

The Arbitrage Pricing Theory Approach to Strategic Portfolio Planning

NOTICE: THIS MATERIAL
BE PROTECTED BY
LAW (TITLE 17, U.S.C.)

The well known Capital Asset Pricing Model asserts that only a single number—an asset's "beta" against the market index—is required to measure risk. Arbitrage Pricing Theory asserts that an asset's riskiness, hence its average long-term return, is directly related to its sensitivities to unanticipated changes in four economic variables—(1) inflation, (2) industrial production, (3) risk premiums and (4) the slope of the term structure of interest rates. Assets, even if they have the same CAPM beta, will have different patterns of sensitivities to these systematic factors.

The central focus of portfolio strategy is the choice of an appropriate pattern of sensitivities. This choice will depend upon the economic characteristics of the beneficiary of the portfolio's income—in particular, upon whether it is more or less concerned with the economic factor risks than the broad average of investors. An organization, for example, may expect to spend less on food than the average investor. To the extent that food prices coincide with general inflation and are somewhat independent of productivity risk, the organization's optimal portfolio could have a lower inflation beta and a higher productivity beta than the broad-based average has.

Such concerns have tactical implications. The organization unconcerned about inflation in agricultural prices would also wish to skew its portfolio holdings out of this sector. Conversely, a sensitivity to energy costs might lead it to skew its holdings in the direction of the energy sector. Implementation of the chosen strategy might be carried out directly by the fund itself, or it may be delegated to select investment managers who follow established investment policy guidelines.

THE ARBITRAGE PRICING THEORY (APT) has now survived several years of fairly intense scrutiny.¹ Most of the explanations and examinations have taken place on an advanced mathematic and econometric level, which means that few persons outside academia have had the time to read them.² Nevertheless, APT has gained the notice of the investment community, and their curiosity will no doubt grow considerably during the next few years as the logical appeal and, more important-

ly, practical implications of APT become apparent. This article aims to accelerate the process by providing an intuitive description of APT and discussing its merits for portfolio management.

Richard Roll is Allstate Professor of Finance and Insurance at the Graduate School of Management of the University of California, Los Angeles. Stephen Ross is Edwin Beinecke Professor of Finance at the School of Organization and Management, Yale University.

This article was originally presented at the Fall Seminar of the Institute for Quantitative Research in Finance at Colorado Springs, Colorado, on October 11, 1983.

1. Footnotes appear at end of article.

The Intuitive Theory

At the core of APT is the recognition that only a few systematic factors affect the long-term average returns of financial assets. APT does not deny the myriad factors that influence the daily price variability of individual stocks and bonds, but it focuses on the major forces that move aggregates of assets in large portfolios. By identifying these forces, we can gain an intuitive appreciation of their influence on portfolio returns. The ultimate goal is to acquire a better understanding of portfolio structuring and evaluation and thereby to improve overall portfolio design and performance.

The Influence of Systematic Factors

The returns on an individual stock in, say, the coming year, will depend on a variety of anticipated and unanticipated events. Anticipated events will be incorporated by investors into their expectations of returns on individual stocks and thus will be incorporated into market prices. Generally, however, most of the return ultimately realized will be the result of unanticipated events. Of course, change itself is anticipated, and investors know that the most unlikely occurrence of all would be the exact realization of the most probable future scenario. But even though we realize that some unforeseen events will occur, we do not know their direction or their magnitude. What we *can* know is the sensitivity of asset returns to these events.

Asset returns are also affected by influences that are not systematic to the economy as a whole, influences that impinge upon individual firms or particular industries but are not directly related to overall economic conditions. Such forces are called "idiosyncratic" to distinguish them from the systematic factors that describe the major movements in market returns. Because, through the process of diversification, idiosyncratic returns on individual assets cancel out, returns on large portfolios are influenced mainly by the systematic factors alone.

Systematic factors are the major sources of risk in portfolio returns. Actual portfolio returns depend upon the same set of common factors, but this does not mean that all large portfolios perform identically. Different portfolios have different sensitivities to these factors. A portfolio that is so hedged as to be insensitive to these factors, and that is sufficiently large and well-proportioned that idiosyncratic risk is diversified away, is essentially riskless.

Because the systematic factors are the primary sources of risk, it follows that they are the principal determinants of the expected, as well as the actual, returns on portfolios. The logic behind this view is not simply the usual economic argument that more return can be obtained only by bearing more risk. While this line of reasoning certainly contains a great truth, its appeal comes more from Calvin than from Adam Smith. There is a far simpler reason why the expected return on a portfolio is related to its sensitivity to factor movements.

The logic is the same as that which leads to the conclusion that two three-month Treasury bills or two shares of GM must sell for the same price. Two assets that are very close substitutes must sell for about the same price, and nowhere in the entire economy are there any closer substitutes than two financial assets that offer the same return. Two portfolios with the same sensitivity to each systematic factor are very close substitutes. In effect, they differ only in the limited amount of idiosyncratic, or residual, risk they might still bear. Consequently, they must offer the investor the same expected return, just as the two Treasury bills or the two shares of the same stock offer the same expected return.

At this point, a bit of mathematics is probably desirable, if not inevitable. Given what we have said so far, it is possible to see that the actual return, R , on any asset—be it a stock, bond or portfolio—may be broken down into three constituent parts, as follows:

$$R = E + b f + e, \quad (1)$$


where

- E = the expected return on the asset,
- b = the asset's sensitivity to a change in the systematic factor,
- f = the actual return on the systematic factor, and
- e = the return on the unsystematic, idiosyncratic factors.

Equation (1) merely says that the actual return equals the expected return plus factor sensitivity times factor movement plus residual risk.

As we have noted, however, there is more than one systematic factor. There are several important ones, and if all of them are not represented, then our understanding of how the capital market works is inadequate. Our basic equation, then, must be expanded to in-

Figure A Returns and Factor Three



corporate multiple systematic factors.

Empirical work suggests that a three or four-factor model adequately captures the influence of systematic factors on stock market returns. Equation (1) may thus be expanded to:

$$R = E + (b_1)(f_1) + (b_2)(f_2) + (b_3)(f_3) + (b_4)(f_4) + e. \quad (2)$$

Each of the four middle terms in Equation (2) is the product of the returns on a particular economic factor and the given asset's sensitivity to that factor.

What are these factors? They are the underlying economic forces that are the primary influences on the stock market. Our research has suggested that the most important factors are (1) unanticipated inflation; (2) changes in the expected level of industrial production; (3) unanticipated shifts in risk premiums and (4) unanticipated movements in the shape of the term structure of interest rates. We will elaborate on this result later. Right now, our task is to show that there is a simple relation between the factor sensitivities of an asset— b_1 , b_2 , b_3 and b_4 —and the asset's expected return, E .

Factor Sensitivity and Asset Returns

Figure A shows a hypothetical plot of Equation (2) using the third factor as an example and holding factors one, two and four at zero. The figure shows the straight-line relation between actual realized returns and movements in factor three for a particular asset. A more sensitive asset—i.e., one with a larger value for b —would have a steeper line, indicating that factor three has a greater influence on its return. Conversely, the plot for an asset with a lower b would be closer to the horizontal; its return would be less affected by movements of the third factor. There is, in fact, nothing to prevent a sensitivity from being negative. If this were the case, then a rise in the factor would cause this asset's price to fall.

Note that a factor return of zero ($f = 0$) does not mean the actual return will be zero. The actual return will in this case equal the expected return, E . The factor movements represented by f are unanticipated. Any anticipated changes have already been incorporated into the expected return on the portfolio, E . Thus f stands for the deviation of the actual factor return from its expected return. When it is zero, actual factor movements have been just as was expected, and actual portfolio returns will be just what investors had expected. Put simply, if there are no surprises in factor movements, then there can be no surprises in portfolio returns.

Figure B illustrates the relation that must hold between expected return, E , and sensitivity, b . Here point A represents a riskless asset, perhaps a short-maturity bond, with an expected return, r , of 15 per cent. Points B and C represent two stocks with, respectively, expected returns of 20 and 35 per cent and sensitivities of one and two.

A portfolio that is evenly divided between the bond A and stock C will have a return that is a simple average of the returns of the two constituent assets:

$$\begin{aligned} E &= 1/2 \times 15\% + 1/2 \times 35\% \\ &= 25\%. \end{aligned}$$

The sensitivity of this portfolio will also be halfway between the sensitivities of bond A and stock C:

$$\begin{aligned} b(3) &= 1/2 \times 0 + 1/2 \times 2 \\ &= 1. \end{aligned}$$

This portfolio is plotted as point P in Figure B.

Figure B Expected Return and Exposure





Figure C Actual Returns: Stock B vs. Portfolio P



Note that P lies directly above stock B. Consider what this means. A portfolio of bond A and the higher risk stock C has the same sensitivity to systematic factor risk as stock B. But, although the portfolio has the same sensitivity as stock B, it has a higher expected return—25 per cent, versus an expected return of only 20 per cent for stock B.


More importantly, no matter what value factor three happens to take, the portfolio's return will dominate that of stock B. Figure C displays the actual returns on the portfolio and on stock B in relation to the factor three return. Regardless of the outcome (and remember that the actual outcomes cannot be known in advance), portfolio P does 5 per cent better than stock B. The situation presented is the very same sort of arbitrage opportunity that would occur in the bond market if two Treasury bills with the same maturity sold at different yields. It is the same sort of situation that foreign exchange traders exploit when the dollar/mark price differs from what a dollar could buy if it were first exchanged for marks. In well-functioning capital markets, such opportunities exist only momentarily, until they are closed by traders whose reward comes from eliminating such gaps.

When this arbitrage takes place, with investors reducing their holdings of stock B and covering themselves by purchasing the portfolio, the price of stock B falls and that of stock C rises. At the lower price, stock B becomes more attractive relative to stock C. This process terminates only when the portfolio and stock B offer the same expected return. In fact, as in the foreign exchange market or in the bond market, the process works sufficiently rapidly that a gap would probably be too fleeting for an outside investor even to notice. Arbitrage opportunities will no longer exist only when all three assets in Figure B lie on the same line; in any other case, there will always be another portfolio that beats (or is beaten by) one of the assets, no matter what unanticipated developments come to pass.

Figure D plots the line on which all three assets must fall. As we have drawn it, there is a direct positive relation between the expected return, E , on any portfolio or individual asset and its risk sensitivity, $b(3)$. The slope of this line measures the market price of this type of risk.

In Figure D, the price of risk for factor three is displayed as the difference between the expected

Figure D Equilibrium Expected Returns



ed return at a sensitivity of one and the riskless return. As the riskless rate is 15 per cent and assets with a factor three sensitivity of one have a 24 per cent return, the market price of risk is 9 per cent ($24\% - 15\%$). This means that any asset with a $b(3)$ of one—i.e., any asset whose return rises or falls by 1 per cent whenever the third factor rises or falls by 1 per cent—will have an expected return 9 per cent above the riskless return of 15 per cent. An asset that is more sensitive will have a higher expected return; for example, the return for an asset with a b of two is 33 per cent ($15\% + 2 \times 9\%$). In other words, the price of risk for factor three of 9 per cent is the rate at which the investor is rewarded for assuming a unit of sensitivity to movements in this factor.

In summary, the expected return on any asset is directly related to that asset's sensitivity to unanticipated movements in major economic factors. If we let E_3 stand for the return on a portfolio with a sensitivity of one to factor three (E_3 equals 24 per cent in the example of Figure D), then the total expected return (E) on the portfolio may be computed as:

$$E = r + (E_1 - r)(b_1) + (E_2 - r)(b_2) + (E_3 - r)(b_3) + (E_4 - r)(b_4). \quad (3)$$

This equation simply states the relationship we have proved: The expected return on any asset, E , exceeds the riskless return, r , by an amount equal to the sum of the products of the market prices of risk, $E_i - r$, and the sensitivities of the asset to each of the respective factors.

Examining the Factors

We have defined sensitivities as the responses of asset return to unanticipated movements in economic factors. But what are these factors? If we knew them, we could measure directly the sensitivities of individual stocks to each. We could, for example, attribute a particular fraction of the observed price movements in a given stock to movements in the economic factor.

Unfortunately, this is much more difficult than it sounds. To begin with, any one stock is so influenced by idiosyncratic forces that it is very difficult to determine the precise relation between its return and a given factor. At a more practical level, we have so much more data available on individual stock returns than we have on broad economic factors that this approach would be very inefficient. It would be a bit like attempting to see what happens to the yield on a Phoenix Power and Light bond when the money supply changes. A much better approach would be first to determine the impact of an index of municipal bond yields on the Phoenix bond; this can be done with considerable accuracy. We can then see how sensitive bond yields as a whole are to money supply changes. The sensitivity of the Phoenix bond to the money supply can then be determined as the product of these two sensitivities, each of which can be measured with some precision.

The biggest problem in the measurement of sensitivities, however, is separating unanticipated from anticipated factor movements. The b s measure the sensitivity of returns to *unanticipated* movements in the factors. By just looking at how a given asset relates to movements in the money supply, we would be including the influence of both anticipated and unanticipated changes, when only the latter are relevant. Anticipated changes are expected and have already been incorporated into expected returns. The unanticipated returns are what determine the b s, and their measurement is one of the more important components of the APT ap-

proach.

What economic factors relate to unanticipated returns on large portfolios? As noted above, empirical research indicates that the following four economic factors are relevant:³

- (1) unanticipated changes in inflation,
- (2) unanticipated changes in industrial production,
- (3) unanticipated changes in risk premiums (as measured by the spread between low grade and high grade bonds), and
- (4) unanticipated changes in the slope of the term structure of interest rates.


It is possible, of course, to think of many other potential systematic factors, but our research has found that many of them influence returns only through their impact on the above four factors. The money supply, for example, is an important variable, but it is not as good a yardstick against which to measure sensitivities, because most of the influence of unpredicted money supply changes is captured by the other variables. For instance, the change in interest rates on a Friday (from before the money supply announcement to after) is an adequate measure of the surprise in the announcement.

It's hardly surprising that the variables listed above were found to be important determinants of market returns. They appear in the traditional discounted cash flow (DCF) valuation formula. Two of them—changes in industrial production and unanticipated inflation—are related to the numerator in the DCF formula, i.e., to the expected cash flows themselves. Expected industrial production is a proxy for the real value of future cash flows. Inflation enters because assets are not neutral; their nominal cash flow growth rates do not always match expected inflation rates.

The other two variables would seem intuitively to be more related to the denominator in the DCF formula—i.e., to the risk-adjusted discount rate. The risk premium measure is an amalgam of investor attitudes toward risk-bearing and perceptions about the general level of uncertainty. The term structure of interest rates enters because most assets have multiple year cash flows and, for reasons relating to risk and time preferences, the discount rate that applies to distant flows is not the same as the rate that applies to flows in the near future.

These variables make intuitive sense, and it also makes sense that they are indeed "systematic." Every asset's value changes when one of

Figure E Sensitivities to Productivity and Inflation Risks



these variables changes in an unanticipated way. Thus investors who hold portfolios that are more exposed to such changes—i.e., that contain assets whose bs are higher on average—will find that their portfolios' market values fluctuate with greater amplitude over time. They will be compensated by a higher total return in the long run, but they will have to bear up under more severe reactions to bear markets.

Strategic Portfolio Planning

No "off-the-shelf" approach to strategic planning is appropriate for all investment funds any more than one size of suit fits all customers. Below, we outline some general considerations that figure into the determination of investment goals.

The Structuring Decision

Traditionally, portfolio strategy is perceived as the choice of the proper mix of stocks and bonds (with real estate and other assets occasionally included). Every portfolio has its own pattern of sensitivities to the systematic economic factors. Stocks as a group and bonds as

another group have different sensitivities to systematic risks, hence the traditional approach may offer a rough solution to the choice of the optimal pattern of risk exposure. But the results can be improved significantly by examining the sensitivity of *each asset* to systematic risks.⁴

The first problem facing the architect of the fund's investment strategy is that of determining the most desirable exposure to systematic economic risks. Altering the mix of stocks and bonds in the portfolio will certainly affect the amount and type of risk exposure, but so will nearly every other purchase and sale decision. The strategist must first choose the desired level of exposure, then appropriate transactions can move the fund toward that desired position.

For example, assume that two of the empirically relevant exposures—to the general level of risk tolerance and to the term structure of interest rates—are held constant and that we are interested in the choice of exposure to inflation risk and to industrial production risk. In Figure E, the horizontal axis depicts the sensitivity, or "exposure," of a portfolio to inflation risk. The vertical axis plots the same portfolio's exposure to production risk. We will refer to these sensitivities as the inflation and productivity "betas," respectively.

The betas measure the average response of a portfolio or an asset to unanticipated changes in the respective economic factors. For example, a portfolio with an inflation beta of one will tend to move up and down by 1 per cent in response to a 1 per cent unanticipated rate of inflation. A beta greater than one, say an inflation beta of 1.5, means that the portfolio's returns are magnified by inflation, with a 1 per cent unanticipated inflation leading to a 1.5 per cent additional return on the portfolio. Similarly, if beta is less than one, unanticipated inflation has a less than proportional impact on the portfolio's returns. A portfolio with a beta of 0.5 will show a 0.5 per cent increase in return for every 1 per cent unexpected inflation. And a portfolio with a beta of zero will, on average, be unaffected by unanticipated inflation. Of course, many assets actually have negative betas and tend to do worse than expected when inflation is greater than expected. A utility stock with an inflation beta of -0.3 loses 0.3 per cent of return for each 1 per cent unanticipated inflation.

In Figure E, point A depicts a large investment fund with an inflation sensitivity of about 0.7 and a production sensitivity of 0.4. Is this a

usual or an unusual pattern of sensitivities? There is no way to answer this question without referring to some landmarks.

One obvious landmark is the origin, O—the point at which both betas are zero. A portfolio at this point would be affected by neither unanticipated inflation nor by changes in expected industrial production. This may seem to be desirable, but it is not necessarily so. For one thing, such a portfolio offers no insurance against unexpected inflation risk; when inflation is greater than anticipated, this portfolio will, on average, not respond. Perhaps more importantly, there is a tradeoff between return and risk exposure. Moving a portfolio to O, where it will not respond to changes in inflation or to productivity, will have an impact on average return.

Point U represents unit sensitivity to both economic factors. A portfolio located at U will increase in value by 1 per cent with either a 1 per cent unexpected inflation or a 1 per cent increase in expected industrial production. The expenditures of an investment portfolio such as a pension fund are probably exposed to the risk of inflation in an adverse way; unanticipated increases in inflation will, on average, increase expenditures. The inflation sensitivity of a portfolio at U will help to offset this. Industrial production, however, could tell a different story. Declines in industrial production will generally be associated with increases in unemployment, which in turn will place greater economic burdens on individuals and corporations. In addition, productivity changes will be associated with changes in the relative prices of the goods and services purchased by the plan sponsor and its beneficiaries, and these may also be adverse. But, rather than helping the fund to insure against these risks, a portfolio with a productivity beta of one actually magnifies them. When industrial production turns down, so too does the return on the portfolio. Whether or not point U is attractive depends upon the particular situation of the fund.

Point B represents the typical pattern of sensitivities for a portfolio of long-term government bonds. Notice that it has a negative beta with inflation and a slightly positive beta with productivity. Investments in bonds are subject to significant adverse inflation effects and are also somewhat sensitive to productivity (although to a far lesser extent than equities). Productivity sensitivity is larger for corporate bonds than for

governments, for obvious reasons.

Point S is the location of a broad-based market index of large, listed stocks. Although this is a useful reference point, it would be wrong to ascribe too much importance to it. The right choice of a pattern of sensitivities for a given fund depends upon a variety of considerations unique to that fund and to the markets in which its beneficiary is a buyer, and these will not generally result in choosing the market index of stocks. The market index should not be ignored, but neither should it be worshipped. It is simply a useful landmark on the horizon, a signpost that is a guide in unfamiliar territory.


APT and the CAPM

We now have the necessary apparatus to relate the well known Capital Asset Pricing Model (CAPM) to APT. The CAPM asserts that only a single number—the CAPM “beta” against a market index—is required to measure risk. As Figure F illustrates, the CAPM beta measures the distance along a ray from the origin through S, where a broad-based market index is located. We assume that portfolio S is the market index used in computing CAPM betas; it could be any of the commonly used indexes, such as the S & P 500.

Portfolio S has a CAPM beta of 1.0 (by construction). Another portfolio, such as Qa, which is located halfway along the ray between O and S, has a CAPM beta of one-half. Similarly, Qb has a CAPM beta of two, because it is twice as far from the origin as S itself. Note that the CAPM beta of any portfolio can be measured by its distance along the ray *relative* to the market index S. The CAPM beta is a relative risk measure.

But there are many portfolios that are not on the ray OS. For instance, portfolios such as B, A and U, all of which have certain desirable properties, are located in the productivity-inflation space *off* the CAPM ray. What are their CAPM betas? It turns out that there are entire families of portfolios with a particular value of the CAPM beta whose members are not on the ray. The dashed lines in Figure F show some of these families. For example, portfolio A is in the family whose CAPM beta is 0.7; but so are all the portfolios along the dashed line that passes through A. There are portfolios in this family, that have no inflation risk (such as Aa) and there are portfolios with no productivity risk (such as Ab). All of them have CAPM betas of

Figure F CAPM and APT Betas



0.7. We doubt very much, however, that most investment managers and clients would regard them as equally desirable.

If S happens to be a mean-variance efficient portfolio, a so-called “optimized” portfolio, then all portfolios whose CAPM betas are the same will have equal returns on average over time. In this sense, the CAPM beta measures the overall desirability of an asset as perceived by the average investor in the marketplace. Even in this case, however, it is not necessarily true that a particular individual or client will consider all portfolios with the same expected return equally desirable. For example, portfolios Aa and Ab in Figure F might have the same long-term expected return, but they are exposed to far different types of risk and neither is preferable for all funds.

Finally, there is usually no reason to think that a particular portfolio such as S, even though it is a broad index such as the S & P 500, is itself optimized. If it is not optimized, then

portfolios A, Aa and Ab will not have equal expected returns, even though they do have equal CAPM betas. Recent empirical evidence has shown unequivocally that most of the commonly used market indexes are not optimized portfolios. Under this condition, the CAPM beta is not even a reliable indicator of expected return and, as we have already seen, it is virtually worthless as a measure of the type of risk to which the portfolio is exposed.

Now consider fund A, located on the exposure terrain with an inflation sensitivity of 0.7 and a productivity sensitivity of 0.4. What should be the strategy for fund A? How should it go about making its strategic investment decision? To put the question another way, where should the fund go to in Figure E? Should it move closer to S, the stock market index? Should it be somewhere between B and S, divided between bonds and stocks? Just "choosing between bonds and stocks" limits the fund to a position along a line between B and S. The strategic decision is clearly much broader than this.

The appropriate choice of risk exposure depends upon the uses to which the income generated by the fund is to be put. Just as different individuals choose to live in different places, different investment funds will choose different patterns of risk exposure.

Analyzing Portfolio Strategy

To choose the optimal pattern of risk sensitivities and move to the best position in Figure E, we must examine the economic situation of the sponsors and the beneficiaries of the fund. To argue that there is one best strategy for everyone—such as "buying the market"—is simply wrong. In the case of pension systems, we might assume that the principal goal is to serve the interests of the beneficiaries by meeting the promised pension benefits with a minimum of additional taxes (if the plan is public) or of corporate contributions (if it is private), but this goal structure may not be appropriate in all cases (for example, for a nonprofit institution, such as a university).

The economic situations confronting the sponsor and the beneficiaries are determined by the markets within which they operate and by the uses to which they put funds. The sensitivity of prices in these markets to overall inflation, for example, is an important determinant of the proper investment policy. The location of the

organization is important as well. A company that employs blue-collar workers in the Los Angeles area has a different pattern of expenditures than a white-collar service firm located in New York.

Although organizations do not constitute a homogeneous group, they all share broad economic concerns. The key questions involve (a) their patterns of expenditures, (b) their other sources of income, and (c) the economic conditions they will face. These questions can be answered by detailed economic study. In the case of a company, for example, central questions would be, "what are its products, its costs and its prospects? How sensitive is it to the business cycle?" In the case of a museum, the study might begin with an examination of the markets for antiquities. How have these markets behaved, and what plans does the museum have for new acquisitions? Also of great importance is the need to meet current and forecast expenditures of a more prosaic sort, such as those related to maintenance and security. Such a study must be continually updated if the fund is to respond to changes in the economic environment and to changes in the goals and operations of the organization. But, even before the initial study is concluded, it will have important implications for strategic portfolio decisions.

Given an economic profile of the organization, one can begin to structure the overall risk exposure of the portfolio. Expenditures on major commodity groups—on salaries and materials, say—should be compared with the general expenditure pattern in the country as a whole. For example, suppose that the organization spent less on food and relatively more on travel than the average investor. The higher expenditure on travel would render it more vulnerable to energy costs than the typical investor, whereas the lower expenditures on foodstuffs would make it less exposed to food prices.

At the strategic level, these considerations will influence the optimal pattern of risk exposure. To the extent that food prices coincide with general inflation, for example, the optimal portfolio could be less hedged against inflation—i.e., could have a lower inflation beta than a broad-based market average has. Similarly, to the extent that food prices tend to be somewhat independent of productivity risk, the organization could accept a higher sensitivity to productivity risk than a broad-based average has. By bearing more risk in this dimension, the

portfolio could expect a higher return.

The influence of these kinds of considerations on the idiosyncratic risk of specific industry groupings has tactical implications. If the organization is unconcerned about inflation in agricultural prices, it would also wish to skew its portfolio holdings out of this sector. Similarly, a sensitivity to energy costs might lead it to skew its portfolio holdings in the direction of the energy sector. An organization will wish to hold a pattern of investments tailored to its own needs. Its optimal portfolio will therefore have a pattern of investments that is modestly skewed from the broad-based market index owned by the average investor.

It should be emphasized that tactical portfolio adjustments can be accomplished without reducing the average return on the portfolio. The strategic decisions determining the level of exposure to systematic economic factors influence the average return, but the tactical decisions can be made without any sacrifice of portfolio return, because they deal merely with idiosyncratic risk.

Implementing the Strategy

To implement the chosen strategy, the fund may direct the investments itself, or it may select investment managers who will follow established investment policy guidelines. The adoption of the APT approach to strategy has implications for the choice and the evaluation of investment managers. If the strategy dictates that investments should be made in particular sectors, then it would be natural to look for managers who specialize in these sectors.

More generally, managers implicitly tend to choose portfolios that have particular patterns of sensitivities to the economic factors. One manager might, for example, focus on high price-earnings ratio companies, so that his portfolio has a characteristic pattern of sensitivities. Another might be heavily invested in utilities, and this would result in a different pattern of sensitivities. The investment strategy for the portfolio as a whole may be implemented by choosing a portfolio of managers in such a way that pooling them together results in the desired pattern of sensitivities. If, for example, Manager A's portfolio typically has an inflation beta of two and manager B's portfolio has an inflation beta of one, then a desired inflation of beta 1.4 for the overall strategy could be achieved by placing \$0.40 with manager A for every \$0.60

given to manager B.

Of course, the complete manager evaluation issue is more complicated than this. Given that a manager has a certain pattern of risk exposures we also want to know whether he or she accomplishes this in the least costly fashion and with the least amount of idiosyncratic risk. This is the subject of performance evaluation, which is well-developed in the APT framework but is beyond the scope of this article.

Finally, a fund's choice of investments will generally be constrained by legal and other considerations. Typical of such constraints is the requirement that all investments be of a certain grade or from an approved list, or that the investments include bonds or equities from a particular issuer. The APT approach to strategy is particularly well-suited to these situations; because of its flexibility, it can be adapted to special situations when many traditional approaches cannot.

For example, suppose that the portfolio is constrained to hold a significant portion of its investments in the equities of the bonds of a particular company or government agency. For two related reasons it will generally be the case that this constraint is binding, in the sense that the fund would rather reduce its holdings of this security. First, the large holding subjects the fund to a substantial amount of idiosyncratic risk and, second, the fund may already be implicitly subject to much of the risk associated with the issuer.

The total risk of this security, however, can be substantially mitigated if the remainder of the portfolio is explicitly selected to offset its influence. If the security in question has a lower than desired sensitivity to inflation risk—e.g., a beta of 0.6 when the desired beta is 0.9—then the influence of the holding on the inflation exposure of the portfolio may be countered by choosing alternative investments with inflation betas in excess of 0.9. As a result, however, the fund may be subjected to idiosyncratic risk, which would not be a problem if the constraints were absent.⁵

Summary

The APT approach to the portfolio strategy decision involves choosing the desirable degree of exposure to the fundamental economic risks that influence both asset returns and organizations. This focus differs from that of traditional investment analysis and is ideally suited to the

management of large pools of funds.

Choosing the optimal degree of risk exposure requires an understanding of the level of risk exposure of the organization. Optimally, the pattern of risk exposure in the fund will balance the organization's current level of risk exposure. The fund should be positioned to hedge the organization against the economic uncertainties it faces.

Implementing this strategy may involve either choosing managers according to their typical pattern of exposure to economic risks and their ability to offer excess returns with low idiosyncratic risk or by choosing assets directly according to estimates of their exposure characteristics and relying upon diversification to remove idiosyncratic risk. The former approach is "active APT," whereas the latter approach may be quasipassive, inasmuch as systematic exposure is planned and implemented but there is no attempt at selection based on anticipated abnormal returns.

Footnotes

1. Arbitrage Pricing Theory was originated by Stephen A. Ross in "The Arbitrage Theory of Capital Asset Pricing," *Journal of Economic Theory*, December 1976, pp. 341-360. Theoretical refinements have been made by the following: Gregory Connor, "A Factor Pricing Theory for Capital Assets" (Working paper, Northwestern University, 1981); Gur Huberman, "A Simple Approach to Arbitrage Pricing Theory," *Journal of Economic Theory*, October 1982, pp. 183-191; Nai-fu Chen and Jonathan E. Ingersoll, Jr., "Exact Pricing in Linear Factor Models with Finitely Many Assets: A Note," *Journal of Finance*, June 1983, pp. 985-988; Philip H. Dybvig, "An Explicit Bound on Deviations from APT Pricing in a Finite Economy," *Journal of Financial Economics*, December 1983, pp. 483-496; Mark Grinblatt and Sheridan Titman, "Factor Pricing in a Finite Economy," *Journal of Financial Economics*, December 1983, pp. 497-507; Robert Stambaugh, "Arbitrage Pricing with Information," *Journal of Financial Economics*, November 1983, pp. 357-369;
2. Empirical testing with equities is described in the following: Richard Roll and Stephen A. Ross, "An Empirical Investigation of the Arbitrage Pricing Theory," *Journal of Finance*, December 1980, pp. 1073-1104; Nai-fu Chen, "Arbitrage Asset Pricing: Theory and Evidence" (Dissertation, Graduate School of Management, University of California at Los Angeles, 1981); Mark R. Reinganum, "The Arbitrage Pricing Theory: Some Empirical Results," *Journal of Finance*, May 1981, pp. 313-321; Patricia Hughes, "A Test of the Arbitrage Pricing Theory" (Working paper, University of British Columbia, August 1981); Lawrence Kryzanowski and Minh Chan To, "General Factor Models and the Structure of Security Returns," *Journal of Financial and Quantitative Analysis*, March 1983, pp. 31-52; and Stephen J. Brown and Mark I. Weinstein, "A New Approach to Testing Asset Pricing Models: The Bilinear Paradigm," *Journal of Finance*, June 1983, pp. 711-743. Tests with Treasury bills are presented in George Oldfield, Jr. and Richard J. Rogalski, "Treasury Bill Factors and Common Stock Returns," *Journal of Finance*, May 1981, pp. 337-350. Issues of testability are discussed by Jay Shanken, "The Arbitrage Pricing Theory: Is it Testable?" *Journal of Finance*, December 1982, pp. 1129-1140; Philip H. Dybvig and Stephen A. Ross, "Yes, the APT Is Testable" (Working paper, Yale University, June 1983); and Gunter Franke, "On Tests of the Arbitrage Pricing Theory" (Working paper, Universitat Giessen, 1983).
3. See Nai-fu Chen, Richard Roll and Stephen A. Ross, "Economic Forces and the Stock Market" (Working paper, University of California at Los Angeles, 1983).
4. The systematic, and idiosyncratic, risks at the heart of the APT approach to investment strategy have been identified in technical econometric work. For assets on which data are available, the pattern of exposure of each asset is known.
5. In some countries other than the U.S., constraints on investments are so stiff as to preclude achievement of the desired overall pattern of systematic risk exposure.