The Journal of Portfolio Management

VOLUME 22 NUMBER 2 WINTER 1996

# When Do Bond Markets Reward Investors for Interest Rate Risk?

*Empirical evidence and implications for long-term investors.*

Antti Ilmanen

# When Do Bond Markets Reward Investors for Interest Rate Risk?

*Empirical evidence and implications for long-term investors.*

Antti Ilmanen

**ANTTI ILMANEN** is a vice president in fixed-income research at Salomon Brothers Inc in New York (NY 10048).

Fixed-income portfolio managers pay considerable attention to risk/return trade-offs. They devote most of their effort, however, to developing interest rate risk measures and risk control tools. They have not analyzed the other side of the trade-off, the potential reward, nearly as much.

In this article, we will examine *what historical data can tell us about the reward for interest rate risk.* We show empirical evidence from several bond markets and discuss the implications for interest rate risk management.

Investors often measure interest rate risk by the volatility of bond returns over some investment horizon, say, a month. Monthly returns of long-term bonds are highly variable, while the returns of one-month bonds are quite stable. We define the realized bond risk premium as the excess return of a long-term bond over a one-month bond. Most of the monthly excess return is unexpected — this is the risk. A small part of it is *expected* — this is the reward for interest rate risk; that is, the compensation that investors require and the market offers for bearing the long-term bond's volatility.

It is hard to dissect the realized bond risk premium into expected and unexpected parts, because market expectations are not observable. On the one hand, if the expected bond risk premium is constant over time, we can estimate it by the historical average premium. On the other hand, if the expected premium varies over time, a historical average is not the best estimate of the near-term expected premium.

If we find that excess bond returns can be predicted with available information, we can use the predictable part as our estimate of the expected bond risk premium. For example, if a bond's yield equals its near-term expected return, the yield difference between long and short bonds (the term spread) is an exact measure of the expected bond risk premium.

We can evaluate the empirical validity of using yields as proxies of expected bond returns by studying whether the term spread can forecast the subsequent realized bond risk premium. Many studies find a significant positive relation, which often is taken as evidence that the expected bond risk premium varies over time.[1]

**EXHIBIT 1**
**AVERAGE CASH AND BOND RETURNS, 1978-1993**

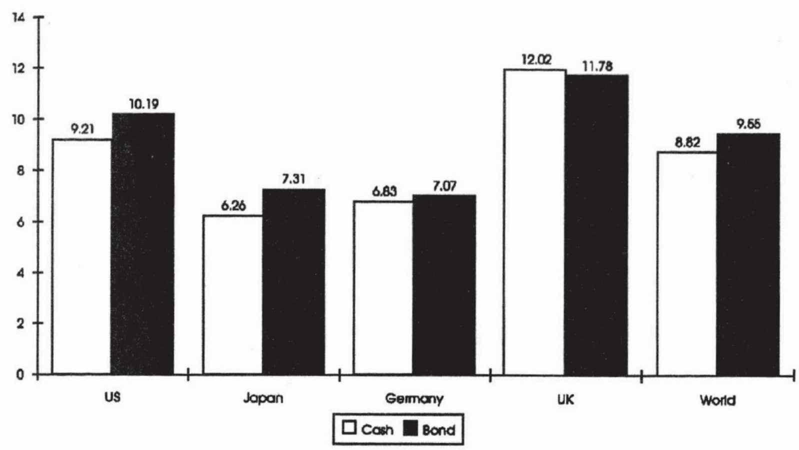


We can draw four main conclusions in our examination of the behavior of international bond markets:

- On average, the historical returns of long bonds have not been much larger than those of short bonds.
- Investors can forecast excess bond returns by using variables such as the term spread, the real bond yield, and past stock market performance. Although the average reward for interest rate risk is small, at times the reward is large. We also offer an economic explanation for the time variation in the expected bond risk premium.
- If we can identify periods when the expected bond risk premium is abnormally high or low, investment strategies that adjust portfolio duration over time should outperform static-duration strategies. Such dynamic strategies would have earned average annual returns between 1978 and 1993 that were several hundred basis points higher than static-duration strategies, with less risk.
- These findings have implications for choosing the benchmark duration and for managing interest rate risk (actively or passively). Short-term traders are not likely to take advantage of these empirical regularities; the proposed strategies are better suited for long-term investors.

## AVERAGE BOND RISK PREMIUM

Exhibit 1 reports the annualized average returns of long-term government bonds and one-month Eurodeposits in the United States, Japan, Germany, the United Kingdom, and "the world" between January 1978 and June 1993.[2] The bond data are from the Salomon Brothers World Bond Index$^{SM}$ (government bonds whose maturities exceed five years). The returns are compounded continuously and measured in local currency terms. The realized bond risk premium is the difference between the bond return and the one-month Eurodeposit rate; it may be interpreted as the currency-hedged bond risk premium for any international investor.[3]

Exhibit 1 shows that long-term bonds earned, at most, 105 basis points more in average annual return than one-month assets. These average return differences are not nearly statistically significant in any country — *the market offered meager reward, if any, for bearing interest rate risk.*

These results are based on one fifteen-year period. If the sample period began a few years later than 1978, the average premium would be higher, but then the sample would cover an exceptionally disinflationary period.

To alleviate the problem of sample-specific results, we plot in Exhibit 2 the 120-month rolling average excess bond return in the United States from 1952-1961 to 1983-1992 (the five- to ten-year maturity Treasury bond return minus the one-month Treasury bill return). The average was very low, often negative, during the rising inflation in the 1960s and 1970s. The average has been significantly positive only recently, when the sample covers the disinflationary period.

Even when we examine historical data over

**EXHIBIT 2**
**120-MONTH ROLLING EXCESS RETURN (AND RETURN) OF FIVE- TO TEN-YEAR U.S. TREASURY BONDS**

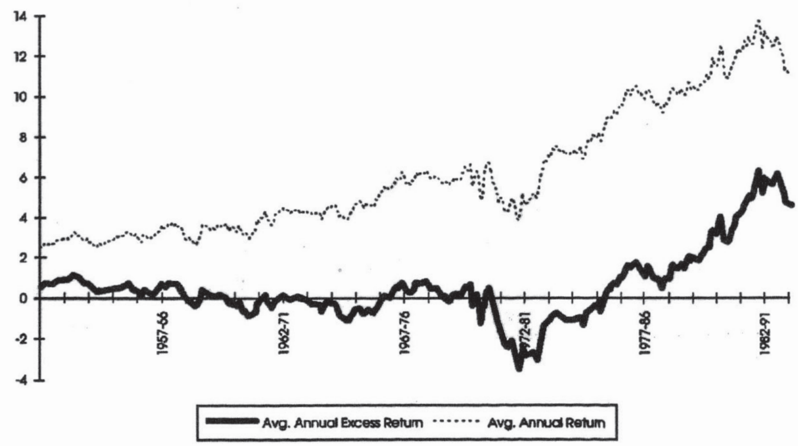


required real return and thus led to higher yields and lower bond prices. Since this major reevaluation is behind us, bondholders can now "enjoy" the higher expected premium.

Why should investors care about these results? The average bond risk premium is relevant mainly for choosing the benchmark duration. A positive expected bond risk premium is often stated as a reason for lengthening the benchmark duration. While the evidence is weak, our view is that expected bond returns do increase with duration, albeit quite slowly. Our preliminary analysis in Ilmanen [1995a] suggests that the reward-to-risk ratio is best at short durations.

forty years, we can make few conclusions because the systematic rise in inflation in the 1960s and 1970s and decline in the 1980s dominate the observed bond returns. We cannot reasonably interpret these realized ten-year average bond risk premiums as reflections of *expected* bond risk premiums when they seem to be driven mainly by *unexpected* inflation. A historical sample average value is a good estimate of the expected value only if the sample is long enough that the unexpected news in each period cancels out over time.

Despite the weak statistical evidence, intuition suggests that the expected bond risk premium should be positive. Long-duration bond returns are more volatile than short-duration bond returns, and we would expect investors to demand some incremental reward for the higher risk.[4] The positive average slope in the yield curve over time also may be seen as evidence of a positive average premium. Moreover, the future expected bond risk premium is more likely to resemble the high average premium over the past ten years than the low average premium over the past forty years.

Until the 1960s, bond risk was low, as was the expected bond risk premium investors demanded. The riskiness of the bond markets increased substantially in the 1970s and 1980s. The interest rate volatility and the correlation between bond and stock markets rose, a trend that probably increased the expected bond risk premium. While rising inflation caused most of the world bond market's poor performance in the 1970s and early 1980s, part of the poor performance reflects a *reevaluation of the market's riskiness*, which increased the

## TIME VARIATION IN EXPECTED BOND RISK PREMIUMS

The historical average premium may be a good estimate of the *long-term* expected bond risk premium, but it is a suboptimal estimate of the *near-term* expected premium, unless the expected premium is constant over time. Recent analysis shows that the expected bond risk premium is not constant but varies over time with changing economic conditions — that is, we can identify periods when the expected bond risk premium is abnormally high or low.[5]

Is there any reason to think that expected risk premiums vary over time? Many academics claim that stock and bond returns are not predictable. Standard models in finance, such as the one-period Capital Asset Pricing Model or the Liquidity Premium Hypothesis of interest rates, assume that expected risk premiums are constant, an assumption made for modeling convenience. In the real world, there are intuitive reasons that expected risk premiums should vary with economic conditions.

Exhibit 3 is a good starting point. It shows the average excess stock and bond returns at varying stages of the U.S. business cycle, using data between 1953 and 1992. For this exhibit, we computed monthly excess returns in seven-month "windows" around each business cycle trough, as defined by the National Bureau of Economic Research, and took an average of all these monthly excess returns over the nine cycles experienced in this period. We then did the same for each

**EXHIBIT 3**
**AVERAGE BUSINESS CYCLE PATTERN OF EXCESS STOCK AND BOND RETURNS IN THE U.S., 1953-1992**

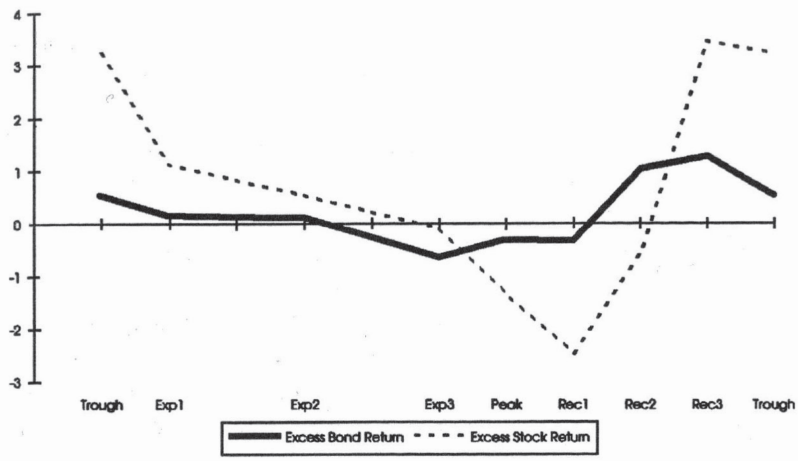


business cycle peak. Finally, we allocated each month to one of six subsamples, depending on whether it occurs in the first, second, or third part of an expansion (trough-to-peak) or in the first, second, or third part of a contraction (peak-to-trough), and then computed the average monthly excess returns for every subsample.

Exhibit 3 reveals some general findings that more formal analysis can confirm. Markets are forward-looking. For example, the excess stock return tends to be highest before a recession officially ends in a trough. Movements in the excess bond return appear to precede those in the excess stock return near cyclical turning points. Most important, both excess stock and bond returns are high near business cycle troughs (after recession) and negative near peaks (after expansion).

Exhibit 3 describes realized premiums, but Fama and French [1989] have shown that *expected* stock and bond returns also are high near business cycle troughs and low near peaks.

What is the *economic explanation* for the cyclical variation in expected risk premiums? The rational answer is that either the amount of risk or the general risk aversion level varies over the business cycle. If this is true, the average investor will not increase risk exposure in periods of abnormally high expected premium, because the higher reward is just a fair compensation for the higher risk or risk aversion.

Marcus [1989] and Sharpe [1990] argue that wealth-dependent risk aversion may cause the time variation in expected risk premiums. Investors are more risk-averse when their current wealth is low relative to their past wealth. The higher risk aversion makes them demand larger compensation for holding risky assets, such as stocks and long-term bonds, which in turn leads to higher expected risk premiums. Conversely, higher wealth near business cycle peaks makes investors less risk-averse, so they bid down the expected risk premiums (by bidding up the asset prices). Such behavior can explain the observed cyclical patterns in expected premiums.

The conjecture of wealth-dependent risk aversion guides our choices of variables that should forecast bond returns. Our main variable reflects the recent performance of the stock market, and it is motivated by a model in which investors have a positive subsistence level. When investors' wealth declines toward the subsistence level, their risk aversion increases (since they cannot afford to risk falling below that level). Investors want to switch from risky assets to riskless assets, a desire that reduces risky asset prices and increases the required market risk premium.

Our empirical proxy for the overall risk aversion level is "*inverse stock*," the ratio of past wealth (stock market level) to current wealth, where recent levels of wealth have greater weight in "past wealth" than do distant levels of wealth. A high ratio (that is, depressed stock market) should reflect a high risk aversion level today and predict a high risk premium in the near future.

The other two variables, *the real bond yield* and *the term spread*, may be good proxies for the expected bond risk premium. A long-term bond yield reflects, first, the market's expectations of future short rates and, second, the expected bond risk premium. By subtracting the recent year-on-year inflation rate or the current one-month interest rate (which may reflect the first part) from the bond yield, we hope to obtain an enhanced proxy for the second part.

We stress that our forecasting variables are financial market data that are publicly known at the beginning of the forecasting month. Empirically, these variables are better predictors than macroeconomic variables, perhaps because the latter are less accurately measured and less timely. While market-related variables are forward-looking, contemporaneous economic data describe history, and with a publication lag.

Extending our analysis beyond the U.S. market, we predict monthly excess bond returns in four countries (and in "the world") between January 1978 and June 1993. We use both country-specific (local) predictors — the term spread, real bond yield, and inverse stock — and their GNP-weighted averages, which we call global predictors. Ilmanen [1995b] describes these variables and the data sources in detail.

## CORRELATIONS BETWEEN FORECASTING VARIABLES AND SUBSEQUENT BOND RISK PREMIUMS

Exhibit 4 shows the correlations between the realized world bond risk premium and various global predictors. A correlation coefficient measures how closely two series are related. Its possible values range from −1 to 1, with 1 indicating a perfect positive correlation, −1 indicating a perfect negative correlation, and 0 indicating lack of any correlation. The conventional view that bond risk premiums cannot be predicted using publicly available information implies that all these correlations should be very close to zero.

Exhibit 4 displays first the predictive ability of two potentially useful forecasting variables, the current long-term rate and the historical twelve-month yield volatility. The correlations are almost zero, so these variables do not exhibit any forecasting ability. In contrast, *the correlation between each of our three predictors and the subsequent monthly world bond risk premium is 0.23-0.25.* The expected bond risk premium really appears to vary over time; it is abnormally high when real bond yields and term spreads (the overall proxies for expected bond returns) are high and when "inverse stock" is high (that is, after a stock market decline, when risk aversion should be high).

These signs are consistent with our economic story. An alternative explanation for our empirical findings is that they reflect systematic forecast errors by investors, perhaps related to wealth-dependent sentiment or naive inflation forecasts.

Each of our three variables can forecast next month's bond risk premium. Combining the information in these predictors — by running a multiple regression — may lead to even better forecasts. A predictive regression splits the realized bond risk premium into an expected part (fitted value) and an unexpected part (residual). The last column in Exhibit 4 shows the correlation between the realized bond risk premium and its expected part, which is based on a multiple regression of the realized bond risk premium on three global predictors. The correlation is 0.38 — using the information in all three predictors does seem to improve the forecasts.[6]

The correlations reported in Exhibit 4 are higher than those found in other published literature, and they are statistically significant.[7] Yet portfolio managers should view any "exciting" empirical findings about asset market behavior with a healthy dose of skepticism. We will examine the economic significance of these findings and address questions about robustness to transaction costs, to risk adjustment, and to out-of-sample estimation.

Exhibit 5 summarizes the information from several multiple regressions; the detailed results appear in Exhibit 6. We regress each country's realized bond risk premium first on the three country-specific predictors and then on the three global predictors.

Exhibit 5 displays correlations between the realized and the expected bond risk premiums. We can see that the predictability patterns are reasonably similar across countries and that the global variables seem to be better predictors than

**EXHIBIT 4**
**CORRELATION OF THE REALIZED WORLD BOND RISK PREMIUM WITH GLOBAL PREDICTORS, 1978-1993**

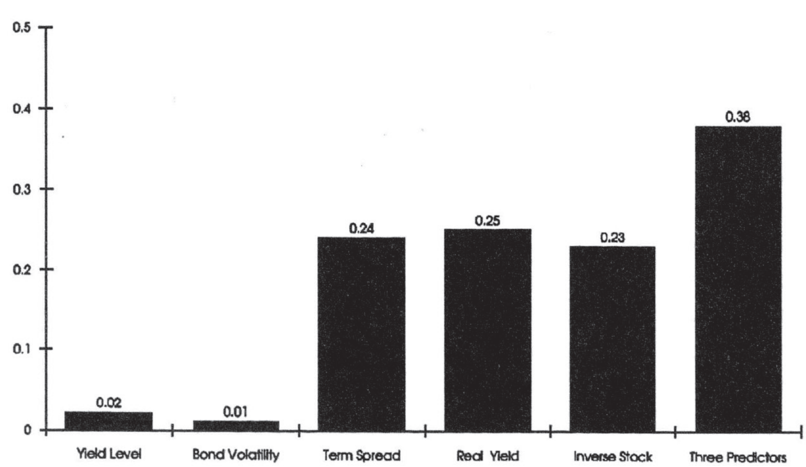

**EXHIBIT 5**
**CORRELATION OF EACH REALIZED BOND RISK PREMIUM WITH LOCAL AND GLOBAL PREDICTORS, 1978-1993**

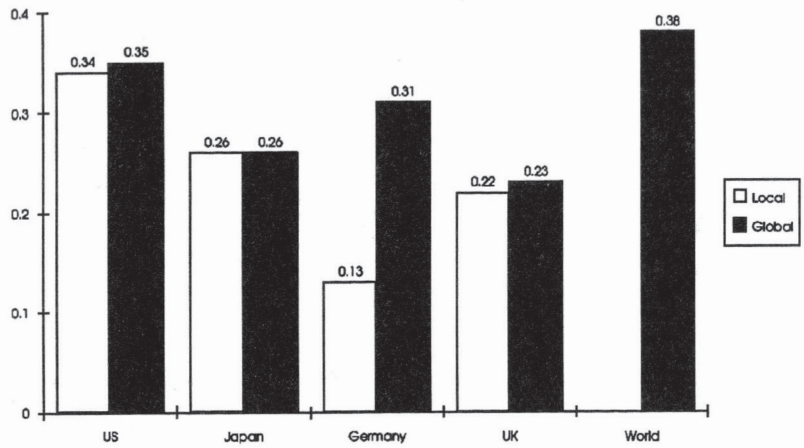


the country-specific variables. These findings may reflect the high level of integration between these markets — expected risk premiums are driven more by worldwide economic conditions than by national economic conditions.

Exhibit 7 plots the expected bond risk premiums (fitted values from a regression of the realized bond risk premiums on the three global predictors) in the United States, Japan, Germany, and the United Kingdom. The expected bond risk premiums seem to move almost fully synchronously over time. The common variation reflects the worldwide time variation in the reward for bearing interest rate risk.

The time series patterns in Exhibit 7 also appear economically reasonable. After negative expected bond risk premiums in the early years, the series peaks at 1982 (worldwide recession), 1984 (real yields peak during an "inflation scare"), 1987 (after the stock market crash), and 1990 (U.S. recession and the Gulf crisis).

## EXPLOITING PREDICTABILITY BY DYNAMIC INVESTMENT STRATEGIES

Investors are more concerned about the financial significance of our predictability findings than their statistical significance. We examine the performance of a dynamic investment strategy that exploits the forecasting ability of the global predictors and compare it with a static "stay in bonds" strategy for each country. Both strategies may be viewed as zero-investment "arbitrage" positions that are financed with

**EXHIBIT 6**
Regressing Monthly Bond Risk Premiums on Global and Local Predictors, January 1978-June 1993

| | U.S. | | Japan | | Germany | | U.K. | | World | |
|---|---|---|---|---|---|---|---|---|---|---|
| | b | (t) | b | (t) | b | (t) | b | (t) | b | (t) |
| World: | | | | | | | | | | |
| Term Spread | 0.06 | (0.20) | 0.14 | (0.76) | 0.19 | (1.01) | 0.13 | (0.47) | 0.15 | (0.70) |
| Real Yield | 0.46 | (3.50) | 0.12 | (1.66) | 0.14 | (1.68) | 0.21 | (1.33) | 0.30 | (3.39) |
| Inv. Stock | 9.92 | (3.77) | 5.23 | (2.78) | 4.57 | (3.29) | 7.60 | (2.54) | 8.02 | (4.11) |
| $R^2$ | 10.7% | | 5.2% | | 7.9% | | 3.7% | | 13.4% | |
| Local: | | | | | | | | | | |
| Term Spread | 0.25 | (1.35) | −0.20 | (−1.28) | 0.01 | (0.07) | −0.10 | (−1.08) | — | |
| Real Yield | 0.22 | (2.15) | 0.29 | (3.05) | 0.11 | (1.23) | 0.17 | (2.09) | — | |
| Inv. Stock | 9.76 | (2.97) | 2.81 | (2.77) | 1.50 | (1.65) | 6.25 | (2.18) | — | |
| $R^2$ | 10.2% | | 5.0% | | 0.2% | | 3.1% | | — | |

Regressions of the realized monthly bond risk premium in the U.S., Japan, Germany, the U.K., and the "world" (their GNP-weighted average) on two predictor sets. Local predictors are the country-specific term spread (the difference between long-term bond yield and short-term rate), real yield (the difference between long-term bond yield and lagged year-on-year inflation rate) and inv. stock (the inverse of exponentially weighted past stock market performance). Global or "world" predictors are GNP-weighted averages of these local predictors. The exhibit reports each regression slope coefficient b and its t-statistic (adjusted for heteroscedasticity in bond returns), as well as the coefficient of determination ($R^2$), which is adjusted for the degrees of freedom.

the one-month Eurodeposit rate.

The static strategy involves always holding the long-term bond (market); thus, the average premium equals the difference between the average returns of bond and cash in Exhibit 1. The dynamic strategy involves holding one unit of the bond if its predicted premium (based on an in-sample regression with three global predictors) is positive and none of the bond if it is negative. The position is rebalanced at the beginning of each month. Note that a third strategy, "stay in cash," gives each period a zero excess return because this strategy involves rolling over one-month Eurodeposits.

Exhibit 8 shows that *the dynamic strategies outperform both the "stay in cash" and "stay in bonds" strategies.* The average annual outperformance ranges from 188 basis points (difference between the dynamic strategy and "stay in bonds" strategy in Japan) to 541 basis points (difference between the dynamic strategy and the "stay in cash" strategy in the U.S.). These return differences are statistically and financially significant.

Even with this level of financial significance, a skeptic could still ask:

1. Is the outperformance consistent over the sample?
2. Do transaction costs offset the high returns?
3. Does the dynamic strategy offer superior returns after risk adjustment?
4. Is the superior performance really predictable ex ante, or was it just benefit of hindsight?
5. Is the past performance likely to prevail in the future, or will it be "arbitraged away"?

### Consistency

By plotting the relative performance of these strategies we can assess whether the outperformance is stable over time. Exhibit 9 shows the return differences for the world bond market.

The dark line shows the cumulative premium of the dynamic strategy. Because the premium of the "stay in cash" strategy is always zero, this line also shows the relative performance of the dynamic strategy versus the one-month rate. Any time the dark line declines, the dynamic strategy underperforms the cash market.

The other solid line shows the cumulative premium of the "stay in bonds" strategy. The dashed line shows the relative performance of the dynamic strategy versus "stay in bonds." The relative performance of the dynamic strategy over the cash market increases gradually through most of the sample period.

In contrast, most of the dynamic strategy's out-

**EXHIBIT 7**
**PATH OF EXPECTED BOND RISK PREMIUMS IN FOUR COUNTRIES**

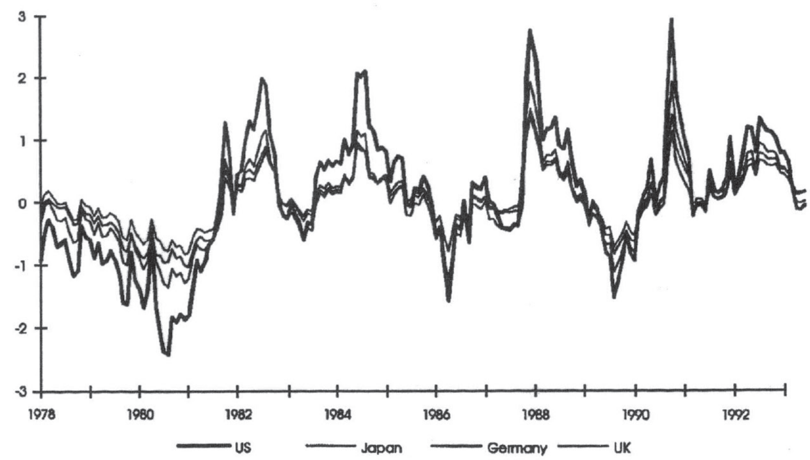


**EXHIBIT 8**
**AVERAGE PREMIUMS OF STATIC AND DYNAMIC BOND MARKET STRATEGIES, 1978-1993**

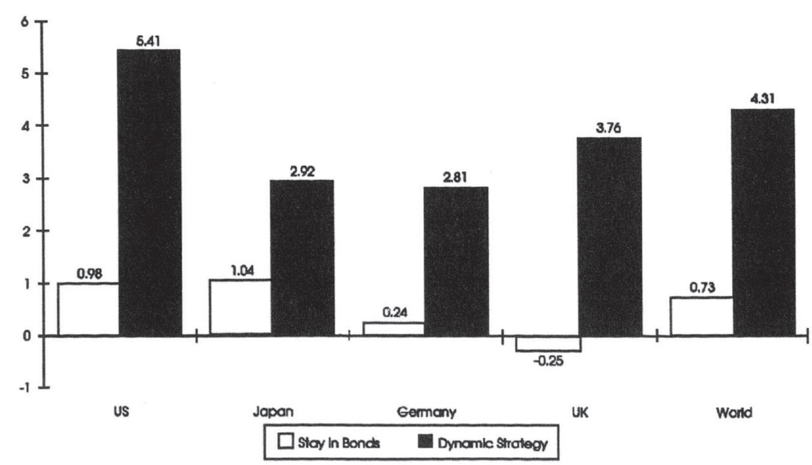

**EXHIBIT 9**
**RELATIVE PERFORMANCE OF A STATIC AND A DYNAMIC STRATEGY WITH WORLD BONDS**

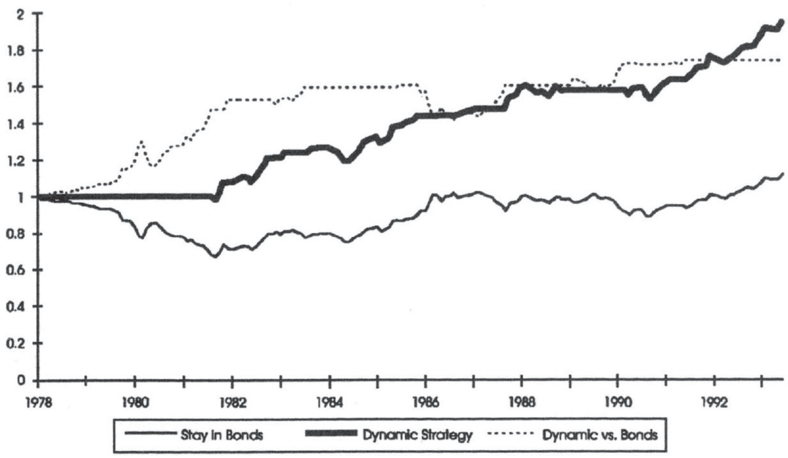


performance relative to bonds occurred before 1982 when the dynamic strategy avoided bonds during the long bear market. (Because this dynamic strategy can, at best, match the bond market performance during a bull market, its performance during the disinflationary 1980s and 1990s must be deemed quite good.)

We also examine the stability of the predictive regressions. The correlations reported in Exhibits 4 and 5 are estimated for the whole sample period 1978-1993. If the correlation between the realized bond risk premium and a predictor variable is stable across shorter subsamples, the relation is less likely to be sample-specific or spurious.

Exhibit 10 shows the thirty-six-month rolling correlations between the realized world bond risk premium and each of the global predictor variables. The correlations are at their highest in the early 1980s and in the early 1990s, near cyclical turning points. Although only "inverse stock" has positive correlation in each subsample, all the rolling correlations are positive in most subsamples — a comforting degree of stability.

### Transaction Costs

In many models that are designed to exploit market opportunities, some apparent anomalies disappear when transaction costs are taken into account. In this case, the opportunities survive after transaction costs are included because the predicted bond premiums move quite slowly over time, and the simple investment strategy requires very infrequent trading. One can confirm from Exhibit 7 that the dynamic strategy switches between the bond and money markets on average fifteen times during the period (that is, the predicted bond risk premiums cross the zero line about once a year). The consequent transaction costs amount to only a few basis points annually.

### Risk-Adjusted Returns

The risk-adjusted performance of the dynamic strategy is even better than its unadjusted performance. Exhibit 11 measures risk in three ways: volatility, downside risk, and "being wrong." The dynamic strategy offers less volatile returns than the "stay in bonds" strategy because it has, at most, the same level of volatility as the bonds, and the rest of the time it is invested in cash.

Various downside risk measures, however, show best the attractiveness of the dynamic strategy.[8] This strategy tries to avoid periods of negative returns — and Exhibit 11 and the histogram in Exhibit 12 show that it has mostly succeeded. The "historical shortfall prob-

**EXHIBIT 10**
**ROLLING 36-MONTH CORRELATION BETWEEN THE REALIZED WORLD BOND RISK PREMIUM AND EACH GLOBAL PREDICTOR**

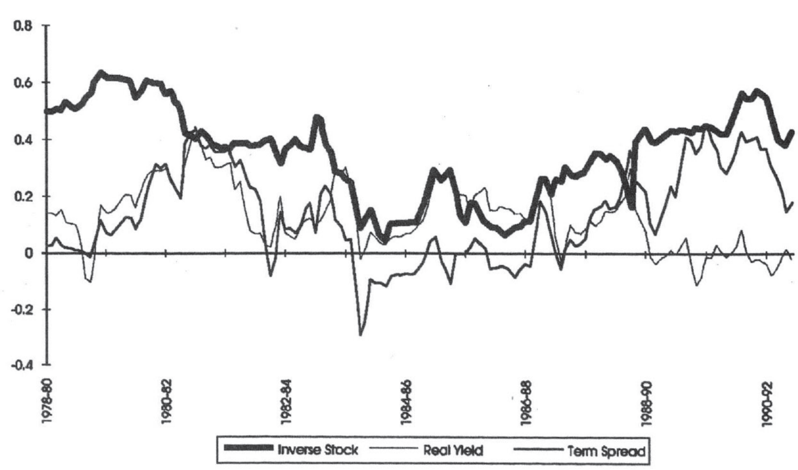

**EXHIBIT 11**
Performance of a Static and a Dynamic Strategy, 1978-1993

| | Stay in Bonds Premium | | | | |
|---|---|---|---|---|---|
| | U.S. | Japan | Germany | U.K. | World |
| Mean | 0.98 | 1.04 | 0.24 | −0.25 | 0.73 |
| Volatility | 10.12 | 6.03 | 5.37 | 9.73 | 6.86 |
| T-Statistic | 0.38 | 0.68 | 0.18 | −0.10 | 0.42 |
| Sharpe Ratio | 0.10 | 0.17 | 0.04 | −0.03 | 0.11 |
| Shortfall Frequency | 0.49 | 0.40 | 0.45 | 0.52 | 0.52 |
| | **Dynamic Strategy Premium** | | | | |
| | U.S. | Japan | Germany | U.K. | World |
| Mean | 5.41 | 2.92 | 2.81 | 3.76 | 4.31 |
| Volatility | 6.37 | 3.85 | 2.81 | 5.75 | 4.40 |
| T-Statistic | 3.34 | 2.99 | 3.94 | 2.57 | 3.86 |
| Sharpe Ratio | 0.85 | 0.76 | 1.00 | 0.65 | 0.98 |
| Shortfall Frequency | 0.19 | 0.15 | 0.16 | 0.19 | 0.18 |
| Frequency of Forecasts with Correct Sign | 0.65 | 0.66 | 0.62 | 0.60 | 0.67 |

This exhibit examines the monthly premiums (excess returns over the one-month rate) of a static and a dynamic strategy in four countries and "the world." The static strategy involves always staying in long-term bonds. The dynamic strategy involves buying one unit of long-term bonds if the expected bond risk premium is positive and none if it is negative. The expected bond risk premium is the (in-sample) fitted value from a regression of the realized bond risk premium on three global predictor variables. The exhibit reports the annualized mean and standard deviation of each premium, the t-statistic against the null hypothesis that the average monthly premium is equal to zero, the annualized Sharpe ratio (sample mean over standard deviation), the shortfall frequency (observed frequency of negative premiums, or the underperformance relative to the one-month rate), and — for the dynamic strategy — the frequency of forecasts with correct sign (observed frequency of predicting positive bond risk premium when the realized premium is positive plus observed frequency of predicting negative bond risk premium when the realized premium is negative).

ability" (defined here as the frequency of negative realized monthly risk premiums) for the global dynamic strategy is 18%.

That is, the dynamic strategy with the world bond market underperforms the cash market in only 18% of the months in the sample. The corresponding number for the "stay in bonds" strategy is 52%. Of course, the dynamic strategy spends many months avoiding the bond market and earning zero premium.

A third way to evaluate the riskiness of the dynamic strategy is to compute how often it is wrong about the market direction (up versus down). Exhibit 11 reports the "frequency of forecasts with a correct sign," which tells how often the expected bond risk premium (based on an in-sample regression with three global predictors) has the same sign (+/−) as the subsequent realized monthly bond risk premium. The expected premium has the correct sign in 60%-67% of the months in the sample (in various countries).

Exhibit 13 shows a scatter plot of the realized and expected world bond risk premiums. In this case, 67% of the observations are in the upper-right and lower-left quadrants where the realized and expected premiums have the same sign.

The lower-right quadrant contains 18% of the 186 monthly observations, those causing negative premiums and "shortfall" for the dynamic strategy (the predicted premium was positive and the bond was bought, but the realized premium was negative). In contrast, the 15% of the observations in the upper-left quadrant make the dynamic strategy underperform relative to the bond market but not relative to the cash market (the predicted premium was negative and no bond was bought, so the realized positive premium was missed).

### Ex Ante Predictability

Overfitting or data mining is a serious concern. Generations of academic and professional researchers have examined financial market data, searching for interesting regularities. Because the sample size is limited, this collective effort may result in spurious findings (patterns that appear systematic but really are sample-specific).

One can guard against spurious findings by, first, limiting the number of predictor variables and choosing them based on economic reasoning. Therefore, we use only a few well-motivated predictors. Second, one can examine the subperiod stability of the findings (see Exhibits 9 and 10). Third, one can conduct out-of-sample analysis.

Our results are based on in-sample estimation, which assumes that investors know in advance the statistical relation between the predictors and the subsequent bond returns. This approach probably exaggerates the potential profitability of the dynamic strategy. A more realistic approach is to forecast the next month's bond risk premium using only available historical data (since 1978) in the regression and to update the forecast each month using new data.

**EXHIBIT 12**
**HISTOGRAM OF THE REALIZED MONTHLY PREMIUMS OF A STATIC AND A DYNAMIC STRATEGY WITH WORLD BONDS, 1978-1993**

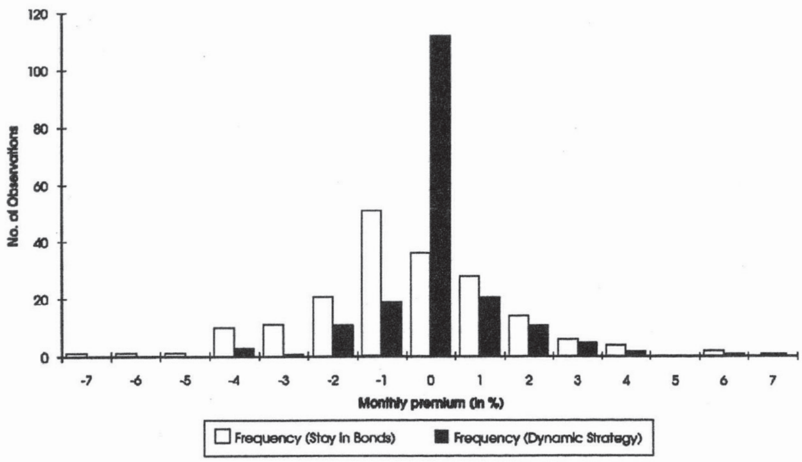


We conduct such an out-of-sample analysis with data starting from January 1990 and ending in April 1994. (We updated the data series for this analysis, adding ten months after June 1993.) Exhibit 14 shows that the dynamic strategy works quite well out-of-sample. It clearly outperforms the "stay in cash" strategy, and it marginally outperforms the "stay in bonds" strategy. This is a good achievement, given that the 1990s have been an exceptionally good time for holding bonds.

While the dynamic strategy may not offer quite as high returns in the future as it has in the past, we expect it to continue to outperform the static strategies. Its risk reduction ability may be even more important than the superior return. Moreover, we have been conservative, and presented the results of a very simple dynamic strategy. Another strategy that also uses the information in the size of the expected bond risk premium, not just its sign, would have performed even better.

## Future Performance

If the superior returns are caused by a time-varying risk premium, the pattern will prevail even if the market is efficient. If the predictability reflects market inefficiency, various impediments to arbitrage still exist.

The dynamic strategy is by no means riskless. It does improve the odds of correctly predicting the market direction, but it predicted the wrong direction in 33%-40% of the months in-sample. The failure rate may well be closer to 50% in the future.

Interestingly, the frequency of correct forecasts increases with the forecast horizon. For traders and arbitrageurs with short horizons, the strategy works "too slowly." The strategy captures business cycle-related variation in expected returns, and its performance should be evaluated using several years of data. Even if the strategy works well in the long run, it may underperform static strategies for a long enough time for a trader to lose his or her job.

Exhibit 15 plots the realized world bond risk premium and its expected value based on out-of-sample estimation; it shows that the dynamic strategy would have missed the bond market rally in the second half of 1993, and would have been long for part of the bear market in spring 1994. Traders whose performance is measured frequently may find even an attractive long-horizon trade too risky. Therefore, the strategy is best suited for institutions that believe in it and that can take the *long-term view.*

**EXHIBIT 13**
**SCATTER PLOT OF THE REALIZED AND EXPECTED MONTHLY WORLD BOND RISK PREMIUM, 1978-1993**

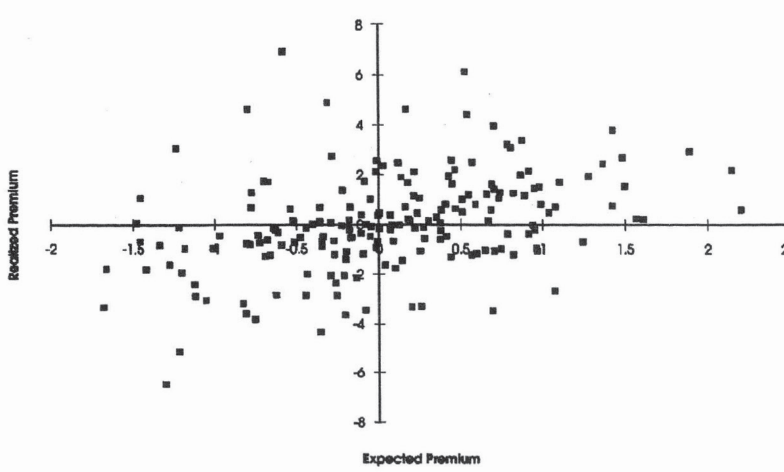

**EXHIBIT 14**
**AVERAGE PREMIUMS OF A STATIC AND A DYNAMIC (OUT-OF-SAMPLE) STRATEGY IN THE 1990s**

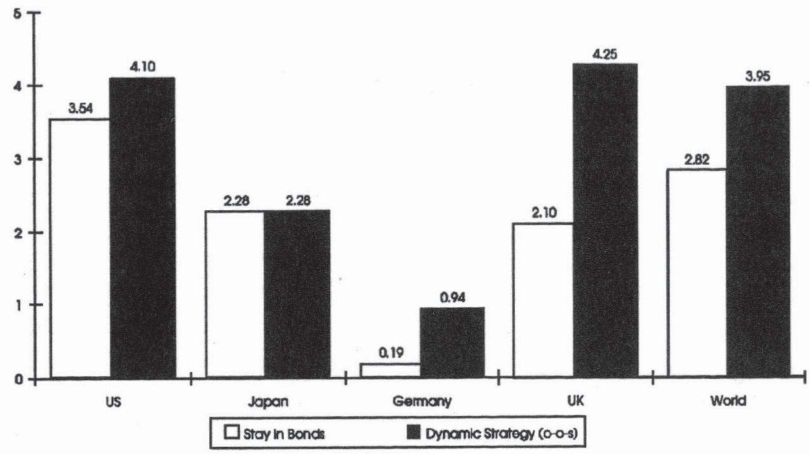


Of course, frequent performance measurement may also make the dynamic strategy too risky for institutional portfolio managers. Even if the strategy has less objective risk than a static strategy, it may have more subjective risk for a portfolio manager. Many managers feel that if they deviate much from their peer group strategies, the downside risk is greater (job loss) than the upside potential. These concerns may produce "herd" behavior and an aversion to innovative strategies.

Even if the predictability is eventually arbitraged away, this may be a slow process. It would probably disappear first from the U.S. market and only later from the other bond markets of the world. Although tactical asset allocation strategies (market timing based on quantitative, no-judgment forecasts) have increased in popularity, only a small minority of market participants use them.

**EXHIBIT 15**
**REALIZED MONTHLY WORLD BOND RISK PREMIUM AND ITS OUT-OF-SAMPLE PREDICTED VALUE IN THE 1990s**

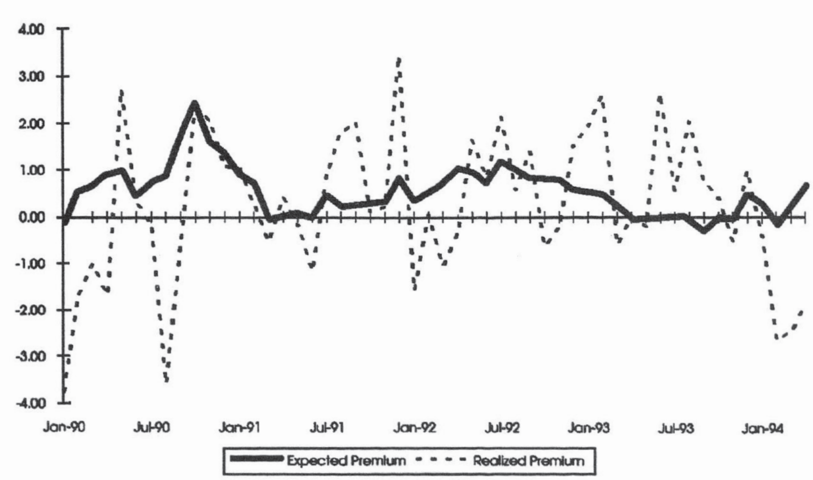


## IMPLICATIONS FOR LONG-TERM INVESTORS

We show that the average reward for bearing interest rate risk is small, but the expected reward varies with economic conditions. Excess bond returns can be forecasted using variables such as past stock market performance, real bond yield, and the term spread. Dynamic strategies that exploit this predictability earn higher returns than passive static-duration strategies, with less risk. These findings indicate that active management of interest rate exposure can enhance portfolio performance.

Of course, the no-judgment model-based forecasting approach is very different from a traditional forecasting approach, which uses fundamental economic analysis or technical analysis in a subjective way. Empirical evidence suggests that the latter approach has not produced exceptional results, at least on average. Perhaps because of some institutional rigidity and/or lack of discipline, the market-timing activities of institutional investors have rarely led to a consistent outperformance over passive strategies.

Investors can implement the dynamic adjustment of interest rate exposure in one of two ways. They can adjust the duration of a benchmark against which portfolio managers are measured, or they can set aside a small part of the portfolio for this strategy, and invest the remainder as before. The simplest implementation method is: Each month, buy long-term bonds if the expected premium is positive; stay in cash if it is negative. This strategy would lead to very infrequent trading, as Exhibit 7 shows, but it also may cause intolerably large jumps in the portfolio duration.

The bond risk premium forecasts may also be used more efficiently than in the simplest method. The performance of the dynamic strategy should be better if investors also use information about the size of the expected premium — not only its sign. That is, a large predicted premium implies a longer duration than a small predicted premium. Another possibility is to overweight (currency-hedged) positions in countries that have the highest expected bond risk premiums.

Whichever implementation is chosen, some key decisions should be made *in advance* — and not changed on the basis of subjective market views or on the strategy's recent performance. These include: What level of estimated expected bond risk premium triggers a trade? How long a duration and how large a position should be taken at each level? How often should the model be reviewed, for example, to add new predictors?

In choosing the implementation, investors face the familiar risk/reward trade-off. More conservative implementation implies reduced potential for return enhancement and risk reduction benefits. This choice should depend on investor risk tolerance and confidence in the recommended strategy.

## ENDNOTES

This article is based on the author's Ph.D. dissertation at the Graduate School of Business at the University of Chicago. The author thanks his advisors, Eugene Fama and Kenneth French, for helpful discussions. He also thanks Ray Iwanowski, Thomas Klaffky, Martin Leibowitz, John Liew, and Anthony Lynch for their comments and Peg Pisani for editorial assistance. All opinions expressed are those of the author and not those of Salomon Brothers Inc.



[1]Fama and French [1989] and Jones [1992] examine the relation between the term spread and subsequent bond returns in the United States, while Mankiw [1986] and Ilmanen [1995b] examine the same relation in many countries.

[2]The "world" variables are the gross national product-weighted average of the local variables in the four countries, and in Canada and France. The results would be similar if we used bond market capitalization weights because the GNP weights and the market weights are quite similar (roughly, U.S. 50%, Japan 25%, and so on).

[3]A currency-hedged bond investment consists of 1) the purchase of an unhedged foreign bond, and 2) the sale of a foreign exchange forward contract. Investors tend to roll over one-month forward contracts, selling the bond's expected end-of-month value back to their home currencies, leaving them exposed only to a small residual currency risk. The "cost" of the currency hedging equals the differential between the home and the foreign one-month Eurodeposit rates, guaranteed by covered interest arbitrage. Thus, the return of a hedged foreign bond equals its local currency return plus the Eurodeposit rate differential. Consequently, the hedged foreign bond return minus the investor's home-currency one-month rate equals the foreign bond's local currency return minus the foreign one-month rate.

[4]One can counter this argument by pointing out that many bond market participants (pension funds, life insurance companies) view the long-term rate as the riskless rate because it better matches their long-term liabilities. If one-half of the market considers the short rate to be riskless, while the other half considers the long rate to be riskless, the expected bond risk premium may be zero. Only if the latter group is a minority should the expected bond risk premium be positive.

[5]It is better to study bond risk premiums than bond returns, because the former reflect pure reward for interest rate risk while the latter also include the one-month rate (reward for time). Moreover, this one-month rate is known at the beginning of the month and is thus fully predictable, while the bond risk premium should not be predictable at all if the expected premium is constant.

[6]Another way to interpret this number is to recall that the square of the correlation coefficient, $R^2$, measures how much of the variation in the dependent variable can be explained by the variation in the independent variables. Thus, the three global variables together predict 14% (= $0.38^2$) of the monthly variation in the realized world bond risk premium. The other 86% is related to unexpected events during the month.

[7]The correlations may appear low to readers who are used to looking at regressions that try to explain the behavior of a series using *contemporaneous* variables (unexpected news). Because unexpected events have a major impact on each period's bond returns, one should not expect to see very high correlations in our predictive regressions, in which we use only variables that are known at the time of forecasting. In fact, one should expect to see zero correlations if the expected bond risk premium is constant.

[8]Downside risk measures are discussed in Leibowitz and Henriksson [1989] and in Harlow [1991].

[9]We examined this issue using more U.S. data (1953-1992). The three predictors (inverse stock, real bond yield, and term spread) forecast the sign of the bond premium correctly for 60% of the months, 64% of the (overlapping) quarters, 66% of the years, and 70% of the three-year periods in the sample.

## REFERENCES

Fama, E.F., and K.R. French. "Business Conditions and Expected Returns on Stocks and Bonds." *Journal of Financial Economics*, Vol. 25 (1989).

Harlow, W.V. "Asset Allocation in a Downside-Risk Framework." *Financial Analysts Journal*, Vol. 47, 5 (1991).

Ilmanen, A. "Does Duration Extension Enhance Long-Term Expected Returns?" Salomon Brothers Inc, July 1995a.

——. "Time-Varying Expected Returns in International Bond Markets." *Journal of Finance*, June 1995b.

Jones, I.E. "Can a Simplified Approach to Bond Portfolio Management Increase Return and Reduce Risk?" *Journal of Portfolio Management*, Winter 1992.

Leibowitz, M.L., and R.D. Henriksson. "Portfolio Optimization with Shortfall Constraints." *Financial Analysts Journal*, Vol. 45, 2 (1989).

Mankiw, N.G. "The Term Structure of Interest Rates Revisited." *Brookings Papers on Economic Activity*, 1986.

Marcus, A.J. "An Equilibrium Theory of Excess Volatility and Mean Reversion in Stock Market Prices." Working paper #3106, National Bureau of Economic Research, 1989.

Sharpe, W.F. "Investor Wealth Measures and Expected Return." In *Quantifying the Market Risk Premium Phenomenon for Investment Decision Making*. Charlottesville, VA: The Institute of Chartered Financial Analysts, 1990.



# Acknowledgments

In addition to our Editorial Advisory Board, the following have been most helpful in providing reviews of manuscripts submitted to us in recent months:

Michael Rosenberg (Merrill Lynch)
David Beaglehole (Goldman Sachs)