Short-Horizon Inputs and Long-Horizon Portfolio Choice

Use simulations to avoid a mismatch.

William N. Goetzmann and Franklin R. Edwards

Despite its well-known limitations, mean-variance optimization is commonly used to make long-term portfolio allocation decisions. See Hakansson [1972, 1974], Jorion [1985, 1986, 1990], Michaud [1989], Broadie [1991], and Best and Grauer [1991], for instance.

In this article we examine the implications of using as Markowitz model inputs the means, standard deviations, and correlations derived from short-term (or short-horizon) asset class returns as opposed to long-term (or long-horizon) asset class returns. More specifically, we simulate long-term returns using a linear model that allows stock, bond, and bill returns to be autocorrelated, and then use these returns to form efficient portfolios.

Our results indicate that for investors with multiple-year investment horizons summary statistics based upon short-term (annual) returns can be grossly misleading for making asset allocation decisions. In particular, the autocorrelation structure of asset class returns strongly influences both the variance and the interclass correlation of assets, and, as a consequence, changes the composition of the efficient frontier.

Our principal finding, therefore, is that investors with long-term horizons should consider alternative methods of estimating inputs to the mean-variance model when making portfolio decisions.

METHODOLOGY

The mean-variance framework is a single-peri-

WILLIAM N. GOETZMANN is associate professor of finance at the Yale School of Management in New Haven (CT 06520).

FRANKLIN R. EDWARDS is the Arthur F. Burns professor of finance and economics at the Columbia University Graduate School of Business in New York (NY 10027).

od model in which returns are assumed to be normally distributed (see Markowitz [1952]). Under the assumption that asset returns are serially independent, short-term returns (such as annual returns) can be used to estimate the statistical characteristics of long-term returns. That is, the annual variance of the natural log of one plus the annual asset return should be an unbiased estimate of one-tenth the variance of the ten-year log return.[1]

This assumption, however, may not be valid for stock returns (see Fama and French [1988] and Poterba and Summers [1988], for instance). It is certainly not true for bond and Treasury bill returns (see Hakansson [1972]). Indeed, as we show, the latter exhibit an extraordinary level of autocorrelation. Holton [1992] also demonstrates the degree to which long-term volatility may be misjudged by considering the uniformly compounded monthly variance to be representative of long-horizon risk.

Kandel and Stambaugh [1987] show that vector autoregression (VAR) can be used to capture the effect of temporal dependencies in asset returns when making long-horizon forecasts. The VAR is an econometric method that generalizes autoregression to multiple data series: It estimates the lagged relationship of a set of variables (for instance, several asset classes). Campbell and Shiller [1988], Goetzmann and Jorion [1993], Hodrick [1992], and Nelson and Kim [1993], among others, use VAR to simulate distributions of long-horizon asset class returns for the purpose of testing the predictability of stock returns.

In this article, we use the VAR technology instead to simulate long-horizon inputs to the mean-variance model. In a manner similar to Ibbotson and Sinquefield [1976], we use short-term returns to bootstrap empirical distributions of long-horizon asset class returns.

The autocorrelation structures of short-term asset class returns are explicitly incorporated into simulations of long-horizon returns by modeling them using vector autoregression in the following manner:

$$R_t = A'R_{t-1} + \varepsilon_t \qquad (1)$$

where R, A, and ε are matrixes. The columns of R are time series observations of log returns to the S&P 500, long-term government bonds, and Treasury bills over the period 1926 through 1991, and the matrix A is the estimate of the VAR coefficients.[2]

The estimated coefficients and the errors, ε, from the VAR are used to generate simulated returns to an investment in each asset class in the manner:

$$R_t^* = \hat{A}' R_{t-1}^* + \hat{\varepsilon}_t^* \qquad (2)$$

where $R_0^* = R_0$. In other words, the initial value for the bootstrap is the actual set of log returns for the four series in 1926, and ε_t^* is the bootstrapped error vector, drawn with replacement from the rows of the VAR error matrix.

By drawing the errors as a vector of four, one for each asset series, the contemporaneous covariance structure of the asset classes is preserved. For comparison, we also draw errors from a normal distribution, whose means, variances, and covariances match those of the error matrix ε. This semiparametric approach is used to show the degree to which our results depend upon the deviations of the model errors from normal.[3]

We do not take into consideration here uncertainty about the parameters estimated in the vector autoregression itself. Goetzmann and Jorion [1993] and Jorion [1985] point out that this estimation uncertainty represents an additional element of risk.

We perform this bootstrap procedure 1,000 times, creating 1,000 simulated ten-year return histories. The resulting joint distributions of stock, bond, and bill log returns for the ten-year horizon are then used as inputs to the mean-variance framework, under the constraint that investment weights in each asset class are either zero or positive and add to 100%.

RESULTS

Exhibit 1 reports the coefficients estimated by the VAR as well as by OLS regression. Stock returns are not well-explained by past values of bonds and bills. Bond returns are explained by past values of T-bill returns, and, conversely, T-bill returns have a significant relationship to bonds.

When the analysis is performed on the risk premiums — that is, upon the stock returns minus the T-bill returns and the bond returns minus the T-bill returns — both the VAR and the regression results show that no relationships exist. Thus, the intertemporal dependencies that exist are clearly due to the time series behavior of T-bill returns.

Exhibit 2 reports the means, standard deviations,

EXHIBIT 1
Vector Autoregression Coefficients and OLS Coefficients for Stocks, Bonds, and T-Bills Estimated from the Vector Autoregression

R_t	VAR Coefficients for Lagged Variable R_{t-1}			Regression Coefficients for Lagged Variable R_{t-1}		
	Stocks	Bonds	Bills	Stocks	Bonds	Bills
Stocks	0.023	0.517	−0.063	0.023	0.517	−0.062
				(0.128)	(0.339)	(0.815)
Bonds	−0.028	0.032	0.882	−0.030	0.032	0.905
				(0.047)	(0.123)	(0.296)
Bills	0.013	−0.073	0.963	0.021	−0.072	0.983
				(0.006)	(0.017)	(0.041)
	Tests on Equity Risk Premiums and Bond Risk Premiums					
	Stocks-Bills	Bonds-Bills		Stocks-Bills	Bonds-Bills	
Stocks-Bills	0.00	0.00		0.008	0.601	
				(0.127)	(0.325)	
Bonds-Bills	0.00	0.00		−0.056	0.111	
				(0.050)	(0.128)	

Notes: VAR and OLS estimates based upon the times series of log returns to the S&P Index, a portfolio of long-term government bonds, and Treasury bills of six-month maturity. Regression estimates are OLS coefficients and standard errors, using the row value as the dependent variable and the set of column values as the independent design matrix.

and correlations of the simulated ten-year log returns and the risk premiums, as well as estimates based upon one-year time series returns under the assumption of serial independence. The standard deviations of the Treasury bills are increased significantly by the VAR model: They have nearly the same volatility as do long-term government bonds. (This result is not due to the fact that the riskless asset at the ten-year horizon is a zero-coupon ten-year government bond. The government bond portfolio maintains its long-term duration by rebalancing monthly, so nothing in the portfolio reaches maturity in the tenth year of the simulation.)

The simulations using normal errors yield similar results, indicating that only a small portion of the increased volatility may be attributable to outlying observations, or to deviations of the errors from normality. Further, with respect to the risk premiums, the inputs are not affected. Clearly, it is the serial correlation of Treasury yields that is causing the long-horizon inputs to differ from the short-horizon inputs.

Equally interesting are the differences in the correlation matrixes. While stocks have a slight negative contemporaneous correlation to T-bills, over the long term, their log returns appear to be slightly positively correlated. This is not surprising, as most pricing models predict that the riskless rate is a component of expected stock returns.[4] On a short-horizon basis, variations in the equity risk premium dominate variations due to the riskless rate, but on a long-horizon basis the riskless component of returns manifests itself. When the riskless return is subtracted out, the correlations are unaffected.

Exhibits 3 and 4 show the composition of the efficient frontier using the two different sets of inputs: short-horizon returns and long-horizon returns generated with the bootstrap methodology. Each portfolio on the frontier assumes that the investor rebalances annually to maintain fixed weights over the entire ten-year horizon. One result is that the bootstrapped inputs have relatively little impact on the high-risk, high-return portion of the frontier — the location of the all-stock portfolio is virtually unchanged.

The major difference between the two frontiers occurs for the minimum-variance portfolio. Inputs generated by the bootstrap result in a minimum-variance portfolio composed of 50% bonds and 50% bills. In contrast, inputs taken from the time series of short-horizon returns result in a minimum-variance portfolio of 10% bonds and 90% bills. Not only do the bootstrapped inputs increase the minimum achievable risk,

EXHIBIT 2
Annualized Statistical Inputs for Stocks, Bonds, and Bills Based upon VAR Bootstrap of Long-Horizon Returns, VAR Simulation with Normal Errors, and Short-Horizon Summary Statistics for the Period 1926-1991

VAR Bootstrap	Mean	Standard Deviation	Stocks	Bonds	Bills
S&P Returns	0.098	0.207	1.00		
Long-Term Government Bonds	0.043	0.061	0.34	1.00	
Treasury Bill Returns	0.036	0.060	0.19	0.08	1.00
VAR Simulation with Normal Errors			Stocks	Bonds	Bills
S&P Returns	0.089	0.186	1.00		
Long-Term Government Bonds	0.048	0.051	0.31	1.00	
Treasury Bill Returns	0.036	0.056	0.19	0.17	1.00
Summary Statistics for Annual Log Returns			Stocks	Bonds	Bills
S&P Returns	0.096	0.199	1.00		
Long-Term Government Bonds	0.044	0.077	0.14	1.00	
Treasury Bill Returns	0.036	0.032	−0.03	0.22	1.00
VAR Bootstrap Statistics for Risk Premiums			ERP	BRP	
Equity Risk Premiums	0.059	0.210	1.00		
Bond Risk Premiums	0.008	0.082	0.11	1.00	
Summary Statistics for Annual Log Risk Premiums			ERP	BRP	
Equity Risk Premiums	0.060	0.208	1.00		
Bond Risk Premiums	0.008	0.081	0.19	1.00	

Notes: All VAR estimates and summary measures performed on logs of one plus the annual returns. VAR bootstrap procedures are described in Equation (2). For purposes of comparison, annualized returns are calculated by dividing the decade returns by ten. Annualized standard deviations are calculated by multiplying decade standard deviations by the square root of 10. Bootstrap based upon 1,000 simulations.

EXHIBIT 3
LONG-HORIZON INPUT FRONTIER
STOCKS, BONDS, BILLS

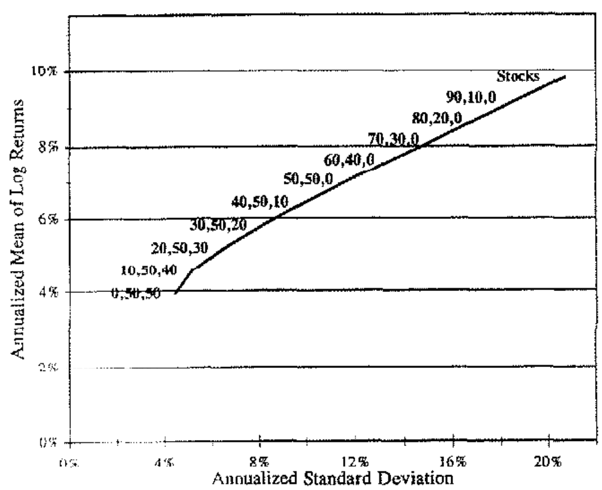

EXHIBIT 4
SHORT-HORIZON INPUT FRONTIER
STOCKS, BONDS, BILLS

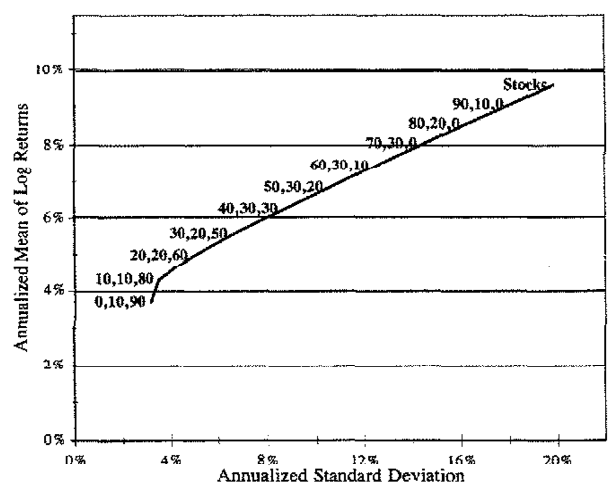

but the slightly higher correlation across assets also reduces the curvature of the frontier.

These results must be qualified by the fact that the simulations are conditional upon point estimates of the VAR coefficients, and the assumption of stationarity of all underlying parameters. Further, the period 1926-1991 is chosen for purposes of example, not for purposes of making the best long-horizon forecast of ten-year asset returns. Investors wishing to use this technique should consider further simulations that perturb the underlying parameters: means, standard deviations, correlations, and VAR coefficients.

CONCLUSION

Simulations of long-horizon returns clearly indicate that mean-variance inputs based upon short-horizon return statistics can lead to incorrect conclusions regarding the composition of the minimum-variance portfolio. This result complements other research regarding the relationship between investment horizons and deviations of returns from a random walk. For instance, Lee [1990] and Butler and Domian [1991] both demonstrate that stocks are more attractive to long-term investors when the time structure of returns is taken into account.

In our research we find that the investment horizon and the temporal structure of returns matter for the opposite end of the efficient frontier as well. The time series behavior of T-bills and bonds dramatically influences the composition of the minimum-variance portfolio and efficient portfolios close to it. In general, Treasury bills are more volatile and are more highly correlated with other assets once their time series behavior is taken into account. As components of long-term returns to stocks and bonds, they also influence the interrelationships among other assets.

The significance of these results for investors depends in large measure upon their portfolio choice criterion. Our results, however, clearly suggest that caution should be exercised when using nominal asset class returns as inputs to the Markowitz model.

In particular, if investors select portfolios on the basis of real expected returns, or on the basis of the mean and variance of the risk premium over Treasury bills, Treasury bills will constitute a major proportion of the minimum-variance portfolio. Similarly, if investors optimize over the mean and variance of the difference between assets and liabilities, the minimum-variance portfolio might be dedicated to assets matching the cash flows of liabilities.

This approach is predicated on the assumption that investors can accurately identify both their investment horizon and the timing of future cash needs. In the absence of such knowledge, there is no riskless asset, and, consequently, the Markowitz model cannot be applied with precision. When investor preferences are expressed over the mean and variance of portfolio wealth for a multiple-period horizon, even though the horizon is uncertain, the time series behavior of returns becomes critical for the prediction of the variances and the correlation matrix to be used as inputs to the Markowitz model.

ENDNOTES

The authors thank Mark Broadie, Yasushi Hamao, Roger Ibbotson, Philippe Jorion, and Paul Kaplan for helpful conversations. They also thank an anonymous referee for useful comments and Ibbotson Associates for their data. Goetzmann's research was supported by a grant from the Futures Center, and summer support from the Columbia Business School.

[1]In discrete time, that is, $\ln(1 + r_t)$, where r is the total annual return for that asset in time t. In fact the ratio of the annual variance to the multiyear variance has been used as a general test of serial dependence by Poterba and Summers [1988] among others.

[2]The VAR coefficients are estimated using the S-PLUS statistical package. For a description of the estimation algorithm, see Box and Jenkins [1976]. The VAR allows for as many lags as one might want, although we find that the first lag captures virtually all the temporal structure. This is determined by applying the Akaike information criterion (AIC), which chooses the lag k that balances the reduction in error variance achieved by including additional lags against the loss of information as a function of the degrees of freedom.

[3]This semiparametric approach also allows for estimating parameters for the VAR and the error matrix using different time periods. This is useful for modeling changing variances as in Hodrick [1992], and for investigating the performance of asset classes for which only a short time period of returns may be obtained.

[4]Ibbotson and Sinquefield [1976] break returns into components that include the real rate and the risk premiums for stocks and bonds, in order to make their bootstrap consistent with asset pricing models and the assumption that equity risk premiums follow random walks.

REFERENCES

Best, M., and R. Grauer. "On the Sensitivity of Mean-Variance Efficient Portfolios to Changes in Asset Means: Some Analytical and Computational Results." *Review of Financial Studies*, 4 (1991), pp. 315-342.

Box, G.E.P., and G.M. Jenkins. *Time Series Analysis: Forecasting and Control.* Oakland, CA: Holden-Day, 1976.

Broadie, M. "Computing Efficient Frontiers using Estimated Parameters." Working Paper, Columbia University, 1991.

Butler, K., and D. Domian. "Risk, Diversification and the Investment Horizon." *Journal of Portfolio Management*, Spring 1991, pp. 41-47.

Campbell, J.Y., and R.J. Shiller. "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors." *Review of Financial Studies*, 1 (1988), pp. 195-228.

Fama, E.F., and K.R. French. "Permanent and Temporary Components of Stock Prices." *Journal of Political Economy*, 96 (1988), pp. 246-273.

Goetzmann, W., and P. Jorion. "Testing the Predictive Power of Dividend Yields." *Journal of Finance*, 48, 2 (June 1993), pp. 663-679.

Hakansson, N. "Fallacy of the Log-Normal Approximation to Optimal Portfolio Decision-Making Over Many Periods: Comment." *Journal of Financial Economics*, 1 (1974), pp. 95-96.

——. "Mean-Variance Analysis in a Finite World." *Journal of Financial and Quantitative Analysis*, 7, 4 (1972), pp. 1873-1880.

Hodrick, Robert. "Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement." *Review of Financial Studies*, 5 (1992), pp. 357-386.

Holton, Glyn A. "Time: The Second Dimension of Risk." *Financial Analysts Journal*, November-December 1992, pp. 38-45.

Ibbotson, R., and R. Sinquefield. "Stocks, Bonds, Bills and Inflation: Simulations of the Future 1976-2000." *Journal of Business*, 49, 1 (1976), pp. 22-33.

Jorion, Philippe. "Bayes-Stein Estimation for Portfolio Analysis." *Journal of Financial and Quantitative Analysis*, September 1986, pp. 279-292.

——. "International Portfolio Diversification with Estimation Risk." *Journal of Business*, July 1985, pp. 259-278.

——. "Portfolio Optimization in Practice: An Application to International Diversification." Working Paper, Columbia University, 1990.

Kandel, S., and R.F. Stambaugh. "On Correlations and Inferences About Mean-Variance Efficiency." *Journal of Financial Economics*, 18 (1987), pp. 61-90.

Lee, Wayne Y. "Diversification and Time: Do Investment Horizons Matter?" *Journal of Portfolio Management*, Spring 1990, pp. 21-32.

Markowitz, H. "Portfolio Selection." *Journal of Finance*, 7, 1 (1952), pp. 77-91.

Michaud, Richard. "The Markowitz Optimization Enigma: Is 'Optimized' Optimal?" *Financial Analysts Journal*, 45, 1 (1989), pp. 11-21.

Nelson, C., and M. Kim. "Predictable Stock Returns: Reality or Statistical Illusion?" *Journal of Finance*, 48 (1993), pp. 641-661.

Poterba, James M., and Lawrence H. Summers. "Mean Reversion in Stock Prices, Evidence and Implications." *Journal of Financial Economics*, 22 (1988), pp. 27-59.