Developments in Asset Allocation Modeling

LUIS M. VICEIRA

The "traditional" approach to designing policy portfolios assumes that expected returns, risk, and real interest rates do not change over time so that short-term and long-term risk properties of asset returns are the same. Thus, target asset allocations are the same regardless of investment horizon and remain constant over time. The "modern" approach, in contrast, recognizes that expected returns, risk, and real interest rates may change over time. This view creates a wedge between the short-term and long-term risk properties of asset returns and implies that target allocations may vary with investment horizon and over time. One implication of this view is that short-term bonds may not be the "safe asset" for long-term investors.

In this presentation, I will discuss the design of policy portfolios for long-term investors.[1] To do so, I will begin by looking at the traditional and modern approaches to policy portfolio design. I will then address return patterns over time and the other side of the coin, risk.

For investors with longer investment horizons, the "safe" investment, T-bills, is not very safe because it leaves investors with an exposure to changing real interest rates.

Traditional Design

In the traditional approach to the design of policy portfolios, the investor starts by formulating a set of capital market assumptions (e.g., beliefs about risk premiums, interest rates, and risk itself) and then assumes that the risk–return trade-off is constant over time. In other words, the investor assumes that bonds have, say, an expected return of 2 percent in excess of the return on Treasury bills and that this expected excess return will remain constant in the future, regardless of economic conditions.

Based on that assumption, the investor formulates the policy portfolio with a target asset allocation by seeing how much each asset class contributes to

[1] Note that this presentation is based on Campbell and Viceira (2005).

Luis M. Viceira is an associate professor of business administration at Harvard University, Cambridge, Massachusetts.

expected excess returns relative to risk in the policy portfolio. When that step is done, the investor can decide on the specific allocations. One way to formalize this analysis is to use the mean–variance approach because it indicates what kind of tilts the capital market assumptions are introducing into the target policy portfolio. The investor can then adjust these allocations depending on other constraints and for "risk outside the model."

The key point is that in the traditional approach, no distinction is made between short-term and long-term risk properties of asset returns. One set of numbers summarizes risk at all horizons. And the underlying belief is that the target asset allocation should be constant over time and independent of investment horizon.

Figure 1 is an illustration of this traditional approach, with expected return on one axis and risk on the other. Suppose an investor is trying to allocate among stocks, bonds, and cash. The investor formulates his or her capital market assumptions, which then indicate where these three asset classes are on the return–risk spectrum. Based on that analysis, the investor considers all portfolios that combine stocks and bonds so as to find the optimal mix. Then, depending on the degree of risk aversion, the investor includes cash—perhaps none for the risk tolerant, perhaps a good bit for the risk averse. The point is that the policy portfolio is independent of the investment horizon and that cash is seen as the safe asset regardless of the investment horizon.

Figure 1 Risk–Return Trade-Off, Traditional View

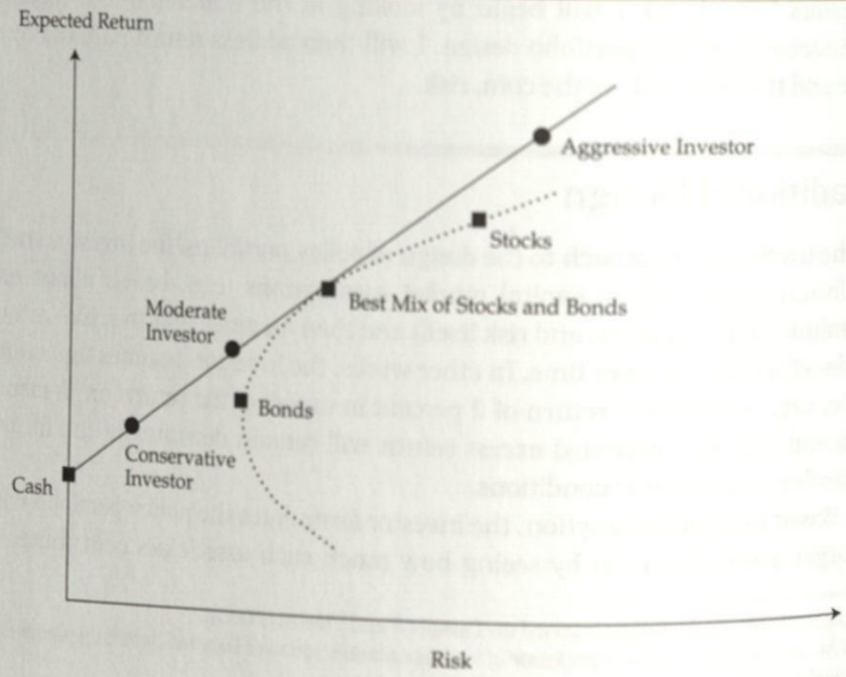

Modern Design

Some institutions have been experimenting with policy portfolio design, what I call the "modern approach." Some large endowments, large foundations, and large pension funds have adopted a practice by which they revise periodically their capital market assumptions. Accordingly, they also revise the target allocations for their policy portfolios. These revisions are based on their belief that investment opportunities, real interest rates, and expected returns on equities and bonds change over time. Thus, they change their policy portfolios over time too. Note that these institutions are making gradual changes in their policy portfolios. It is not a tactical asset allocation program, which calls for high-frequency trading in and out of asset classes. They are periodically revising their capital market assumption based on current market conditions and introducing gradual changes to their policy portfolios.

Return Patterns

Because these changes in the policy portfolio depend on changes in expectations about returns, it is important to look at history to see how these patterns change over time. Changes in nominal dividend growth and nominal earnings growth are generally driven by changes in inflation and have been stable over 10-year periods in real, inflation-adjusted terms. For example, dividend growth per share (DPS) in the S&P 500 Index over the past 10, 30, or 84 years has been, in nominal terms, 4.0 percent, 5.8 percent, and 4.4 percent, respectively, as shown in **Table 1**. It was on the high side for the past 30-year period because this period captures the high inflation of the late 1970s and early 1980s. In contrast, in recent times, it has been in the 3.0–3.5 percent range. It is the same story for EPS. In nominal terms, it has been, respectively, 6.7 percent, 6.5 percent, and 5.0 percent for the same 10-, 30-, and 84-year periods. Earnings may be negatively affected in the short term, but they always end up catching up with inflation. So, when projecting future dividends or earnings growth,

Table 1 Dividend and Earnings Growth of S&P 500, Periods Ending 2004

Period	DPS	EPS
10 years	4.0%	6.7%
30 years	5.8	6.5
84 years	4.4	5.0

Figure 2 S&P 500 Price/10-Year Average Earnings, 1881–2005

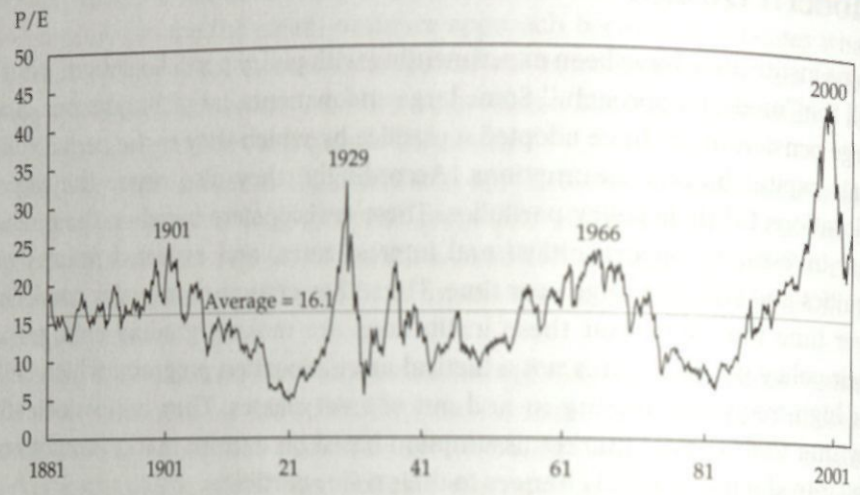

an investor can reasonably use long-term historical values, unless the investor believes things have changed significantly from the past.

P/E multiples also need to be examined when making capital market assumptions. **Figure 2** shows the ratio of the price of the S&P 500 to the 10-year moving average of earnings from 1881 to 2005. By this measure, the P/E of the S&P 500 averaged a little more than 16 times in this period. Big jumps from this average occurred in 1901, 1929, 1966, and 2000, but in general, the P/E moved slowly around the average.

If you are someone who thinks that the P/E multiple of the market will revert to its 100-year average, then you must ask yourself what has to happen for it to revert in, say, 10 years. The answer is that either the denominator must grow (i.e., earnings) or the numerator must fall (i.e., prices). So, earnings would have to grow at their long-term average of 5 percent a year, or 60 percent in total, over the next 10 years while stock prices stay flat. Alternatively, prices would have to fall by 40 percent, as in 1929, 1966, and 2000, to get back to the average. Of course, a combination of growth in earnings and a fall in prices could also happen. This type of thinking is precisely what modern portfolio policy design has in mind.

The dividend/price (D/P) multiple, shown in **Figure 3**, is another factor to be examined. For the 1881–2005 period, it averaged 4.5 percent. And again, by and large, it has hovered around the average, although it is currently at roughly 1.7 percent. When making predictions, one needs to decide whether D/P will remain below its historical average or will mean revert, with resulting implications for dividend and price movements.

Figure 3 S&P 500 10-Year Average Dividend/Price, 1881–2005

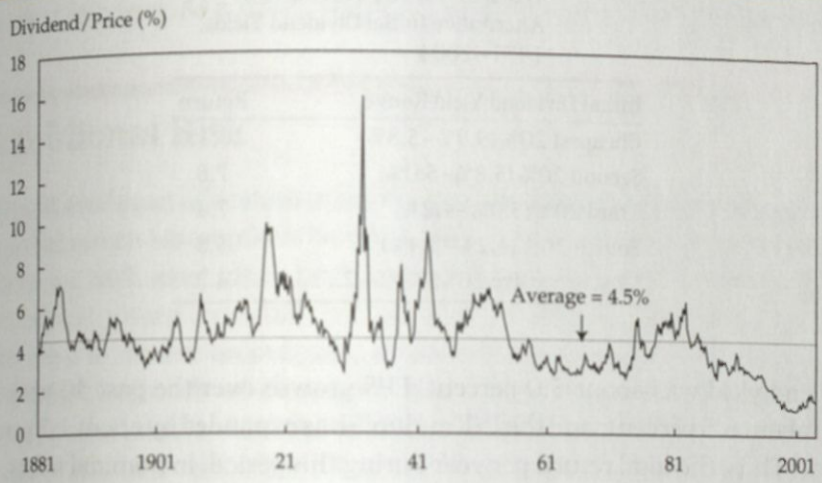

Table 2 Future 10-Year Real Rates of Return When Stocks Are Purchased at Alternative Initial P/E Multiples, 1871–2004

Initial P/E Range	Return
Cheapest 20% (5.6×–10.1×)	11.6%
Second 20% (10.1×–12.7×)	8.1
Third 20% (12.7×–14.9×)	6.8
Fourth 20% (14.9×–17.9×)	4.1
Most expensive 20% (17.9×–26.6×)	4.7

In fact, the level of multiples provides important information about prospective returns. **Table 2** reports the subsequent 10-year real return on the S&P 500 for different current P/E ranges. On average, when the market is cheap (P/E range of approximately 6 to 10 times), it does very well for the next 10 years. When the market is expensive (P/E range of approximately 18 to 27 times), it does not do well. Right now, the market is in the most expensive category, unfortunately. The same pattern holds for dividend yields, or D/P, shown in **Table 3**. That is, when markets are expensive based on dividend yields, returns over subsequent periods are below average.

This type of analysis forces investors to consider how their expectations for risk and return and how their allocations to various asset classes might change over time depending on how expensive or cheap the markets are. This analysis can be extended in the following manner. Consider just three projections based on what has happened over the past 30 years: 1975 to 2004. In 1975, the

Table 3 Future 10-Year Real Rates of Return When Stocks Are Purchased at Alternative Initial Dividend Yields, 1871–2004

Initial Dividend Yield Range	Return
Cheapest 20% (9.9%–5.8%)	10.1%
Second 20% (5.8%–5.0%)	7.8
Third 20% (5.0%–4.2%)	7.4
Fourth 20% (4.2%–3.4%)	5.8
Most expensive 20% (3.4%–2.7%)	4.3

dividend yield was about 5.0 percent. EPS growth over the past 30 years has been about 6.5 percent, and the P/E multiple has expanded by about 1.7 percent a year. Thus, the total return per year during this period, in nominal terms, was about 13.2 percent.

An investor could repeat this exercise by looking at the current yield on stocks, which is around 1.7 percent, and inserting his or her own views about the other two components of expected stock returns—EPS growth and changes in P/E multiples. For example, one could reasonably assume that the kind of hyperinflation seen in the late 1970s and early 1980s will not happen any time soon. So, reducing expectations for nominal growth of EPS from the historical 30-year average of 6.5 percent to 5.0 percent may be reasonable.

The critical question at this point becomes: What will happen with P/Es? If the investor thinks multiples will remain at their current level of about 26 for the next 10 years, then the return expectation is 6.7 percent (5.0 percent earnings growth plus 1.7 percent dividend yield). But if the investor thinks multiples will double again, then about 2.0 percent a year from multiple expansion needs to be added as a source of return. And finally, if the investor thinks multiples will contract down to their historical level, then multiples contraction will be a subtraction from the expected return. For example, if the institution expects that multiples will go from 26 to 16, then it must subtract about 1.7 percent a year from returns on an annual basis over the next 30 years, which would produce an expected return of 5.0 percent. If that is the investor's view, then it would probably make sense to reduce the allocation to equities. Note that this action does not mean that the investor is trying to time the market on a tactical, short-term basis. Rather, the investor is simply trying to base his or her actions on a plausible scenario for expected returns in the foreseeable future.

In a world of low expected returns, managing the risk exposure becomes more important than ever. The institutions I referred to previously manage their risk exposures by slowly moving their target asset allocations. Basically, they believe that there is some low-frequency, slow-moving asset return predictability. That is, they are not timing the market in the sense of predicting what it will

do over the next month, but they may have a sense of what will happen over the next 10 years. As a result, they change their asset allocations accordingly.

Conditional Risk

In traditional asset allocation analysis, investors use constant, or unconditional, expected return assumptions based on long-term historical patterns for risk and return for each asset class. In the model described herein, investors establish conditional return expectations that may be different for each time period. This view of return, accordingly, forces investors to take another look at the traditional approach of measuring risk as the unconditional volatility and correlation of asset returns. Specifically, if there is predictability in asset returns, then not all of the unconditional volatility of asset returns should count as risk because part of this volatility is caused by predictable changes in asset returns. Risk should be measured by the conditional volatility of asset returns at the investor's horizon. For example, depending on the investment horizon, stocks may be less or more risky and bonds may be more or less risky. Similarly, cash may be riskless at certain horizons and risky at others. The point is that what we call a "term structure of the risk–return trade-off" exists, as described in Campbell and Viceira (2005), and accordingly, the optimal policy portfolio (the target policy portfolio) might be different across investment horizons.

To capture horizon effects on risk, we used a simple, flexible statistical model that essentially describes the dynamic behavior of asset returns. We used a well-known model from time-series econometrics—a first-order vector autoregression model, or VAR(1) model—and applied it to asset allocation analysis. The idea is simple. Start with the set of asset classes that you are interested in. Then, add a set of variables that you think is relevant for forecasting returns, such as P/E and D/P, the level of interest rates, or the slope of the yield curve. These are all variables that seem to have some forecasting power for stock and bond returns. Once you have this information, construct forecasts of returns by regressing each asset class return on its own lagged value, the lagged values of other asset returns, and the lagged values of the return-forecasting variables. The model is then used to extract the volatility and correlation structure of asset returns at different horizons. The technical details are included in an appendix to Campbell and Viceira (2005).

Figure 4 plots the annualized standard deviation of real (or inflation-adjusted) returns for horizons between 1 year and 50 years for four investment strategies. The first is an all-equities strategy, and the second represents a five-year constant-duration bond portfolio strategy. The third strategy is to be fully invested in T-bills. The final strategy is to buy and hold bonds with maturities equal to the investment horizon. Notice that when the risk–return trade-off is

Figure 4 Annualized Volatility of Real Returns across Investment Horizons

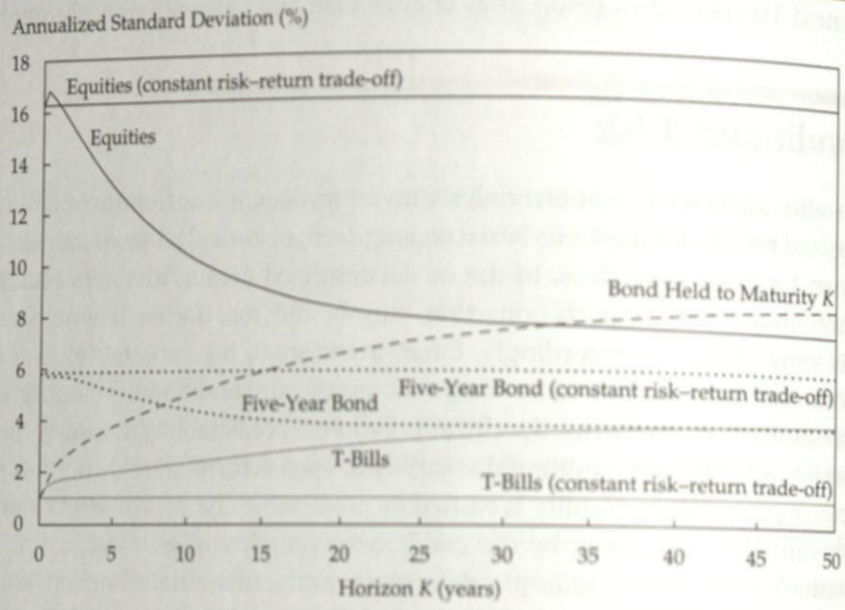

constant (i.e., when future returns are not predictable and interest rates are constant), the annualized standard deviation of real returns is independent of investment horizon and should be the same at short and long horizons. Under the traditional approach to asset allocation, one would use this constant level of risk, which corresponds to the horizontal lines in the plot. But this figure tells a completely different story. Over a short-term horizon, equities have a 16 percent annualized standard deviation, but for a 20-year horizon, they have approximately an 8 percent standard deviation. This level of risk is still quite significant and fully contradicts the claim that stocks are riskless in the long term. But it does represent a reduction in risk at long horizons relative to short horizons. This reduction in risk is caused by stock return predictability, or mean reversion. Thus, the gradual tendency of stock returns to exhibit mean reversion reduces the risk of stocks for investors who have long-term horizons relative to those who have short-term horizons.

Figure 4 also shows that volatility for a five-year constant-duration bond strategy declines as the horizon increases, but not as much as for equities. Two partially offsetting effects are responsible for this pattern. First, a steepening of the yield-curve forecasts positive bond returns in the future, which creates long-term mean *aversion*, or an amplification of volatility at long horizons. Second, an increase in nominal interest rates causes bond prices to decline immediately. But over the long term, nominal interest rates exhibit a mean-reversion pattern, which ultimately reduces volatility. This mean-reversion effect dominates the

mean-*aversion* effect of yield-curve steepening, resulting in this pattern of slowly declining volatility in bond returns. The net effect is that bonds are less risky at long horizons.

Compare this pattern with that of T-bills, for which annualized standard deviation increases with the investment horizon. Because this result is completely counterintuitive, it must be explained. Certainly, T-bills provide a good hedge against inflation risk because their short-term maturity implies that their rates adjust frequently and quickly reflect changes in expected inflation. But the downside of this frequent adjustment is that T-bills do not protect against unexpected declines in real interest rates. The reinvestment risk of T-bills is actually quite pronounced at long horizons and causes the volatility of T-bill returns to increase from 1 percent a year for short horizons to 4 percent a year for long horizons. As tangible proof, consider the situation of those who retired in the early 1980s with their investment funds in short-term investments yielding in the high teens. These pensioners, who have a 20- to 25-year horizon, have come to realize their exposure to fluctuations in real interest rates the hard way because significant declines in real rates (and not just inflation) during the 1990s and early 2000s hampered their ability to spend out of the principal of their investments.

A similar analysis explains the risk of buying and holding a nominal bond to maturity. With bonds, investors receive regular coupons. But the purchasing power of those coupons depends on how inflation moves in the meantime. Buying and holding bonds subjects an investor to inflation risk. Figure 4 shows that over long horizons, inflation risk is quite pronounced.

As can be seen in **Figure 5**, an interesting correlation pattern exists between equities and a constant-duration bond portfolio for varying investment horizons. At short horizons, the correlation of stocks and bonds is quite low. But for medium investment horizons, the correlation becomes quite large. Perhaps even more surprisingly, for very long investment horizons, it declines significantly. The question is: What is driving these changes? Over the intermediate term, the correlation increases dramatically because of the effect of changes in the discount rates, which tend to move stocks and bonds in the same direction. When nominal short-term interest rates increase, bond returns fall at once but stock returns react with some delay, which explains the low short-term correlation of stocks and bonds. Over an intermediate-term business cycle, however, stocks respond in the same manner, causing a significant increase in their correlation. Over long horizons, economic growth has a far greater impact on stock returns than discount rates. As a result, the return patterns diverge again.

A similar case can be made for the correlation between stock returns and inflation. Over short horizons and intermediate horizons, stocks react negatively

Figure 5 Correlation between Equities and Five-Year Bond Portfolio across Investment Horizons

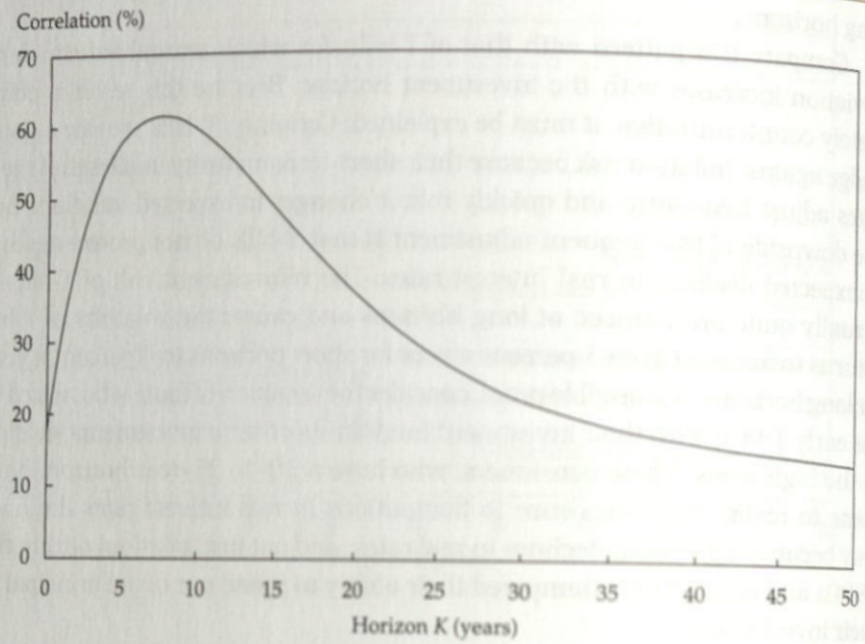

to increases in inflation risk. But over longer horizons, as earnings and dividends adjust for the effects of inflation, stocks become a good inflation hedge.

Thus, for investors with longer investment horizons, the "safe" investment, T-bills, is not very safe because it leaves investors with an exposure to changing real interest rates. The safe asset for a long-term investor thus becomes a bond that matches his or her investment horizon and whose coupons and principal adjust with inflation. In the United States, Treasury Inflation-Protected Securities (TIPS) probably constitute a good proxy for this safe asset. At the same time, the "risky" investment, stocks, appears to be less risky for two reasons. First, long-term volatility in stock prices is less than short-term volatility. Second, stocks provide a better hedge against inflation.

Conclusion

I conclude with some caveats about my analysis as well as some final thoughts. First, the caveats. The dynamic properties of stock and bond returns are extremely difficult to estimate accurately. Thus, asset allocation recommendations under any asset allocation approach (traditional or modern) are sensitive to how the model characterizes future movements in stock and bond returns. Consequently, investors should be aware of this uncertainty and should trim back

extreme positions in stocks and bonds that may be suggested by a particular model.

Nevertheless, empirical evidence shows that the volatility and correlation structure of asset returns can change across investment horizons. These patterns have important implications for portfolio allocations among cash (or T-bills), stocks, and long-term bonds. Real interest rate risk (or reinvestment risk) tilts the composition of minimum variance portfolios at long horizons toward long-term inflation-indexed bonds and away from T-bills. The large positive correlation of bond and stock returns at intermediate horizons and the declining volatility of stock returns at long horizons create bias in the composition of risky portfolios toward stocks and away from long-terms bonds at long horizons. Investors thus might want to adopt policy portfolios whose target allocations vary with their investment horizon and with changing long-term capital market conditions.

This presentation comes from the 2006 Financial Analysts Seminar held in Evanston, Illinois, on 16–21 July 2006.

References

Campbell, John Y., and Luis M. Viceira. 2005. "The Term Structure of the Risk–Return Trade-Off." *Financial Analysts Journal*, vol. 61, no. 1 (January/February):34–44.

——. 2006. "Strategic Asset Allocation for Pension Plans." In *Oxford Handbook of Pensions and Retirement Income*. Edited by Gordon L. Clark, Alicia H. Munnell, and J. Michael Orszag. Oxford: Oxford University Press.

More from Luis M. Viceira

This section presents the speaker's lively question and answer session with the conference audience.

Question: Proposed accounting changes are predicted to make year-over-year liabilities much more volatile for U.S. corporations. Should corporations be looking at one-year time horizons based on that change?

Viceira: If a corporation is sponsoring a pension fund, it has a very long horizon and, therefore, needs to think accordingly. A properly chosen 100 percent allocation to fixed income will certainly hedge liabilities that are derived from accrued pension obligations. But this means that the corporation will now have to fund out of operations future pension obligations as they accrue. These obligations typically grow with productivity and inflation, and over long horizons, stocks tend to grow with productivity and inflation.

It is not entirely clear to me that stocks should be completely discarded from the investment portfolio of a traditional pension fund. Moreover, I am not sure that adopting a 100 percent fixed-income portfolio in a traditional defined-benefit pension fund is in the best interests of the shareholders of the corporation sponsoring the plan, as Bill Sharpe and Jack Treynor pointed out 30 years ago.[2]

Question: For correlations modeled out 50 years, doesn't that put a lot of stress on the data period for which you are making your estimations?

Viceira: Absolutely. That's why we look more to the 20- to 25-year horizon. I don't think we can see much beyond that.

Question: Have you looked at the long-term volatility of TIPS?

Viceira: In my view, TIPS are the truly riskless asset. If you are an investor with a 10-year horizon, you can buy a 10-year TIPS, and that TIPS will save you from an inflation risk that can be quite large over 10 years. It also saves you from real interest rate risk because you now fix the real interest rate. So, this is truly the safe asset—not cash. But keep the trade-off in mind. TIPS have much lower yields.

Question: What questions would you ask your clients to determine what their actual risk preferences are in the context of the time horizon concept?

Viceira: You need to understand what their liabilities are and what their needs are. For someone who is 60 years old and thinking about retiring now, we know objectively that he or she has a 20- to 25-year horizon. How much does this person want to leave

[2] For pension fund investing, please refer to Campbell and Viceira (2006).

to his or her children, and how much does he or she intend to spend? If the client doesn't need it all, his or her risk tolerance increases immediately.

Question: Should you be looking at expected earnings rather than trailing realized earnings?

Viceira: You have to base your analysis on projected earnings. The problem is how to project earnings 10 years into the future. I have found that regardless of how you analyze earnings history, it is a remarkably consistent story. Over long periods, earnings consistently grow at a 3 percent real, 5 percent nominal, rate.

So, yes, you want to use earnings projections, but the best projections should be guided by the historical experiences just cited. If you think that our future experience will be different from our past, then by all means use your projected earnings.