ASSET ALLOCATION: COMBINING INVESTOR VIEWS WITH MARKET EQUILIBRIUM

FISCHER BLACK AND ROBERT LITTERMAN

FISCHER BLACK is a Partner in Goldman, Sachs & Co.

ROBERT LITTERMAN is a Vice President in the Fixed-Income Research Department at Goldman, Sachs & Co.

Investors create global bond portfolios for a variety of reasons: to diversify interest rate risk, to manage yield, to control exposure to foreign currencies, and to enlarge the universe of possible trading opportunities. This article describes a new approach to international asset allocation of fixed-income securities.

We show how to construct portfolios by choosing the optimal weights to invest in assets in each country and the optimal degree of hedging of currency exposure, given the investor's views for interest rates and exchange rates. While our approach brings several new features to the traditional asset allocation problem, its most innovative contribution is to allow investors to compare their outlook for currencies and interest rates with expected returns generated by an International Capital Asset Pricing Model (ICAPM) equilibrium.

The simple idea that expected returns ought to be consistent with market equilibrium, except to the extent that the investor explicitly states otherwise, turns out to be of critical importance in making practical use of the model. In particular, it allows investors to specify views in a much more flexible way than otherwise would be permitted.

For example, rather than requiring investors to specify views about absolute returns on every asset, our approach allows investors to specify as many or as few views as they wish — views with different degrees of confidence and views about relative returns on different assets. This use of the expected returns associated with asset market equilibrium as a reference point for investors is a unique feature of the model. Much of our article focuses on this aspect of our approach.[1]

Another advantage to our approach is that it jointly determines the optimal allocations of bonds into different countries and the most desirable currency hedges. It is fashionable today to recognize that in multi-currency bond portfolios there is no reason to have the foreign currency exposure equal to the bond allocation.[2] The usual implication that is drawn — that the two decisions can be made sequentially — is just as wrong, however. The weights on bonds and currencies certainly need not be identical, but because returns on currencies and bonds are generally highly correlated, the optimal weights on both should be determined in a joint optimization.

Finally, our approach includes many choices for its user, not only in expressing views but also in estimating the risks from historical data and in stating objectives and constraints. These choices give the approach a flexibility in addressing the different types of objectives and constraints that different international fixed-income investors face.

One innovative use of the model is to take an investor's actual portfolio and derive the "implied views," that is, the set of expected returns that would make the investor's portfolio optimal. We will come back to these points, but let us first focus on the most important innovation in our approach: the use of the ICAPM equilibrium.

USING EQUILIBRIUM EXPECTED RETURNS IN ASSET ALLOCATION

We make the standard assumption that the objective of the portfolio manager is to maximize total return, measured in a specific currency and possibly relative to a benchmark portfolio, for any given level of risk. This mean-variance approach has proved to be a valuable framework for addressing investment decisions.

Our approach goes beyond the standard mean-variance analysis, however, in that we incorporate a key insight provided by the ICAPM: In equilibrium, prices of bonds and exchange rates will adjust until investors worldwide are comfortable holding the outstanding supply of assets with a given degree of currency hedging. This equilibrium[3] provides a useful reference point that allows investors to generate balanced portfolios reflecting both their views about currencies and interest rates and the equilibrating forces of supply and demand in international asset markets.

We feel this innovative use of the ICAPM equilibrium not only is theoretically sound but also provides an important practical advantage in solving the asset allocation problem. As we illustrate below, it significantly ameliorates the usual tendency of mean-variance models to map seemingly reasonable views into what appear to be extremely unbalanced portfolios.

We begin by considering the role that a quantitative asset allocation model ought to play in helping the portfolio manager make the asset allocation decisions. The first step in the portfolio allocation process is for the manager to formulate views about the relative attractiveness of different assets. These must then be translated into a portfolio that will perform well, in some sense, if the views are correct. The asset allocation model should play a key role at this stage by helping managers map their views into an optimal portfolio.

The Standard Mean-Variance Approach

There are many possible aspects of the asset allocation problem that we can address by formulating the problem in different ways. For example, by making specific assumptions about investor preferences, by considering a dynamic allocation strategy, or by including various constraints, transactions costs, and so on, we can bring different aspects of the problem into focus. The optimization problem in particular becomes very difficult in a dynamic context when there are transactions costs.

In this article, however, we concentrate on a relatively simple version of the problem, focusing attention on a key aspect of the asset allocation decision that has limited the practical application of quantitative models. That aspect is the role that investor views play in formulating the basic trade-off between risk and expected return.

The mean-variance framework that we adopt assumes that investors have views that they want to incorporate into expected returns for different assets, and that they want to use historical returns to estimate their risks. It is our impression from talking to many portfolio managers that this is their usual approach. Given these inputs, our model then finds the frontier of optimal portfolios at that time.

A portfolio lies on the optimal portfolio frontier if, for a given level of risk, the weights in different asset classes generate the maximal expected return among all feasible portfolios. An investor who is risk-averse should choose a portfolio on the frontier, because any portfolio not on the frontier is dominated by one on the frontier that offers at least as much expected return with less risk. In choosing one of the portfolios along the frontier, the investor should pick the one that best balances the desire for additional return against aversion toward additional risk.

In theory, an asset allocation model should be an indispensable tool for the portfolio manager. In practice, portfolio managers have found asset allocation models difficult to use and the portfolios they generate often unacceptable. To explain the advantages offered by our model, we need to consider the difficulties that have historically hindered use of these types of models.

The main problem that portfolio managers have found when using mean-variance models is that the mapping between views and optimal portfolios can be very sensitive to small changes in the views.[4] This is particularly true when the portfolio includes a balanced weighting for various assets whose returns are highly correlated, as is likely to be the case in a multi-currency bond portfolio.

To illustrate this problem, let us consider the portfolio in Table 1, which lies on the optimal portfolio frontier when the expected returns are as shown.

We construct the expected returns in such a way that the portfolio includes weights on bonds in different countries that match the proportion of the total outstanding market value of government debt in those countries.[5] We also construct the expected returns so that the degree of currency hedging in the optimal portfolio is 80%.

These expected returns are, in fact, equilibrium values implied by the ICAPM when people in different countries have the same average risk tolerance as in Black [1989]. We choose these values for this example because the optimal portfolio is then balanced in the sense that it matches the market portfolio weights.

Now consider what happens when we modify these expected returns just a bit. Unfortunately, it turns out that a nicely balanced optimal portfolio such as the one displayed here is, in a sense, rather precariously balanced. Small changes in the expected returns can lead to large reallocation of weights to different assets.

Table 2 shows the optimal portfolio when we raise the annualized expected returns on German unhedged bonds by just ten basis points (bp).[6] We posit that this increase in expected returns is generated by an appreciation of the deutsche mark relative to the dollar, so we assume that the returns on hedged German bonds do not change.

Given this very slight modification in expectations, the model's optimal portfolio jumps to an extremely long position in deutsche marks, financed by U.S. dollars that have been hedged into French francs. How can this extreme sensitivity arise?

Consider a simple two-asset example. Suppose we are a U.S. dollar-based investor considering optimal weights on German bonds versus French bonds. The returns on these assets have had a historical correlation of about 0.9 from the point of view of a U.S. dollar investor. Unless we specify a very carefully balanced view of future returns (which does not necessarily mean equal expected returns) in these two assets, the model will find a combination of weights — long one of these assets and short the other — that together can be expected (based on our specified views and the historical covariances) to generate higher returns and lower risks than a balanced portfolio with similar weights in each country.

Table 3 shows, for this simple two-asset problem, the optimal portfolio weights on German bonds and

TABLE 1 ■ Equilibrium Expected Returns and the Optimal Portfolio (percent)

	U.S.	Japan	Germany	France	U.K.	Canada	Australia
Equilibrium Expected Returns on Bonds (August 20, 1990)							
Unhedged Bonds	9.77	9.81	9.37	9.18	9.93	9.60	8.30
Bonds Hedged into $		8.98	9.21	9.03	9.65	9.55	8.19
Optimal Portfolio Weights							
Cash	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Bonds	53.6	18.3	9.4	7.1	6.3	4.4	1.0
Forward Contracts		14.6	7.5	5.6	5.1	3.5	0.8

TABLE 2 ■ Slightly Modified Expected Returns and the New Optimal Portfolio (percent)

	U.S.	Japan	Germany	France	U.K.	Canada	Australia
Slightly Modified Expected Returns on Bonds (August 20, 1990)							
Unhedged Bonds	9.77	9.81	9.47	9.18	9.93	9.60	8.30
Bonds Hedged into $		8.98	9.21	9.03	9.65	9.55	8.19
New Optimal Portfolio Weights							
Cash	-56.8	0.0	56.9	0.0	0.0	0.0	0.1
Bonds	57.7	17.0	8.7	9.0	5.5	2.0	0.0
Forward Contracts		15.2	0.0	71.4	5.0	1.9	0.0

TABLE 3 ■ Optimal Portfolios in a Simple Two-Asset Example

		Three-Month Expected Yield Change in France (Basis Points)			
		-6	-4	-2	0
Three-Month Expected Yield Change in Germany (Basis Points)	-4	**3.8**, *12.0* 58/42	**3.9**, *12.5* 90/10	**4.3**, *13.3* 124/-24	**5.1**, *14.2* 153/-53
	-2	**3.5**, *11.9* 26/74	**3.3**, *12.0* 58/42	**3.4**, *12.5* 91/9	**3.8**, *13.2* 122/-22
	0	**3.6**, *12.2* -8/108	**3.0**, *11.9* 24/76	**2.7**, *12.0* 56/44	**2.8**, *12.5* 87/13
	2	**4.8**, *13.6* -73/173	**3.5**, *12.8* -41/141	**2.6**, *12.2* -8/108	**2.0**, *11.9* 22/78

Key: **return** *volatility*
% Germany bonds / % French bonds

French bonds and the expected portfolio returns and volatilities. The columns and rows of the table show different views, expressed as expected changes in yields on ten-year maturity benchmark bonds in each country over a three-month horizon. Small changes in the views imply extreme changes in the optimal weights.

For example, expectations of no change in German yields and a 6 bp drop in French yields lead to an optimal portfolio long 108% French bonds financed by going short 8% in German bonds. An increase of just 4 bp in the outlook for French yields and a decline of 4 bp in the outlook for German bonds leads to an optimal portfolio long 124% German bonds and financed by going short 24% in French bonds.

Even this simple two-asset example illustrates a number of the difficulties and complexities that arise in using the standard mean-variance optimization model. Consider, for example, the relationship between anticipated yield changes in the two countries and the portfolio weights.

In forming the optimal portfolio weights, the model balances the expected returns of the two assets relative to their domestic short rates. Thus, the translation of yield changes into these expected excess returns involves both the durations of the bonds and the level of short rates. Beyond these calculations to get excess returns, the optimal weights also depend on the volatilities of returns of each security and the covariance terms. Of course, the complexities only expand when we consider more assets.

What this exercise illustrates is that the model and the portfolio manager are not likely to look at expected returns or portfolio weights in the same light. On the one hand, portfolio managers look at these different portfolio weights and find them extreme relative to their notions of a balanced portfolio. The model, whose assessment is based on the high correlations of historical returns, finds that the risks of the extreme portfolio weights are not that great and therefore does not treat them as unreasonable.

On the other hand, while expected yield changes that differ by just a few basis points may not appear significantly different to the portfolio manager, to the model such a difference in expected returns may indeed represent a significant break from past behavior.

Note that in markets where returns are largely unpredictable, as they must be in liquid assets such as bonds and currencies,[7] reasonable variations in expected returns must be small — in fact, much smaller than historical variations in actual returns. Thus, the model may be correct in recognizing that relative to historical correlations, changes of a few basis points in mean returns are large differences that logically imply significant reallocation of portfolio weights in order to optimize the risk-return trade-off.

Nonetheless, we think it is unrealistic to expect portfolio managers to have refined their views to the degree that they can recognize the extent to which their current views do or do not fall outside of historical norms, especially when such regions may only be a few basis points wide.

The problem of portfolio managers' uncertainty about what are reasonable expected returns is compounded by the fact that in the standard approach to asset allocation models, there is no way to incorporate the *degree* of uncertainty about the expected returns. Because of this, the standard model has no way to distinguish a firmly held view from a weak view, and it will react to both equally strongly.

Given this sensitivity of optimal portfolio weights to expected returns, what can investors do when attempting to use quantitative asset allocation models? One approach model users have taken in the past is to put relatively tight constraints on the weights given to different asset categories.[8] Such constraints can certainly lead to more reasonable looking and performing portfolios, but they also impair the validity of the solution.

The model is designed to balance the return-to-risk ratio across all assets, but if many portfolio weights are forced against constraints, the balance is lost. If the constraints are not real, but rather are imposed in order to generate reasonable looking portfolios, then they clearly reflect some inadequacy of the model, and there is no justification for accepting the constrained solution. Unless we understand why the model would rather choose a different portfolio when unconstrained, we should not be satisfied with a solution that it gives when it is constrained.

Models that include transactions costs can also avoid unbalanced portfolios in an artificial way. By setting the transactions costs high enough, the user can force the model to find an optimal portfolio that is close to the current portfolio. For the same reasons that apply to the artificial use of constraints, we cannot recommend the use of transactions costs to exclude unbalanced portfolios.

The Equilibrium Approach

The equilibrium approach to asset allocation offers a better solution to the problem. The fundamental advantage of the equilibrium is that it points the investor toward a reasonable region for expected returns.

A basic deficiency in the standard approach to mean-variance optimization is the way that it requires investors to specify their views. Because the optimal portfolio weights are highly sensitive to the specification of expected returns, the requirement that investors formulate views for all assets and currencies in an absolute sense, rather than relative to a reasonable point of reference, is a severe problem that has to be addressed directly.

Recall that we showed that a 10 bp change in expected returns on the deutsche mark led to an extreme change in the weights of an optimal portfolio. Given this type of sensitivity of optimal portfolios to expected returns, it is basically impossible for investors, on their own, to specify expected returns on a large number of assets and currencies in a way that will lead to an unconstrained portfolio being at all reasonably balanced.

Thus, we have a good practical reason to point investors toward expected returns that lead to balanced portfolios. There is also a sound theoretical justification.

It is natural to expect market forces to push the prices of bonds and exchange rates to levels where the expected returns are consistent with optimal portfolios approximating the market capitalization weights. This is the basic implication of the capital asset pricing model of Sharpe (1964) and Lintner (1965), and its extension to the ICAPM.

It is natural to define a balanced portfolio as one whose returns can be expected to correlate highly with the market portfolio. If that is what we mean by "balanced," then it turns out, not surprisingly, that expected returns that are close to the equilibrium returns (in a sense that can be made precise)[9] lead to portfolios that are well-balanced.

To fix the intuition of the ICAPM equilibrium approach, consider the demand for an asset such as an Australian government bond. Two features of this asset make it interesting: 1) its returns have a very low correlation with those of other government bonds, and 2) there is a very small quantity of it available relative to other government bonds.

One of the reasons investors diversify into assets around the world is to reduce the total risk of their portfolios. From this perspective, the Australian government bonds, with their low correlations of returns, can make a valuable contribution toward lowering risk.

Unfortunately, the bonds are in short supply, and it is impossible for all investors to hold a significant weight in Australian bonds. What has to happen, in equilibrium, is that the prices of Australian bonds have to be bid up to the point where investors around the world are comfortable holding the small outstanding supply. This price adjustment process can stop only when prices are so high that the expected returns on Australian bonds are low enough relative to other securities that investors will willingly forgo the benefits of diversification.

How low are those expected returns? This is a complicated calculation involving the entire set of covariances of returns. We would not expect portfolio managers to be able to figure this out themselves. Our equilibrium returns provide this useful reference.

Using equilibrium returns also gives investors a much more flexible approach toward specifying their own views. In particular, it is probably reasonable to assume that portfolio managers have, on average, only a few strong opinions at a given time. Our equilibrium approach allows portfolio managers to specify as many or as few views as they wish. Moreover, they can specify relative strengths for each view. Finally, we allow the investor to specify expected returns in a relative rather than an absolute sense.

For example, we allow the investor to specify a view such as "asset A will outperform asset B by X." This is in contrast to the standard approach, which requires views in the form, "asset A will return X, asset B will return Y, asset C will return Z, etc."

In that standard approach, investors are required to specify a complete set of expected returns, and no distinction is possible between the strongly and the weakly held views. Small differences in some of the expected returns for which an investor does not have a strong opinion could easily lead, for reasons explained above, to sharp differences in the optimal portfolio.

The equilibrium approach fosters this flexibility by making the assumption that mean returns ought to be close to the equilibrium except to the extent explicitly stated otherwise. When we incorporate this basic assumption into our approach, then we can formulate views in a much more general way: as one or more statements about either relative or absolute returns.

As an example, we may want to state that French bonds will outperform German bonds by 5 bp over the next three months. Or we may have a view that Australian bonds will return 20 bp above equilibrium (even though we do not know exactly what the equilibrium return is). It is also quite natural in this context to associate different degrees of certainty with different views.

Having formulated views in this way, we can then find the set of expected returns most consistent with 1) our stated views and their different degrees of confidence, and 2) the equilibrium returns. The expected returns formulated this way are more likely to lie close to the equilibrium returns — and therefore to reflect the investor's views in a balanced portfolio — than those generated by requiring the investor to state a full set of expected returns with no reference to the equilibrium.

An Example

We can best illustrate the differences between the equilibrium approach and the standard approach in the context of an example. Suppose our strongest current view is that Japanese government bond yields will rise. We also feel U.S. Treasury yields will rise, but we do not have a view about European and Australian yields. With respect to currencies, we expect a depreciation of the Canadian dollar with respect to the U.S. dollar and a small appreciation of sterling with respect to the deutsche mark and French franc.

We have tried to describe these views in a way that a portfolio manager would typically formulate them. Now consider the difference between our approach and the way we express these views in the standard model, where we must specify an expected return for every asset.

We give in Table 4 an example of the expected returns that we might use in a standard model to try to capture the views described above. Although we have tried to do this in a sensible way, we have to admit at the outset that without the equilibrium reference it is by no means obvious how we should translate our views (expressed in the terms in which we conceive them) into a complete set of expected returns.

To construct the returns shown in Table 4, our ad hoc method was first to convert the views into yield changes over a three-month period (+10 for Japan, +5 for the United States, 0 elsewhere) and annualized percent changes in spot exchange rates (foreign currency per U.S. dollar) relative to the current forward rates (-1% for Canada, +0.5 for the United Kingdom, -0.5 for Germany and France, 0 elsewhere). We then converted those changes into their implied returns for currency-hedged and unhedged bonds.

In contrast, the equilibrium approach allows us to specify views using language that is closer to the way we have formulated them above. Applying our model to this example, we can describe our outlook as follows:

VIEWS EXPRESSED IN THE EQUILIBRIUM APPROACH

View

1. Bonds in Japan will underperform the equilibrium returns by 100 bp (on an annualized basis).
   Confidence = 50%
2. U.S. bonds will underperform the equilibrium returns by 50 bp.
   Confidence = 40%
3. Canadian currency will decline 100 bp relative to the U.S. dollar.
   Confidence = 30%
4. Sterling will outperform the deutsche mark by 50 bp.
   Confidence = 20%
5. Sterling will outperform the French franc by 50 bp.
   Confidence = 20%

The confidence percentages indicate the degree of certainty that we have in each individual view, on a scale

TABLE 4 ■ Ad Hoc Expected Returns Constructed From Views (percent)

	U.S.	Japan	Germany	France	U.K.	Canada	Australia
Unhedged Bonds	7.36	4.75	8.01	7.68	6.08	5.22	8.36
Bonds Hedged into $		4.75	8.51	8.18	5.58	6.22	8.36

from 100% (very strong) to 0% (very weak). In Table 5, we show three sets of expected returns that arise from placing different overall weights on the views versus the equilibrium returns.[10]

By putting different overall weights on the views, we can map out a path of expected returns and associated optimal portfolios that range from the market portfolio to one that incorporates the views exactly. We use such a range of weights so that if the optimal portfolio associated with the higher weights turns out to be too extreme, we can consider a portfolio associated with returns arising from placing more weight on the equilibrium.

ALTERNATIVE OPTIMAL PORTFOLIOS

Let us now consider the optimal portfolios associated with these two ways of specifying the same sets of views. Table 6 shows the portfolio that is optimal with views expressed according to the standard approach.

Notice that these weights might be considered somewhat extreme. Of course, portfolio managers would probably not accept such an allocation but would start adding constraints until they found some acceptable portfolio weights. Other managers might "play" with their views, hoping to find expected returns that would lead to a more balanced portfolio, or they might simply specify weights directly and not use the model at all.

Now consider the optimal portfolios shown in Table 7, which arise from the different combinations of the investor's views with the equilibrium expected excess returns. As this table makes clear, adjusting the weights on the views and the equilibrium leads to a whole spectrum of portfolios from which investors can choose, depending on the degree of confidence they have in their views.

TABLE 5 ■ Expected Returns Generated from the Equilibrium Approach (percent)

	U.S.	Japan	Germany	France	U.K.	Canada	Australia
100% Weight on Investor Views							
Unhedged Bonds	9.27	8.62	8.56	8.58	9.94	7.76	8.40
Bonds Hedged into $		7.98	8.63	8.66	9.51	8.76	8.56
75% Weight on Investor Views							
Unhedged Bonds	9.50	8.85	9.15	9.03	9.86	9.15	8.37
Bonds Hedged into $		8.22	8.92	8.80	9.43	9.25	8.23
50% Weight on Investor Views							
Unhedged Bonds	9.62	9.16	9.25	9.11	9.90	9.39	8.35
Bonds Hedged into $		8.47	9.03	8.89	9.51	9.38	8.20

TABLE 6 ■ Optimal Portfolio Weights Based on the Standard Approach (percent)

	U.S.	Japan	Germany	France	U.K.	Canada	Australia
Cash	-51.4	3.4	0.0	0.0	8.6	0.0	0.0
Bonds	-13.1	0.0	50.5	0.4	0.5	0.9	60.5
Forward Contracts		0.0	54.7	44.0	0.0	90.8	50.8

TABLE 7 ■ Optimal Portfolio Weights Based on the Equilibrium Approach (percent)

	U.S.	Japan	Germany	France	U.K.	Canada	Australia
Portfolio Based on 100% Weight on Investor Views							
Cash	-19.5	2.1	0.0	0.0	17.5	0.0	0.0
Bonds	49.5	0.0	12.7	0.2	4.8	10.8	21.8
Forward Contracts		0.0	5.5	20.4	0.0	118.4	20.2
Portfolio Based on 75% Weight on Investor Views							
Cash	-7.0	3.7	0.0	0.0	3.4	0.0	0.0
Bonds	53.7	0.0	13.4	0.0	5.6	9.6	17.5
Forward Contracts		0.0	15.4	0.0	0.0	21.6	16.6
Portfolio Based on 50% Weight on Investor Views							
Cash	-0.9	0.2	0.0	0.4	0.4	0.0	0.0
Bonds	54.8	4.0	13.1	0.0	6.0	8.4	13.5
Forward Contracts		0.0	14.1	0.0	0.0	10.3	13.0

None of these portfolios is as extreme as the portfolio in Table 6 — which is to be expected, because we move away from the equilibrium only to the extent that the views require it. Using the standard model, we might move away from equilibrium (in this case toward Australian bonds) simply because in our ad hoc approach we did not recognize how low the expected returns on Australian bonds had to be to keep us from wanting to take large positions in them (of course, the nature of the ICAPM equilibrium makes it very reasonable for us to expect the returns to be that low).

While the portfolio that puts 100% weight on views may be a little extreme, there are clearly other portfolios — associated with returns incorporating less weight on views — that do look reasonably balanced. By allowing adjustment of the weights on views in this way, the equilibrium approach removes the need to use artificial constraints on portfolio weights in order to generate balanced portfolios.

A Demonstration of the Equilibrium Approach

The key assumption of the equilibrium approach — use of the equilibrium expected excess returns *except to the extent explicitly stated otherwise* — is a subtle but extremely useful concept. To illustrate the power of this idea further, let us revisit the example that we used at the outset (see Tables 1 and 2) to demonstrate the sensitivity of portfolios to changes in expected returns.

We raised the expected return on deutsche marks by 10 bp above the equilibrium value and found that the optimal portfolio moved from a balanced portfolio to one with extreme positions. We will again make an adjustment to expected returns, but this time by expressing a view in the equilibrium context. Rather than raise the expected return on deutsche marks and hold all other expected returns fixed, we express the view that deutsche mark returns will be 10 bp above the equilibrium.[11] The result is shown in Table 8.

In the equilibrium context, expressing this view does not lead to an unbalanced portfolio. The difference is that rather than hold all other expected returns fixed, we find the expected returns that are most consistent with both this view and the market equilibrium.

In this example, given the high correlation of returns between deutsche marks and French francs, that most consistent set of expected returns includes a higher expected return on French francs, such that the optimal portfolio remains well-balanced.[12]

OTHER FEATURES OF THE MODEL

We have focused specifically on our approach to combining views with the ICAPM equilibrium because we feel it is critical to the successful use of this type of analysis. Several other distinguishing features of the model are also worthy of mention, however.

Implied Views

One feature of our model is its ability to identify the investor views that are implied by a given portfolio. The usual steps in using the model are to specify views, estimate covariances, and solve for the optimal portfolio. Using the model this way defines a mapping from expected returns to optimal portfolios.

But another approach that is often useful — especially for an investor who has undefined views but who has a current portfolio — is to turn the mapping around. In other words, the investor starts with the portfolio and lets the model solve for the expected returns for which the portfolio is optimal.

We call these expected returns the "implied views" of the portfolio. The implied views can then form the basis for a careful formulation of a set of explicit

TABLE 8 ■ Combined View With Market Equilibrium and the New Optimal Portfolio (percent)

	U.S.	Japan	Germany	France	U.K.	Canada	Australia
Combined View With Market Equilibrium Expected Returns on Bonds (August 20, 1990)							
Unhedged Bonds	9.77	9.85	9.47	9.28	9.99	9.60	8.33
Bonds Hedged into $		8.97	9.21	9.04	9.65	9.55	8.19
New Optimal Portfolio Weights							
Cash	0.1	0.0	0.0	0.0	0.0	0.0	0.0
Bonds	53.7	18.3	9.3	7.0	6.3	4.3	1.0
Forward Contracts		14.4	5.5	5.6	5.0	3.5	0.8

views that more adequately reflect the investor's current investment strategy.

Use of Daily Data in Estimating Risks

Another feature is the way we estimate the risk profile of the international bond portfolio. We treat the variances and covariances of bond and currency returns in a manner analogous to the way options traders treat volatility when they price interest rate and currency options. That is, we regard them as unobservable quantities that change continuously through time. We use the historical returns to estimate the current values of these moving targets.

There are many ways to do so, just as there are many ways to estimate empirical volatilities in pricing options. We can characterize these different approaches by the frequency of data used and by differential weighting of more recent versus older observations.

When we think about the variances and covariances as unobserved time-varying quantities, it becomes clear that more recent observations of returns contain more information about current values than do older observations. It is also clear that we should use data gathered as frequently as is practical, such as daily data, because these observations will have more relevant information than will a similar sample of observations gathered at less frequent intervals.

There are several problems that have discouraged users of other international asset allocation models from employing daily data. The first, and perhaps most serious, problem in using daily data in this context is dealing with unsynchronized time intervals: International markets do not open and close at the same times around the world, so the daily returns are not recorded at the same times.

If the problem is ignored, this lack of synchronization of returns induces a downward bias in the estimation of correlations across markets. In our approach we do use daily data, but we avoid the biased estimates by explicitly considering the exact times of trading in each market and therefore the actual overlaps that occur in the daily return data.

For example, we take into account that Tuesday's returns on U.S. bonds will correlate not only with Tuesday's returns on Japanese bonds, for which there are eleven hours of overlap, but also with Wednesday's returns in Japan, for which there are thirteen hours of overlap.[13] We use a weighted sum of both effects in estimating the correlation coefficient for returns in the two markets.

In the case of U.S. and Japanese bonds, the overlap of returns from Tuesday and Wednesday in Japan with Tuesday's returns in the United States is almost the same, so we give nearly equal weight to both sets of returns in estimating the covariances.

In contrast, the returns in London on Tuesday and Wednesday also are correlated with U.S. returns on Tuesday, but the Wednesday correlation reflects only the five hours of price changes that occur in New York after the close of trading in London, while the correlations in the Tuesday returns are based on price changes reflecting nineteen hours of overlap.

In this case, we give much more weight to the Tuesday returns in estimating the covariances. We choose the exact weights that will maximize the efficiency of our estimated covariances.

Treatment of Holidays

Another problem with using daily data arises out of the large number of different holidays in different countries. This makes the treatment of missing data a serious dilemma. Fortunately, it turns out that the method we use for correcting for lack of synchronous returns lends itself to missing data as well.

For example, if the U.S. bond market is closed on Monday, we form the total return from the close on Friday through the close on Tuesday. The correlation we want to estimate between returns in the United States and Japan will manifest itself in a relationship between the U.S. bond returns reflecting price changes from Friday's close to Tuesday's close and the Japanese bond returns for Monday, Tuesday, and Wednesday. By knowing the hours of overlap of each of those returns with the U.S. returns, we can weight them correctly in the covariance estimation.

Specification of the Objective

Another feature of our model is its flexibility in specifying the investor's objective and constraints. While the basic objective is always to maximize expected return for a given level of risk, we feel that it is important to recognize that for many portfolio managers, risk and return are measured not in an absolute sense but rather relative to a specific benchmark such as a particular international bond index. For this reason, we allow both risk and return to be measured relative to any benchmark portfolio that the investor wishes to specify.

In any application, the optimization generates not just one optimal portfolio, but one portfolio for every level of risk. In practice, the investor has to search along the optimal portfolio frontier to find the portfolio that best balances the trade-off between risk and return. We allow several approaches to that choice, such as selecting a particular level of risk, maximizing the ratio of return to risk, or specifying a utility function.

Types of Constraints

In discussing our equilibrium approach, we warned against using *artificial* constraints on portfolio weights to get around the problem of unreasonable looking portfolios. Of course, there are often *real* constraints on portfolio weights or on the degree of currency hedging that can be employed. We therefore include the capability of maximizing an objective subject to constraints on the maximum and minimum weights in each country, on the degree of currency hedging, and on the degree of leverage.

Yield Versus Expected Return

We recognize in addition that many portfolio managers face an explicit or implicit trade-off between yield and expected return. For this reason, we also allow the portfolio manager to specify a constraint on the yield of the portfolio. By adjusting this constraint and looking at the resulting portfolios, the portfolio manager can easily investigate the nature of this trade-off.

Adjustments to Equilibrium Expected Returns

One final feature of our asset allocation approach is something we call "adjustments to equilibrium expected returns based on market conditions." These adjustments are movements of expected returns away from the equilibrium values, based on such current conditions as relative current yields in different countries, different steepnesses of yield curves, and variations in forward discounts. In each case we give investors the freedom to specify the degree to which they expect the market condition to affect returns.

For example, one common view among investors is that a large forward discount, such as that associated with sterling relative to the U.S. dollar, is likely to lead to higher excess returns from holding that currency. An investor holding that view can specify a gain parameter, such as 5 bp of additional excess return (relative to the excess returns of the ICAPM equilibrium) for each currency per percent of forward discount in that currency.

These adjustments are intended not to substitute for a careful consideration of current market views but rather to provide a standard accounting for the effects of certain types of market conditions that regularly play a role and yet change in magnitude daily.

CONCLUSION

We have presented a new international asset allocation model that brings several innovations to the asset allocation problem. Its most interesting feature lets us combine and compare our outlook for currencies and bond markets with expected returns generated by an ICAPM equilibrium. This provides an elegant solution to the major problem of the standard mean-variance approach to asset allocation: the tendency for small variations in expectations to produce dramatic changes in the optimal portfolio.

The standard approach typically takes a set of bond yield and currency forecasts, calculates the expected returns on each asset, and then — using some measure of the riskiness of assets and their interrelations — finds the portfolio with the highest expected return for a given level of risk. The ICAPM recognizes that, in equilibrium, bond yields and exchange rates will adjust until investors worldwide are comfortable holding the outstanding assets. The equilibrium returns are those that would make an investor comfortable holding a market capitalization portfolio, i.e., that would make the market capitalization portfolio optimal.

In exactly the same way, the model can take an investor's actual portfolio and derive the "implied views" — the set of returns that would make the investor's portfolio optimal. Comparing actual views with implied views of an investor's portfolio is a powerful way of revealing flaws in investment strategy.

This new asset allocation approach — which allows the investor to combine particular expected returns with the equilibrium returns — lets investors derive portfolios that appear balanced and reflect their views, without resorting to arbitrary constraints on portfolio composition (although real constraints may be incorporated in the solution). A strength of the model is its emphasis on using recent daily data to assess the riskiness of assets and their interrelationship.

The model also lets portfolio managers incorporate their market views in a manner that approximates the way they actually think about their outlook. It uses relative rather than absolute returns and takes into account the degree of confidence with which investor views are held.

ENDNOTES

This article is based on a report by the same title published by Goldman, Sachs & Co.

The model described in this article is the new Goldman, Sachs international asset allocation model. The authors acknowledge the many helpful suggestions of Casper Bowden, Shahid Chaudhri, Jeremy Hale, Kevin Keane, Peter Knez, Tim Knowles, Kaoru Ono, Scott Pinkus, Scott Richard, Ken Singleton, Richard Thomasson, and Jeff Wecker.

[1]The closest precedent to this approach is in Treynor and Black [1973], where the authors show how to combine a market portfolio with an "active portfolio" consisting of securities for which there is an "appraisal premium."

[2]See, for example, Rosenberg [1990] and Thomas [1990].

[3]Described in Adler and Dumas [1983], Solnik [1974], and most recently by Black [1989].

[4]Academic researchers in finance are well aware of this problem, although they have not yet come to agreement on the appropriate solution. For a recent treatment, see Green and Hollifield [1990]. Academics have tended not to focus on the formulation of views by investors but rather on the use of statistical techniques to estimate expected returns from historical data. Formally, however, both historical data and investor views can be treated as sources of information about future returns. The problem in either case is basically the same: how best to use the available information to generate a distribution for expected returns. Our solution to the problem can be stated as the use of an informative Bayesian prior, whose distribution reflects the ICAPM equilibrium.

[5]All tables display portfolios from a U.S. dollar currency perspective in terms of percentage weights in cash, bonds, and forward contracts in different countries. The weights on bonds refer to benchmark ten-year maturity bonds in that country. A weight on forward contracts refers to a holding of foreign exchange contracts that investors can use to hedge the currency risk in holding foreign bonds by fixing an exchange rate at which they will exchange foreign currency into dollars on a future date. We adopt a convention that the weight shown for forward contract holdings is always positive. Instead of showing a negative weight on forward contracts, we show the active choice to increase currency exposure as a positive holding of foreign cash. We show the financing of a levered portfolio as a negative weight on the domestic (U.S. dollar) cash position. In practice, of course, the same effective exposures to interest rate and currency risk could be accomplished through a variety of other types of assets, such as swaps, forward rate agreements, futures contracts, or options. Our model treats all portfolios with the same local exposure to interest rate and exchange rate changes as equivalent in terms of the risk-return trade-off.

[6]Up to this point we have not specified how we will pick one of the frontier portfolios as "optimal." In order to do this we need somehow to specify a set of preferences between risk and expected return. Throughout the examples in this paper we use a utility function to determine the frontier portfolio that we call "optimal." The extreme sensitivity of portfolio weights to expected returns that we focus on here is itself not sensitive to how we make this choice, as long as it reflects a reasonably smooth trade-off between risk and expected return.

[7]For illiquid assets, predictable returns may persist simply because transactions costs would be too high to make it profitable for investors to take advantage of the expected change in price. For precision, note also that it is the excess returns above the risk-free rate that are unpredictable.

[8]See, for example, Frost and Savarino [1988].

[9]To make this mapping precise, we must measure distance between two vectors of returns using a metric that is based on the inverse of the covariance matrix of returns.

[10]Formally, we generate expected excess returns as the mean of a posterior distribution that incorporates the information in the equilibrium and the investor's views. We assume a prior distribution based on the ICAPM equilibrium, with the covariances of the prior proportional to the historical covariances of excess returns. The constant of proportionality reflects the relative weight given to the equilibrium versus the views. We translate individual views into "observations" on a set of linear combinations of the expectations for individual securities, again with covariances proportional to historical values (here the constants of proportionality reflect the relative degrees of confidence in the individual views), and combine the results with the prior in a procedure that Theil [1971] describes as "mixed estimation" (pp. 346-352 and 670-673). Note that when we adjust the overall weight on views versus equilibrium, putting 100% weight on the views does not in itself imply that we ignore the equilibrium, but rather that we have chosen the expected returns from a subspace of expected returns defined by linear restrictions incorporated in the views. Only when the number of views is equal to the dimensionality of the expected return vector and 100% weight is given to the views does the equilibrium have no impact.

[11]In the Bayesian approach, we have to specify a degree of confidence. We have indicated 100% confidence in this view. If we specify less confidence in the view, the portfolio will move even closer to the original balanced portfolio consistent with the equilibrium. The point we are making here is that the equilibrium approach leads to balanced portfolios that reflect our views — even when we have strong confidence in those views.

[12]Of course, if we actually do want to specify that the relative return of marks over francs will be 10 bp higher than the equilibrium differential, that is a very distinct view that would imply a significant reallocation of the portfolio. An advantage of the equilibrium approach is that it makes this type of distinction clear and allows investors to distinguish exactly which of these very different views they want to express.

[13]We measure returns for a given day on the basis of the change in closing prices from the previous trading day to the given day. The relevant overlap is then based on the entire time between closes, during which new information will affect the returns, rather than just the overlap of trading hours (of which there may be none). We make the assumption that information about asset values is generated at a constant rate twenty-four hours each day. A more complicated model — for example, one that assumes that information about each asset value is generated at a higher rate during the trading hours of that asset — would lead to a different weighting scheme.

REFERENCES

Adler, Michael, and Bernard Dumas. "International Portfolio Choice and Corporate Finance: A Synthesis." *Journal of Finance,* Vol. 38 (June 1983), pp. 925-984.

Black, Fischer. "Universal Hedging: How to Optimize Currency Risk and Reward in International Equity Portfolios." *Financial Analysts Journal,* Vol. 45, No. 4 (July/August 1989), pp. 16-22.

Frost, Peter A., and James E. Savarino. "For Better Performance: Constrain Portfolio Weights." *Journal of Portfolio Management,* Fall 1988, pp. 29-34.

Green, Richard C., and Burton Hollifield. "When Will Mean-Variance Efficient Portfolios Be Well Diversified?" Working Paper, Graduate School of Industrial Administration, Carnegie-Mellon University, February 1990 (revised).

Lintner, John. "The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets." *Review of Economics and Statistics,* Vol. 47, No. 1 (February 1965), pp. 13-37.

Rosenberg, Michael R. "A Framework for Formulating International Fixed-Income Strategy." *Journal of Portfolio Management,* Summer 1990, pp. 70-76.

Sharpe, William F. "Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk." *Journal of Finance,* Vol. 19, No. 3 (September 1964), pp. 425-442.

Solnik, Bruno H. "An Equilibrium Model of the International Capital Market." *Journal of Economic Theory,* Vol. 8 (August 1974), pp. 500-524.

Theil, Henri. *Principles of Econometrics.* New York: John Wiley & Sons, 1971.

Thomas, Lee R. "Bring Discipline to Global Asset Allocation." *Global Investing,* Spring 1990, pp. 22-33.

Treynor, Jack L., and Fischer Black. "How to Use Security Analysis to Improve Portfolio Selection." *The Journal of Business,* Vol. 46, No. 1 (January 1973), pp. 66-86.