



Fixed Income Research

GLOBAL FIXED INCOME PORTFOLIO STRATEGY

June 1998

Using the Black-Litterman Global Asset Allocation Model: Three Years of Practical Experience

Andrew Bevan
Kurt Winkelmann

Acknowledgments

We wish to acknowledge the comments and suggestions of Joanne Hill, Robert Litterman, Robert Litzenberger, Tom Macirowski, Scott McDermott, Victor Ng, and Scott Pinkus. Over the past three years, we have also received valuable feedback on our portfolio process from the Global Economics Group.

*Andrew Bevan
London
44-171-774-1168*

*Kurt Winkelmann
London
44-171-774-5545*

Andrew Bevan is an Executive Director in the Economic Research Group at Goldman Sachs International in London.

Kurt Winkelmann is an Executive Director in the Fixed Income Research Department at Goldman Sachs International in London.

Editor: Ronald A. Krieger

Copyright © 1998 by Goldman, Sachs & Co.

This material is for your private information, and we are not soliciting any action based upon it. This report is not to be construed as an offer to sell or the solicitation of an offer to buy any security in any jurisdiction where such an offer or solicitation would be illegal. Certain transactions, including those involving futures, options, and high yield securities, give rise to substantial risk and are not suitable for all investors. The material is based upon information that we consider reliable, but we do not represent that it is accurate or complete, and it should not be relied upon as such. Opinions expressed are our current opinions as of the date appearing on this material only. While we endeavor to update on a reasonable basis the information discussed in this material, there may be regulatory, compliance, or other reasons that prevent us from doing so. We and our affiliates, officers, directors, and employees, including persons involved in the preparation or issuance of this material, may, from time to time, have long or short positions in, and buy or sell, the securities, or derivatives (including options) thereof, of companies mentioned herein. No part of this material may be (i) copied, photocopied, or duplicated in any form, by any means, or (ii) redistributed without Goldman, Sachs & Co.'s prior written consent.

This material has been issued by Goldman, Sachs & Co. and/or one of its affiliates and has been approved by Goldman Sachs International, which is regulated by The Securities and Futures Authority, in connection with its distribution in the United Kingdom and by Goldman Sachs Canada in connection with its distribution in Canada. This material is distributed in Hong Kong by Goldman Sachs (Asia) L.L.C., in Japan by Goldman Sachs (Japan) Ltd., and in Singapore through Goldman Sachs (Singapore) Pte. This material is not for distribution in the United Kingdom to private customers, as that term is defined under the rules of The Securities and Futures Authority; and any investments, including any convertible bonds or derivatives, mentioned in this material will not be made available by us to any such private customer. Neither Goldman, Sachs & Co. nor its representative in Seoul, Korea, is licensed to engage in the securities business in the Republic of Korea. Goldman Sachs International and its non-U.S. affiliates may, to the extent permitted under applicable law, have acted upon or used this research, to the extent it relates to non-U.S. issuers, prior to or immediately following its publication. Foreign-currency-denominated securities are subject to fluctuations in exchange rates that could have an adverse effect on the value or price of, or income derived from, the investment. In addition, investors in certain securities such as ADRs, the values of which are influenced by foreign currencies, effectively assume currency risk.

Further information on any of the securities mentioned in this material may be obtained upon request, and for this purpose persons in Italy should contact Goldman Sachs S.I.M. S.p.A. in Milan, or at its London branch office at 133 Fleet Street, and persons in Hong Kong should contact Goldman Sachs (Asia) L.L.C. at 3 Garden Road. Unless governing law permits otherwise, you must contact a Goldman Sachs entity in your home jurisdiction if you want to use our services in effecting a transaction in the securities mentioned in this material.

Using the Black-Litterman Global Asset Allocation Model: Three Years of Practical Experience

Contents

Executive Summary

I. Introduction.....	1
II. An Overview of Our Process	2
III. Calculating Equilibrium Returns	2
IV. Setting the Weight-on-Views and Confidence Levels.....	4
V. Setting Target Risk Levels.....	5
VI. Optimization and Risk Decomposition	7
VII. How Have We Done?	8
VIII. Conclusions	10
Appendix A: Interpreting the Published Portfolios	12
Appendix B: Practical Issues in Developing a Process	13
Bibliography	15

Executive Summary

The Black-Litterman model, developed at Goldman Sachs in the early 1990s, provides a framework for combining investor views with a global capital market equilibrium. Its purpose is to help investment managers determine an optimal portfolio allocation for specific classes of assets in a manner consistent with their market views. With this model, we can calculate optimal portfolio weights by using volatilities and correlations across asset classes.

For the past three years, we have been publishing optimal global fixed income portfolios using the Black-Litterman framework — both to offer portfolio advice to our clients consistent with the views of our economists and to illustrate how the model can be used to solve practical investment management problems. This paper provides a summary of our experiences in using the model for investment strategy.

Following an overview of our “investment process,” we first explain how we set the key parameters in the Black-Litterman framework. This involves a discussion of how we use the model to observe the equilibrium returns in global capital markets and then blend the equilibrium returns with our own views to provide a set of expected returns. We explain how we determine the weight and confidence levels on our own views relative to the equilibrium. Next we discuss risk control and optimization. We describe the process we follow to set tracking error risk and Market Exposure (a statistical measure of a portfolio’s sensitivity to market moves). Finally, we discuss our performance over the three-year period and consider how the same framework can be applied to other fund management issues.

It turns out that in the aggregate, our published portfolio has outperformed its benchmark over the last three years. However, we focus on our performance in this paper only to illustrate how the Black-Litterman framework can be used for designing investment strategies. Clearly, the performance relative to the benchmark will in large part reflect the accuracy of our views. At the same time, however, it is at least as important to consider the impact of the risk control mechanisms used in the model. We have constructed our own portfolio against a global government bond index, but we can easily apply the approach to other asset classes and other benchmarks. Designing risk-controlled portfolios is likely to become increasingly important in global fixed income markets with the advent of European economic and monetary union and with the growth of local currency emerging government bond markets.

Using the Black-Litterman Global Asset Allocation Model: Three Years of Practical Experience

I. Introduction

The basic investment management problem is to simultaneously maximize performance and manage risk: Investors determine risk-controlled allocations to specific asset classes that make best use of the information at their disposal. Quantitative investment theory holds out the potential of providing a systematic framework for solving this problem. At the heart of this framework is an asset allocation model, which practitioners can use to find portfolio allocations that best reflect their investment views. They can determine optimal portfolio weights by using the volatilities and correlations across asset classes.

However, many practitioners have not fully incorporated this framework into their investment management processes, for two reasons: First, this approach can lead to dramatic swings in the optimal portfolio weights for small changes in investment views. Second, the optimal portfolio weights in the traditional mean/variance framework often seem to be taking risk positions that appear to be at odds with the strongest investment views. The Black-Litterman model (see Black and Litterman, 1990, and Black and Litterman, 1991) was developed to provide a systematic resolution to these problems.¹

A central feature of the Black-Litterman framework is the notion that investors should take risk where they have views, and correspondingly, they should take the most risk where they have the strongest views. In the Black-Litterman framework, all expected returns are viewed as a blend of a set of equilibrium returns (reflecting a neutral reference point) and an actual set of investor views (which should differ from the equilibrium returns). Consequently, the problem facing the practitioner is to determine the weight given to the actual views.

In February 1995, we started publishing optimal global fixed income portfolios using the Black-Litterman framework. These portfolios, which reflect both our economists' forecasts and the equilib-

rium returns, appear in our monthly publication, *Global Fixed Income Asset Allocation*. We had two purposes in starting this publication: First, we wanted to provide portfolio advice to our clients consistent with the views of our Global Economics Group. Second, we wanted to illustrate how the Black-Litterman model could be used to solve practical investment management problems.

In addition to using the equilibrium concept to find optimal portfolio weights, our published portfolios have incorporated risk control as an explicit component of the investment management process. In particular, we have used the concepts of Hot Spots™ (see Litterman, 1996) and Market Exposure (see Litterman and Winkelmann, 1996) to identify the risk distribution and directional bias in our portfolios.²

This paper provides a summary of our experiences in using the Black-Litterman model for investment strategy. In the aggregate, over a three-year period, our published portfolio has outperformed its benchmark. However, the primary purpose of this paper is not to discuss performance. Rather, our published portfolio (see Appendix A) provides an illustration of how the Black-Litterman framework can be used for designing investment strategies. While we have constructed our portfolio against a global government bond index, the approach can be easily applied to other asset classes and other benchmarks. *Indeed, designing risk-controlled portfolios is likely to become increasingly important in global fixed income markets as the effects of European economic and monetary union (EMU) begin to take hold.*

The paper is organized as follows: In the next section we give an overview of our "investment process." In Sections III and IV we explain how we set the key parameters in the Black-Litterman framework. In Sections V and VI we discuss risk control and optimization, and in Section VII we review our actual performance over the three-year period. Section VIII presents our conclusions and considers some open issues.

¹ A list of references appears at the end of this report, on page 15.

² Hot Spots is a trademark of Goldman, Sachs & Co.

II. An Overview of Our Process

As discussed in the Introduction, in the Black-Litterman framework, expected returns are viewed as a combination of a specific set of investor views with a neutral reference point. By extension, the framework has implications for risk control. As a result, successful implementation of the Black-Litterman model requires careful consideration of the key parameters influencing both expected returns and risk control.

Exhibit 1 Optimization Procedure

Step	Action	Purpose
1	Calculate equilibrium returns.	Set neutral reference point.
2	Determine weightings for views.	Dampen impact of aggressive views.
3	Set target tracking error.	Control risk relative to benchmark.
4	Set target Market Exposure.	Control directional effects.
5	Determine optimal portfolio weights.	Find allocations that maximize performance.
6	Examine risk distribution.	Determine whether risk is diversified.

The process that we have adopted is explicit in its determination of these parameters. It has six principal steps, which we summarize in Exhibit 1. First, we find the neutral reference point as the returns associated with a global capital market equilibrium. Considered from a different perspective, the asset returns from our neutral reference point are those required for a representative investor to hold the global capitalization-weighted portfolio.

Our second step is to determine how much weight to put on the neutral reference point returns relative to our actual views on asset returns. In this step, we also set a relative weighting for each of our individual views. The effect of this step is to dampen the impact on the portfolio composition of particularly aggressive views.

In our third and fourth steps, we set target risk levels. We choose our portfolio weights relative to the characteristics of a benchmark portfolio. Consequently, the two dimensions of risk for which targets are set are the tracking error (i.e., projected volatility of performance differences) and Market Exposure

(i.e., the directional bias of the portfolio relative to the benchmark).

Our fifth step is to find an optimal portfolio that maximizes expected return subject to the risk constraints. Finally, we consider the risk decomposition of the optimal portfolio to determine whether it satisfies our diversification requirements and whether the sources of risk are consistent with our most strongly held views. In the event the portfolio does not satisfy these requirements, we return to the second step and recalibrate both the weight assigned to the equilibrium returns and the relative weights given to individual views. We will discuss each of these six steps in more detail in the sections that follow. Also, in Appendix B, we consider some practical issues involved in distinguishing between tactical and strategic portfolio positioning.

III. Calculating Equilibrium Returns

This section describes in more detail how we find equilibrium returns. In the Black-Litterman framework, expected returns are viewed as a blend of equilibrium returns and an actual set of investor views. The equilibrium returns can be interpreted as the long-run returns provided by the global capital markets. Under this interpretation, the equilibrium returns represent the information that is available through the capital markets. Investor views, by contrast, correspond to the interpretation of information that is unique to the individual investor. Thus, expected returns are a blend of the information available through the capital markets and information unique to a specific investor. As the mixture of sources of information changes, expected returns will also change.

While equilibrium returns represent a useful neutral starting point, one immediate problem is that they are unobservable. However, with a few straightforward assumptions, we can easily infer equilibrium returns from other data that are readily available.

A natural way to proceed is as follows: We can start by assuming that if asset markets are in equilibrium, a representative investor would hold some proportion of the global capitalization-weighted portfolio. This assumption provides us with one readily observable piece of information. We can work from the observable capitalization weights to equilibrium

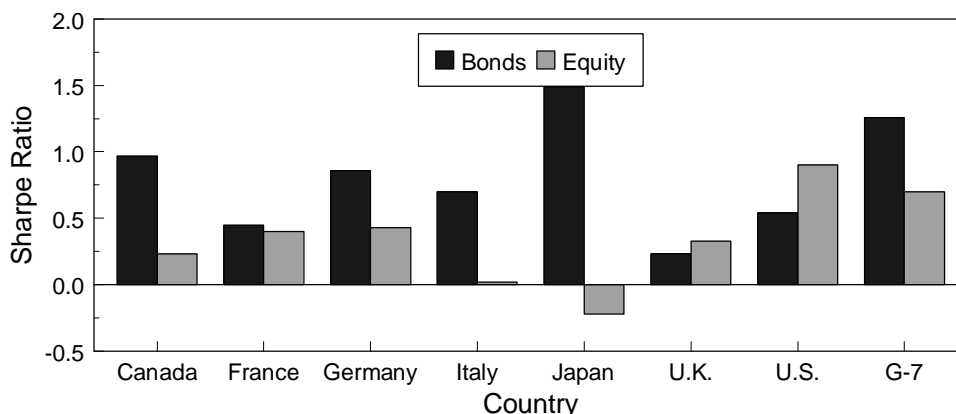
Exhibit 2 Benchmark Characteristics

Market	Weight	Duration
Australia	0.74	4.43
Austria	0.98	4.13
Belgium	2.32	4.80
Canada	3.73	5.17
Denmark	1.56	4.27
France	7.19	5.12
Germany	8.06	4.21
Italy	6.51	3.82
Japan	18.34	5.80
Netherlands	2.96	4.84
Spain	2.83	3.75
Sweden	1.58	3.87
U.K.	6.64	5.93
U.S.	36.58	5.02

returns by calculating this portfolio's volatility and then finding the asset returns that are consistent with a target Sharpe Ratio (the ratio of projected portfolio excess returns over cash to portfolio volatility).

Since our problem is a global fixed income asset allocation problem, our capitalization-weighted portfolio consists of the sector allocations for the Goldman Sachs government bond indexes. These indexes give the capitalization weights for the one- to three-year, three- to seven-year, seven- to 10- (or 11-) year, and greater-than-10- (or 11-) year maturity sectors in 13 government bond markets. For markets where index data are not available, we use a short-

Exhibit 3 Historical Sharpe Ratios January 1988 – October 1997



Hedged excess returns
G-7 Sharpe Ratios are for cap-weighted portfolios
Sharpe Ratio over all G-7 assets is 1.05

and a long-maturity bond. Representing the global fixed income markets in this way allows us to capture the effects of country allocation, overall Market Exposure, duration within each country, and steepness of the curve in each market. Exhibit 2 shows an illustration of the capitalization weights.

After finding the capitalization weights, we need to estimate a covariance matrix. At Goldman Sachs, we estimate covariance matrices using daily data. Volatilities and correlations are determined with weighted averages of daily squared returns. In this calculation, the weight depends upon the rebalancing horizon (see Litterman and Winkelmann, 1998, for more detail on our covariance matrix estimation procedures). For the purposes of our Global Fixed Income Asset Allocation portfolio, we assume a weight that is consistent with a two- to three-month rebalancing horizon. We combine the covariance matrix with the capitalization weights to find the volatility of the capitalization-weighted portfolio.

Once we have determined the cap-weighted portfolio's volatility, we calibrate asset returns to produce a Sharpe Ratio of 1.0. We calibrate the model in this way for two reasons: first, a Sharpe Ratio of 1.0 can be interpreted equivalently as a one-standard-deviation event. When returns are distributed normally, a one-standard-deviation event occurs approximately 66% of the time. Since one-standard-deviation events occur with such a high frequency, we can be reasonably sure that they will have been observed historically, with the corresponding implication that the equilibrium returns are not dependent on events that have few historical precedents.

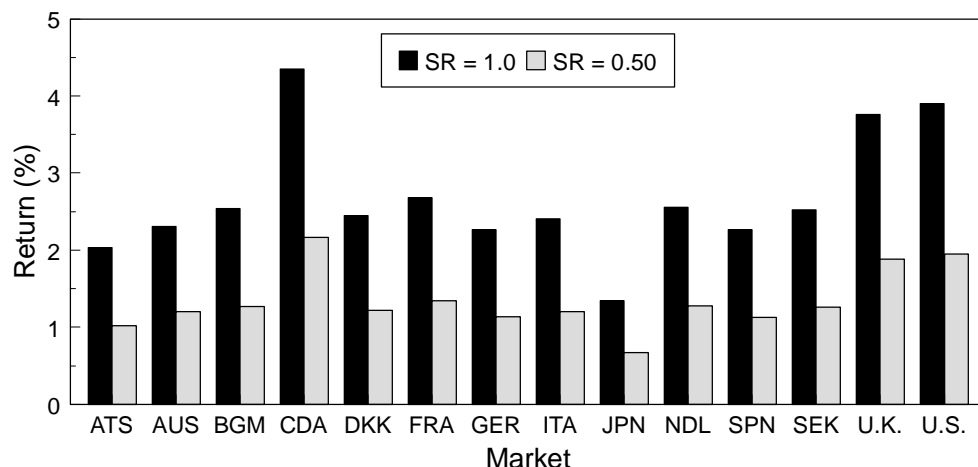
The second reason we calibrate the model to produce a Sharpe Ratio of 1.0 is because that figure is roughly consistent with historical experience. Exhibit 3 shows the historical Sharpe Ratios for 17 portfolios from January 1988 through October 1997. The portfolios consist of the Group of Seven equity and bond

markets individually, a capitalization-weighted portfolio of the G-7 equity markets, a capitalization-weighted portfolio of the G-7 bond markets, and a capitalization-weighted portfolio of the combined G-7 equity and bond markets. As the chart illustrates, historical Sharpe Ratios have ranged between -0.2 and 1.5, with an average over all 17 portfolios of 0.8. The last figure is roughly consistent with our assumption of an equilibrium Sharpe Ratio of 1.0.

Exhibit 4 shows two sets of equilibrium returns. The two sets of returns correspond to projected Sharpe Ratios of 1.0 and 0.50. *It is important to note that the equilibrium returns are not the result of an econometric forecasting exercise.* Instead, the equilibrium returns represent the idea of a set of long-run returns that are consistent with market clearing. One clear role for econometric forecasting models is to provide indications about short-term movements around the long-run equilibrium.

As previously discussed, the Black-Litterman approach combines equilibrium returns with an explicit set of views. Expected returns can, in some sense, be interpreted as a complicated weighted average of the neutral (or equilibrium) returns and an investor's views. In the next section, we will discuss the procedure that we follow for determining how much weight to put on the equilibrium returns and how to set the relative weights for each specific view.

Exhibit 4
Equilibrium Expected Returns



Returns are excess over local cash

IV. Setting the Weight-on-Views and Confidence Levels

In the next stage of the optimization process, we determine how much weight to put on the neutral reference point (equilibrium) returns relative to an explicit set of investor views. For the latter, we will take our own views as an example, using the interest rate and currency forecasts provided by the Goldman Sachs Economic Research Group in London (we will not discuss the forecasting procedure itself).

The blending of individual market views with the equilibrium returns is an important step in our methodology. There are two major reasons for following this procedure. The first reason is that by referring to the covariance matrix of historical returns (implicit in calculating the equilibrium), we ensure greater consistency across our own views. In this sense, the equilibrium returns help to serve as a macro constraint on our forecasts. The second reason is that, in damping extreme views relative to the implied equilibrium, the methodology produces more balanced portfolios than are typically produced from an unconstrained mean-variance optimization.

In determining the “weight-on-views” (WOV), we consider the projected excess return on the portfolio relative to the benchmark. In particular, we look at the portfolio's projected “information ratio” (IR). We define the IR as the forecast excess return on the portfolio over the benchmark, divided by its tracking error, where the latter is one standard deviation of excess returns relative to the benchmark (see the next section for a discussion of tracking error risk).

As such, the IR can be considered as the forecast excess returns on the portfolio expressed in terms of standard deviations, based on the covariance matrix of historical returns. We can interpret the IR as a measure of the information contained in the investor's views. In our process, we select the WOVS to produce an IR of no greater than 2.0, on grounds that returns of

greater than two standard deviations in excess of the benchmark would be statistically unlikely to occur.

From this we can see how the WOV serves, in effect, partly as a constraint at the macro level on our own views. An optimization based on the unconstrained views often will lead to a projected IR significantly in excess of 2.0. This would usually be the result of one or more extreme views relative to the others. Controlling the WOV usually leads to more-balanced portfolios. If left unconstrained, the portfolio will overweight the views with the largest risk-adjusted returns. By combining our forecasts with the implied equilibrium returns, we dampen the influence of extreme views. Typically, in order to keep the IR at 2.0 or lower, we have attached a WOV of 50–70%.

Having determined the WOV and the set of expected returns to be used in the optimization, we then consider constraints at the micro level on the individual market views, which we refer to as “confidence weightings.” We use the model to determine the conditional probability of the projected excess returns on each asset and currency. In other words, using the covariance matrix of historical returns for N assets, we determine the probability of observing the return on the N th asset, conditional on the return forecast for the $N-1$ assets. In the same way that we impose the constraint at the macro level that we would not want the projected IR of the portfolio to be in excess of 2.0, we attempt also to ensure that the individual market views are not greater-than-two-standard-deviation events.

We constrain the views on projected returns on the individual exposures by allocating confidence weightings. We attach either a Low, Medium, or High confidence weighting to each of the separate views. If, for example, the forecast excess return on Australian bonds is a three-standard-deviation event relative to the other views, then we would give a Low confidence weighting to the Australian bond forecast. This has the effect of using less of the information in the actual view, and more of the information in the equilibrium returns.

Why would we dampen the influence of a particular view? The principal reason is because our statistical analysis is telling us that this view is unlikely to prove correct, based on the historical behavior of

asset returns. Thus, putting less emphasis on a particular view lets us incorporate the information that gave rise to that view, and at the same time produce a portfolio whose performance is not overly dependent on an asset return that is statistically unlikely to occur.

An alternative to using the WOV and confidence weightings to constrain the projected excess returns would be to revisit the currency and interest rate forecasts. Indeed, this occurs in the process of producing our own optimal portfolio; the model is used, to a certain extent, as part of an iterative process in helping to build the forecasts of the Economic Research Group. The intention is not to suppress the deliberate expression of strongly held views, however. After all, as we will discuss in Section VI, it is the identification of these key views that we use to make sure that the portfolio is taking risk where we are comfortable. Rather, we use WOV and confidence intervals to ensure greater statistical consistency across the bond and currency forecasts.

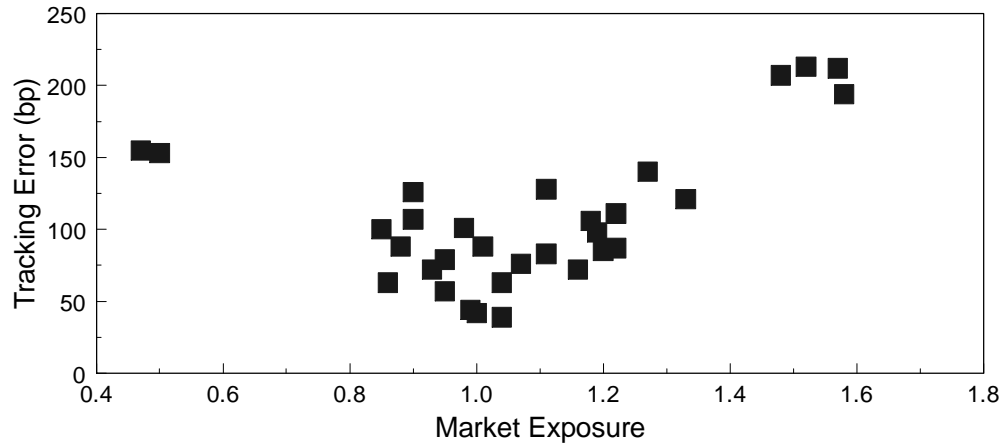
V. Setting Target Risk Levels

After finding expected returns, we then set target risk levels. Since we construct our optimal portfolio relative to a benchmark, we consider all of our risk measures as risks relative to the benchmark. The two risks that we care most about are the tracking error and the Market Exposure. (See Litterman and Winkelmann, 1996, for a detailed discussion of both of these measures.)

Tracking error measures the volatility of the portfolio and benchmark performance differences. Since tracking error is a standard deviation, we can give it a natural probabilistic interpretation. For example, a tracking error of 100 basis points tells us that (roughly) 66% of the time, the actual performance of the portfolio should be within 100 bp of the benchmark's performance, irrespective of whether the benchmark return is positive or negative.

To measure the directional bias of the portfolio, we look at the Market Exposure. A portfolio's Market Exposure tells us how responsive the portfolio's performance is for given levels of benchmark performance, with the caveat that all other influences are assumed to be constant. For example, a Market Ex-

Exhibit 5 Survey Risk Characteristics



Source: GS Global Fixed Income Survey, August 1997

posure of 1.2 tells us that if the benchmark has a 1.0% return, then the portfolio will have a 1.2% return, other variables held constant. When the Market Exposure is greater than 1.0 the portfolio has a bullish bias, and when the Market Exposure is less than 1.0 the portfolio has a bearish bias.³

³ Practitioners sometimes look at other types of risk measures. For example, portfolio managers might look at the Value-at-Risk (VaR) of a portfolios. The VaR shows the size of the potential loss for specific probability choice. For example, suppose that we set the probability level at 1.7% and we find that our portfolio will lose 200 bp or more relative to its benchmark at this probability level. In this case, our VaR is 200 bp.

When assets have a symmetric payoff, VaR has a natural interpretation in terms of tracking error. Continuing with the example, when asset returns follow a Normal distribution, the probability level of 1.7% corresponds to two standard deviations. Thus, the portfolio has a VaR that corresponds to a two-standard-deviation event. Since tracking error is expressed in terms of one standard deviation, we find that the VaR of 200 bp corresponds to a tracking error of 100 bp. We can do the same kind of calculation for other choices of probability and VaR whenever the assets have symmetric payoffs.

A second popular risk measure is called “downside risk.” The basic idea behind downside risk is to find the expected value of portfolio returns less than some critical level. For example, if the critical value is set at zero, we can calculate both the expected value of all portfolio returns less than zero and the probability that the portfolio will have negative portfolio returns. As with VaR, when the distribution of asset returns is symmetric, downside risk has equivalent interpretations in terms of tracking error. Also in common with VaR, when payoffs are not symmetric (i.e., when options are present), downside risk does not have an equivalent interpretation in terms of standard deviations.

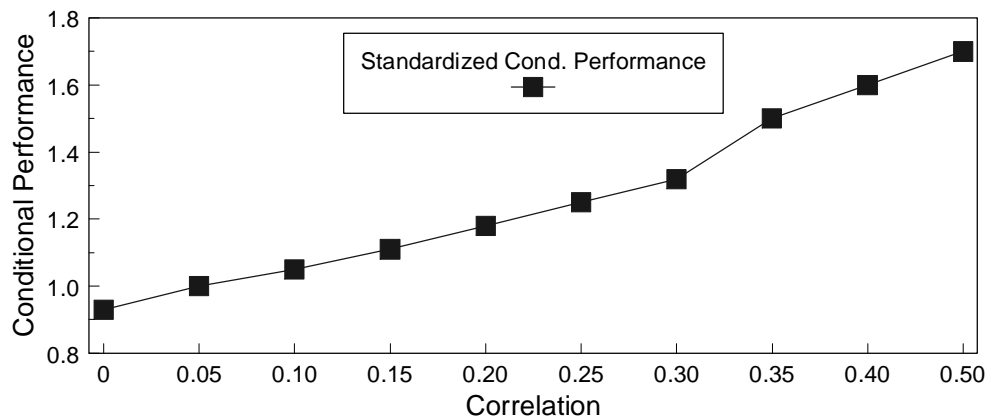
How do we set our two risk parameters? Let’s look first at the Market Exposure. When our views on market direction are relatively neutral, we constrain the optimal portfolio to have a Market Exposure of 1.0. Imposing this constraint means that our optimal portfolios reflect our curve views and country views but are neutral in any overall directional sense. By contrast, when we have

strong directional views, we let the optimizer find the Market Exposure number. For example, when our views are strongly bullish, the optimizer will select a portfolio whose Market Exposure is greater than one.

Our tracking error figures have been set to produce a portfolio with a tracking error of roughly 100 bp. While we also analyze portfolios with 200 bp of tracking error, our focus is on the lower risk portfolio. The reason for this is that in our experience, the higher risk optimal portfolios have tended to concentrate risk in a smaller number of positions. In other words, we have found that a tracking error of 100 bp has been more appropriate to give adequate expression to our views.

We also consider actual industry practice in choosing our tracking error figures. Recent experience shows that a tracking error of 100 bp has been consistent with the practice of global fixed income fund managers. Exhibit 5 illustrates this point. The chart plots the tracking error and Market Exposure for participants in a survey of global fixed income portfolio managers. As shown in the chart, our tracking error of 100 bp is approximately the median tracking error for the survey participants.

Exhibit 6
Correlation and Performance



Constant one-standard-deviation rally in Italian 10-year
Constant one-standard-deviation sell-off in German 10-year

VI. Optimization and Risk Decomposition

Once our target risk levels have been set, we find portfolios that maximize expected return. We then take the optimal portfolio weights and identify the major sources of risk (see Litterman, 1996), on the premise that the major sources of risk should also be consistent with the most strongly held views. In our process, we typically look for optimal portfolio weights that satisfy three criteria. First (as discussed above), we want to make sure that the risk exposures in the optimal portfolio correspond to positions where we want to take risk. Second, we typically impose an upper bound of 20% of the risk from any particular position: That is, any currency or bond market position should contribute no more than 20% of the tracking error. Finally, we try to ensure that the tracking error is evenly balanced between our bond and currency positions.

We look for a balanced distribution of risk for one principal reason: Spreading the risk across several positions means that we can exploit the power of diversification. When the risk in a multi-asset portfolio is concentrated in one or two positions, the implication is that all other portfolio management decisions are not likely to have a significant impact on the portfolio's performance. (In this case, the investment manager could be well advised to consider implementing the optimal portfolio through options strategies, thereby limiting the potential losses). In practice, of course, we are likely to have more than

one or two strongly held views that we would like to see expressed in the portfolio.

Why would the optimal portfolio concentrate risk in one particular position? Simply put, the optimal portfolio will allocate its risk to the most aggressively held views. Thus, when one or two views are much more aggressive than the remaining views, the optimal portfolio will concentrate its risk in those positions.

It is important to keep in mind that views can be aggressive in two senses: relative to their own history (as captured by the volatility) and relative to all other views. In the first case, the portfolio concentrates its risk in positions that are driven by views that are not likely to prove correct in the context of their own history. In the second case, the portfolio concentrates its risk in positions based on views that are not likely to prove correct if all other views are borne out, assuming that the structure of the covariance matrix remains unchanged. An implication of the second point is that views also have implications for correlations. We illustrate this point in Exhibit 6.

Exhibit 6 plots the correlation between German and Italian 10-year bond returns against the conditional Z-Score for Italian bonds.⁴ The assumption underlying the chart is that German bonds will sell off and Italian bonds will rally. The size of the respective sell-off and rally is set equal to one standard deviation. As the chart illustrates, given a correlation

⁴ We calculate the conditional Z-scores as follows: Suppose that we have N assets and (for illustrative purposes) we have views on all N . For each of the N assets, we can compare the actual view with the view that is most likely given the other $N-1$ views. We find the most likely views by using the covariance matrix to project the N th asset on the $N-1$ assets and then substitute in the views on the $N-1$ assets to find the conditional return on the N th. We then find the difference between the actual view and the most likely view. To make comparisons across assets, we take the ratio of the difference to the conditional volatility.

between German and Italian bond returns of 0.50, if a one-standard-deviation sell-off in German bond returns occurs, then a one-standard-deviation rally in Italian bonds is a very unlikely event (as measured by the conditional Z-score). Indeed, the correlation between Italian and German bond returns must be -0.05 in order for the conditional Z-score to be zero.

One natural way to correct for the implicit view on the correlation matrix is to dampen the views — in other words, take less-extreme views. In fact, this is exactly the role of the weight-on-views and confidence parameters. Decreasing the weight-on-views or assigning a Low confidence to a particular view has the effect of pulling the views back toward the equilibrium views. From a correlation perspective, decreasing the weight-on-views (or assigning a Low confidence) has the effect of pulling the correlation that makes the expected returns statistically likely more consistent with the actual covariance matrix. The net result is that the ultimate portfolio provides a better diversification of the portfolio's risk — i.e., less-aggressive risk positions are taken.

The risk distribution has important implications for comparing the optimal portfolio with alternative portfolio weights. For example, it is often heard that portfolios with the same risk level as an optimal portfolio, but interior to the efficient frontier, may have desirable properties. These properties can be summarized as lower exposure to adverse moves in a particular set of asset prices.⁵

Let's look at these desirable properties in more detail. The idea that a point that is interior to the frontier is less exposed to potentially adverse moves in some asset prices can be interpreted as saying that risk is better diversified than in the optimal portfolio. As discussed above, the risk distribution in the optimal portfolio depends on the relative strength of the investment views: The most risk is taken where views are the most aggressive. Thus, the desirable properties of suboptimal portfolios are a result of the optimal portfolio seeking the maximum leverage

from very aggressive investment views. Using the equilibrium in a systematic way to dampen the impact of aggressive views has the effect of providing a portfolio that is well diversified in a risk sense and expressing views that are internally consistent.

VII. How Have We Done?

In this section, we consider the performance of our portfolio strategy recommendations. Clearly, the performance relative to the benchmark will in large part reflect the accuracy of our views. At the same time, however, it is at least as important to consider the impact of the various control mechanisms used in the Black-Litterman model. That is because the procedures that we have outlined above are designed to help in the process of risk management. Independently of the accuracy of our views, it is of interest to know whether these control procedures have worked effectively.

We have been publishing regular updates to our portfolio recommendations since the beginning of 1995. We calculate the relative performance of the portfolio on a monthly basis by adjusting the over- or underweighting in each market for the recommended duration relative to the index. Similarly, we calculate the relative foreign exchange performance as the product of the actual performance of the currency and the relative overweighting (or underweighting) in that currency.

For two reasons, these results should be regarded as only an approximate measure of portfolio performance. First, the calculation of returns is based on the portfolios at the beginning of each month, though the actual portfolio revisions were often made at different times. Second, no allowance is made for short-term trading views that have been expressed in other regular commentaries.

⁵ The observation that points interior to the efficient frontier have potentially desirable properties is sometimes used to justify so-called "scenario-dependent" optimization.

Last year, our base case portfolio underperformed the benchmark by 26 bp (through end-December 1997). In 1996, the portfolio outperformed the benchmark by 83 bp, and in 1995, the portfolio's return exceeded that of the benchmark by 103 bp. The average over the three-year period has been 52 bp. Given that we aim for a projected tracking error of 100 bp, the ex-post information ratio (i.e., performance relative to risk) has been 0.52.

Monthly Returns: Exhibit 7 shows the portfolio's monthly excess returns over the benchmark. The chart further decomposes total relative performance into its fixed income and currency components. Exhibit 8 shows a history of standardized monthly performance for close to three years. As we discuss below, this history provides some evidence of the efficacy of our risk control procedures

In spite of a bearish outlook for global bond markets at the outset of 1995, the portfolio outperformed the benchmark over the year as a whole, for three reasons. The first reason was that, although bearish views predominated in the major markets at the outset of 1995, we switched to a bullish near-term outlook for both the Japanese and European bond markets and extended the duration of the portfolio to longer than that of the index for April–August. The second reason is that we remained appropriately underweighted in the United States relative to Europe and were mostly close to index weight in Japan. The third reason is that we adopted a bullish outlook on the U.S. dollar

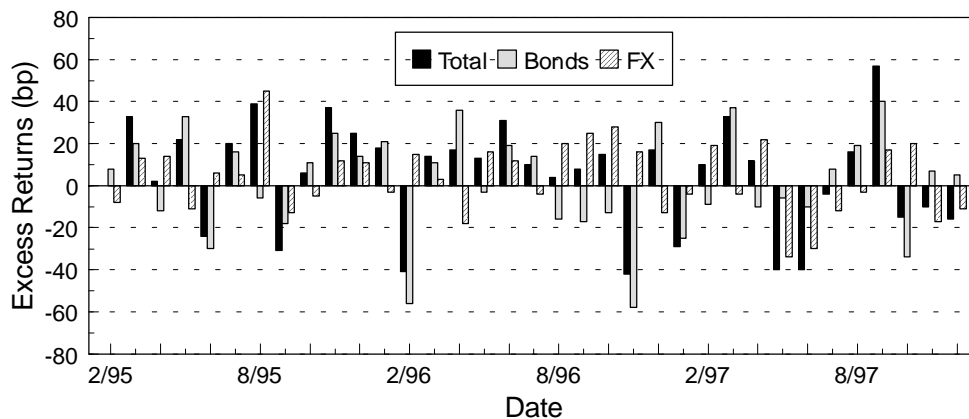
against both the yen and European currencies from February, and the portfolio then remained long the dollar, mostly in terms of European currencies, for the rest of the year.

In 1996, the dominant source of outperformance relative to the benchmark was our currency positioning. Cumulatively, our bullish stance on the U.S. dollar for the bulk of 1996 provided 106 bp of the portfolio's relative outperformance. In only four of the 12 months were our foreign exchange positions at odds with actual events in the currency market. The figures also indicate that the portfolio's bond positioning acted as a drag on performance: The bond component of the portfolio underperformed the fixed income component of the benchmark by roughly 23 bp. This is because the portfolio was not positioned for the global bond rally in the latter part of the year. In the period up until the end of July, the fixed income portion of the portfolio outperformed the benchmark by 42 bp on the benchmark, while from August through December, relative performance was -65 bp.

Over the first eight months of 1997, the portfolio's performance was dominated by the incorrect bearish stance that we adopted, for the most part, on overall market direction. The duration of the portfolio was short relative to the benchmark for the four months ending in July. In August, we set a constraint that the bond exposures should have a Market Exposure (or Beta) coefficient of 1.0, and we retained this until December, when we lengthened the Market

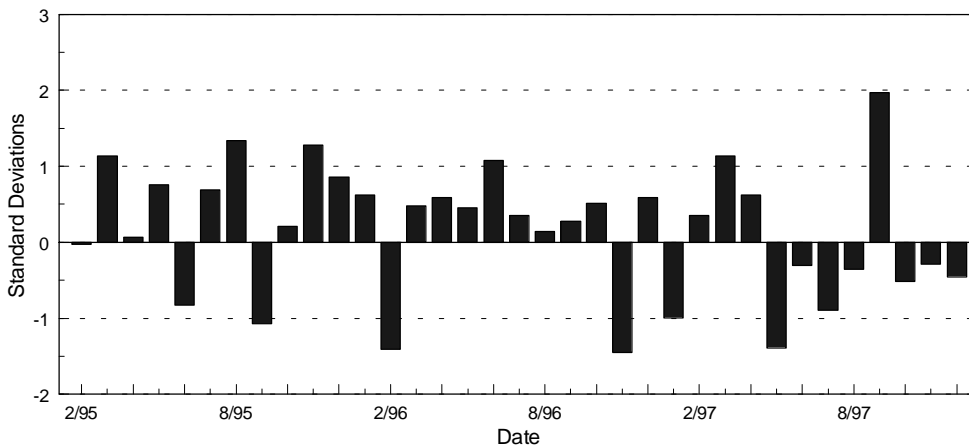
Exposure of the hedged bond holdings to 1.08. For the year as a whole, the bond component of the portfolio outperformed the benchmark by 12 bp, while the currency component underperformed by 39 bp.

Exhibit 7
Monthly Performance
February 1995 – December 1997



GS Global Fixed Income Asset Allocation
Low Risk Portfolio/Three-Month Horizon
Predicted Tracking Error = 100 bp

Exhibit 8
Standardized Monthly Performance



GS Global Fixed Income Asset Allocation
 Low Risk Portfolio/Three-Month Horizon
 Predicted Tracking Error = 100 bp

Risk Analysis: Clearly, even after taking into account the control mechanisms detailed above and the approach that we take to blending expected excess returns with the equilibrium, the ex-post performance relative to the benchmark will depend heavily on the accuracy of our views. We would make two observations, however. First, we should compare the performance with an unconstrained optimization. This would be one way of measuring the efficacy of the model. Second, it is at least as important to consider whether the risk control parameters applied in the model work efficiently. We can explore the latter issue by referring to Exhibit 8.

Recall that Exhibit 8 shows the entire history of the portfolio's performance on a standardized basis. We calculated the figures in the chart by taking the actual monthly performance and dividing by the predicted monthly tracking error. Standardized performance provides a method for evaluating the risk control elements of the portfolio construction process; our objective is to eliminate very large fluctuations in monthly performance.

To illustrate, for an annualized tracking error of 100 bp, we would anticipate a monthly tracking error of roughly 29 bp. That is, if our risk measurement and control procedures are accurate, then 66% of the time we would expect the actual monthly performance of the portfolio to be within 29 bp of the performance of the benchmark. Similarly, 34% of the time we would expect monthly relative performance to exceed 29 bp.

As Exhibit 8 shows, in 22 of the 33 months (the portfolio started in February 1995) we had actual performance within 29 bp of the benchmark. By contrast, in only one of the 33 months was our actual performance outside the two-standard-deviation band. On balance, then, we can conclude that our portfolio process is succeeding in its approach to risk control.

VIII. Conclusions

This paper has focused on our use of the Black-Litterman model to develop portfolios that best reflect our investment views. We have discussed the procedures that we use to set all of the Black-Litterman parameters and shown how the choice of these parameters affects the structure of risk-controlled optimal portfolios. Finally, we have discussed the performance of our particular portfolios. Two features of our actual performance stand out: First, we have outperformed the benchmark over the three-year period. Second, and in our view of equal importance, our portfolio's performance has been enhanced by using the risk-control characteristics inherent in the Black-Litterman model.

We have not discussed procedures for actual implementation of the optimal portfolios. The output from our exercise is allocations, sector exposures (e.g., the one- to three-year sector in Italy), and durations in each of the markets in our optimization problem. Investors clearly cannot purchase the individual sectors and must implement the portfolio through the purchase of actual securities. The natural solution to this problem is to use a fitted curve to identify cheap securities in each of the relevant sectors. Investors would choose weights in the individual securities to match the allocations and durations produced by the optimization process.

The discussion in this paper has been oriented toward selecting an optimal asset allocation for a portfolio whose performance is measured against a global bond index. However, the approach can be extended to other applications. For instance, one popular investment strategy is to allow international investing versus a domestic benchmark, an investment style that has been labeled “opportunistic international investing.” For example, French investors may purchase Japanese government bonds, even though their benchmark is a French government bond benchmark. Alternatively, investors with a global bond benchmark may take exposure in emerging markets on an opportunistic basis (we have included Poland and South Africa in our recommended portfolio since January 1998). Two key aspects of our framework can provide insight into this process:

- First, since investors are taking risk in international positions because they expect international markets to outperform domestic markets, the risk control dimensions of our framework provide a means for ensuring that risk is well balanced.
- Second, the identification of aggressive views (both absolute and relative) helps investors determine the probability of international outperformance contingent on their views of domestic performance.

Active managers add value by using information to take risk positions that deviate from their benchmarks. We have shown through a practical example how the Black-Litterman model can be used in this process. The model has provided an effective means to combine a particular set of investor views (i.e., investor information) with the returns given by a capital market equilibrium (i.e., the market’s information). As the effects of risk positions become more pronounced — e.g., through the influence of EMU — the ability to efficiently use information to produce risk-controlled portfolios becomes more acute. Our framework provides one means to this end. ■

Appendix A: Interpreting the Published Portfolios

In Exhibit 9, we provide an example of a recently published portfolio taken from *Global Fixed Income Asset Allocation*. Drawing on the discussion in the various sections above, we will illustrate how to interpret the portfolio recommendations and the risk analysis.

In the “Bonds” section of the table, the first column shows the market-capitalization weights of the benchmark, using Goldman Sachs government bond indexes. Alongside this, in the second column, we show the recommended bond weightings derived from our optimization.

In the “Duration” section, the first column is the duration of the benchmark sectors. The second column shows the recommended duration. At the foot of the table, we can see both the aggregate duration of the benchmark and the aggregate duration of the

portfolio. We would caution against using duration as a measure of the portfolio’s risk. Our preferred measure is the Market Exposure of the portfolio, as discussed in Section V of this report but not shown here.

The “Forex Overlay” column shows the additional units of foreign currency to be purchased or sold, which — combined with the bond exposures — will produce the total currency exposure recommended in the optimal portfolio. The total currency position is shown in the “Net Allocation” column. The short or long currency exposures of the portfolio can be found by comparing this column with the benchmark in the first column.

Finally, the “Risk” section shows the marginal contribution to the tracking error risk of the portfolio arising from the individual bond and currency exposures relative to the benchmark. (For a detailed discussion of the contribution to risk, see Litterman, 1996.)

Exhibit 9 Recommended Bond and Currency Weightings

GS Global Bond Portfolio - One Month Horizon (Dollar Based)

	Bonds (%)		Duration (years)		Forex Overlay (%)	Net Allocation (%)	Risk*(%)	
	Index	GS Recommended	Index	GS Recommended		GS Recommended	Bonds	Currency
US	35.9	11	5.3	10.5	23	34	27.1	-
Canada	3.4	0	5.6	0.0	3	3	3.6	0.0
Australia	0.7	2	4.3	6.5	-1	1	2.3	0.0
Japan	17.7	12	5.8	9.1	6	18	1.0	0.0
Euroland	32.5	39	4.7	5.0	-2	37	16.7	12.0
UK	6.9	4	6.3	10.2	0	4	0.2	0.1
Denmark	1.5	18	4.9	3.6	-16	2	17.0	0.0
Sweden	1.4	8	4.3	4.1	-7	1	12.1	0.0
Poland	-	5	-	1.9	-5	0	7.4	0.0
South Africa	-	1	-	5.0	-1	0	0.5	0.0
TOTAL	100	100	5.3	5.9	-	100	87.9	12.1

*Percentage contributions to tracking error risk of the portfolio. Errors due to rounding.

Source: *Global Fixed Income Asset Allocation*, Goldman Sachs International, June 1, 1998

Appendix B: Practical Issues in Developing a Process

In the past three years, we have attempted different approaches to distinguish between strategic and tactical portfolio positioning. We constructed a strategic portfolio to draw a distinction between the bond holdings that reflect our longer-term views (based on our interest rate and currency forecasts) and portfolio positions that represent shorter-term tactical views. Strategic portfolio positions should, in principle, exhibit little change on a month-to-month basis (until, of course, our fundamental views change). By contrast, tactical trading positions can change frequently.

Our procedure was relatively straightforward. Using our Low Risk (100 bp of tracking error) portfolio derived from our three-month horizon views, we identified combinations of deviations from the benchmark portfolio that gave rise to an information ratio (projected performance divided by portfolio risk) that was about 80% of the optimal portfolio's information ratio. Among the set of portfolios that satisfied this condition, we then selected the portfolio that satisfied two additional conditions. First, we looked to minimize the number of deviations from the benchmark, and second, we looked for the portfolio whose tracking error (risk relative to the benchmark) was closest to the risk of the Low Risk optimal portfolio. We then assumed that all other bond positions in the portfolio were held at index weight and at index duration.

Using this framework, we then introduced tactical trades to overlay the strategic portfolio. The idea was to retain a portfolio that best represented the key elements of our three-month horizon interest rate and currency forecasts but that at the same time would allow for flexibility to express near-term trading views. In this way, we intended to more closely replicate actual practice in the fund management industry.

However, we encountered a number of difficulties in our attempt to distinguish between tactical and strategic portfolio positioning. An obvious problem arose when our near-term trading views differed from the views expressed in our three-month horizon forecasts. For example, although the

forecasts might look for a rise in U.S. interest rates over a three-month horizon, that should not preclude expressing a trading view that is bullish for the very near term.

A second related difficulty was one of presentation. Readers failed to understand that portfolio allocations are not always intuitive based on any given set of expected returns. The problem is compounded when we distinguish between trading views and three-month horizon forecasts. Trading views are developed in isolation and take no account of the relative volatility of returns across asset classes or the correlation structure. Thus, it would be perfectly possible — and not necessarily inconsistent — to express a bullish view on an individual market and then to be underweighted in the portfolio. Nevertheless, such positioning served to generate confusion.

Partly intending to resolve the difficulties that arose in attempting to distinguish tactical and strategic portfolios, we decided to explicitly develop both portfolios that would be more consistent with our trading views and portfolios based on our strategic views. In the case of the former, we developed a methodology that would ensure consistency with our trading views. In addition, we continued to publish a strategic portfolio using our three-month horizon forecasts, though not employing the approach to limiting the portfolio to a small number of exposures, as set forth above.

In the first step of developing our portfolios, we identified the small number of trading views in the bond markets that we held most strongly. We then generated a set of one-month horizon excess returns on two-year and 10-year bonds consistent with these views. Using the Black-Litterman model to produce an optimal bond portfolio, we then continued with the procedure described in Section IV to use weight-on-views and confidence levels.

In addition, as a further control procedure, we usually looked for the aggregate currency positions to contribute between 40% and 60% of the overall risk (measured as tracking error), with bond positions contributing the remaining risk. Moreover, we typically looked for any individual position to contribute no more than 20% of the aggregate risk. Finally, again to ensure consistency with trading views, we

used the risk analysis to ensure that we were taking risk where we had the strongest views.

At the same time as making this distinction between tactical and strategic views, we also separated the bond and currency allocation decisions. The motivation was twofold. First, it allowed for better tractability; in other words, it avoided problems that had arisen where it was not possible to be sure whether it was the bond view or the currency view that was driving the allocation. Second, it more closely resembled industry practice.

In making the distinction between the bond and currency allocation decisions, we first optimized the bond holdings relative to a fully hedged benchmark. We next considered the currency exposure separately as an overlay. We assume that currency positions are taken through outright positions in the forward market.

Bibliography

Black, Fischer, and Robert Litterman, *Asset Allocation: Combining Investor Views With Market Equilibrium*, Goldman, Sachs & Co., Fixed Income Research, September 1990.

Black, Fischer, and Robert Litterman, *Global Asset Allocation With Equities, Bonds, and Currencies*, Goldman, Sachs & Co., Fixed Income Research, October 1991.

Litterman, Robert, *Hot Spots™ and Hedges*, Goldman, Sachs & Co., Risk Management Series, October 1996

Litterman, Robert, and Kurt Winkelmann, *Managing Market Exposure*, Goldman, Sachs & Co., Risk Management Series, January 1996.

Litterman, Robert, and Kurt Winkelmann, *Estimating Covariance Matrices*, Goldman, Sachs & Co., Risk Management Series, January 1998.