. In the cross-sectional model framework, factor returns are parameters to be estimated and the loading variables are assigned. For example, if the first factor represents a certain industry factor, the *loadings* β_{ik} for stocks fully in this industry will be assumed to be one and otherwise zero. For these stocks within the same industry, the above regression model becomes

$$r_i = f_1 + \sum_{k=2}^{K} f_k \beta_{ik} + e_i . \qquad (6)$$

Therefore, all stocks within this industry are assumed to respond similarly to the industry shocks. But stocks within an industry might react to industry shocks very differently. Assuming the same sensitivities (reactions) of stock returns to risk factors simply introduces systematic error into the model. And it limits the prediction capacity of such a risk model and makes benchmarking and hedging inappropriate.

A time series version of the three-factor model has drawn much attention in recent literature. An extensive bibliography can be found in Chan, Karceski, and Lakonishok (1999). The time series version of a multiple factor model is specified as:

$$r_{it} = \sum_{k=1}^{K} f_{kt} \beta_{ik} + e_{it} , \qquad (7)$$

where f_{kt} is the return of factor k at time $t = 1, \ldots, T$, β_{ik} the loading of stock $i = 1, \ldots, N$ on factor k, and e_{it} the disturbance term. In such a framework, β_{ik} is estimated on a stock-by-stock basis and, therefore, it is different across stocks. Security specific exposures to risk factors produce obvious advantages over a cross-sectional approach. Specifically, the systematic error associated with the cross-sectional model is no longer an issue and the resulting loading estimates provide correct characteristics of stock return sensitivities to factors. Consequently, benchmarking, hedging, and many other investment decisions are more reliable and meaningful. Further, the factor returns in a time series model are pre-calculated and the factor exposures are computed on a stock-by-stock basis, and hence are independent of the selected universe. By contrast, risk estimates from a cross-sectional regression significantly depend on the selected estimation universe, an issue if the user is concerned with only a small subset of the total universe.

A disadvantage of time-series regression models in a large universe is that some stocks will, purely by chance, have extreme loading estimates. These outliers can skew risk forecasts and result in optimized portfolios with undesirable characteristics. For example, in a long-short portfolio optimization problem, the result might be a "dumbbell" type portfolio with concentrations in highly positive and highly negative beta stocks. Fortunately, the problem can be overcome with statistical techniques that "shrink" extreme loadings towards more reasonable levels or Bayesian methods. Generally speaking, cross-sectional factor models are better suited to explain cross-sectional differences in stock returns while, as suggested by Fama and French (1993), time series factor models are more appropriate for capturing risk.

THE USE AND MISUSE OF RISK MODELS

Risk models are powerful tools for predicting future risk and understanding the past. However, there are many challenges in the practical application of risk models. The examples that follow are based on our experience with portfolio managers and traders using risk models for: (1) **Measurement**, i.e., prediction of future portfolio volatility; (2) **Analysis**, i.e., an understanding of the risk exposures of the portfolio; and (3) **Evaluation** of past performance. The organization of the section follows these three functions: We begin with issues in forecasting volatility (horizons mismatch), then describe challenges in interpreting and using risk models (creating zero-beta hedged portfolios, global v. regional model output, analysis of risk scenarios on an out of sample basis), and finally conclude with an example of evaluation (performance analysis).

A. Measurement

Forecasting Risks When Horizon Varies

Most users of risk models use monthly risk models typically constructed using 60 months (five years) of prior data. Although this choice appears dictated by custom, it is worth questioning given that risk model users are not alike. Indeed, fundamental portfolio managers might hold their positions for months at a time, while hedge funds and traders might have much shorter horizons of weeks or days. Can the needs of all users are served by risk models with a fixed horizon? EXHIBIT 2 presents the annualized short-term risk versus long-term risk of S&P 500 index from January 1985 through February 2002. At the beginning of each month, the short-term risk is measured by the standard deviation of the daily returns within the month and the long-term risk measured by the standard deviation of the daily returns within the next six months.

EXHIBIT 2: Short-term versus long-term risk of S&P 500

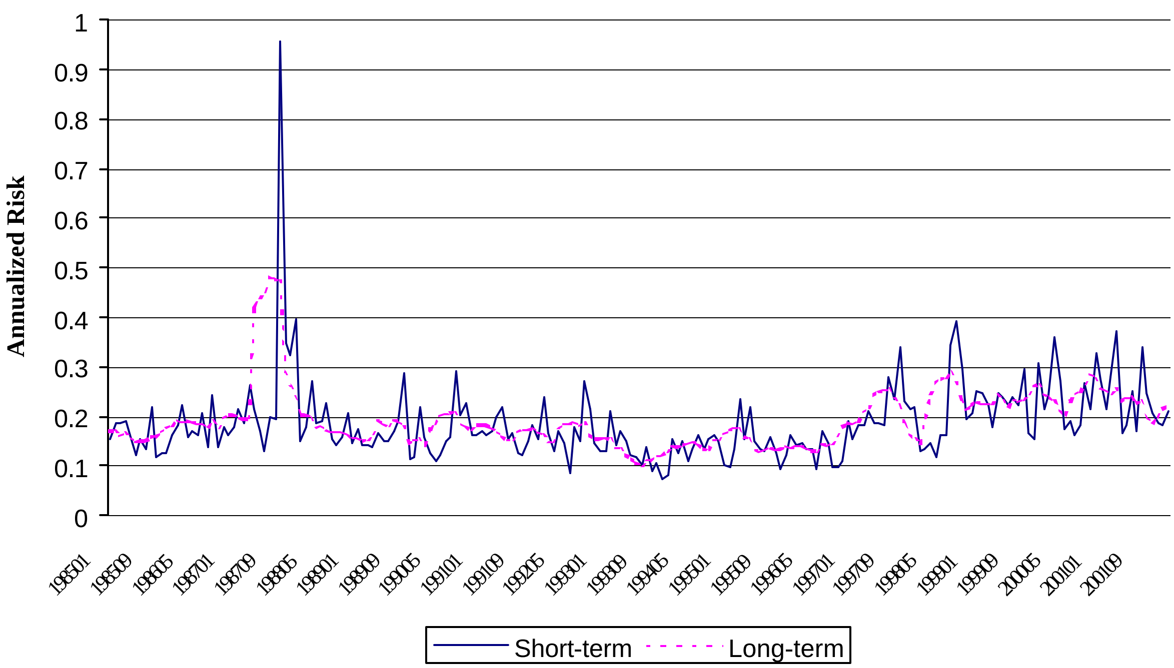

Not surprisingly, the discrepancy between short- and long-term volatility can be dramatic, and long-term volatility is significantly less variable than short-term volatility. This example highlights the importance of considering investment horizon or holding period in forecasting risk.

In addition to the risk level discrepancy, there are big differences between the portion of risk attributable to stock specific events versus those arising from market movements. In general, a longer horizon risk model will "explain" more risk in terms of systematic movements represented by style factors and subgroup factors because the behavior difference will only prevail in a long term.

Indeed, the exposures to risk factors can exhibit dramatic differences according to the horizon over which we forecast volatility. As an example, consider EXHIBIT 3, which provides the risk values and their breakdown of S&P 500 using Monthly and Daily Risk Models (MRM and DRM, hereafter) produced by ITG Inc. Briefly, these models are time-series stock-specific models with market, sector, and industry factors in addition to the size and value factors suggested by Fama and French (1993). Both models perform well over their forecast horizons, but as shown in EXHIBIT 3, there are major differences in the factor exposures.

EXHIBIT 3: Risk Decomposition of S&P 500 index

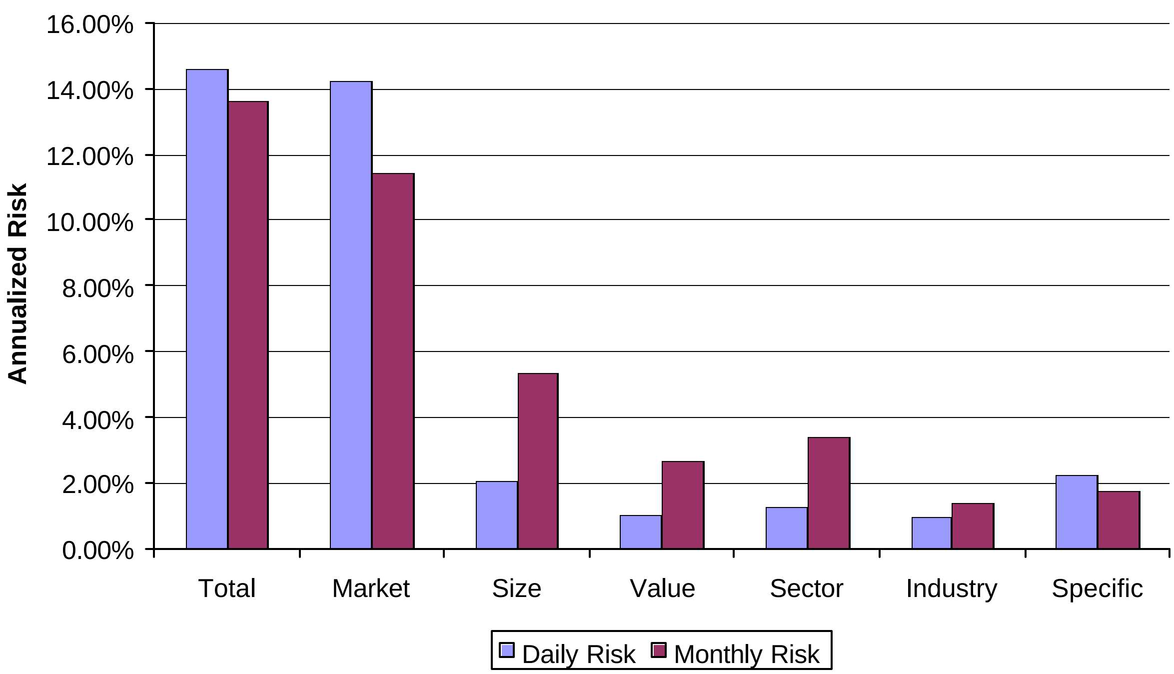

In particular, the market factor is more significant in the Daily Risk Model. Size and value factors play important roles in the long term rather than in the short term, as discussed above and consistent with the findings in the literature. With such significant differences of risk decomposition in a monthly risk model versus a daily risk model, a user should carefully choose the right model with appropriate horizon for their specific investment needs.

Risk models are often used to create hedged portfolios, and the horizon mismatch problem is particularly acute for long-short portfolios. For illustration, we take three stocks (CRAY, IBM, and DELL) and estimate their market betas using daily returns (last 60 trading days) and monthly returns (last 60 months). The estimated betas are used to construct dollar neutral and market neutral portfolios of the three stocks, every day from June 1, 2001 through May 31, 2002. The absolute values of the realized daily portfolio returns are presented in EXHIBIT 4. Using the daily beta to hedge market risk is generally more effective for these small portfolios, especially in very volatile periods. We conclude that a forward-looking forecast of risk needs to match the user's intended trading horizon.

EXHIBIT 4: Daily Hedged versus Monthly Hedged

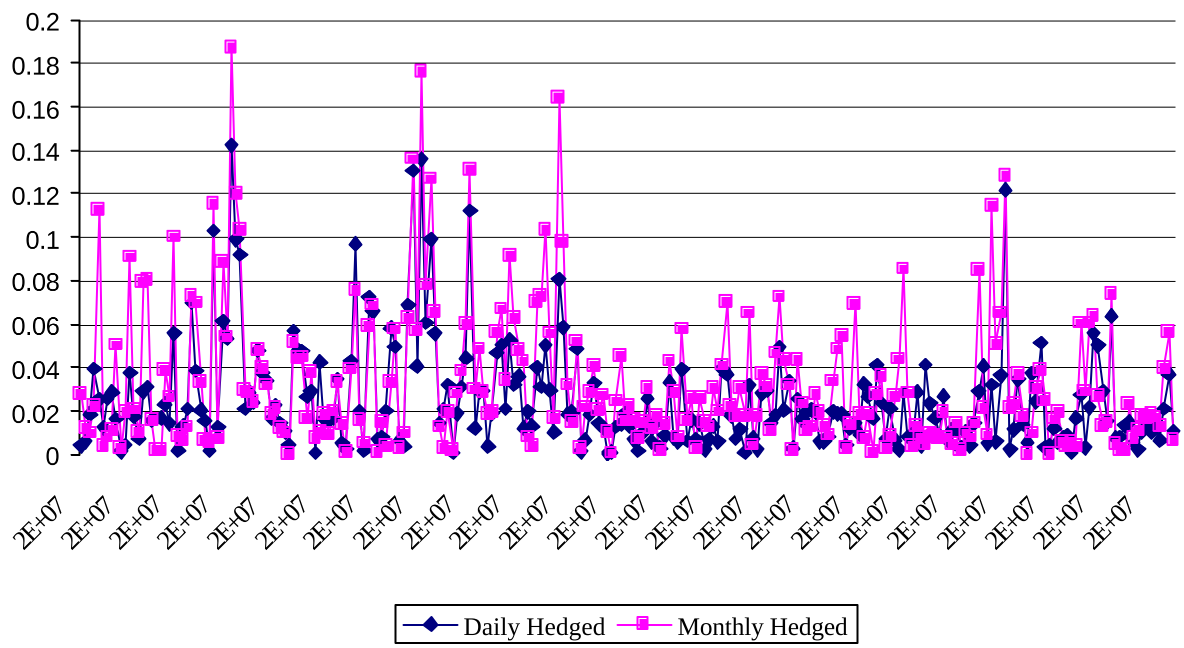

A Zero-Beta Hedge

Risk models are often used to create hedged portfolios. But while a long-short portfolio might have zero-beta with respect to a particular factor, it could well turn out to actually have exposure to that factor depending on how the underlying risk model was constructed. For example, Dell Computer Corp. (DELL) and Gateway Inc. (GTW) are in the same industry, i.e., computer hardware. Cross-sectional risk models assume that the exposures to the industry factor are the same. Indeed, we know that some cross-sectional risk models assigned both stocks industry loadings of 100% on Computer Hardware in 2001 and 2002. Time-series models, by contrast, do not make this assumption.

If the cross-sectional approach is valid, a dollar-neutral portfolio of the two stocks should be associated with about the same level of risk as the beta-hedged portfolio constructed by using the industry exposures estimated using historical data. To test whether this is the case, we use weekly samples from January 2001 through December 2001 to estimate industry exposures and then construct the beta-hedged portfolio. The realized risk is then computed based on the samples from January of 2002 to December of 2002, which is measured as the standard deviation of the realized portfolio returns. The realized risk values are reported in EXHIBIT 5. For this pair of stocks, we find the dollar-neutral portfolio has 28.19% higher risk than the beta-hedged, a substantial difference that illustrates the importance of correctly estimating industry exposures.

EXHIBIT 5: Dollar-Neutral versus Beta-Hedged

Portfolio	Dollar-neutral	Beta-Hedged
Realized Risk	0.1113	0.0869

EXHIBIT 6 plots the absolute value of the realized returns of the dollar-neutral portfolio, DELL against GTW, and the beta-hedged portfolio for out-of-sample periods.

EXHIBIT 6: Dollar-neutral versus beta-hedged

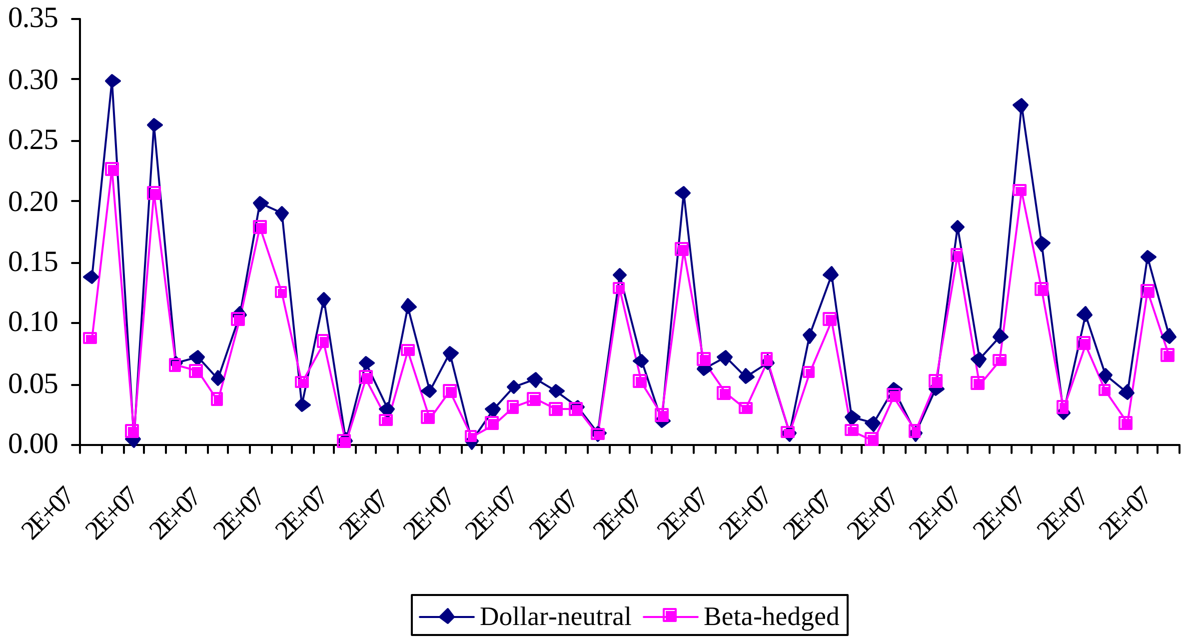

The beta-hedged portfolio has clearly lower risk especially for certain highly volatile periods. The example illustrates the importance of correctly understanding factor exposures when constructing hedged portfolios. To the extent that a risk model makes simplifying assumptions (i.e., that all industry exposures are identical), the results are potentially quite misleading, and open the manager up to unexpected factor risk. This example illustrates that models are not alike; it is worth the effort to understand how a risk model is constructed before using it for important portfolio decisions.

B. Analysis

Out-of-Sample Analysis

Given the variety of risk models available, how does a user select the model that best suits their needs? This section provides develops an out-of-sample testing methodology that can also be applied to a variety of other practical problems. Risk models are supposed to capture volatility, so the focus should be on second moments, not returns. Accordingly, we develop here statistic – ***the risk ratio test*** – to examine the forecasting ability of a risk model for any randomly chosen portfolio, or for some portfolios of special interest. Denote the forecasted variance using information up to period t as fv_t and let fs_t denote the square root of fv_t, i.e., forecasted risk. Denote the realized variance by v_t and the realized risk as s_t. If a risk model works well, then the ratio of actual to forecast risk, constructed as an average of over a significant number of out-of-sample periods, will tend to 1 as sample size increases. Define:

$$Risk \quad ratio \quad statistic = \sqrt{\frac{p}{2}}\frac{1}{T_1}\sum_{t=1}^{T_1}\frac{s_t}{fs_t} = \sqrt{\frac{p}{2}}\frac{1}{T_1}\sum_{t=1}^{T_1}\frac{|r_t - E[r_t]|}{fs_t}, \tag{8}$$

where T_1 is the number of out-of-sample periods. The statistical properties of (8) can be derived straightforwardly from asymptotic theory. This statistic can also be used to evaluate different portfolios or hedging strategies.

We apply the risk ratio test to the MRM for Russell 2000 stocks. In the long-only case, we randomly generate 100 portfolios with 100 names in each portfolio. In the long-short case, we randomly generate 100 dollar-neutral portfolios with 100 names for both long and short positions. EXHIBIT 7 presents our results. The second row provides the mean statistic across all portfolios and the third row reports the t-statistic for the null hypothesis that the ratio equals one.

EXHIBIT 7: Risk Ratio Statistic

Portfolio	Long-Only	Long-Short
Risk Ratio Statistic	1.0455	0.9738
t-statistic	0.9827	-0.3193

While the model under-predicts long-only portfolio risk and over-predicts long-short portfolio risk, the forecasting errors are statistically insignificant.

EXHIBIT 8 plots the risk ratio statistic across portfolios. Interestingly, the risk ratio statistic varies more across portfolios for long-short portfolios. Stock-specific risk is relatively more important in a long-short portfolio than in a well-diversified long-only portfolio. Therefore, we would conclude that it is more difficult to predict stock-specific risk.

EXHIBIT 8: Risk Ratio Statistic Across Portfolios

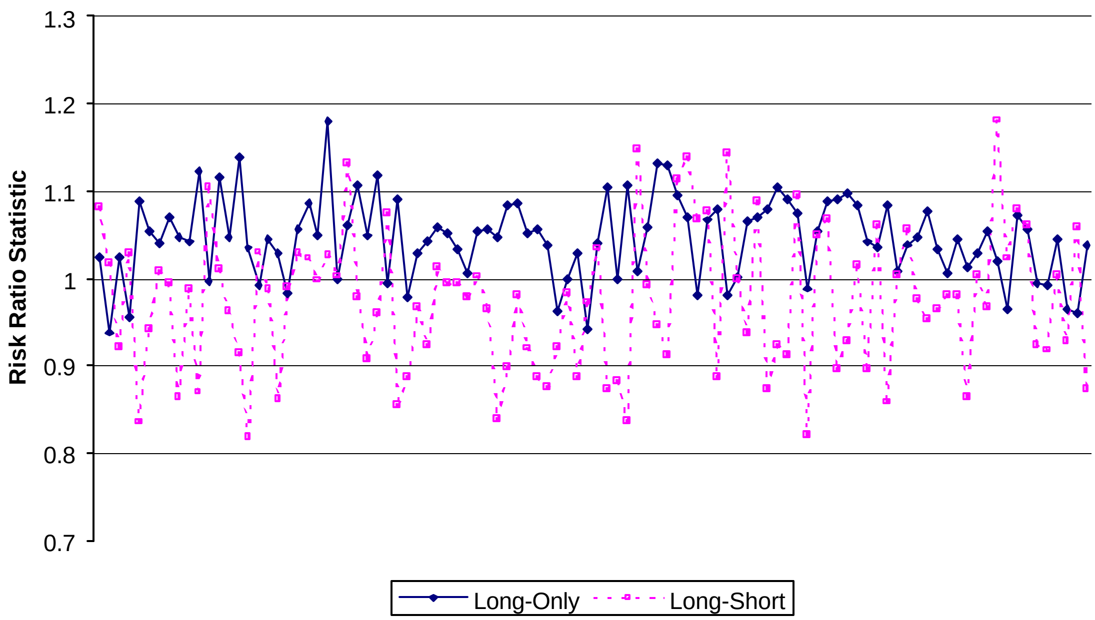

Portfolio Optimization

Index fund managers traditional attempt to track their benchmark portfolio by its matching the fundamental characteristics. Increasingly, this objective is achieved using factor risk models. Intuitively, a portfolio with *exposures* closely matching that of the benchmark might track the benchmark well, provided the idiosyncratic risk of the portfolio constituents is appropriately controlled. Hedge fund managers also use tracking error minimization to construct long-short portfolios. See Lo (2001) for a discussion of risk management issues pertaining to hedge funds. In this case, the "benchmark" is a preselected long portfolio or short portfolio, and a portfolio optimizer is used to determine the optimal hedge. To evaluate the power of factor models in forming optimized portfolios, we form optimal long-only and long-short portfolios using sample covariance and factor models, and then compare their out-of-sample performance. Specifically, at the beginning of each month starting January 1993 through December 2001, we randomly select 100 stocks. We estimate the variance matrix using historical returns (i.e., the sample variance matrix), using the previous five years of data. We also construct the variance matrix using MRM (estimated for each month using the previous five years) for each month in the same period.

We then construct optimal long-only and long-short portfolios using the estimated matrix. For the long-only case, we form minimum variance portfolios. Let S denote the estimated variance matrix using historical data. We form a minimum variance portfolio by minimizing $w'Sw$, where w is 100×1 vector of portfolio weights, we constrain portfolio weights to be non-negative with an upper limit 10%. We repeat this process using the estimated variance matrix from the factor model. The realized risk of the optimal portfolio is measured by the annualized standard deviation of the realized returns of the portfolio for the next 12 months, re-balanced based on price changes. The best variance matrix estimate will manifest itself in producing optimized portfolio with the lowest realized risk.

For the long-short portfolios we adopt a similar procedure. Long-short equity investing is the most common hedge fund strategy by far, accounting for 38 percent of hedge fund assets. We mimic a popular strategy, namely picking short positions following a proprietary stock selection model and selecting long positions to hedge risk. Essentially, the weights of the long positions are obtained by minimizing the overall portfolio risk. We randomly select 100 stocks to form an investment universe. Among these 100 stocks, 20 stocks are chosen at random for short positions with random weights satisfying a 10% upper limit. The optimal portfolio is obtained by minimizing the risk of the long-short portfolio. Again, the realized risk of the optimal portfolio is measured by the annualized standard deviation of the realized returns of the long-short portfolio for the next 12 months, re-balanced to maintain the same weights.

EXHIBIT 9 reports the annualized realized risk values of different optimal portfolios. Column 2 depicts the risk values of optimal portfolios obtained by using a factor model (the MRM); Column 3 details the risk values of the optimal portfolios obtained by using the historical sample variance matrix; and Column 4 displays the risk values of corresponding to equally weighted portfolios. For the long-only strategy, optimization clearly reduces portfolio risk. Indeed, by using a factor model, the risk of an equally weighted portfolio is cut almost in half. The reduction in risk by using a factor model for long-short portfolios is less dramatic, about 12%, consistent with the results in Chan, Karceski, and Lakonishok (1999).

EXHIBIT 9: Annualized Risk of Optimal Portfolios

Variance matrix estimate	Factor Model	Sample Estimate	Equally-Weighted
Long-Only Portfolio	0.0826	0.0940	0.1585
Long-Short Portfolio	0.1387	0.1583	0.1403

Controlling Exposure Distribution

In the long-short examples discussed above, a short position is hedged by selecting long positions from a particular universe. This practice assumes that a portfolio can be characterized by its exposures to risk factors and specific risks. However, portfolio exposures to risk factors are a weighted average of individual stock exposures, and therefore do not fully reveal the information on the exposures of the underlying stocks. A related point is made by Clarke, de Silva, and Wander (2002), who compare risk and asset allocation strategies. In particular, information on the exposure distributions across the stocks is not retained. Exhibit 10 below provides an illustrative example of a histogram of factor exposure for a given universe. As depicted, the majority of stocks have no exposure to the factor, but there is considerable dispersion over the universe.

EXHIBIT 10: Exposure

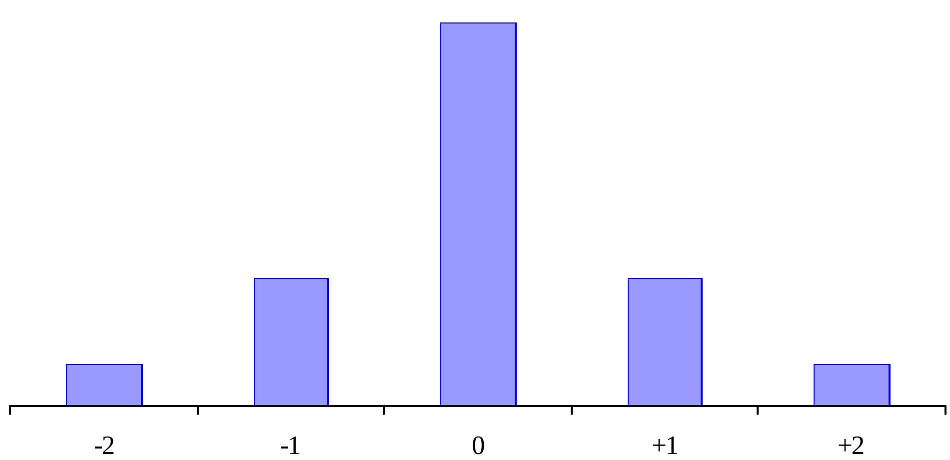

Now consider, say, a portfolio manager who wants to hold 20 stocks that has the same factor exposures as the universe as a whole. This "zero beta" portfolio can be constructed in many ways, for example, with 20 stocks selected from the universe whose exposure is 0 or with 10 stocks each from the extreme factor exposure categories of –2 and +2. But clearly these two portfolios are not alike. The first portfolio is overly concentrated in "zero" exposure stocks while the second portfolio, which is comprised only of extreme beta stocks, is subject to risks arising from estimation error. A portfolio with dispersed exposure distribution across stocks behaves quite differently than one with less dispersion, even though the portfolio exposures are the same. To construct optimal hedging portfolios that match the *distribution* of risk exposures, we need to match not only the value-weighted exposures of long and short portfolios but also the value-weighted positions with exposure values belonging to certain pre-specified intervals. In the example above, we would select stocks whose distribution is similar to the overall universe in Exhibit 10. For larger portfolios, this goal is difficult if not impossible to achieve without the use of powerful portfolio optimizers that can handle integer constraints.

Risk Analysis for Traders

Portfolio managers traditionally are the primary users of risk models. Increasingly though, traders are using risk models to manage their positions, formulate trading strategies, and create optimized trading lists. This trend is fueled by the growth of portfolio trading, increased sophistication by both buy- and sell-side traders, and automation. High frequency (daily or weekly), stock-specific (time-series) risk models are especially useful for trading applications where parameter estimates must be current, horizons are typically short, and idiosyncratic risk is often undiversified. This section describes some interesting new uses of risk analysis in the context of trading, focusing primarily on portfolio traders. Many of our examples, however, apply equally well to single-stock traders.

In large mutual fund companies, orders originate from several different portfolio managers and are aggregated by the trading desk at the open. The collective actions of several managers might result in a trade list with significant sector or industry risks. Even if the trade list at the start of the day is balanced by sector or industry, this can change if a trader executes a large-block (e.g., in a crossing system or through the block-brokerage "upstairs" market), suddenly creating an "active" bet. Risk models can be used to assess the risks of the entire trading list across all orders sent to the trading desk, and identify those securities that are most risky. (From the viewpoint of the desk, a portfolio manager's buy order actually represents a "short" position, since the desk is to deliver the shares to the manager.) Risk models have a particularly useful feature in this context, namely the ease with which we can compute the effect on risk of an incremental change in holdings. Specifically, the **Marginal Contribution to Risk** (MCR) is the first derivative with respect to holdings, computed for each security in the trade list. Here, risk can be either total risk or active risk, i.e., risk relative to a benchmark portfolio. The MCR indicates which orders pose the greatest risk, and can be an invaluable aid in prioritizing execution.

Risk models are valuable in pre-trade analysis, where they are used to determine optimal trading strategy. In particular, institutions typically have orders that are large relative to available liquidity, necessitating that they be broken up over time. Indeed, some institutional orders are traded over periods of weeks. Trading too quickly might result in market or price impact costs (Madhavan, 2002) that can substantially erode investment performance. But trading slowly increases the risk of an adverse price movement before the order is filled; the resulting in delay costs can also reduce alpha substantially. Traders must thus balance opportunity costs (risk) against impact costs (liquidity) in determining their optimal strategy. While traders often have a good sense of current liquidity (i.e., volumes, depth), quantifying opportunity costs is much trickier. One use of risk models is to measure these risks andbalance the two cost elements to determine the optimal trading strategy. Formally, the problem to be solved is a dynamic one: find a sequence of trades $q_1, \ldots, q_T$ (where q_t is vector representing the signed volumes in each stock in the program in period t) over T periods to minimize a weighted average of the expected price impact costs (pi) and the corresponding opportunity costs. With opportunity costs proportional to the volatility (standard deviation) of the realized proceeds, the specific objective is

$$Min_{\{q_t\}} C = l\, pi(q_1,...,q_T) + (1 - l) s\, (q_1,...,q_T)\,, \tag{9}$$

where λ is the weight on market impact cost and s is the volatility of realized proceeds. (Higher values of λ yield more passive optimal strategies.) Complicating matters is the fact that through price dynamics, a trade in period 1 affects all future prices. This necessitates a solution using stochastic dynamic programming. Models of this type are discussed further in Madhavan (2002), and are now commonly implemented by sophisticated traders.

This type of pre-trade analysis can be performed routinely, but is especially valuable for portfolio transitions and rebalancing. Consider, the annual reconstitution of the Frank Russell Company's equity indexes, which are widely used as performance benchmarks for investment managers. At the end of June each year, the Frank Russell Company reconstitutes its indexes based on market capitalization at the end of May. To avoid the price impacts associated with index additions and deletions, some portfolio managers benchmarked to these indexes trade ahead of the reconstitution date. But this can involve tracking error risk, because the differences between the indexes immediately before and after reconstitution can be significant from a risk perspective. In 2002, the active risk between the pre- and post-reconstitution Russell 2000 was over 2.2%, a significant figure for an index manager. Exhibit 11 shows the breakdown of this "active" risk for the Russell 1000 and 2000 indexes. The majority of active risk is due to exposure to market and value/growth factors. This analysis is valuable for a trader considering establishing or liquidating positions based on additions and deletions to these indexes prior to actual reconstitution.

EXHIBIT 11: Active Risk

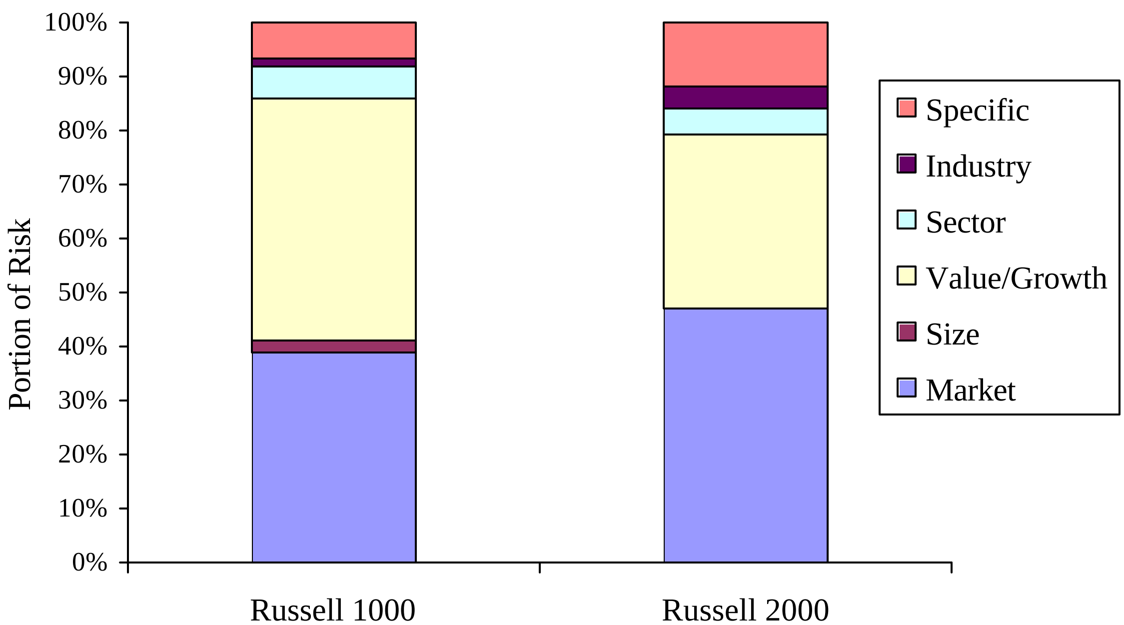

Portfolio transitions also provide a good illustration of the use of risk models in developing strategy. Transitions can be extremely costly because of the costs incurred in entering into and exiting from large positions. For this reason, portfolio transitions are often executed over long periods of time, sometimes days or even weeks. Simple strategies to mitigate risk arising from unfavorable price movements over the trading horizon – sector or industry neutrality – are a common solution to this problem. However, simple risk controls might not be *optimal* when balanced against other portfolio considerations such as controlling trading costs. Further, such strategies simply might not be *feasible* given common portfolio constraints such as short sale restrictions or cash balance restrictions. One solution is to combine a portfolio optimizer with a pre-trade cost predictors and a high-frequency risk model. The optimizer can generate optimal "waves" of trades to control risk while minimizing trading costs and satisfying other portfolio constraints such as tracking error, benchmark concentration, or industry sector exposures.

Finally, an important use of risk models is to evaluate guaranteed principal bids, where a dealer guarantees execution of a list at a benchmark price (e.g., the close) for a fixed commission. We can think of a principal bid as comprised of two elements: the cost of trading the list as an agent, plus a premium that represents compensation to the dealer for bearing principal risk, as in (9) above. The second term is directly related to the risk of the portfolio under bid. Using a risk model together with a pre-trade cost predictor allows a trader to evaluate bids relative to the cost of trading the list on an agency basis or directly. Similarly, dealers can use these tools to formulate their bids.

C. Evaluation

Risk models can also be used to assess performance looking backward. Consider, for example, a portfolio manager whose return in a given quarter is 3.4%. This raises the natural question: What is the source of this performance? Traditionally, the question is answered by simple stratification analysis, i.e., a breakdown of performance by market capitalization, market, currency, relative to some benchmark index. Although stratification analysis is helpful, it does not help us understand what risks the manager took in the form of factor exposures to achieve this return, nor does it explain whether, on a risk adjusted basis, performance was positive.

Risk models can be used to assess factor exposures and compute risk-adjusted returns. Performance analysis is typically motivated by plan sponsor concerns regarding the nature of risks taken by investment managers to achieve their returns. However, performance analysis is also valuable to investment managers because it is a tool to help them monitor and improve performance. It helps ensure that active positions are properly compensated, and avoid unnecessary risks. Let a_p represent performance, after controlling for risk factors, where:

$$r_p - r_f = a_p + \sum_{i=1}^{k} b_{ip} F_i + e_p, \quad (10)$$

where r_p is the portfolio realized return, r_f is the risk-free rate, a_p is the intercept, F_i is the factor excess return, b_{ip} is the factor loading on factor i and e_p is a stochastic disturbance term.

Further understanding of the portfolio manager's performance can be gained by juxtaposing the portfolio's return to each individual factor against the benchmark return. In the case described above, the analysis shows that the return, while strongly positive, fell short of the benchmark return of 5.8%, a difference of –2.4%. EXHIBIT 12 below illustrates the breakdown by factor type for the portfolio and benchmark returns.

EXHIBIT 12: Return Attribution

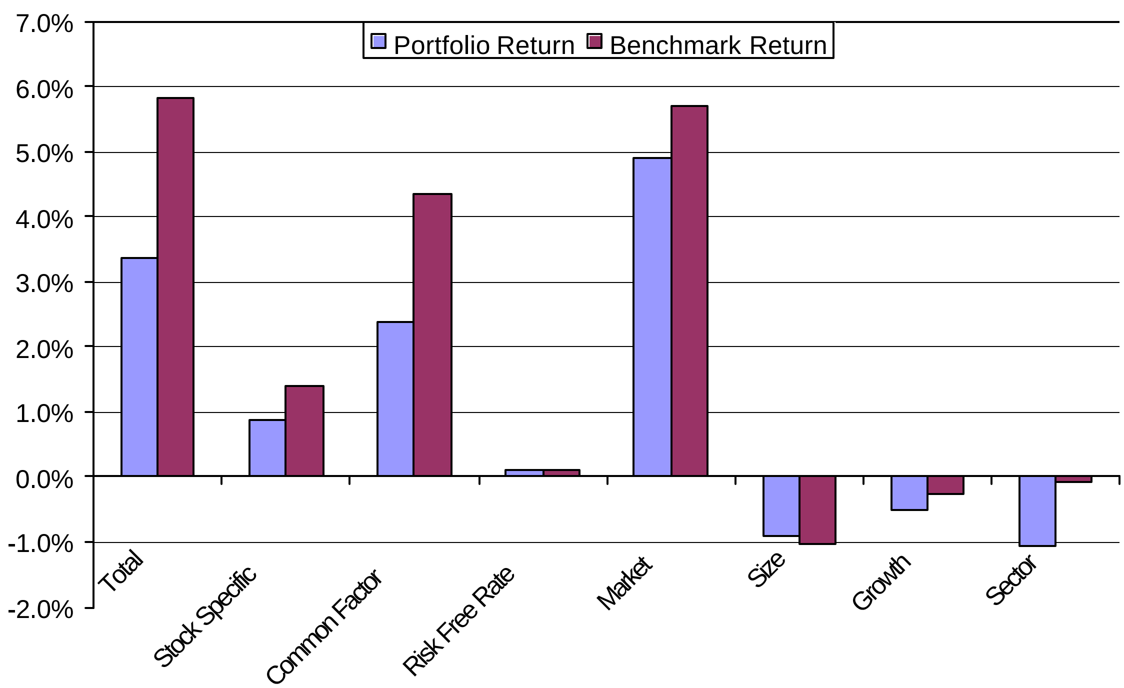

In this case, much of the underperformance came from the fact that the manager's stock selections underweighted the market factor, which was strongly positive in the period. Certain sectoral bets (underweighted in Technology and overweighted in Capital Goods and Energy) also reduced performance relative to the index. Further analysis shows that although the raw shortfall relative to the index was –2.4%, the risk-adjusted shortfall (a_p) is much smaller, only –0.6%, because the actual portfolio was much less risky than the benchmark. This example shows the power of a risk-based attribution analysis over naïve stratification approaches.

CONCLUSION

Risk models are powerful tools to measure, analyze, and evaluate risk. These tools are superior to naïve historical forecasts or simple heuristics in many dimensions, most importantly in solving the dimensionality problem. As shown here, risk models are not alike, and users need to be aware of how differences in their construction can affect their processes. Risk models can be applied to a variety of practical problems. A new and interesting set of applications concerns trading, especially in a portfolio context, and we expect further developments in this area as traders recognize the value of high-frequency risk models to their operations. While risk models are powerful, their complexity can also pose challenges in implementation and interpretation, as seen here. Particular care must be taken in construction of hedged portfolios and in matching the risk forecast to the holding period. Risk analysis, as it stands today, is as much an art as a science.

REFERENCES

Chan, Louis K. C. and Josef Lakonishok. "Are the Reports of Beta's Death Premature?" *Journal of Portfolio Management*, 19, Summer 1993, pp. 51-62.

Chan, Louis K. C., Jason Karcerski and Josef Lakonishok. "On Portfolio Optimization: Forecasting Covariances and Choosing the Risk Model." *Review of Financial Studies*, 12 (1999), pp. 937-974.

Clarke, Roger, Harindra de Silva, and Brett Wander, "Risk Allocation versus Asset Allocation," *Journal of Portfolio Management*, 29, Fall 2002, pp. 9-30.

Connor, Gregory and Robert A. Korajczyk. "Performance Measurement with the Arbitrage Pricing Theory." *Journal of Financial Economics*, 15 (1986), pp. 373-394.

Connor, Gregory, Robert A. Korajczyk, and Oliver Linton. "The Common and Specific Components of Dynamic Volatility." Finance Working Paper No. 311 (2002), Kellogg School of Management, Northwestern University.

Fama, Eugene F. and Kenneth R. French. "Common Risk Factors in the Returns on Stock and Bonds." *Journal of Financial Economics*, 33 (1993), pp. 3-56.

Jones, Christopher S. “Extracting Factors from Heteroskedastic Asset Returns.” *Journal of Financial Economics*, 62 (2001), pp. 293-325.

Lo, Andrew. “Risk Management for Hedge Funds: Introduction and Overview,” *Financial Analysts Journal* 57, November/December 2001, pp.16-33.

Madhavan, Ananth. “Market Microstructure: A Practitioner’s Guide.” *Financial Analysts Journal*, 58, September/October 2002, pp. 28-42.