

The Mean–Variance–Optimization Puzzle: Security Portfolios and Food Portfolios

Kenneth L. Fisher and Meir Statman

The Markowitz mean–variance–optimization framework presents a puzzle. Although it is the standard model of portfolio construction, investors rarely use it and, when used, it is constrained so much that portfolios reflect the constraints more than the optimization. Why do investors dislike unadulterated optimized mean–variance portfolios? The typical answer focuses on biases introduced by estimation errors. More important is the fact that investors' goals are quite different from mean–variance optimization. People care about more than cost and nutrition when they construct their diets: They also care about palatability. Similarly, investors care about more than expected returns and variance as they construct their securities portfolios: They also care about the investment equivalent of palatability. We use an analogy between security portfolios and food portfolios to explore the nature of the gap between intuitively appealing portfolios and mean–variance–optimized portfolios.

The Markowitz mean–variance–optimization framework presents a puzzle. As Green and Hollifield (1992) wrote, mean–variance optimization plays an important role in finance and the properties of mean–variance–efficient portfolios are central in both static and dynamic models of asset prices. Yet the practical application of mean–variance analysis has proved problematic. Portfolios constructed using historical means, variances, and correlations do not appeal to investors' intuition. They consist of few of the many available assets, and the weights of the assets included in the portfolios are extreme—large long or large short positions. Investors are suspicious of portfolios with extreme weightings, although many like the idea of mean–variance optimization.

As a result of these suspicions, mean–variance optimization is implemented with extensive sets of constraints such as the prohibition of short positions and the assignment of maximum weights for long positions. The constraints enforce portfolios that appeal to investors' intuition.

Do optimized portfolios constructed with historical means, variances, and correlations display extreme weightings because historical parameters are not good estimates of the expected parameters? This view is common in the literature. For example,

Michaud (1989) argued that mean–variance optimizers are, in effect, "estimation–error maximizers," shifting too much portfolio weight to assets that, because of estimation errors involved in using historical parameters as estimates of expected parameters, show high returns, low variances, and low correlations with other assets.

If estimation errors inherent in historical parameters are the problem, then better estimation is the solution, and many people have offered better methods for estimation. Green and Hollifield, however, found that the extreme weightings that make portfolios unappealing to investors are, in fact, inherent in efficient mean–variance portfolios constructed with proper estimates. Eliminating estimation errors would not eliminate extreme weighting and would not make efficient mean–variance portfolios intuitively appealing. Given their results, Green and Hollifield recommend that investors abandon their intuition about the features of desirable portfolios and accept optimized mean–variance portfolios.

The purpose of this article is to explore investors' intuition about features that make portfolios appealing. A recommendation that investors concern themselves only with expected returns and variances and abandon their intuition about features that make portfolios appealing is equivalent to a recommendation that people concern themselves only with cost and nutrition when choosing a diet and abandon their intuition about palatability, variety, prestige, and conformity to culture. This advice is neither good nor is it advice investors would follow.

Kenneth L. Fisher is chairman, CEO, and founder of Fisher Investments. Meir Statman is a professor in the Department of Finance at the Leavey School of Business and Administration at Santa Clara University.

The intuition of investors combines good and bad. The good is that intuition reflects the investors' true objectives, objectives that might well be different from mean-variance optimization. Vegetarians avoid meat even if it is nutritious and low in cost, and some investors avoid junk bonds even if junk bonds are high in expected return and low in risk. The bad part is that intuition is likely to lead investors to inferior portfolios, portfolios that they themselves would identify as inferior when presented with alternative portfolios. The diet of the typical vegetarian is probably not the lowest-cost diet that conforms to vegetarian rules and nutrition demands. Similarly, the typical portfolio of an investor who wants to avoid junk bonds is probably not the one with the highest expected return among all portfolios that exclude junk bonds and have the desired level of risk. If mean-variance optimization is useful, it is useful in the task of uncovering portfolios investors themselves accept as superior to portfolios they choose by intuition alone. We build on the foundation of behavioral portfolio theory by Shefrin and Statman (1995) and on Statman (1982).

ESTIMATION ERRORS

The common view in the literature is that the unintuitive nature of optimized mean-variance portfolios constructed with historical mean variances and covariances is the result of estimation errors. Specifically, historical parameters might be poor estimates of expected parameters. Consequently, efforts are directed at estimation techniques that might reduce estimation errors (see, e.g., Frost and Savarino 1986, 1988; Black and Litterman 1992; Chopra, Hensel, and Turner 1993; and Chopra 1993).

Frost and Savarino (1988), using the capital asset pricing model (CAPM) equation, in which expected returns of stocks are determined by beta rather than by historical returns, estimated expected returns of stocks. They found that mean-variance optimization leads investors to plunge into a small subset of the 200 stocks from which investors in their study can choose.

Frost and Savarino attributed the tendency to plunge to biases in the estimation of expected returns, even though they began with a technique designed to eliminate such biases. Then, rather than offering a better way to estimate expected returns, they resorted to the brute force of constraints by imposing maximum and minimum values on the allocation that each stock can receive in the portfolio. In particular, they disallowed short sales.

Chopra, Hensel, and Turner compared optimized mean-variance portfolios in which expected

returns, standard deviations, and correlations are taken as their historical averages over a period of five years to optimized mean-variance portfolios using Stein estimators. Stein estimators are designed to move estimates of historical returns closer to their expected values. The assets in the Chopra, Hensel, and Turner study are indexes of stocks, bonds, and cash for the United States, as well as Canada, Germany, Japan, the United Kingdom, and Australia.

Chopra, Hensel, and Turner used three versions of Stein estimation. In the first, the expected returns of each stock index are estimated as the global mean of historical returns of all stock indexes. Similarly, the expected return of each bond index is estimated as the global mean of historical returns of all bond indexes. The second version of Stein estimation follows the first, but in addition, correlations between each pair of stock indexes are estimated as the global mean of all correlations between stock indexes. Correlations between each pair of bond indexes are estimated as the global mean of the pairs, and correlations between each pair of stocks and bond indexes are estimated as the global mean of the pairs. The third version of Stein estimation follows the second, but in addition, variances are estimated as the global means of their groups.

The authors then compared the performance of their Stein estimation portfolios to the performance of a "passive" portfolio constructed without benefit of a mean-variance optimizer. The passive portfolio was 60 percent in stocks and 40 percent in bonds, and country allocations were equal within each asset group. In all, Chopra, Hensel, and Turner examined five portfolios: an "active" portfolio for which expected returns and other parameters were taken as their historical values over the preceding five-year period, the three versions of Stein estimation portfolios, and a passive portfolio. They examined the returns of the five portfolios over a subsequent 72 months and found that the passive portfolio provided somewhat lower monthly returns than the Stein estimation portfolios, but the passive portfolio had a much lower standard deviation of returns.

Chopra, Hensel, and Turner noted that the Sharpe ratio represents a reasonable method for comparing the risk-adjusted performance of alternative portfolios and concluded that the passive portfolio is best when judged by the Sharpe ratio. Nevertheless, they recommended mean-variance optimized portfolios using Stein estimation rather than the passive portfolio.

Good estimates of expected returns, variances, and covariances are always better than bad

estimates, but good estimates will not bridge the gap between optimized mean–variance portfolios and intuitively appealing portfolios. As noted by Green and Hollifield, the gap is unlikely to be the result of errors in estimation. Moreover, even if the gap were caused by errors in estimation, such errors can never be eliminated. At best, they can be made smaller. Errors come from two sources. First, we do not know with certainty the nature of the equilibrium that governs expected returns, variances, and correlations. Second, even if we were willing to assume a particular equilibrium model, estimates will always deviate from true values because we rely on the available observations of samples rather than on the unavailable observations of the population.

The problem of estimation errors would be minor if small changes in the estimation of the parameters led to small changes in the composition of optimized mean–variance portfolios. The fact is, however, that optimized mean–variance portfolios are extremely sensitive to small changes in the estimation of the parameters (see Best and Grauer 1991). This extreme sensitivity indicates that mean–variance optimization is unlikely to be saved by reasonable attempts at better estimation. To glimpse the degree of sensitivity of mean–variance-optimized portfolios to changes in the parameters, compare two sets of mean–variance-optimized portfolios in which estimates are based on the same data set.

From The Vanguard Group, we obtained the monthly returns of their index funds, or the indexes that the funds mimic, for the 15-year period from January 1980 through December 1994. We used

index numbers for periods when Vanguard index funds were not available and adjusted index returns to reflect current expenses of Vanguard index funds. We also used the monthly returns of The Vanguard Group’s money-market (prime) fund for the same period.

Consider two sets of estimates of the parameters. In the first, the annual set, all the parameters are estimated from annualized returns. In the second, the monthly set, expected returns are estimated from annualized returns but standard deviations and correlations are estimated from monthly returns. Consider the returns of the Vanguard Total Stock Market fund, an index fund that mimics the Wilshire 5000 Index, a broad index of U.S. stocks. The mean return in the annual set over the 1980–94 period was 14.51 percent, and the annualized standard deviation of returns in the annual set was 13.87 percent. The annualized standard deviation in the monthly set, however, was 15.60 percent. The correlation between the annual returns of the Vanguard Total Stock Market fund and the annual returns of the European Portfolio was 0.62. The correlation between their monthly returns, however, was 0.65. Data on returns, standard deviations, and correlations of the Vanguard funds are presented in Table 1.

Optimized mean–variance portfolios with standard deviations of 6 percent, 10 percent, and 44 percent are described in Table 2. Note that, although differences between monthly and annual estimates of the mean–variance parameters are small and have only minor effects on estimates of expected returns, they lead to huge differences in the composition of the portfolios. For example, although the 20 percent

Table 1. Estimates of Mean Returns, Standard Deviations, and Correlations of U.S., European, and Pacific Stock, Bond, and Money-Market Funds

| Fund | U.S. Fund | European Fund | Pacific Fund | Bond Fund | Money-Market Fund |
|---------------------------------------|-----------|---------------|--------------|-----------|-------------------|
| Mean annual return | 14.51% | 16.01% | 19.50% | 10.75% | 8.41% |
| <i>Annualized standard deviation</i> | | | | | |
| Annual data | 13.87 | 22.86 | 29.91 | 8.69 | 3.85 |
| Monthly data | 15.60 | 17.38 | 23.03 | 7.32 | 1.04 |
| <i>Correlation with U.S. fund</i> | | | | | |
| Annual data | 1.00 | 0.62 | 0.30 | 0.39 | 0.02 |
| Monthly data | 1.00 | 0.65 | 0.48 | 0.35 | 0.15 |
| <i>Correlation with European fund</i> | | | | | |
| Annual data | | 1.00 | 0.59 | 0.39 | –0.22 |
| Monthly data | | 1.00 | 0.62 | 0.42 | 0.00 |
| <i>Correlation with Pacific fund</i> | | | | | |
| Annual data | | | 1.00 | –0.05 | –0.13 |
| Monthly data | | | 1.00 | 0.09 | –0.02 |
| <i>Correlation with bond fund</i> | | | | | |
| Annual data | | | | 1.00 | 0.22 |
| Monthly data | | | | 1.00 | 0.45 |

Table 2. Comparison of Optimal Mean-Variance Allocations in Portfolios with Various Standard Deviations, Annual and Monthly Data

| Fund | Standard Deviation: 6 percent | | Standard Deviation: 20 percent | | Standard Deviation: 44 percent | |
|-----------------|-------------------------------|--------------|--------------------------------|--------------|--------------------------------|--------------|
| | Annual Data | Monthly Data | Annual Data | Monthly Data | Annual Data | Monthly Data |
| United States | 25% | 6% | 101% | 24% | 55% | 225% |
| European | -7 | 2 | -44 | 1 | 0 | -104 |
| Pacific | 13 | 18 | 51 | 61 | 135 | 112 |
| Bond | 29 | 34 | 108 | 130 | 293 | 237 |
| Money market | 40 | 40 | -116 | -117 | -384 | -370 |
| Expected return | 11.5% | 11.7% | 19.3% | 19.8% | 33.6% | 32.1% |

standard deviation portfolio using monthly data calls for a 24 percent allocation to U.S. stocks, the 20 percent standard deviation portfolio using annual data calls for a 101 percent allocation. And although the 44 percent standard deviation portfolio using annual data calls for zero allocation to European stocks, the 44 percent standard deviation portfolio using monthly data calls for a 104 percent short position.

What we learn most from the attempts to use estimation errors as a way to reconcile intuitively appealing portfolios with optimized mean-variance portfolios is that investors assign an extraordinary level of importance to their intuition about portfolios and, to a lesser degree, that investors want “scientific” solutions for their portfolios. Perhaps nothing compares better to the desire to have portfolios that are both intuitively appealing and mean-variance optimized than the desire to have meals that are palatable but also nutritious and economical.

SECURITIES AND FOOD

Portfolios of securities are not the only case in which a gap exists between optimized portfolios and intuitively appealing portfolios. Another case is that of diets, or food portfolios. We might gain insight into the nature of the gap between mean-variance-optimized security portfolios and intuitively appealing security portfolios by examining the gap between optimized food portfolios and intuitively appealing food portfolios.

Investors in the mean-variance-optimization framework care only about the expected returns and the variances of their portfolios. They attempt to minimize the variance of their portfolios for any given expected return. Similarly, ranchers in the feed-optimization framework care only about the nutrition and the cost of the diet they provide their cattle. Ranchers attempt to minimize the cost of a diet for any given level of nutrition.

Candler and Heady (1958) described a rancher who chooses from foods such as alfalfa meal at \$66 a ton, fish meal at \$156 a ton, and distillers solubles at \$92 a ton. Nutrients include fiber, protein, and fat. Alfalfa meal contains 25 percent fiber, 17 per-

cent protein, and 2 percent fat; fish meal contains 1 percent fiber, 60 percent protein, and 7 percent fat. The minimum-cost combination that supplies the required nutrients is found by applying the optimization algorithm of linear programming. It turns out that the optimal portfolio of foods consists of 0.14 ton of alfalfa meal, 0.31 ton of distiller solubles, 0.56 ton of soybean meal, and no fish meal.

Extreme portfolio weightings and high sensitivity to small changes in parameters are not unique to mean-variance optimization. Linear-programming optimization solutions, such as the feed solution, often contain extreme weightings, and the composition of portfolios is sensitive to small changes in the parameters. For example, a small drop in the price of fish meal might lead not only to its inclusion in the diet but also to radical changes in the entire composition of the food portfolio. “So what?” asks the rancher. After all, the food portfolio is on the efficient frontier; it provides sufficient nutrients at minimum cost. But is it palatable? Does it offer variety? Is it prestigious? Does it conform to the prevailing culture?

Palatability, variety, prestige, and culture may not matter to cattle—they offer no protest—but people expect their food portfolios to be more than economical and nutritious. People expect food portfolios also to be palatable, varied, prestigious, and conforming to culture. Mean-variance-optimized portfolios are economical and nutritious because they minimize risk for a given level of expected return⁵, but evidently, such portfolios rank low on the investment equivalents of palatability, variety, prestige, and culture.

What does a human food portfolio that is economical and nutritious look like? Stigler (1945) considered 77 foods, from wheat flour to sirloin steak and strawberry preserves. Each food item has a set of nutrients and a cost. The minimum-cost portfolio turns out to be an “extreme” portfolio: It consists of only 5 food items out of the 77. A moderately active man weighing 154 pounds would have satisfied all his nutritional requirements at an annual cost of \$39.93 (as of 1939). The food portfolio consisted of

370 pounds of wheat flour, 57 cans of evaporated milk, 111 pounds of cabbage, 23 pounds of spinach, and 285 pounds of dried navy beans. Are you ready for a life with this optimized portfolio of foods? If you answered no, you are not alone.

Stigler compares his minimum-cost portfolio of foods with the portfolio of foods that dietitians Carpenter and Stiebeling described as one that “gives the cheapest combination of foods that is desirable to use for an indefinite period.” The Carpenter and Stiebeling food portfolio would have cost \$100 in 1939, more than double the cost of Stigler’s food portfolio. Similarly, the food portfolio that dietitian Rose presented as “an unqualifiedly minimum diet” would have cost about \$115 in 1939, almost three times the cost of Stigler’s food portfolio. The Carpenter and Stiebeling and the Rose food portfolios are clearly below the efficient frontier illustrated in Figure 1.

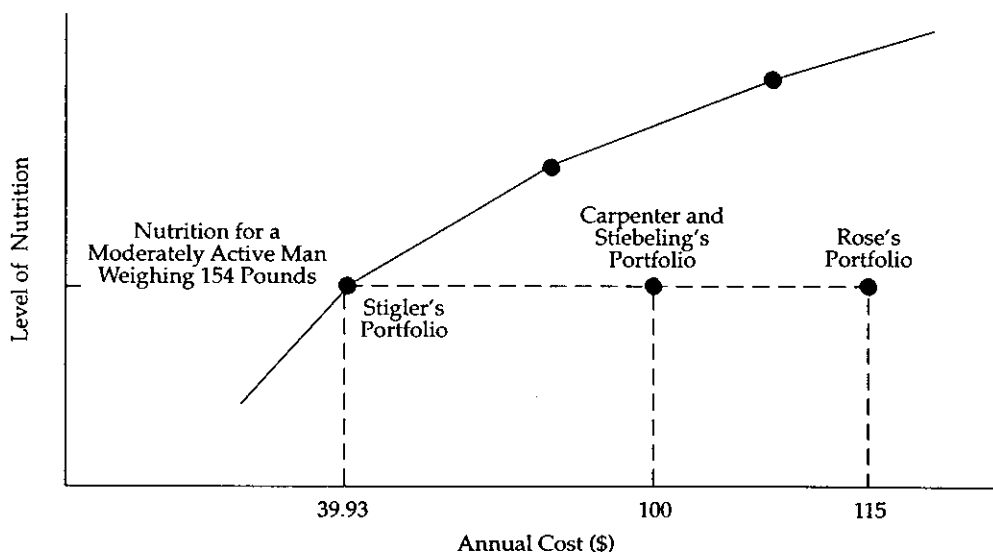
Why do diets proposed by dietitians as cheapest cost so much more than Stigler’s diet? The answer, writes Stigler, is that dietitians take into consideration not only cost and nutritional values of foods but also palatability, variety, prestige, and culture. Primarily on such grounds, he notes, can one explain dietitians’ inclusion of meat and sugar in minimum-cost diets. Stigler does not object to diets that contain meat and sugar; he objects only to attempts to present palatable, varied, prestigious, and cultural diets in the guise of minimum-cost diets. Similarly, we object to attempts to “mas-

sage” the parameters of securities so as to present palatable, varied, prestigious, and cultural security portfolios in the guise of mean–variance–efficient portfolios.

Mean–variance investors follow the farmer in the cattle feed problem in the realization that foods are no more than bundles of nutrients, and because all foods mix in the stomach, features other than nutrition and cost do not matter. But normal people do not judge meals from the perspective of the stomach alone, and normal investors do not judge security portfolios by risk and expected return alone. Normal people commonly associate particular food items such as cereal with particular meals, eaten at particular times of the day. Similarly, normal investors commonly associate particular securities such as money-market funds with particular goals such as liquidity. Normal people commonly have a sugar-laden dessert but insist on a sugar substitute for their coffee. Similarly, normal investors commonly insist on choosing only the safest money-market fund but go for the wildest IPOs.

The Markowitz mean–variance portfolio theory is prescriptive: It instructs investors about the way they *should* construct optimal portfolios. In contrast, the Shefrin and Statman behavioral portfolio theory is descriptive: It describes the way investors *actually* construct portfolios. And the portfolios investors construct are portfolios that appeal to their intuition.

Figure 1. An Efficient Frontier of Food Portfolios



MEAN-VARIANCE AND BEHAVIORAL PORTFOLIOS

To understand the difference between mean-variance portfolios and behavioral portfolios, it is useful to divide the elements in Markowitz's mean-variance framework into two parts. One concerns the importance of viewing a portfolio as a whole rather than as a simple sum of the portfolio assets. Specifically, viewing the portfolio as a whole requires that proper account be taken of the covariances between all assets in the assessment of the risk of each asset. The second part concerns the roles of the expected return and the variance of a portfolio as the sole determinants of the desirability of the portfolio. Those who take issue with the role of portfolio variance as the measure of risk propose other measures, such as semivariance. Those who take issue with the roles of expected return and variance as the sole determinants of the desirability of a portfolio modify these variables or add others such as skewness. Few people, however, argue that portfolios should be assessed asset by asset rather than as a whole and that covariances between assets be overlooked. Yet covariances are commonly overlooked as investors assess the risk of assets and construct portfolios.

Imagine that you are presented with the following two investment positions: a \$10,000 long position in U.S. Treasury bills, and a \$10,000 short position in Standard & Poor's (S&P) 500 futures contracts. Which investment position is riskier? Investors who characterize the second as more risky than the first fail to consider covariances properly unless their portfolios covary negatively with the S&P 500 Index. Note that the covariance of the T-bill position with the S&P 500 Index is zero, while the covariance of the short position in the S&P 500 futures with the S&P 500 Index is negative.

The lesson of Markowitz, a lesson that Sharpe developed into the CAPM, is that the risk of an asset is proportional to the covariance of the asset with the overall portfolio, not to the variance of the asset itself. That lesson is reflected in the structure of beta. The beta of an asset is proportional to the covariance of the asset with the overall "market portfolio," not to the variance of the asset itself. Indeed, the risk of the \$10,000 short position in the S&P 500 futures is lower than the risk of the \$10,000 long position in T-bills because the beta of the short position in the S&P 500 futures is negative, lower than the zero beta of the long position in T-bills.

The fact that the short position in S&P 500 futures contracts is less risky than the long position in T-bills seems counterintuitive to most people

because they assess the risk of an asset by the volatility of that asset and overlook the covariance with the overall portfolio. They compare the high volatility of S&P 500 futures with the low volatility of T-bills and conclude, in error, that the risk of S&P 500 futures is higher.

Markowitz attempted to help investors by framing portfolios as a whole, but typical investors do not follow Markowitz's advice. Instead, as Shefrin and Statman write, behavioral investors build portfolios as pyramids, layer by layer, and covariances among assets are overlooked in the building process. Consider the portfolios of pension funds. Typically, pension funds begin constructing their portfolios with an asset allocation decision that defines the layers, or asset classes, of the portfolio pyramid: so much for growth stocks, so much for small-capitalization stocks, and so much for international stocks. Next, the funds fill each layer with suitable securities. They allocate growth monies among growth money managers, small-cap monies among small-cap money managers, and international monies among international money managers. This layer-by-layer process virtually guarantees that covariances between asset classes will be overlooked.

Institutional investors who invest globally often put securities and currencies in distinct layers of the pyramid. They separate the management of securities from the management of currencies, which they delegate to "overlay" managers. As Jorion (1994) notes, the overlay structure is inherently suboptimal because it ignores covariances between securities and exchange rates. He calculated the efficiency loss that results from overlooking covariances as the equivalent of 40 basis points.

Cases of suboptimal portfolios that result from layer-by-layer construction are plentiful, much beyond Jorion's example. Some investors provide no layer at all to derivatives because they consider derivatives too volatile. The volatility of derivatives can indeed be high, but proper consideration of covariances makes it plain that derivatives need not be associated with high risk. Derivatives can be used to reduce the risk of the overall portfolio just as easily as they can be used to increase it. Other investors who are ready to allocate substantial proportions of their portfolios to bonds so as to keep risk low are reluctant to allocate anything to foreign stocks in the belief that foreign stocks are too risky. Such investors forgo the benefits of global diversification because they fail to understand the nature of covariances between foreign stocks and other assets, as well as the effect of these covariances on the overall risk of their portfolios.

Pyramid portfolios are categorized into layers

in many ways, and the effect of these categories on the characteristics of portfolios is often opaque. Investors whose portfolios have large layers devoted to "value" stocks are not always aware that they are, in fact, tilting their portfolios toward oil companies and banks and away from pharmaceutical and high-technology companies. The definition of layers often leads to the exclusion of securities. Investors who define fixed-income layers of their portfolios as a place for low-risk assets might exclude junk bonds from those layers because they have higher default risk than investment-grade bonds. The same investors might exclude junk bonds from stock layers of their portfolios because junk bonds are bonds, not stocks. Yet an analysis of the portfolios as a whole, taking covariances into account, might show that junk bonds improve the risk and return characteristics of the overall portfolio.

Coordination of the layers so that portfolios are considered as a whole is difficult when plan sponsors allocate the various layers of the portfolios to many money managers. One version of mean-variance optimization compounds the problem. "Tracking-errors optimization" is perhaps the most popular application of mean-variance analysis among money managers. In this framework, returns are replaced by tracking errors—deviations from benchmark portfolios. Money managers find their optimal portfolio on the tracking-error-efficient frontier, a portfolio they expect will beat the benchmark without too much risk of falling below it. Roll (1992) demonstrated that portfolios on the tracking-error-efficient frontier are quite different from portfolios on the return-efficient frontier. Therefore, money managers who optimize in the tracking-errors framework, suboptimize in the returns framework. Clarke, Krase, and Statman (1994) argued that tracking-error optimization is popular because it helps money managers deal with the pain of regret that comes with underperforming the benchmark and with the business risk of losing a client. Ironically, suboptimization is inherent in a major use of optimizers in institutional settings.

The observation that people have inconsistent attitudes toward risk is long-standing. This is the familiar Friedman-Savage observation that many people who buy insurance policies also buy lottery tickets. The inconsistent attitudes toward risk are evident in the pyramid structure of portfolios. Portfolio layers are often associated with particular goals. Indeed, mutual fund companies design their funds for particular layers, and the labels of the layers indicate their goals. Examples include

"income," "growth," and "aggressive growth." Income funds are designed to assure investors that they will not be poor; aggressive growth funds are designed for a chance at getting rich. Many people who buy insurance as if they are risk averse also buy lottery tickets as if they are risk seeking. Similarly, many investors who insist that their income funds contain nothing but Treasury securities also insist that their stock funds contain nothing but aggressive growth stocks.

People who frequent casinos often distinguish between their "own" money, which they brought with them, and "house" money, which they have won at the craps tables or slot machines. They take risks with the house money that they would not take with their own money, even though the distinction between the two kinds of money is entirely in their heads.

Pension funds, as described by Sharpe (1987), play a similar game. When the pension fund is overfunded, the committee might say, "Go for it; be aggressive; we have plenty of protection; the cushion is big." But when the pension fund is underfunded, the committee might say, "Don't take many chances; we are underwater already and need to be conservative." Portfolio insurance is an application of this structure. Users of portfolio insurance take more risk when they are ahead, playing with house money, and scale back risk when they fall behind, playing with their own money.

INVESTOR PREFERENCES

Tversky and Kahneman (1986) showed that frameworks affect choices. People often choose one alternative when the problem is framed one way but choose another alternative when the problem is framed differently. In particular, they showed that typical subjects who choose optimal, or "dominating," portfolios when the portfolio is framed in a "transparent" frame choose suboptimal, or "dominated," portfolios when the problem is framed in an "opaque" frame. Transforming a problem from an opaque frame to a transparent one is a difficult mental task.

To understand the difference between opaque and transparent frames, consider again the two investment positions: a \$10,000 long position in U.S. Treasury bills or a \$10,000 short position in S&P 500 futures contracts. Is the second more risky than the first? Posed this way, the question frames the choice in an opaque frame. Investors who have substantial long S&P 500 positions in their portfolios might overlook the negative covariance between their existing long S&P 500 position and the short S&P 500 position offered in the second option and conclude that the second is more risky

than the first. The opaque frame makes it difficult to see that when judged in the context of the overall portfolios, the short S&P 500 position is, in fact, less risky than the long T-bill position.

If an opaque frame is the problem, transformation into a transparent frame is the solution. If so, investors who choose the long T-bill position as the less risky one when the problem is framed in the opaque frame will choose the short S&P 500 position when the problem is framed transparently. But what if a transformation from an opaque frame into a transparent frame does not lead investors to change their minds? What if investors who are presented with portfolios in a transparent frame still prefer the dominated portfolio over the dominating one?

Imagine that you use a mean-variance optimizer to analyze the portfolios of a group of investors. The optimizer transforms the problem from an opaque frame into a transparent frame as it uncovers covariances that have been overlooked and inconsistent attitudes toward risk. If you find that the portfolios you analyze are dominated, you can construct mean-variance-efficient portfolios with higher expected returns and lower variances by replacing some stocks with others. But the investors may refuse to switch from their dominated portfolios to mean-variance-efficient portfolios. They explain that the mean-variance-efficient portfolios contain stocks of "sin" companies, companies that produce cigarettes or alcohol, and as "socially responsible" investors, they will have none of them.

Are portfolios of socially responsible investors truly dominated? Some might say so. Others, however, would realize that portfolios of investors who use social responsibility criteria might be no less optimal than the portfolios of investors who care only about risk and expected returns. The two groups of investors have different goals, and each investor group optimizes relative to its own goals.

One part of the mean-variance framework relates to the role of covariances in the construction of portfolios. Mean-variance investors take proper account of covariances, but many perfectly normal investors overlook covariances as they construct their portfolios. The other part of the mean-variance optimization framework relates to the nature of investor preferences. Even those who agree that investors should consider covariances do not necessarily agree that investors should have mean-variance preferences. For example, mean-variance-efficient portfolios are not optimal for investors who want to maximize the likelihood of long-term survival. The preferences of such investors are captured in a logarithmic utility function. Moreover, some—perhaps most—investors have preferences that go

beyond expected returns and risk. A preference for stocks of socially responsible companies is one example.

There is always a place for debate, persuasion, and education about goals, but in the end, people should be left alone to choose their goals, even if those goals conflict with mean-variance goals. If mean-variance optimization techniques are useful, they are useful only in helping investors frame portfolios in a transparent frame so that obscure covariances become clear and better portfolios—better in the eyes of investors—become apparent. To understand the issue, consider again portfolios of food.

Smith (1959) begins his discussion of diets with the observation that the passage of time since the publication of Stigler's article in 1945 did little to make Stigler's minimum-cost diet popular. Evidently, people care about more than cost and nutrition. Smith modifies Stigler's analysis by adding palatability to cost and nutrition as elements in the objective function.

The baseline for Smith's analysis is a "midget model," in which—as in Stigler's model—palatability is set aside and only cost and nutrition are considered. His set includes 73 foods and 12 nutritional requirements. The optimized portfolio, a portfolio that would feed a family of three for a four-week period, costs \$28.33 and consists of only six foods: milk, margarine, carrots, potatoes, pork, and flour. He concedes that the diet is a bit dull.

Palatability is taken into account in Smith's "large model." The constraints in the large model force the diet to provide some of the common pleasures of dining by setting minimum quantities for commodities that are widely consumed. The optimized large model contains 57 foods and costs \$43.58, quite a jump from the \$28.33 of the midget model. The large model diet, however, is much more palatable than the limited one; it includes coffee, meats, eggs, fish, ice cream, cakes, and other foods that make dining pleasurable.

Can consumers benefit from the solution of the large model? Possibly. The solution of the large model seems to dominate diets that people choose by intuition alone. Smith notes that the large-model diet costs only 17 cents a meal (as of 1953), although according to the Michigan State University Survey, even families that spent 20 cents a meal did not have fully nutritious diets. Will consumers accept the large-model diet as one that dominates their current diet? Perhaps. If families overlooked palatable and nutritious minimum-cost diets because they are presented in an opaque frame, presentation in a transparent frame might sway them to change their diets. If so, the optimization procedure

has practical benefit. Similarly, the benefit of the mean-variance optimization is in presenting opaque-frame portfolios in transparent frames, but the choice among transparent-frame portfolios is best left to investors.

CONCLUSIONS

Michaud (1989) argued that mean-variance optimizers are little used because mean-variance-optimized portfolios are unintuitive. He traces the nature of those portfolios to the fact that mean-variance optimizers are "estimation-error maximizers"; they shift too much weight to assets that, because of estimation errors, show high returns, low variances, or low correlations with other assets. In contrast, we traced the unintuitive character of mean-variance-optimized portfolios to the observation that they do not reflect investors' preferences.

Securities portfolios are like diets. There is much to be said for portfolios that are on the mean-variance-efficient frontier. Similarly, there is much to say for diets that are low cost and nutritious. But people care about more than cost and nutrition when they select diets. They also care about palatability, variety, prestige, and culture. Similarly, investors care about more than risk and expected return when they select securities portfolios. They care about the investment equivalents of palatability, variety, prestige, and culture. Mean-variance objectives simply do not reflect the objectives of

investors as they choose portfolios.

Green and Hollifield (1992) found that errors in the estimation of expected returns, variances, and correlations do not explain the unintuitive nature of mean-variance-optimized portfolios. They call for investors to abandon their intuitive portfolios in favor of mean-variance-optimized portfolios. We disagree. Optimization techniques should fit the goals of investors, rather than dictate goals to investors. Mean-variance optimization can help "socially responsible" investors uncover the best socially responsible portfolios, but it cannot tell investors whether social responsibility goals are worth pursuing.

Today's diets are different from the diets of decades ago, and today's portfolios are different from the portfolios of decades ago. Many decry "home bias," the tendency of investors to concentrate their holdings in their own domestic stocks. But at one time, even U.S. stocks were considered too exotic for U.S. investors. Lowenstein (1996) reports that when Charles Dow initiated the Dow Jones Industrial Average in 1896, prudent investors bought bonds, not stocks. Stocks were then the playthings of raiders and speculators. U.S. stocks are now considered tame investments. International stocks are considered more risky than U.S. stocks, and stocks in emerging markets are still at the frontier of risk. Tastes, intuition, and goals, whether in food or securities, can be debated and pushed, but they can be pushed only so fast.¹

NOTES

1. We thank Guss Souter for Vanguard data and Michelle Hawthorne, Julie Sauter, and Paras Trivedi for assistance. We thank Peter Bernstein, Van Harlow, Mark Kritzman, Hersh

Shefrin, and Drew Starbird for comments. Meir Statman acknowledges financial support from the Dean Witter Foundation.

REFERENCES

Best, Michael J., and Robert R. Grauer. 1991. "On the Sensitivity of Mean-Variance-Efficient Portfolios to Changes in Asset Means: Some Analytical and Computational Results." *Review of Financial Studies*, vol. 4, no. 2 (Summer):315-42.

Black, F., and R. Litterman. 1991. "Asset Allocation: Combining Investor Views with Market Equilibrium." *Journal of Fixed Income*, vol. 1, no. 2 (September):7-18.

———. 1992. "Global Portfolio Optimization." *Financial Analysts Journal*, vol. 48, no. 5 (September/October):28-43.

Candler, Wilfred, and Earl Heady. 1958. *Linear Programming Methods*. Ames, IA: Iowa State University Press.

Chopra, Vijay K. 1993. "Improving Optimization." *Journal of Investing*, vol. 2, no. 3 (Fall):51-59.

Chopra, Vijay K., Chris R. Hensel, and Andrew L. Turner. 1993. "Massaging Mean-Variance Inputs: Returns from Alternative Global Investment Strategies in the 1980s." *Management Science*, vol. 39, no. 7 (July):845-55.

Chopra, Vijay K., and William T. Ziemba. 1993. "The Effect of Errors in Means, Variances, and Covariances on Optimal Portfolio Choice." *Journal of Portfolio Management*, vol. 19, no. 2 (Winter):6-12.

Clarke, Roger G., Scott Kruse, and Meir Statman. 1994. "Tracking Errors, Regret and Tactical Asset Allocation." *Journal of Portfolio Management*, vol. 20, no. 3 (Spring):16-24.

Frost, Peter A., and James E. Savarino. 1986. "An Empirical Bayes Approach to Efficient Portfolio Selection." *Journal of Financial and Quantitative Analysis*, vol. 21, no. 3 (September):293-305.

- . 1988. "For Better Performance: Constrain Portfolio Weights." *Journal of Portfolio Management*, vol. 15, no. 1 (Fall):29–34.
- Green, Richard C., and Burton Hollifield. 1992. "When Will Mean-Variance Efficient Portfolios Be Well Diversified?" *Journal of Finance*, vol. 47, no. 5 (December):1785–1810.
- Jorion, Philippe. 1994. "Mean-Variance Analysis of Currency Overlays." *Financial Analysts Journal*, vol. 50, no. 3 (May/June):48–56.
- Kallberg, Jarl G., and William T. Ziemba. 1984. "Mis-Specification in Portfolio Selection Problems." In *Risk and Capital: Lecture Notes in Economics and Mathematical Sciences*, edited by G. Bamberg and A. Spremann. New York: Springer-Verlag:74–87.
- Lowenstein, Roger. 1996. "Charles Dow's Vision Was Well above the Average." *Wall Street Journal* (January 8):C1.
- Michaud, Richard O. 1989. "The Markowitz Optimization Enigma: Is 'Optimized' Optimal?" *Financial Analysts Journal*, vol. 45, no. 1 (January/February):31–42.
- Roll, R. 1992. "A Mean/Variance Analysis to Tracking Error," *Journal of Portfolio Management*, vol. 18, no. 4 (Summer):13–23.
- Sharpe, William F. 1987. "The Risk Factor: Identifying and Adapting to the Risk Capacity of the Client." In *Asset Allocation for Institutional Portfolios*, edited by Michael D. Joehnk. Homewood, IL: Dow Jones Irwin:35–45.
- Shefrin, Hersh, and Meir Statman. 1995. "Behavioral Portfolio Theory." Working paper, Santa Clara University.
- Smith, Victor E. 1959. "Linear Programming Models for the Determination of Palatable Human Diets." *Journal of Farm Economics*, vol. 31, no. 2 (May):272–83.
- Statman, Meir. 1982. "Fixed Rate or Index-Linked Mortgages from the Borrower's Point of View: A Note." *Journal of Financial and Quantitative Analysis*, vol. 17, no. 3 (September):451–57.
- Stigler, George J. 1945. "The Cost of Subsistence." *Journal of Farm Economics*, vol. 27, no. 2 (May):303–14.
- Tversky, A., and D. Kahneman. 1986. "Rational Choice and the Framing of Decisions." *Journal of Business*, vol. 59, no. 4 (October):Part 2, S251–78.