

The Next Best Thing to Knowing Someone Who is Usually Right

Gary Smith

Department of Economics

Pomona College

Claremont CA 91711

Joseph Steinberg

Department of Economics

Pomona College

Claremont CA 91711

Robert Wertheimer

Department of Economics

Pomona College

Claremont CA 91711

The Next Best Thing to Knowing Someone Who is Usually Right

Abstract

Mean-variance analysis is widely used for portfolio allocation decisions. The use of historical data for the inputs may be inferior to using informed estimates that reflect one's beliefs about the current financial environment. In this paper we show that portfolios based on expert opinion can outperform portfolios based on historical data, and that even better performance can be achieved by taking into account regression to the mean.

Mean-variance analysis is one of the most widely used techniques for allocating wealth among different assets or asset classes (Tew, Reid, and Witt 1991; Thorley 1995). The parameters needed for mean-variance analysis are the expected values and standard deviations of the asset returns and the correlations among the asset returns. Typically, these parameters are estimated from observed historical returns.

An alternative approach is to work with genuine informed predictions about asset returns. This has the appealing virtue of identifying a portfolio allocation that is consistent with the investor's views—putting one's money where one's mouth is. To illustrate this approach, we use the Federal Reserve's Livingston Survey which asks professional economic forecasters to predict future values of the S&P 500 index and long-term Treasury bond yields.

Optimal Portfolios

Asset allocation can be modeled as a utility maximization problem in which an investor chooses a portfolio that maximizes expected utility, taking into account the probability distribution of the asset returns. In mean-variance analysis, expected utility depends on the expected value and standard deviation (or variance) of the portfolio return: Choose the portfolio

weights ω_i to maximize the portfolio mean $\mu = \sum_{i=1}^n \omega_i \mu_i$ for a given value of the portfolio variance

$$\sigma^2 = \sum_{i=1}^n \sum_{j=1}^n \omega_i \omega_j \mu_i \mu_j \rho_{ij} \text{ subject to } \sum_{i=1}^n \omega_i = 1 \text{ where}$$

n = number of assets

ω_i = fraction of portfolio invested in asset i ($\omega_i \geq 0$ if no short selling)

μ_i = expected return on asset i

σ_i = standard deviation of return on asset i

ρ_{ij} = correlation between returns on assets i and j

The efficient frontier is the upper part of a hyperbola in the mean-standard deviation plane, with the smallest variance at the apex of the hyperbola. The Markowitz frontier (1959) identifying those portfolios with the highest possible expected return for a given standard deviation is that part of the hyperbola above the minimum-risk portfolio. If there is a risk-free asset, Tobin's Separation Theorem (1958) states that the optimal risky portfolio is at the tangency of a ray extending from the safe return to the Markowitz frontier.

If the probability distribution is known, then the portfolio allocation will indeed be optimal. If the means, variances, and correlations are estimated with error, because of limited data or because the parameters vary over time, then the portfolio allocation will not be optimal and the expected value of utility will be lower than the investor thinks and lower than the maximum value with the optimal portfolio.

In Figure 1, the solid frontier is the true opportunity locus based on the actual probability distribution of the asset returns; if this true frontier were known, the investor would choose Portfolio A. The dotted frontier is the investor's (incorrect) estimate of the frontier, and the investor chooses Portfolio B based on this frontier. Portfolio C shows this selected portfolio's actual mean and standard deviation. Jorion (1986) defines the loss in expected utility due to estimation error as the percentage reduction in expected utility resulting from the choice of Portfolio C instead of A.

The utility loss is due to estimation error. For example, the incorrectly estimated frontier might be based on the frequency distribution of historical returns, which are an imperfect guide to the probability distribution of future returns. If investors have useful information about the true probability distribution beyond what might be gleaned from observed historical returns, then this information might be used to reduce the utility loss.

Historical Estimates

There is some evidence that Markowitz portfolios based on historical data can beat the

market, in apparent contradiction of the efficient market hypothesis. For example, Cohen and Pogue (1967) used ten years of historical data to estimate the means, standard deviations, and correlation coefficients of widely traded stocks and, from these, identified portfolios on the Markowitz frontier. They then calculated the performance of various portfolios in succeeding years and found that Markowitz-frontier portfolios dominated randomly selected portfolios.

However, investors need not restrict their anticipation of the future to a simple averaging of the past. Markowitz procedures can also be usefully employed to describe one's beliefs about the future and, then, to select a portfolio that reflects these beliefs. For example, the fact that, over the last 80 years, long-term Treasury bonds have yielded an average of 5 percent and stocks 10 percent is not a persuasive reason for assuming that the expected values of the returns on bonds and stocks during the coming year are 5 percent and 10 percent. Perhaps long-term bonds are currently priced to yield 12 percent and we think that bond prices are equally likely to go up or down. If so, we should use 12 percent, not 5 percent, as the expected return. If stock prices have reached what we consider unreasonable levels relative to dividends, earnings, and corporate assets, we are not compelled to assume an expected 10 percent return on stocks this year. If we expect stock prices to drop, we should express that belief by using a negative value for the expected return—which will then lead to a portfolio that reflects our beliefs.

Similarly, the past volatility of bond and stock prices can be helpful in gauging uncertainty but, here too, need not be applied mechanically. The clearest example is Treasury bill rates. In the Great Depression, these were but a fraction of a percent; in the inflationary 1970s, T-bill rates soared to double-digit levels. Over the past 80 years, the average value has been about 4 percent and the standard deviation about this mean has been 3 percent. But this annual variation in the level of Treasury bill rates over the past 80 years is not the same as our uncertainty about the return over the coming year. There is no uncertainty at all about the one-year return on a one-year Treasury bill! If a \$10,000 Treasury bill sells for \$9,400, we are certain to make \$600, a

percentage return of $\$600/\$9400 = 6.4$ percent, and we should use an expected return of 6.4 percent with a standard deviation of 0 percent.

Similarly, the standard deviations for the returns on stocks and long-term bonds should reflect our uncertainty, not their historical volatility. If the Federal Reserve Chair announces that the Fed will hold Treasury rates rock steady by purchasing or selling, as needed, unlimited quantities of Treasury bonds at fixed prices—as it did for several years in the 1940s—then the implied constancy of bond prices means that there is virtually no uncertainty about the return on Treasury bonds. A low standard deviation on bond returns reflects our confidence in the success of the Fed's policy. If the Fed Chair instead announces that the Fed intends to ignore interest rates, a large standard deviation can be used to describe our uncertainty about bond prices and about the rate of return to be realized by investing in bonds.

Anticipated correlations among bond and stock returns need not be a mechanical replication of the past either. Bond and stock returns were negatively correlated in the 1960s, but positively correlated in the 1970s. The correlation coefficient used for portfolio management should reflect the sources of our uncertainty. Bond returns depend on interest rate surprises, while stock returns are affected by surprises in both interest rates and the economy.

The 1960s were a recession-free decade, propelled by increases in private and government spending, and the strong economy increased borrowing and interest rates. The rise in interest rates reduced bond prices; but stock prices increased as the negative effects of high interest rates were overwhelmed by the strong economy. This coexistence in the 1960s of rising interest rates and expanding corporate profits caused a negative correlation between bond and stock prices.

Other historical periods have been dominated by the Federal Reserve raising interest rates to cool the economy and slow inflation, or lowering interest rates to stimulate the economy when it seemed on the verge of collapse. A rise in interest rates that brings an economic recession will cause bond and stock prices to fall together; a decline in interest rates that brings an economic

boom will cause bond and stock prices to rise together. In such circumstances, bond and stock prices are positively correlated.

The question we need to ask for the future is, if interest rates rise unexpectedly, do we think it will be because the economy is expanding or because the Fed is using high interest rates to contract the economy? The former implies a negative correlation between bond and stock prices, the latter a positive correlation.

Bayesian Estimates

In the 1700s, a reformist Presbyterian minister named Thomas Bayes introduced a theoretical framework for updating an experimenter's beliefs by combining the experimenter's understanding of a scientific phenomenon with observed data. The experimenter's initial understanding is called the *prior* belief and it refers specifically to beliefs held prior to analyzing the current data. The revised belief that results from using observed data to update the prior is called the *posterior* belief. Bayes' Theorem for the conditional probability about an unknown and continuous quantity θ takes the form

$$p[\theta | x_1, x_2, \dots, x_n] = \frac{p[x_1, x_2, \dots, x_n | \theta] p[\theta]}{\int p[x_1, x_2, \dots, x_n | \phi] p[\phi] d\phi}$$

where $p[\theta]$ is the prior density, $p[\theta | x_1, x_2, \dots, x_n]$ is the posterior probability density conditional on the observed data x_1, x_2, \dots, x_n , and $p[x_1, x_2, \dots, x_n | \theta]$ is the likelihood function for the data conditional on θ . We can view Bayes' theorem as describing how to either (a) reverse conditional probabilities (going from the probability of observing the data given the unknown parameter value to a probability distribution for the parameter value given the observed data); or (b) revise one's beliefs in light of the data (going from a prior probability distribution to a posterior probability distribution).

The Bayesian framework is well-suited for mean-variance analysis because it provides a

logically consistent apparatus for combining prior beliefs about future asset returns with historical data on asset returns. Bayesian inference allows an investor to maximize expected utility with respect to a more appropriate probability density function that conditions prior beliefs on observed data. The parameter values used in Bayesian-influenced mean-variance analysis will be a weighted average of those derived from the historical data and those that reflect one's priors, with the weights determined by the precision or "tightness" of the priors.

Bayesian inference can be used to modify historical data so that they reflect investors' prior beliefs about future returns, thus producing posterior estimates that combine the historical data and prior values. Several authors, such as Jorion (1986), Pastor (2000) and McClatchey and VandenHul (2005) illustrate empirically the advantages of adopting a portfolio selection method that incorporates Bayesian inference. James-Stein estimators have also been used to shrink individual asset expected returns toward an overall mean (Michaud 1989; Chopra, Hinsel and Turner 1992) when there are a priori reasons for believing that a group of assets should have similar expected returns.

There has been relatively little work on modified covariance matrixes because it is quite difficult to derive estimators without quite drastic and unrealistic assumptions (Press 2003). For example, Elton, Gruber, and Padberg (1976) and Elton, Gruber, and Ulrich (1978) assume that the correlation between any two asset returns is equal to the historical average correlation among asset returns. On the other hand, Chopra and Ziemba (1991) argue that errors in the estimated means are far more serious than are errors in the estimated variances and covariances.

Some authors, like Frost and Savarino (1986), and Jorion (1986) use an empirical Bayes method in which prior beliefs, instead of being assumed known, are derived from the sample data. Frost and Savarino, for example, analyze a sample of 25 U. S. stocks. To specify their priors, they assume that all 25 stocks have identical expected returns and variances, and that all correlation coefficients between any two stocks are the same. The prior values of these universal

parameters are simply assumed to be the historical averages for all stocks. Bayes' Theorem is then used to meld these historical overall averages with the historical means, variances, and correlations for the individual stocks. In this way, the individual values are shrunk toward the overall averages.

Pastor (2000) and Wang (2005) combine estimated means based on historical returns with the expected returns estimated from the Capital Asset Pricing Model (CAPM), while Black and Litterman (1992) argue that investors can combine their subjective views with the CAPM's expected returns. An alternative approach, which we investigate here, is to combine historical data with genuine expert beliefs by analyzing predicted bond and stock returns. The forecasters who make such predictions presumably take into account the historical data as well as their own subjective beliefs. Therefore, we interpret their forecasts as posterior predictions. However, such beliefs might be profitably modified by taking into account the implications of regression to the mean.

Regression Towards the Mean

The educational testing literature provides a well-established framework for explaining regression to the mean in a sequence of cross-section data (Kelley 1947; Lord and Novick 1968). A person's *true score* is the statistical expected value of his or her score on a test. It is assumed that a person's observed score on the test differs from the true score by an independent and identically distributed *error score*. Looking at a test taken by a group of people, there is a distribution of true scores and observed scores. Those who score the highest are likely to have positive error scores because it would be unusual for someone to score below their true score and still score highest on a test. Since a score that is high relative to the group is also likely to be high relative to this person's true score, this person's score on another test is likely to regress toward the mean score.

This framework is applicable to predictions of economic variables, such as the growth rate of a

company's earnings. Because the predicted growth rate deviates from the probabilistic expected value of the growth rate ("the true growth rate"), a predicted growth rate that is high relative to a group of companies is also likely to be high relative to that company's true growth rate. It is unlikely that the company predicted to grow the fastest in 2000 has a negative error score, with predicted growth below the expected value of its growth rate.

We can consequently anticipate regression toward the mean when comparing the predicted and actual values of economic variables. For example, when a company's predicted growth rate is above the predicted growth rate of other companies or above its own trend-adjusted average growth rate, its growth rate is likely to regress to the mean. Thus Keil, Smith, and Smith (2004) show that the accuracy of analysts' earnings growth forecasts for a cross section of companies can be improved by shrinking their forecasts toward the mean forecast. Dorsey-Palmateer and Smith (2006) find that the accuracy of a sequence of predicted changes in interest rates can be improved by shrinking them toward a mean of zero.

Here, we will analyze predicted bond and stock returns. Looking at expert forecasts made at different times, a natural assumption is that when the predicted returns are higher (or lower) than the average predictions made at other times, these predictions are more likely to be too high (or too low) relative to the actual expected value of the return. That is, if the average predicted return has historically been 5% and predicted return this period is 15%, this 15% predicted return is more likely to be too high than too low. Therefore, the expected value of the return may be closer to 5% than is the forecast.

A Model of Regression Towards the Mean

Let the actual return Y at some point in time have a probability distribution with expected value μ ,

$$Y = \mu + \epsilon, \quad E[\epsilon] = 0 \quad (1)$$

and assume that the expert forecast return X is equal to the expected value of the actual return, with an independent and identically distributed error score ϵ

$$X = \mu + \epsilon, \quad E[\epsilon] = 0 \quad (2)$$

All of these random variables and parameters are for a given time period but we have suppressed the time subscript for notational simplicity. Even if each period's predicted return is unbiased, its accuracy can be improved by shrinking it toward the mean of the predicted returns made at other points in time!

If we knew the values of μ , we could use Equation (2) to make unbiased predictions of the forecasts. However, we are interested in the reverse question: inferring the expected value of the return from the forecast return. If we had data on μ and X and estimated

$$\mu = \beta + \alpha X + \epsilon$$

by ordinary least squares, the estimated slope will be

$$\hat{\alpha} = \frac{\text{cov}[X, \mu]}{s_X^2}$$

This estimate can be used to shrink the forecasts toward the mean:

$$\hat{\mu} = \hat{\alpha}(X - \bar{X})$$

In a large sample in which the means of the forecast and actual values are approximately equal, we expect:

$$\hat{\mu} = \hat{\alpha}X + (1 - \hat{\alpha})\bar{X} \quad (3)$$

Thus, to predict the expected return, we use $\hat{\alpha}$ to shrink a given period's expert forecast toward the average forecast over several periods.

Implementing the Model

In practice, we do not observe μ and consequently cannot use data for X and μ to estimate α . However, we do observe Y , the actual return, and if we estimate this equation by ordinary least

squares

$$Y = \beta_0 + \beta_1 X + \epsilon \quad (4)$$

the estimated slope will be

$$\begin{aligned} \hat{\beta}_1 &= \frac{\text{cov}[X, Y]}{s_X^2} \\ &= \frac{\text{cov}[X, \beta_0] + \text{cov}[X, \beta_1]}{s_X^2} \end{aligned}$$

If $\text{cov}[X, \beta_0] = 0$, then $\hat{\beta}_1 = \beta_1$ and we can use our estimate of β_1 as an estimate of β_1 , which can then be used to implement Equation (3). The estimated slope of the empirical relationship between the predicted and actual returns can be used to shrink each period's forecast toward the average forecast.

Data

In 1946, a Philadelphia newspaper columnist, Joseph Livingston, began a semiannual survey of business economists' forecasts of macroeconomic variables. The Federal Reserve Bank of Philadelphia began maintaining a data base of the survey responses in 1978 and took over the survey after Livingston died. The Philadelphia Fed mails the surveys to professional economic forecasters from a wide variety of industries, including non-financial businesses, investment and commercial banks, academic institutions and government agencies, in May and November of each year, and the survey results are released to the press in June and December.

We will use the forecasts of the secondary market interest rates on long-term Treasury bonds and the level of the S&P 500 stock market index. Respondents are asked to predict the values of these variables on the last day of the month the survey is released (June and December) and six and twelve months after that date (the last days of the subsequent June and December). For example, for the survey mailed in November 2000 and released in December 2000, the forecasters

were asked to predict the values on the last day of December 2000, June 2001, and December 2001.

The survey began reporting both the predicted long-term Treasury bond yields and the S&P 500 index in 1992. In order to maximize the number of observations, we will use the 6-month forecasts. Respondents were asked to predict the yield on 30-year Treasury bonds prior to the December 2002 survey and the 10-year yield in later surveys. Because there are no actual historical values for the 30-year bond rate after February 2002, we use the empirical relationship between the 20-year and 30-year rates to estimate the value of the 30-year rate in June 2002 and December 2002.

In order to perform mean-variance analysis, we need to calculate the predicted returns implied by the predicted levels of the Treasury yields and S&P 500. We don't know the exact dates when each forecaster completed the survey and these dates surely vary across forecasters. Thus, we cannot calculate the predicted returns over horizons beginning with the survey completion dates. Instead, we compare the predicted values on the last day of the survey release month with the predicted values six months later. For example, for the survey mailed in November 2000 and released at the end of December 2000, we use the difference between the predicted S&P 500 on the last trading day of December 2000 and the predicted S&P 500 on the last trading day of June 2001 to calculate the implied rate of return from purchasing stocks on the last trading day of December 2000, and selling them on the last trading day of June 2001.

For the S&P 500 stock index, the predicted return \widehat{R}_S is calculated by adding the dividend yield to the predicted capital gain

$$\widehat{R}_S = \frac{D_1 + \widehat{P}_1 - \widehat{P}_0}{\widehat{P}_0}$$

where \widehat{P}_0 is the predicted level of the index at the end of the survey release month, \widehat{P}_1 is the

predicted level six months after that, and D_1 is the dividend paid during this six-month interval, for simplicity assumed to be known with certainty and paid at the end of the period. The actual rate of return R_S is calculated analogously using the actual values of the S&P 500 at the end of the survey release date and six months later:

$$R_S = \frac{D_1 + P_1 - P_0}{P_0}$$

In order to calculate the predicted return on a Treasury bond, we assume that it is initially selling for par (with a coupon rate equal to the yield to maturity). If we let \hat{y}_0 be the predicted yield at the end of survey month, \hat{y}_1 be the predicted yield six months after that, $M = 100$ be the maturation value, and $n = 60$ or $n = 20$ depending on whether it is a 30-year or 10-year bond, then $\hat{C} = \frac{\hat{y}_0}{2}100$ is the semiannual coupon and the bond's initial price is

$$\begin{aligned} \hat{P}_0 &= \sum_{t=1}^n \frac{\hat{C}}{(1 + \hat{y}_0)^t} + \frac{M}{(1 + \hat{y}_0)^n} \\ &= \sum_{t=1}^n \frac{(\hat{y}_0 / 2)100}{(1 + \hat{y}_0)^t} + \frac{100}{(1 + \hat{y}_0)^n} \\ &= 100 \end{aligned}$$

The market price of the bond six months later is predicted to be

$$\hat{P}_1 = \sum_{t=1}^{n-1} \frac{(\hat{y}_0 / 2)100}{(1 + \hat{y}_1)^t} + \frac{100}{(1 + \hat{y}_1)^{n-1}}$$

where \hat{y}_1 is the predicted yield to maturity at the time of sale. (We assume that the yield on a 29.5-year bond will be virtually the same as on a 30-year bond.) The implied predicted return on the bond is

$$\hat{R}_B = \frac{\hat{C} + \hat{P}_1 - \hat{P}_0}{\hat{P}_0}$$

The actual returns R_B were calculated in the same way using the actual yields to maturity in place of the predicted yields.

Constructing Portfolios

We treat the semiannual data as independent—consistent with the efficient market hypothesis—and ignore the sequence in which the observations occur by assuming that the investor knows about all the other observations. This is equivalent to picking an observation at random and assuming that all the other observations have already occurred. Thus, when we make a portfolio decision for December 2000, we assume that the investor knows all outcomes other than December 2004.

Our baseline portfolio is based on the historical distribution of returns; for example, the expected values of the stock and bond returns in any six-month forecasting period are assumed equal to the average returns for the other periods in our data set. The Livingston Survey portfolios use the individual forecasts of the S&P 500 and Treasury yield in each six-month period to calculate predicted the bond and stock returns for every forecaster; the means of these predicted returns are used as the expected values in our mean-variance analysis. The regression-to-the-mean portfolio estimates the relationship between the predicted and actual returns, as in Equation (4), using data for all periods other than the current forecasting period, and uses these estimates to implement Equation (3). For example, to adjust the analysts' forecasts made in December 2004 for the bond or stock return during the first six months of 2005, we estimate the slope of the relationship between the forecast and actual values for other periods and use this estimated slope to shrink the December 2004 forecast towards the average predicted return for other periods.

We do not have any information about the forecasters' assessment of the uncertainty of their predictions. We can't use the variation in predictions across forecasters because this only gauges

uncertainty about the expected value, not uncertainty about the actual return. For example, even if all of the forecasters were agreed that the expected value of the stock return is 5%, this would not mean that there is no uncertainty about the actual stock return.

To estimate each period's posterior covariance matrix of bond and stock returns, we first calculate the squared difference between the predicted and actual returns for each 6-month horizon (other than the current period). Interpreting these prediction errors as a measure of uncertainty about returns, we use the covariance matrix of these squared errors for m observations to estimate the posterior covariance matrix of bond and stock returns:

$$\hat{\sigma}_{ij}^2 = \frac{\sum_{t=1}^m (R_i - \hat{R}_i)(R_j - \hat{R}_j)}{m}$$

where $i = j = S$ for the stock variance; $i = j = B$ for the bond variance; and $i = S$ and $j = B$ for the covariance. Because the historical forecasts use the historical means \bar{R}_S and \bar{R}_B as estimates of the expected returns, portfolios based on the historical data alone use the historical covariance matrix:

$$s_{ij}^2 = \frac{\sum_{t=1}^m (R_i - \bar{R}_i)(R_j - \bar{R}_j)}{m - 1}$$

In each case, we use the observed differences between the predicted and actual returns in all other periods to estimate the covariance matrix in each period. The expected returns and covariance matrix are then used to determine the Markowitz frontier and select an optimal bond-stock mix.

To standardize our comparisons, we invoke Tobin's Separation theorem and focus on the optimal risky portfolio to combine with the safe asset. We use the six-month Treasury bill as the risk-free asset, with the risk-free return calculated from the secondary market rates on the last trading days of June and December of each year. Because T-bill rates are reported on a discount

basis, the actual rate of return is calculated from the reported discount rate d as follows:

$$R_0 = \frac{182d}{360 \square 182d}$$

Results

Table 1 shows the predicted and actual stock and bond returns for every six-month period, as well as the six-month T-bill returns. It is apparent that the predicted stock and bond returns are not constant. The predicted stock returns range from -0.3% to 8.8%, while the predicted bond returns range from -1.4% to 4.3%. If the forecasts were simply averages of the historical returns, the predicted returns would not jump around so much from period to period. These forecasters are surely taking current economic conditions into account. It is also clear (and unsurprising) that their forecasts are imperfect. The correlations between predicted and actual returns are -0.38 (two-sided $p = 0.05$) for stocks and 0.08 (2-sided $P = 0.67$) for bonds.

It is especially interesting that the correlation between predicted and actual stock returns is so strongly negative. If the predicted and actual returns were uncorrelated, the predictions would be worthless. If they are correlated (positively or negatively), they may contain useful information. It may be as profitable to know forecasters who are usually wrong as to know forecasters who are usually right.

Table 2 summarizes the returns from the three portfolio strategies. The Sharpe ratio is the difference between the average portfolio return and the average T-bill return, divided by the standard deviation of the portfolio return. The portfolios selected on the basis of the unadjusted Livingston forecasts do slightly better, on average, than do the portfolios based on the historical returns, but the standard deviation of the six-months portfolio returns is substantially higher and the Sharpe ratio is lower. The regression-to-the-mean portfolios have a much higher average return and a substantially higher Sharpe ratio too.

Thus, the adjusted forecasts produced by regressing each Livingston predicted return to the

mean generates asset allocations that are superior to those resulting from portfolios based solely on the historical returns, as they yield higher average returns and higher Sharpe ratios. Even with this relatively small sample, a matched-pair t-test gives a two-sided p value of 0.10. The differences are certainly substantial. Over this 13-year period, an initial investment of \$1 in the portfolios chosen by the historical data would have grown to \$2.99, while an initial \$1 invested in the portfolios chosen by the regression model would have grown to \$4.80. (61% higher)

Conclusion

Historical returns are commonly used to implement the selection of mean-variance efficient portfolios. An alternative is to take into account one's beliefs about the expected returns and variance-covariance matrix of returns so that one's portfolio reflects one's beliefs. If an investor has useful information beyond what is contained in the historical returns, this information can be used to select superior portfolios. On the other hand, one's beliefs may not be as useful as one thinks. Even professional forecasters have a hard time predicting economic variables, particularly in financial markets where the monetary rewards from accurate predictions can create a reasonably efficient market.

In this paper, the predicted returns implied by the Livingston survey are used to generate predicted stock and bond returns which we interpret as posterior forecasts. We also calculate adjusted predictions by regressing the Livingston predicted stock and returns to the historical average predicted returns. These adjusted predictions are used to estimate not only the expected returns, but also the variance-covariance matrix of returns, which are then used to choose mean-variance efficient portfolios. As it turns out, there is a negative correlation between actual stock returns and the stock returns predicted by the Livingston survey. When these professional forecasters are optimistic, the stock market tends to do poorly; when they are pessimistic, the market tends to do well. Regression-adjusted portfolios have higher average returns and higher Sharpe ratios than do portfolios based on unadjusted predictions or on the historical data alone.

References

- Black, F. and Litterman, R. (1992). Global Portfolio Optimization, *Financial Analysts Journal*, September/October, 8-43.
- Chopra, V. K., Hensel, C. R., and Turner, A. L. (1992). Massaging mean-variance inputs: returns from alternative global investment strategies in the 1980s. *Russell Research Commentaries*, May 1-19.
- Chopra, V. K., and Ziemba, W. T. (1991). The effect of errors in means, variances, and covariances on optimal portfolio choice, *Russell Research Commentaries*, November: 1-11.
- Cohen, Kalman J., and Pogue, Jerry A. (1967). An Empirical Evaluation of Alternative Portfolio Selection Models, *Journal of Business*, pp. 166–193.
- Dorsey-Palmateer, R., and Smith, G. (2006). Shrunken Interest Rate Forecasts are Better Forecasts, *Applied Financial Economics*, forthcoming.
- Elton, Edwin J., Gruber, Martin J., and Padberg, Manfred W. (1976) Simple Criteria for Optimal Portfolio Selection, *Journal of Finance*, 31: 1341-1357.
- Elton, Edwin J., Gruber, Martin J., and Ulrich, Thomas J. (1978). Are betas best?, *Journal of Finance*, 33: 1375-1384.
- Federal Reserve Bank of Philadelphia (2005). Livingston Data Documentation.
- Herold, U. and Maurer, R. (2003). How Much Foreign Stocks? Bayesian Approaches to Asset Allocation Can Explain the Home Bias of US Investors, EFA 2003 Annual Conference Paper No. 207.
- Jorion, P. (1986). Bayes-Stein estimation for portfolio analysis, *Journal of Financial and Quantitative Analysis* 21, No. 3, 279-292.
- Jorion, P. (1991). Bayesian and CAPM estimators of the mean: Implications for portfolio selection, *Journal of Banking and Finance* 15, No. 3, 717-727.
- Keil, M., Smith, G., and Smith, M. H. (2004). Shrunken earnings predictions are better

- predictions, *Applied Financial Economics*, 14, 937-943.
- Kelley, T. L. (1947). *Fundamentals of Statistics*, Harvard University, Cambridge, MA, 409.
- Lord, F. M., and Novick, M. R. (1968). *Statistical theory of mental test scores*, Addison-Wesley, Reading, MA.
- Markowitz, H. (1959). Portfolio selection: Efficient diversification of investments,"New York.
- McClatchey, C., and VandenHul, S. (2005). The Efficacy of Optimization Modeling as a Retirement Strategy in the Presence of Estimation Error, *Financial Services Review* 14, 269-284.
- Michaud, R. (1989). The Markowitz optimization enigma: Is "optimized" optimal?, *Financial Analysts Journal*, 45: 31-42
- Pastor, L. (2000). Portfolio selection and asset pricing models, *Journal of Finance* 55, No. 1, 179-223.
- Press, J.(2003). *Subjective and Objective Bayesian Statistics*, John Wiley & Sons, Inc.
- Tew, B. V., Reid, D. W., and Witt, C. A. (1991). The Opportunity Cost of a Mean-Variance Choice, *The Financial Review*, 26: 31-43.
- Thorley, S. R. (1995). The Time-Diversification Controversy, *Financial Analysts Journal*, 51: 68-76.
- Tobin, James (1958). Liquidity Preference as Behaviour Towards Risk, *Review of Economic Studies*, 25, No. 67, pp. 65-86.
- Wang, Z. (2005). A Shrinkage Approach to Model Uncertainty and Asset Allocation, *Review of Financial Studies*, 2005, vol. 18, issue 2, pages 673-705

Table 1 Predicted and Actual Returns

Forecast Date	Predicted Stock Return	Actual Stock Return	Predicted Bond Return	Actual Bond Return	Actual T-bill Return
June 1992	3.028	8.282	3.643	8.547	1.880
Dec 1992	4.640	4.294	3.052	12.927	1.686
June 1993	3.280	8.875	2.048	7.715	1.608
Dec 1993	3.349	-3.488	2.354	-11.759	1.644
June 1994	2.682	1.667	3.542	0.856	2.407
Dec 1994	1.746	23.849	2.930	20.175	3.247
June 1995	0.991	10.839	1.606	12.569	2.780
Dec 1995	3.365	5.072	2.835	-8.802	2.572
June 1996	1.823	16.950	4.247	6.664	2.679
Dec 1996	-0.240	29.842	1.689	1.426	2.657
June 1997	-0.253	2.529	2.571	15.454	2.636
Dec 1997	3.064	16.300	1.170	7.406	2.716
June 1998	0.076	10.424	0.452	10.859	2.615
Dec 1998	0.842	8.760	2.024	-9.721	2.270
June 1999	1.841	11.216	1.057	-3.550	2.508
Dec 1999	1.370	-2.057	1.800	11.302	2.855
June 2000	1.478	-7.162	0.537	9.364	3.107
Dec 2000	5.231	-7.683	4.304	-1.366	2.844
June 2001	4.100	-4.543	2.390	6.802	1.812
Dec 2001	7.156	-19.906	0.505	2.450	0.913
June 2002	8.006	-2.594	-0.517	14.074	0.872
Dec 2002	8.835	13.468	-0.203	4.237	0.616
June 2003	5.340	13.227	-1.384	-3.882	0.488
Dec 2003	7.131	-0.088	-0.791	-0.532	0.508
June 2004	5.919	10.930	-0.301	5.256	0.836
Dec 2004	4.276	2.720	-1.319	4.478	1.290
June 2005	5.636	2.076	-1.162	-1.495	1.665
average	3.508	5.696	1.447	4.498	1.989

Note: The predicted returns are the means of the predicted returns calculated from the Livingston Survey. The predicted and actual returns are for the 6-month period following the forecast.

Table 2 Portfolio Returns, percent

	Historical Data	Livingston Survey	Regress to Mean
Average Return	4.39	4.51	6.33
Standard deviation	7.33	9.32	8.94
Sharpe Ratio	0.33	0.27	0.49

Basis for estimates of assets means, standard deviations, and correlations

Historical Data: historical distribution of returns

Livingston Survey: returns calculated from predicted S&P 500 and Treasury-bond yield

Regress to Mean: Livingston current predicted returns regressed to average predicted returns

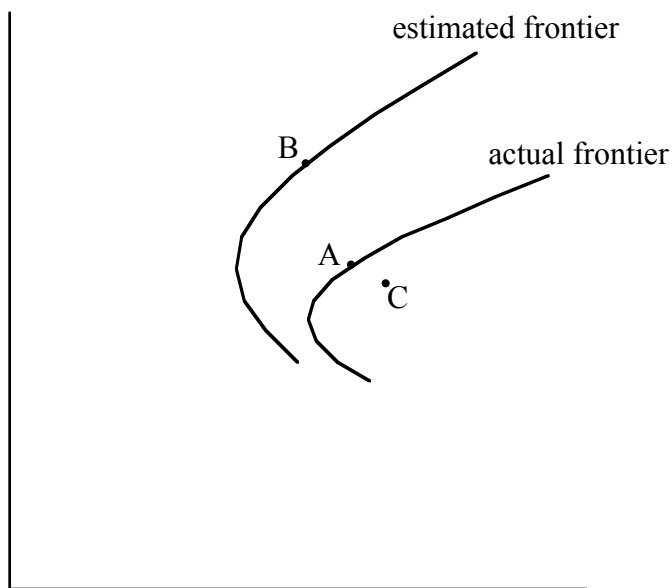


Figure 1 Optimal and Suboptimal portfolios