Risk Management in Asset Management

Gregory Connor*

and

Robert A. Korajczyk**

January 17, 2003

Abstract

Investors are natural risk bearers, in part due to the vast array of risk management tools available to them. These tools allow a risk budgeting process that de-couples the asset allocation and active bets taken in the portfolio. The risk of non-traded assets in the portfolio can be reduced by selective hedging and insurance products. Non-traded assets and a dynamic risk/return tradeoff lead to horizon specific asset allocation. Portfolios should be constructed to account for the systematic shifts in asset liquidity.

Forthcoming in:

The Growth of Risk Management: A History

*Department of Accounting and Finance, London School of Economics, Houghton Street, London WC2A 2AE, United Kingdom; g.connor@lse.ac.uk.

**Kellogg School of Management, Northwestern University, 2001 Sheridan Road, Evanston, IL 60208-2001, USA; r−korajczyk@northwestern.edu.

I. Introduction

The application of risk management techniques to portfolio management crucially depends upon modern portfolio theory, beginning with the seminal contributions of Markowitz (1952), Treynor (1961, 1999), Sharpe (1964), and Lintner (1965). These papers and their extensions are covered elsewhere in this book (see the chapters by Pearson and by Cochrane and Culp) and that foundational material is not repeated here. In section II we discuss risk budgeting and the decoupling of portfolio risk management into three separate components: strategic asset allocation, tactical asset allocation, and security selection. In section III we consider risk-return dynamics and how these affect portfolio risk management. Section IV deals with liquidity problems in portfolio risk management, and Section V concludes.

II. Risk Budgeting Through Strategic Asset Allocation, Tactical Asset Allocation, and Active Deviations from Benchmark Portfolios

The goal of risk management in portfolio management is not to eliminate risk, but to choose which risks to bear and to avoid unnecessary risks. What risks are appropriate for a particular portfolio will depend on the risk preferences of the investor and the role that particular portfolio plays in the investor's overall portfolio strategy. We will make the distinction between the investor's overall/total portfolio and sub-component portfolios. The total portfolio is the ultimate portfolio formed by combining sub-component portfolios. The goal of the investor is to construct a total portfolio from sub-component assets and portfolios that best suits the return requirements and risk aversion of the investor. What might constitute a good sub-component of the portfolio (e.g., a hedge fund manager earning a high risk-adjusted return) may be totally inappropriate as the sole component of the total portfolio because of the high risk involved. Throughout, we treat the portfolio optimization problem and risk budgeting as "flip-sides" to the same coin (also see Scherer (2002)).

The investor should begin by specifying the universe of investment styles, or benchmark portfolios, that make up the set of passive portfolios from which strategic asset allocation is chosen. Time variation in the risk premia or the level of risk of the benchmark portfolios leads to tactical shifts in asset allocation. Finally, the existence of superior or inferior investments (i.e., mispriced assets or active portfolio managers earning abnormal returns) leads to active shifts away from the passive benchmark portfolios. Each step entails managing the tradeoff between extra return and extra risk.

To illustrate optimal portfolio construction, we will use a minor extension of Treynor and Black (1973) to the case of a *k*-style universe. In practice, the choice of the set of styles should be broad enough to encompass the set of risks that are compensated by risk premia, as discussed in the chapter by Cochrane and Culp. These style portfolios make up the set of benchmarks against which portfolio managers and other investments are compared.

To simplify our example, we assume that there are two benchmark portfolios: an aggregate market portfolio and a portfolio of "value" stocks, whose rates of return from period t-1 to period t are denoted $R_{m,t}$ and $R_{v,t}$, respectively (extensions to a greater selection of styles are straightforward but add unnecessary notation). Also let $R_{f,t}$ denote the return on an investment in a one-period riskless asset. In addition to the benchmark style portfolios and the riskless asset we have *n* other investments which might be individual assets/liabilities, mutual funds, hedge funds, derivative positions, and so on (with returns $R_{i,t}$ for $i = 1, 2, \ldots, n$). Returns in excess of the riskless return are denoted $r_{m,t} = (R_{m,t} - R_{f,t})$, $r_{v,t} = (R_{v,t} - R_{f,t})$, and $r_{i,t} = (R_{i,t} - R_{f,t})$. Also let $\mu_{m,t}$, $\mu_{v,t}$, and $\mu_{i,t}$ denote the investors expectations of $r_{m,t}$, $r_{v,t}$, and $r_{i,t}$ and let $\sigma_{m,t}$, $\sigma_{v,t}$, and $\sigma_{i,t}$ denote the investors expectations of the standard deviations of $r_{m,t}$, $r_{v,t}$, and $r_{i,t}$. We can relate the value style portfolio to the market portfolio using the standard market model:

$$r_{v,t} = \alpha_{v,t} + \beta_{v,m} r_{m,t} + \varepsilon_{v,t}. \quad (1)$$

In a CAPM world with no superior information, $\alpha_{v,t} = 0$. However, in a (non-CAPM) world where there are risk premia for additional sources of risk (see the chapter by Cochrane and Culp) or a world in which our investor has superior insights into the mispricing of the value style, $\alpha_{v,t}$ need not be zero. It is either an additional risk premium or abnormal return to value stocks. Similarly, we can relate the additional n potential investments to the factor portfolios (the market and value portfolios in this example) using a multi-factor market model:

$$r_{i,t} = \alpha_{i,t} + \beta_{i,m} r_{m,t} + \beta_{i,v} r_{v,t} + \varepsilon_{i,t}. \quad (2)$$

Since we assume that we have the correct style portfolios, $\alpha_{i,t}$ is the investor's forward looking expectation of the abnormal return on that asset or portfolio. It might be based on the investor's own information or the investor's expectation of the performance that will be obtained by the manager of portfolio i.

We will assume that the investor wishes to maximize the ratio of expected excess return to standard deviation, the Sharpe ratio (Sharpe (1966)), on the investor's total portfolio, p. Let X_m, X_v, X_i ($i = 1, 2, \ldots, n$), and X_f, denote the fraction-s of the net portfolio value (assets less liabilities), or portfolio weights, invested in the market, the value index, asset or portfolio *i*, and the riskless asset. Any wealth not invested in *m*, *v*, or assets/portfolios 1 through *n* will be invested in the riskless asset. Positive (negative) values of *X* correspond to assets/long positions (liabilities/short positions). As in Treynor and Black (1973), for now we assume that there are no restrictions on the portfolio weights and we assume that the non-style returns to the *n* assets are uncorrelated (i.e., $\text{corr}(\varepsilon_{i,t}, \varepsilon_{j,t}) = 0$ if $i \neq j$).

The investor's portfolio has exposure to the second style portfolio, *v*, from two sources: the direct position in the style index and the indirect exposure through assets 1 to *n*. If we define X_v^* to be the total exposure to the value-style index, we have that:

$$X_v^* = X_v + \sum_{i=1}^{n} X_i \beta_{i,v} \, . \tag{3}$$

Similarly, the investor's portfolio has exposure to the market factor portfolio from three sources: the direct position in the market, the exposure to the value-style index and the indirect exposure through assets 1 to *n*. If we define X_m^* to be the total exposure to the market, we have that:

$$X_m^* = X_m + X_v^* \beta_{v,m} + \sum_{i=1}^{n} X_i \beta_{i,m} . \tag{4}$$

The portfolios weights that maximize the reward to risk (Sharpe) ratio are:

$$X_{m,t}^* = \lambda \left[\frac{\mu_{m,t}}{\sigma_{m,t}^2} \right], \tag{5}$$

$$X_{v,t}^* = \lambda \left[\frac{\alpha_{v,t}}{\sigma^2(\varepsilon_{v,t})} \right], \tag{6}$$

$$X_{i,t}^* = \lambda \left[\frac{\alpha_{i,t}}{\sigma^2(\varepsilon_{i,t})} \right]. \tag{7}$$

The constant of proportionality, λ ($\lambda > 0$), reflects the investor's level of risk aversion. More risk-averse investors will have values of λ closer to zero. These portfolio weights embody the investor's strategic and tactical asset allocations as well as the investor's active bets.

The strategic asset allocation is determined by the investor's long run expectations of the returns on the style portfolios, μ_m and α_v, the risk of the market, σ_m, the tracking error of the *v* portfolio, $\sigma^2(\varepsilon_v)$, and the investor's risk aversion (reflected in λ). The tactical asset allocation is determined by the current deviation from the long run expectations of the returns and risks of the style portfolios: $\mu_{m,t} - \mu_m$, $\alpha_{v,t} - \alpha_v$, $\sigma_{m,t} - \sigma_m$, and $\sigma^2(\varepsilon_{v,t}) - \sigma^2(\varepsilon_v)$. Finally, the active positions are larger the higher the abnormal return, alpha, from the active positions, and smaller (in absolute value) the larger the tracking error of the active investment, $\sigma^2(\varepsilon_{v,t})$.

An important insight of the Treynor and Black (1973) analysis is that the tools of risk management allow us to completely decouple the asset allocation and active bets in the portfolio. This is quite different from the manner in which many portfolios are managed. For example, assume that our investor believes that $\alpha_{v,t} = 0$. This implies that the investor does not want any exposure to the value-style portfolio (beyond that inherent in the market index). In many traditional settings this would lead the investor to exclude value managers from the investor's active manager search. What if there exists a value manager that earns a large alpha relative to tracking error? The Treynor and Black analysis implies that the investor should hire that manager ($X_{i,t} > 0$ when $\alpha_{i,t} > 0$) but reverse the exposure to the value-style by taking a short position in the value index (for example, by shorting a futures contract). In essence, through this set of positions we can create a style-neutral portfolio (long the positive alpha style manager and short the style index). This allows us to have our cake and eat it too. We hire active managers on the basis of α relative to tracking error, not on the basis of our asset allocation decision. The risk is then managed by rebalancing the asset allocation to our desired position. Thus, decisions on active bets and decisions on asset allocation can be decoupled. The manager can go "over budget" on style risk with the asset allocation adjusted to stay within budget (i.e., change X_v to get to the desired level X_v^*). This decoupling is referred to as "Portable Alpha" by Arnott (2002).

The above analysis also has important implications for active managers of sub-component portfolios. Active managers, whose clients are using risk management techniques effectively, should not worry about their Sharpe ratio. The Sharpe ratio rewards managers for superior alpha *and* for overall diversification. When the active manager's client is managing risk optimally, the manager does not need to worry about overall diversification, but needs to worry about providing the highest α per unit of tracking error. Thus the total portfolio and the sub-component portfolios should be evaluated by different criteria. The analysis also provides a strong rationale for market- or style-neutral (long/short) portfolios. Not only can the managers specialize in information production, but they are in a better position to adjust the portfolio to remain style-neutral as the composition changes. Some institutions use this type of analysis to reverse engineer the implied α from a manager's position. That is, given the active position taken by the manager in asset i, X_i, we can determine the level of α_i consistent with the position from (7), (Patel (2002)). The manager can then assess whether that level of abnormal return is reasonable.

The intuition of this analysis is clear and robust: α is good, tracking error is bad, and active positions and asset allocation need to be balanced effectively by the risk manager. Relaxing some of the assumptions will not change this intuition, but may change the optimal positions. The assumption that non-style returns are uncorrelated across assets (i.e., corr($\varepsilon_{i,t}$, $\varepsilon_{j,t}$) = 0 if $i \neq j$) can be approximately true, particularly when include a large number of style benchmarks. However, active managers with the same style specialization may have non-zero residual correlation. In practice, this is often handled by choosing only one manager per style, the manager with the highest α-to-tracking error ratio. This is generally sub-optimal, and can be improved upon by a simple mean variance optimization across managers in the sub-style category.

We have placed no restrictions on the portfolio weights. In practice there may be such restrictions. Some portfolios are restricted from entering into short sales. Another common restriction is that the positions in certain assets and/or liabilities are not under the control of the portfolio manager. For example, for a typical investor one of the n assets is that investor's human capital. We can not reduce the portfolio weight in human capital by selling that asset (which would amount to indentured servitude). A second example is a portfolio manager for an insurance company may have liabilities that are dictated by the insurance policies written by the firm (and, hence, fixed from the portfolio manager's perspective). If the assets or liabilities are well-diversified, in the sense of having very little tracking error (e.g., the insurance company's liability portfolio might look like a long-term bond index fund), then we can use the style indices to return to our preferred asset allocation. If assets or liabilities have substantial tracking error, then the investor may choose to specifically hedge that risk by finding assets with negative correlation with $\varepsilon_{i,t}$ (e.g., buying insurance in the human capital example).

The above analysis allows the investor's expectations to change over time, but the investor continues to act myopically, in the sense of solving a one-period mean/variance problem. We now turn to the problem of managing risk while taking into account the dynamics of the risk/return tradeoff.

III. Risk-Return Dynamics and the Planning Horizon

Some investors have very short planning horizons, such as one day for many floor traders and market makers, and other investors have very long planning horizons, such as a century or more for some university endowment funds. As Samuelson (1969) first noted, if asset prices follow a random walk and preferences are logarithmic, then the risk-return trade-off is unaffected by the planning horizon. Very short-term and very long-term investors measure the risk/return trade-off in the same way, without reference to their different planning horizons. If Samuelson's conditions do not hold – asset prices are not a random walk and/or preferences are not logarithmic -- then the planning horizon can affect the measurement or cost-benefit valuation of risk. In addition, non-traded assets, such as human capital, can also lead to interactions between the planning horizon and optimal portfolio choice (Jagannathan and Kocherlakota (1996)).

It is the violation of the random walk assumption, more than the logarithmic preferences assumption, that tends to motivate extensions of the standard portfolio risk management model. It is abundantly clear that (log) prices do not follow a random walk. There are at least two categories of empirical violations: short-term risk dynamics and return predictability.

A. Short-term Risk Dynamics

Asset returns exhibit strong volatility clustering. High-risk days, for example, tend to be followed by high-risk days, and vice-versa for low-risk days. There are many methods for estimating the dynamics of volatility (see the chapter by Duan). To illustrate some of the issues, we use a classic model of volatility clustering, the GARCH(1,1) model. Let $\sigma_{t,t+1}$ denote the variance (conditional on information at time t) for the return of an asset at time t+1. The GARCH(1,1) model assumes that this conditional variance is a linear function of the previous period's conditional variance and the square of the current period's de-meaned return on the asset, z_t:

$$\sigma^2_{t,t+1} = \omega + a\, z_t^2 + b\, \sigma^2_{t-1,t}\,. \qquad (8)$$

Note that the risk of the asset expressed in units of variance not only moves through time, but also depends upon the planning horizon of the investor. A one-period investor can use (8) directly to measure risk. If the investor plans to hold the asset for *k* periods, then it is necessary to generate the conditional forecasts from (8) for each future period, and temporally aggregate the forecasts across the holding period. Performing this calculation (see Bollerslev, Engle, and Nelson (1994) for a general review of GARCH-type models) gives:

$$\sigma^2_{t,t+k} = \omega \frac{1-(a+b)^k}{1-a-b} + \left((a+b)^k - \theta\right)\zeta^2_{t-k,t} + \theta\sigma^2_{t-k,t} \qquad (9)$$

Where $\sigma^2_{t,t+k}$ denotes the conditional variance of the *k*-period return (not to be confused with the k-step-ahead one-period conditional variance) and θ is a function of the parameters in (8). As long as $a + b < 1$ (a necessary condition for covariance stationarity of the GARCH model) the very long-term investor can approximately ignore GARCH effects: the longer the holding period, the smaller the distinction between conditional and unconditional holding-period variance. This turns out to be an important practical consideration: empirically observed GARCH effects are very strong at daily and higher frequency but die out fairly quickly. Correcting risk forecasts for GARCH effects is much more important for short-term than long-term investors. As we will see, the opposite applies to mean-reversion effects, where short-term investors are not affected at all and long-term investors might be.

B. Return Predictability and Long-term Mean Reversion

A prominently discussed anomaly in empirical finance is the presence of "excess volatility" in returns (Shiller (1991)). However, such "excess volatility" is equivalent to time variation, or predictability in asset returns (Cochrane (1991)). Lo and MacKinlay (1988) note that excess volatility plus return stationarity necessarily implies that long-horizon returns have lower proportionate variance than short-horizon returns. This means that a long-term investor and a short-term investor face a different risk-return tradeoff for the same multiple asset opportunity set. The difference in risk between high-risk equities and low-risk cash instruments is proportionately less for the long-term investor. Many analysts have used this finding to propose higher weightings on equities for longer-term holders. This advice has had considerable influence on investment practice in recent decades, particularly in North America.

Intuitively, the long-term investor experiences lower proportionate risk from equities since he/she can "ride out" the short-term price fluctuations due to excess volatility, holding long enough for prices to revert back toward fundamental values. Campbell and Viceira (2001) advocate a more aggressive policy for long-term investors. Excess volatility/mean reversion also implies some small degree of predictability in long-term returns. During “down markets” (defined by some statistical criteria such as yield ratios) the expected return to high-risk equities is higher than during “up markets.” Long-term investors should tilt their asset allocation plan to account for this. This is not strictly speaking a risk management issue since it concerns expected returns rather than risk, but the effects of excess volatility/mean reversion are intimately connected.

C. Other Dynamic Features of Returns and Their Implications for Portfolio Risk Management

There are numerous other dynamic risk-return patterns observable empirically, but only a few will be mentioned here. Over annual horizons, returns seem to exhibit momentum rather than its opposite, mean reversion. This implies that over certain intermediate-length holding periods, variance increases proportionately with the holding period rather than decreasing. Another important empirical finding comes from Campbell, Lettau, Malkiel, and Xu (2001) who show that the proportion of asset-specific risk in total risk (for a typical individual asset) has experienced a secular increase over the last fifty year. This means that portfolio diversification, to a given tolerance level, requires more assets now than it did in earlier decades.

Another important dynamic feature in returns is the decline in kurtosis as the return interval gets longer. Daily and higher-frequency returns have very high positive excess kurtosis. Assuming reasonable limits on return interdependence over time, and the existence of finite higher moments, it follows from the Central Limit Theorem that the excess kurtosis in multiperiod returns will decline towards zero as the measurement interval grows. This strong decline in kurtosis is in fact observed empirically.   If investors care about kurtosis and other higher moments of return, not just variance, in measuring risk, then their evaluation of risk of a given asset can change with the planning horizons due to this affect.

IV. Execution/Liquidity Risk

A common finding is that "paper" (i.e., simulated) portfolios always outperform real portfolios based on the same information. The reason is that the real portfolio incurs execution costs (commissions, price impact, partial executions, etc.) that the "paper" portfolio does not (Treynor (1983)). There is a wide array of execution performance metrics. Many commonly used metrics, such as Volume Weighted Average Price (VWAP), can be easily gamed − especially if the order is large, necessitating working the order (Beebower (1989)). The extensive literature on "market microstructure" provides a set of models (see O'Hara (1995) and Harris (2002)) and a growing empirical literature that can provide benchmarks against which trade execution can be evaluated (e.g., Keim and Madhavan (1995, 1997, 1998) and Breen, Hodrick, and Korajczyk (2002)).

More importantly, there is evidence that there are systematic components to liquidity (Chordia, Roll, and Subrahmanyam (2000), Hasbrouck and Seppi (2001), Pástor and Stambaugh (2001), and Sadka (2002)). Thus, an investor might find that many of the assets in the portfolio are simultaneously difficult to trade. In fact, dramatic systematic shifts in liquidity appear to be an important factor in numerous financial crises (e.g., Edwards (1999)). Sub-component portfolios that incorporate long/short strategies and leverage may have a higher chance of incurring margin calls during periods of low liquidity.

Portfolio construction should take into account the likelihood that assets will need to be traded in low liquidity environments. This is an area where extreme value theory (see the chapter by de Vries and Caserta) or simulation and stress testing (see the chapter by Bhansali) lead to useful insights.

V. Conclusion

There are many important aspects of risk management that, due to space constraints, we have not addressed. The inputs to the optimization process often require the aggregation of information from historical data with prior beliefs about parameters. The risk management process must deal with the associated estimation risk. Some approaches are explicitly Bayesian (e.g., Scherer (2002, Chp. 4)). Others impose constraints on holdings which can have a Bayesian interpretation (Jagannathan and Ma (2002)). Additionally, we have only tangentially addressed non-normality in asset returns, particularly in option held in the portfolio. This non-normality makes stress testing portfolios (see the chapter by Bhansali) all the more important.

Portfolio Optimization and risk budgeting are flip-sides of the same coin. The budget allocated to a particular source of risk depends on the investor's beliefs about the risk/reward tradeoff for that source of risk. The budgeting process should allow opportunistic shifts in the risk budget across investments. The tools of risk management (e.g., derivatives, insurance, leverage, exchange traded funds, etc.) allow a risk budgeting process that de-couples the asset allocation and active bets taken in the portfolio. Non-traded assets and a dynamic risk/return tradeoff lead to horizon specific asset allocation. Finally, portfolios need to be constructed to account for the systematic shifts in asset liquidity.

Bibliography

Arnott, R. D., 2002, "Risk Budgeting and Portable Alpha," *Journal of Investing* 11(2), pp. 15-22.

Beebower, G., 1989, "Evaluating Transaction Cost," in W. H. Wagner *The Complete Guide to Securities Transactions*, (New York: Wiley).

Bollerslev, T., R. F. Engle, D. B. Nelson, 1994, "ARCH Models," in R. F. Engle and D. L. McFadden *Handbook of Econometrics, Volume IV*, (Amsterdam: North Holland).

Breen, W. J., L. S. Hodrick, and R. A. Korajczyk, 2002, "Predicting Equity Liquidity," *Management Science*, 48(4), pp. 470-483.

Campbell, J. Y., M. Lettau, B. G. Malkiel, and Y. Xu, 2001, "Have Individual Stocks Become More Volatile? An Empirical Exploration of Idiosyncratic Risk," *Journal of Finance* 61(1), pp. 1-43.

Campbell, J.Y. and Luis M. Viceira, 2001, *Strategic Asset Allocation: Portfolio Choice for Long-term Investors*, Oxford University Press, Oxford, U.K.

Chordia, T., R. Roll, and A. Subrahmanyam, 2000, "Commonality in Liquidity," *Journal of Financial Economics* 56(1), pp. 3-28.

Cochrane, J., 1991, "Volatility Tests and Efficient Markets," *Journal of Monetary Economics* 27(3), 463-485.

Cochrane, J., and C. L. Culp, 2002, "Equilibrium Asset Pricing and Discount Factors: Overview and Implication for Derivatives Valuation and Risk Management." Chapter in the present volume.

Connor, G., R. A. Korajczyk, and O. Linton, 2002, "The Common and Specific Components of Dynamic Volatility," Working paper #311, Department of Finance, Northwestern University.

Duan, J.-C., 2002, "Modelling Volatility." Chapter ?? in the present volume.

Edwards, F. R., 1999, "Hedge Funds and the Collapse of Long Term Capital Management," *Journal; of Economic Perspectives*, 13(2), pp. 189-210.

Harris, L., 2002, *Trading & Exchanges*, (Oxford, Oxford University Press).

Jagannathan, R., and N. R. Kocherlakota, 1996, "Why Should Older People Invest Less in Stocks Than Younger People?" *Federal Reserve Bank of Minneapolis Quarterly Review*, 26(2), pp. 11-23.

Jagannathan, R., and Ma, T., 2002, “Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraints Helps,” forthcoming, *Journal of Finance*.

Keim, D. B., and A. Madhavan, 1996, “The Upstairs Market for Large-Block Transactions: Analysis and Measurement of Price Effects,” *Review of Financial Studies*, 9(1), pp. 1-36.

Keim, D. B., and A. Madhavan, 1997, “Transactions Costs and Investment Style: An Inter-Exchange Analysis of Institutional Equity Trades,” *Journal of Financial Economics*, 46(3), pp. 265-292.

Keim, D. B., and A. Madhavan, 1998, “The Cost of Institutional Equity Trades,” *Financial Analysts Journal*, 54(4), pp. 50-69.

Lintner, J., 1965, "The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets," *Review of Economics and Statistics* 47(1), pp. 13-37.

Markowitz, H. M., 1952, “Portfolio Selection,” *Journal of Finance* 7(1), pp. 77-91.

O’Hara, M., 1995, *Market Microstructure Theory*, (Oxford, Blackwell).

Pástor, L., and R. F. Stambaugh, 2001, “Liquidity Risk and Expected Stock Returns,” *Journal of Political Economy* forthcoming.

Patel, N., 2002, “Buying and Selling Risk,” *Risk* 15(9), pp. S12-S14.

Samuelson, P. A. “Lifetime Portfolio Selection by Dynamic Stochatic Programming,” *Review of Economics and Statistics* 51(3), pp. 239-246.

Sadka, R., 2002, “Momentum, Liquidity Risk, and Limits to Arbitrage” Working paper, Department of Finance, Northwestern University.

Scherer, B., 2002, *Portfolio Construction and Risk Budgeting*, (London: Risk Books).

Sharpe, W. F., 1964, "Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk,” *Journal of Finance* 19(3), pp. 425-442.

Sharpe, W. F., 1966, “Mutual Fund Performance,” *Journal of Business* 39(1), pp. 119-138.

Shiller, R. J., 1991, *Market Volatility*, (Cambridge, MA: MIT Press).

Treynor, J. L., 1961, “Toward a Theory of Market Value of Risky Assets,” Unpublished manuscript – subsequently published as Treynor (1999).

Treynor, J. L., and F. Black, 1973, "How to Use Security Analysis to Improve Portfolio Selection," *Journal of Business* 46(1), pp. 66-86.

Treynor, J. L., 1983, "Implementation of Strategy: Execution," in J. L. Maginn and D. L. Tuttle (eds), *Managing Investment Portfolios: A Dynamic Process*, pp. 537–572 (Boston: Warren, Gorham & Lamont).

Treynor, J. L., 1999, "Toward a Theory of Market Value of Risky Assets," in R. A. Korajczyk (ed), *Asset Pricing and Portfolio Performance: Models, Strategy and Performance Metrics*, (London: Risk Books).