Tests of conditional mean-variance efficiency of the U.S. stock market [☆]

Charles Engel [a,*], Jeffrey A. Frankel [b], Kenneth A. Froot [c], Anthony P. Rodrigues [d]

[a] *University of Washington and NBER, 1050 Massachusetts Ave, Cambridge, MA 02138, USA*
[b] *University of California, Berkeley and NBER, Cambridge, MA 02138, USA*
[c] *Harvard University and NBER, Cambridge, MA 02138, USA*
[d] *Federal Reserve Bank of New York, New York, NY, USA*

Final version accepted April 1994

Abstract

We test the mean-variance efficiency (MVE) hypothesis using a method that allows conditional expected returns to vary in relatively unrestricted ways. The method takes advantage of the predictability of conditional variances. The data estimate reasonably the price of risk, and the MVE model is valuable in explaining expected equity returns. Nevertheless, we reject the restrictions imposed by MVE on the alternative hypothesis. Unlike with most tests of MVE, we can put an explicit interpretation on the alternative hypothesis – a general linear Tobin portfolio choice model.

Keywords: CAPM; Portfolio balance; Mean-variance efficiency; GARCH

JEL classification: G12

1. Introduction

This paper tests the mean-variance efficiency (MVE) hypothesis using a method that allows conditional expected returns to vary in relatively unrestricted ways. It was first applied in a macroeconomic context in which the "market" portfolio included not only equities, but also money, bonds and real estate. It has since been applied more widely to other portfolios. Recent implementations of this method make use of the predictability of the conditional variances of individual asset returns. [1] Here we apply the method to the U.S. stock market.

The hypothesis of mean-variance efficiency of a given portfolio of assets states a relation between the conditional expected returns on the vector of assets in the portfolio, and the conditional covariance of those returns with a measure of "benchmark" returns. This measure of benchmark returns is a weighted average of the returns on the vector of assets, where the weights are given by the portfolio shares. Early tests of this hypothesis took both the conditional means and conditional covariances of returns to be constant over time.

More recently, tests have allowed for the fact that both the conditional expected returns and the conditional covariance with the benchmark vary over time. The process for expected returns and for the covariance of returns with the benchmark cannot be modeled independently because mean-variance efficiency requires that they be proportional. However, MVE is not a full model. To close the model one needs, in addition to the MVE restriction, either a model of conditional expected returns or a model of the conditional covariance of returns with the benchmark.

A recent influential test of MVE that allows for time-varying moments is Harvey (1989). Harvey assumes that conditional expected returns are linear in observable economic data such as measures of the term premium, junk bond premium or market dividend yield. This model, combined with MVE, then determines the behavior of the covariance of returns with the benchmark.

The approach of this paper is to model the conditional covariance of the returns with the benchmark, and then derive the model of expected returns from the MVE restrictions. It is more natural, we believe, to choose this approach than the opposite approach of Harvey. First, finance texts typically describe an asset-pricing model based on MVE as one in which expected returns are determined by the covariance of the return with the benchmark, and not vice-versa. Second, in practice, applied work has been successful at developing models which help to predict variances of returns on financial assets (e.g., Bollerslev's (1986) GARCH model), while it has been somewhat less successful in developing convincing and

---

[1] See Frankel (1982, 1985a), Frankel and Dickens (1984), Frankel and Engel (1984), and Ferson, Kandel and Stambaugh (1987) for early applications of a constant-variance version of our method. Bollerslev, Engle and Wooldridge (1987), Engel and Rodrigues (1989, 1993), Giovannini and Jorion (1989) and Ng (1991) estimate versions which allows for changing conditional second moments.

robust models to forecast returns. The two approaches are similar in that they test the cross-equation restrictions imposed by MVE while maintaining an auxiliary assumption about the predictability of return moments. The power of the tests of MVE using both our method and Harvey's depends on the predictive ability of this auxiliary model. Since we have considerable ability to forecast covariances, our tests of MVE should be more powerful than tests that depend upon an auxiliary model of expected returns.

Furthermore, our technique nests MVE in a more general, but economically meaningful, theory of portfolio determination. In contrast, most tests of the null hypothesis of MVE have no clear alternative hypothesis. This feature is particularly important because many tests do in fact reject MVE. When one rejects the null, it is crucial to have some idea of the alternative hypothesis. In the central tests below, the alternative to MVE is that investors' portfolio shares are linearly related to expected returns, but that investors' asset demands are not determined in the precise way implied by MVE. The alternative hypothesis is the more general portfolio-balance approach to asset demands that was first introduced by Tobin (1958, 1969) and concerns the relationship between expected returns and the demand for bonds and other assets. Sharpe's (1964) CAPM grew out of an attempt to place more structure on Tobin's portfolio balance model by modeling the behavior of individuals as mean-variance optimizers (as in Markowitz (1952) and Tobin (1958)). However, most modern testing of CAPM has departed from this original context.

The covariance of the vector of asset returns with the benchmark can vary over time for two reasons – the portfolio weights change over time, and the covariance matrix of individual asset returns is time-varying. In this paper, we allow the second moments of the individual asset returns to vary according to a GARCH process. The GARCH model has enjoyed notable success in forecasting conditional variances of asset returns. Alternatively, we could have modeled the variances as depending on observable economic data. [2]

We first estimate a version of the MVE model in which the variances of individual asset returns are assumed to be constant over time. This model still allows for unrestricted variation in the expected returns on equity portfolios. However, the constant variance version of MVE has two significant shortcomings from an empirical standpoint. First, the model has no predictive ability for expected returns. Second, it can be rejected in favor of a version of MVE in which the covariance of the asset returns follows a GARCH process.

---

[2] See for example Bollerslev (1987) and Bollerslev, Engle and Wooldridge (1988) for evidence on the predictability of conditional variances of excess returns. Engel and Rodrigues (1989, 1993) employ methods similar to the ones in this paper on international portfolios, but model the variances as functions of observable economic data such as gold prices, oil prices and the dollar LIBOR.

The GARCH version of MVE does have statistically significant ability to predict stock prices. So, a version of MVE in which market betas vary conditionally both because of changes in asset shares and time variation in the covariances of individual asset returns turns out to have explanatory power and to produce plausible parameter estimates. (The estimated measure of aggregate relative risk aversion is approximately 3.0.) This may be relevant to the recent findings of Fama and French (1992) that betas based on unconditional covariances have no predictive ability once size is included as an explanatory variable. Although we do not address that issue directly, the suggestion of our findings is that the CAPM might have performed better in the Fama and French setting if the betas were conditional on contemporaneous information. Nonetheless, we find the restrictions that this version of MVE imposes on a GARCH version of the Tobin portfolio-balance model are rejected.

Sections 2 and 3 briefly describe the model and the data, respectively. Section 4 tests for the constant-variance MVE and for our specification of GARCH-MVE. Section 5 summarizes our hypothesis tests and offers our conclusions.

2. The model

Mean-variance efficiency implies that the vector of conditional risk premia is a linear combination of the asset shares in the portfolio, with the weights proportional to the conditional variance of asset returns:

$$E_t(r_{t+1}) = \rho \Omega_t \lambda_t, \tag{1}$$

where $E_t(r_{t+1})$ is the expected return above the riskless rate on an $N \times 1$ vector of assets conditional on all information available at time t, Ω_t is the conditional variance of returns between t and $t+1$, λ_t is the $N \times 1$ vector of portfolio weights, with $\sum_{i=1}^{N} \lambda_{t,i} = 1$, and ρ is the price of risk equal to $E_t(m_{t+1})/\mathrm{Var}_t(m_{t+1})$, where m_{t+1} is the return on the aggregate portfolio. If the aggregate stock portfolio is the "market" portfolio, MVE is equivalent to CAPM, and the parameter ρ is to be interpreted as the coefficient of relative risk aversion. We treat ρ as a parameter of tastes. Note that the right-hand side of (1) is equivalent to the risk-adjusted conditional expected return on the aggregate (or benchmark) portfolio,

$$E_t(r_{t+1}) = \beta_t E_t(m_{t+1}), \qquad \text{where} \tag{2}$$

$$\beta_t = \frac{\mathrm{cov}_t(m_{t+1}, r_{t+1})}{\mathrm{var}_t(m_{t+1})} = \frac{\Omega_t \lambda_t}{\mathrm{var}_t(m_{t+1})}.$$

This expression makes it clear that the vector of sub-portfolio β_ts varies both with the shares of assets in the portfolio, λ_t, and the conditional covariance matrix, Ω_t, and thus may move substantially over short time intervals.

Under rational expectations, we can replace the vector of expected excess returns with the actual returns by including a prediction error that is orthogonal to all information at time t:

$$r_{t+1} = \rho \Omega_t \lambda_t + \varepsilon_{t+1}, \tag{3}$$

where $\varepsilon_{t+1} = r_{t+1} - E(r_{t+1})$. The insight in Frankel (1982) was that information about the conditional covariance matrix of returns can be obtained from the error terms, since under MVE:

$$\Omega_t = E_t(\varepsilon_{t+1}\varepsilon'_{t+1}). \tag{4}$$

MVE therefore imposes a set of restrictions that are highly nonlinear in that they constitute proportionality between the coefficient matrix and the variance-covariance matrix of the error term in (2).

To evaluate (4), we must take a position on how Ω_t changes over time. In section 4 below, we assume first that Ω_t is constant and then that it follows a GARCH process. We test the hypothesis that MVE holds against more general alternatives in which investors forecast excess returns as a function of asset shares and past prediction errors (as in the Tobin model).

The portfolio-balance model of Tobin (1958, 1969), representing a general relationship between asset demands and expected returns, can be written as:

$$\lambda_t = B_t E_t(r_{t+1}), \tag{5}$$

where B is an $N \times N$ matrix of coefficients. By inverting the system of equations in (5), we obtain an expression for expected excess returns,

$$E_t(r_{t+1}) = A_t \lambda_t, \tag{6}$$

where $A_t = B_t^{-1}$. This system of equations representing the portfolio-balance model is a generalization of MVE. MVE imposes the restriction that the matrix of coefficients A_t be proportional to the variance of the forecast error, ε_{t+1}.

Hence, MVE can be viewed as the null hypothesis in a test where the alternative hypothesis is the more general unconstrained portfolio balance model. Using ex post returns, (6) can be written:

$$r_{t+1} = A_t \lambda_t + \varepsilon_{t+1}. \tag{7}$$

Although the values of the equities are endogenous variables in an economic sense, they are still uncorrelated with the prediction errors, which under rational expectations are uncorrelated with all information available at time t. [3] We also test the MVE hypothesis above, as well as the more general alternatives, against

---

[3] Note that the N asset shares, $\lambda_{t,1} \dots \lambda_{t,N}$ are perfectly collinear because they sum to 1. This does not pose a problem for the estimation of (7), however, because the equations do not include a constant term.

an even more restrictive null hypothesis: that investors expect conditional excess returns to be zero. The results of our tests are discussed in sections 4 and 5.

Ng (1991) estimates a constrained version of MVE that is in most respects identical to our constrained model. However, her alternative hypothesis differs from ours. She tests the restriction that the intercept term in equation (1) is zero. As mentioned, our test is more in line with Harvey (1989), in the sense that we test cross-equation restrictions. Bollerslev, Engle and Wooldridge (1989) [4] estimate a model much like Ng's for a CAPM that includes treasury bills, bonds and an aggregate stock measure, but do not test the restrictions that CAPM places on the alternative Tobin model.

3. The data

Our tests use monthly stock returns from the New York and American Stock Exchanges from January 1955 to December 1984. To ease the computational burdens in estimating (3) we aggregate the stocks into $N = 7$ industry portfolios. [5]

Table 1 describes the aggregation of stocks into industry portfolios. The returns for each portfolio are value-weighted average returns. The $N \times 1$ vector of portfolio shares, λ_t, is the value of the stocks in the portfolios as a fraction of the total value of all stocks. Because it is desirable to group together equities that have highly correlated returns, we tried to put similar industries into the same portfolio. [6] Stambaugh (1982) aggregates into 20 industries, roughly by type of final output. We further aggregate into 7 industries, combining some of Stambaugh's categories. Table 1 shows Stambaugh's 20 industries, as well as the 7-industry aggregation that we use to perform our maximum likelihood tests of MVE.

The value shares, λ_t, are used to predict excess returns between time t and $t + 1$. The shares are measured monthly from the last day of January 1955 to the last day of November 1984 (359 observations), while the returns are calculated as the dividend plus appreciation over the previous month beginning the last day of February 1955 and ending the last day of December 1984. All returns are nominal excess returns above the return on the one-month Treasury bill recorded by Ibbotson Associates (1986).

---

[4] Others have also used the estimation method in this paper on various portfolios of assets, but have not performed the test performed here. See fn. 1.

[5] It is easy to demonstrate that if the returns on the industry portfolios are computed using the asset shares as weights, that the MVE model of equations (1) or (2) holds for returns on industry portfolios.

[6] On the other hand, we would not want to include together the suppliers of intermediate products and the producers of final output in the same industry. When steel prices rise, the cost of producing autos increases so that it is possible that steel producers' profits rise when auto manufacturers' profits decline.

4. Tests of MVE

A. *Constant variance*

First we estimate a version of equation (3) – the relation between ex-post returns and the covariance of returns implied by mean-variance efficiency – assuming that the covariance matrix of returns, Ω_t, is constant over time. The N equation system (3) must be estimated by maximum likelihood techniques, imposing cross-equation restrictions between the matrix of coefficients in the regressions and the variance matrix of the regression errors. We assume that the ε_t

Table 1
Industry portfolios and S.E.C. codes

Stambaugh's Portfolios	
Industry	*S.E.C. Codes*
1. Mining	10, 11, 12, 13, 14
2. Food and Beverages	20
3. Textile and Apparel	22, 23
4. Paper Products	26
5. Chemical	28
6. Petroleum	29
7. Stone, Clay and Glass	32
8. Primary Metals	33
9. Fabricated Metals	34
10. Machinery	35
11. Appliances, Electric Equipment	36
12. Transportation Equipment	37
13. Miscellaneous Manufacturing	38, 39
14. Railroads	40
15. Other Transportation	41, 42, 44, 45, 47
16. Utilities	49
17. Department Stores	53
18. Other Retail Trade	50–52, 54–59
19. Banking, Financial, Real Estate	60–67
20. Miscellaneous	1, 4, 15–17, 21, 24, 25, 27, 30, 31, 46, 48, 70, 73, 75, 78–80, 82, 89, 99
7 Portfolios (combinations of the 20 portfolios)	
Portfolio	*Industry Portfolios*
1	1, 2, 3, 4, 20
2	5, 7, 8, 9
3	6
4	10, 11
5	12–15
6	16
7	17–19

have a conditional normal distribution. Note that the assumption that Ω is constant is not the same as the assumption of constant betas and expected returns. As we saw in the previous section, even with a constant covariance matrix, the betas, and hence the expected returns on all securities including the aggregate or "market" portfolio, will vary over time.

The constant-variance version of the model is very unsuccessful along two dimensions – it has no predictive power for returns, and it can be rejected in favor of a version of MVE with time-varying covariances. So, for space considerations we do not report the coefficient estimates.

Examination of equations (1) or (3) reveal that for the MVE-constrained model to have power in predicting equity returns, the estimated value for ρ must be significantly different from zero. Our maximum-likelihood estimate of ρ is 2.018, with a standard error of 1.466.

The assumption that the distribution of returns is normal may not be warranted. White (1982) argues under certain regularity conditions, the maximum likelihood estimates of the parameters assuming normality is consistent even if the true distribution is not normal. However, if the distribution of returns is not normal, the sample standard errors calculated under that assumption are not consistent estimates of the true standard errors of the coefficients. Following White (1982) and Amemiya (1985) we calculate a robust standard error for ρ of 1.454.

With either the usual standard error or the robust standard error, we find that ρ is insignificantly different from zero. The constant variance version of MVE does not have statistically significant power for predicting expected returns.

In the next sub-section we describe a version of the MVE in which Ω_t is allowed to vary over time according to a GARCH process. There we shall see that the GARCH-MVE model performs significantly better than the constant variance version.

Before passing onto the GARCH version, we should note that we have not in this case even discussed performing the crucial test of MVE against the more general Tobin asset-pricing model. The constant variance version of MVE performs so poorly along the other dimensions we have discussed that we would not learn much if we found the model were rejected against the Tobin model.

We should, of course, be careful about applying large sample results to a system estimation with so many parameters. Gibbons, Ross and Shanken (1989) demonstrate for a different test of MVE that the true size of the likelihood ratio test (as well as the Wald and Lagrange multiplier tests) can be very different from the nominal size in small samples. They attribute the small-sample problems to the leptokurtosis of the distribution of returns. In the next section, we model the variance matrix as following a GARCH process. Baillie and Bollerslev (1989) find that in the presence of the GARCH model, the sample conditional distribution is less leptokurtotic than the unconditional distribution. There is no guarantee, though, that given any conditional density, the addition of a GARCH process will remove all the excess kurtosis in the unconditional distribution.

B. Time-varying variance

In the estimates reported in section 4.A, we assumed that the return covariance matrix, Ω_t, was constant over time. In this section we turn to a model of time-varying conditional variances.

In simple regression models, the presence of heteroskedasticity often does not affect the consistency of coefficient estimates, although it does cause standard calculations of test statistics to be inconsistent. When the MVE restrictions are imposed, however, changes in variances imply changes in coefficient estimates, which in turn imply changes in expected excess returns. Since Ω_t is time dependent, the coefficient on the asset shares in the constrained model must also move over time. Hence, holding Ω_t constant would lead to inconsistent coefficient estimates.

Inspection of (2) makes it easy to see why it is important to allow for variation in Ω_t. There are two possible sources of variation in expected returns if ρ is constant: changes in asset shares, λ_t, and changes in Ω_t. Suppose, for example, that favorable news about a stock is announced. One could easily think of cases in which the price is pushed up, increasing the stock's share in the aggregate portfolio, even though its expected return is now lower with the news. If the market is mean-variance efficient, this can happen when the riskiness of the asset declines – its own variance falls, or its covariance with other assets decline. But, for the jth asset, this is exactly a change in the jth row of Ω_t. We choose to model variances empirically following a multivariate GARCH process, which is an extension of Bollerslev's (1986) univariate process.

Even with only 7 equations to estimate, the dimension of the GARCH problem can be quite large. For example, if we restrict ourselves to first-order GARCH in which the variances and covariances this period are related only to the squares and cross-products of forecast errors in the previous period and to the elements of one lagged value of the variance matrix, the problem is unmanageably large. There are 28 independent elements in the covariance matrix. If each element were linear related to the 28 lagged squares and cross products of the forecast errors and to each of the 28 elements of the lagged Ω matrix, there would be 1596 parameters to estimate.

Given the complexity of estimating the MVE-GARCH system, and given the limited amount of data, it is helpful to lower the number of GARCH coefficients. Our test of MVE uses a parsimonious version of GARCH, in which the model, has return variances given by

$$\Omega_t = P'P + G\varepsilon_t\varepsilon_t'G + H\Omega_{t-1}H. \tag{8}$$

We treat as parameters the upper triangular matrix P, and the diagonal matrices G and H. Under this formulation, each element of Ω_t is linearly related to its corresponding component in the matrix of cross-products of lagged forecast errors and in Ω_{t-1}. There are only 42 coefficients to estimate. This formulation enforces positive semi-definiteness on the covariance matrix Ω_t. It is a special case of the

Table 2
MVE estimates, GARCH, 7 assets

$$r_{t+1} = \rho \Omega_t \lambda_t + \varepsilon_{t+1}$$
$$\text{var}_t(\varepsilon_{t+1}) = \Omega_t = P'P + G\varepsilon_t\varepsilon_t'G + H\Omega_{t-1}H$$

The estimate of the coefficient ρ:
3.043
(1.604)

The estimate of the upper triangular matrix P:

0.01075	0.01314	0.01078	0.01227	0.01039	0.00716	0.00824
(0.00115)	(0.00072)	(0.00552)	(0.00615)	(0.00430)	(0.00628)	(0.00661)
	0.00625	−0.00006	0.00424	0.00338	−0.00099	0.00102
	(0.00363)	(0.00173)	(0.00135)	(0.00150)	(0.00079)	(0.00133)
		0.01087	−0.00050	0.00011	0.00046	0.00084
		(0.00571)	(0.00084)	(0.00120)	(0.00105)	(0.00226)
			0.00679	0.00134	−0.00045	−0.00020
			(0.00133)	(0.00313)	(0.00122)	(0.00209)
				0.00519	0.00057	−0.00004
				(0.00324)	(0.00176)	(0.00375)
					0.00663	0.00071
					(0.00402)	(0.00153)
						0.00194
						(0.00210)

The estimates of the diagonal elements of G:

0.18916	0.13561	0.25085	0.11798	0.08263	0.22818	0.12647
(0.05636)	(0.10108)	(0.13006)	(0.08464)	(0.14683)	(0.13649)	(0.08398)

The estimates of the diagonal elements of H:

0.94234	0.93786	0.92592	0.95185	0.96264	0.94444	0.97392
(0.08668)	(0.04566)	(0.02410)	(0.02053)	(0.01118)	(0.06946)	(0.02830)

(robust standard errors in parentheses)

two multivariate processes proposed by Baillie and Myers (1991) and an extension to GARCH of the ARCH model of Engel and Rodrigues (1989).

Table 2 reports the results of the estimation with the MVE restrictions imposed on the GARCH system.

Although the coefficients were estimated by maximum likelihood estimates assuming normality, this distributional assumption may not be warranted. We construct normalized residuals using our conditional variance estimates. A multivariate Jarque-Bera test for the joint normality of the seven series of residuals, constructed by comparing the skewness and kurtosis of the residuals with the normal, has a value of 105.60. Under the null hypothesis the test statistic has a χ^2_{14} distribution – hence we strongly reject the normality assumption, and we report the robust standard errors calculated as in White (1982).

There are three hypotheses to test. The first asks whether we can reject the constant-variance MVE model in favor of the GARCH-MVE. A rejection would imply that time-varying variances statistically reduce the distance between the

stock-market portfolio and the mean-variance efficient frontier. The second question is whether we can reject the null hypothesis that the MVE-constrained model has no predictive power for expected returns. If we reject these first two null hypotheses, we can conclude that the GARCH-MVE has significant ability to explain expected returns. Rejection of these two hypotheses would lead us to the other interesting question: can we reject the restrictions implied by MVE on the unrestricted GARCH cum portfolio-balance system in equation (7)?

The constant-variance version of MVE is a special case of the GARCH-MVE model, in which the G and H matrices from (8) are constrained to be zero. This imposes 14 constraints on the GARCH system. The likelihood ratio test statistic is 107.04 and under the null has a χ^2_{14} distribution. We reject the constant-variance restrictions at the 99 percent level. GARCH therefore improves significantly on the constant-variance form of MVE. We can also use a likelihood ratio test to ask whether the GARCH system is overparameterized. That is, would a simple ARCH model of the variance, with the H matrix restricted to be zero, adequately model the behavior of Ω_t? The null hypothesis that H is zero imposes 7 restrictions on the GARCH model. The χ^2_7 statistic for this test is 76.22, so we can strongly reject the simple ARCH formulation in favor of the GARCH model.

Alternatively, the robust Wald statistic from Amemiya (1985) can be used to test the GARCH-MVE against the null of the constant variance version of the mean-variance efficiency hypothesis. Here, the Wald statistic is extremely large (6.362×10^5), and the null is strongly rejected. [7]

The elements of the matrix H are all exceed .9, and are statistically significant individually. The square of these coefficients would serve as a measure of the persistence of the diagonal elements of the variance matrix. There is evidently a great deal of persistence, which in turn implies that the risk premia on these assets are highly serially correlated.

The coefficient ρ is estimated to be 3.04, with a robust standard error of 1.60. If we make the additional assumptions required to obtain CAPM from MVE, ρ has the interpretation of being the coefficient of relative risk aversion. A value of 3 seems quite plausible, and does not imply excessive risk aversion as some other asset pricing models require in order to accord reasonably well with the data (see, for example, Mehra and Prescott (1985)).

Moreover, the fact that the t-statistic ($= 1.90$) is significantly different from zero with a p-value of .03 in a one-sided test, implies that the constrained MVE model with GARCH has statistically significant power in explaining equity returns ex ante. That is, the model is useful in predicting the excess returns on equities.

The next step is then to compare the GARCH model with the MVE constraints

---

[7] A standard Wald test, assuming the distributional assumption is correct, also yields a large chi-square statistic of 3.53×10^6.

imposed on the unrestricted Tobin portfolio balance model given by equation (7). MVE imposes the restriction that $A_t = \rho \Omega_t$, where Ω_t is the conditional variance of r_{t+1}.

The Tobin framework is general enough to incorporate any model in which the demand for the asset is a function of the expected returns. Tobin's (1969) paper discusses how the demand for the asset might be affected by the asset's liquidity, by tax treatment or by government regulation. One could incorporate other effects that lie beyond mean-variance analysis. To illustrate the portfolio-balance framework, we posit that the coefficient matrix contains the mean-variance component, $\rho \Omega_t$, but there is an additional constant matrix which may represent the effects of liquidity or tax treatment on asset demands.

So, we assume that in the unrestricted model, the coefficient matrix A_t evolves according to:

$$A_t = C + \rho \Omega_t. \tag{9}$$

where C is a constant matrix with 49 elements. The MVE constraint, that $A_t = \rho \Omega_t$, imposes 49 constraints on the unconstrained asset demand equations in (7). We test the null hypothesis of the MVE model given by equation (1) with equation (8) describing the evolution of Ω_t against the alternative of equation (7), with equation (9) modeling A_t and equation (8) modeling Ω_t.

Rather than estimate the full-blown unconstrained model, consisting of equations (7), (8) and (9), we test the restrictions that the MVE system, (1) and (8), put on this model using a Lagrange multiplier (LM) test. The LM test is useful in this context because it requires estimation only of the constrained model by maximum likelihood. [8] We use the robust LM statistic proposed by White (1992, p.8). The test statistic is distributed chi-square (49 d.f.) and has a value of 89.38. The constrained MVE model is easily rejected at the 99 per cent level. While the MVE model has predictive power, it can be rejected in favor of a more general portfolio-balance model.

5. Summary of conclusions

Figure 1 provides a graphical summary of our nested hypothesis tests. In the middle of the figure is the model that imposes mean-variance efficiency with GARCH covariances – the model given by equations (1) and (8).
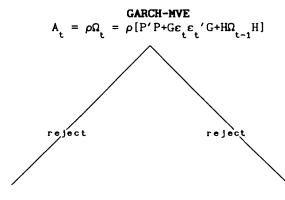
The GARCH-MVE is a useful model. It has significant predictive power for expected excess returns. That is, the model of investor risk-neutrality ($A_t = 0$) can be rejected.

---

[8] Berndt and Savin (1977) show that in linear problems the LM statistic is always smaller than the likelihood-ratio statistic, but do not extend the result to a non-linear test such as the one reported here.

TESTS OF THE MODEL

$$r_{t+1} = A_t \lambda_t + \varepsilon_{t+1}; \quad E_t(\varepsilon_{t+1} \varepsilon'_{t+1}) = \Omega_t$$

Unrestricted, GARCH Coefficients
$$A_t = C + \rho \Omega_t$$

reject

GARCH-MVE
$$A_t = \rho \Omega_t = \rho[P'P + G\varepsilon_t \varepsilon_t'G + H\Omega_{t-1}H]$$

reject   reject

Investor risk-neutrality   **Constant-variance MVE**
$$A_t = 0$$ $$A_t = \rho\Omega = \rho P'P$$

Fig. 1.

It is also apparent that allowing variances to be time-varying significantly improves the explanatory power for the constrained model. In section 4, we reported that the GARCH-MVE model significantly outperformed the constant variance MVE model.

It is interesting that the GARCH version of MVE has power in explaining equity returns. Allowing the covariances of asset returns to be time-varying significantly improves the predictive power of the constrained MVE model. Fama and French (1992) suggest that covariances, as reflected in a measure of an asset's unconditional beta, essentially have no power to predict returns when size is included as an explanatory variable. It seems unlikely that the predictive power in

our model occurs because our betas are correlated with size, because it is unlikely that the GARCH effects are related to firm size. Though we do no formal test of our model with size included as an explanatory variable, it appears that the covariance of asset returns does help predict the mean of asset returns, as CAPM would have it.

Still, we reject MVE in favor of the Tobin portfolio balance model. There are several ways to rationalize this rejection. One would be that the true asset-pricing model is not the CAPM, but rather a multi-factor CAPM, the APT, a version of the intertemporal CAPM, or perhaps a version of the one-period CAPM that allows for more investor heterogeneity in either tastes or information sets.

The portfolio balance model that we posit as the alternative is not very specific. An interesting direction to proceed would be to narrow the behavior allowed under the alternative. Because the Tobin model is so general, we have not ruled out many alternatives to MVE. On the other hand, the Tobin model does not include many of the dynamic elements present in models based on intertemporal optimization. We might interpret the rejection of MVE against the Tobin model as evidence that, while asset returns can be forecast ex ante in part by the asset shares, the relation is not determined by mean-variance optimization. Left for future research is a more specific answer as to why asset shares can predict returns, and whether other intertemporal elements can play a role in addition to asset shares in predicting returns.

Another explanation for the results would rely on the Roll (1977) critique. If the stock market is very unlike the true ''market'' portfolio, we would not expect to find MVE, even if the CAPM holds. [9] Indeed, under this explanation, the asset shares and GARCH processes cannot be accurately observed.

A third explanation of the results would be that the residuals in equation (2) lead to poor measures of the conditional variances. If ''peso problems'' affect stock market returns, the estimated residuals will be biased. Imposing the MVE restrictions only compounds the problems. For example, in the five years following the stock-market boom of August 1982, the market rose at an average annual rate of 22 percent. Few would argue in retrospect that it is possible to obtain from this period ex post, valid measures of ex ante expected risk and return.

One could imagine other reasons as well why the generalized Tobin portfolio balance is better able to describe the moments of a given sample of asset returns than the MVE model. The approach taken in this paper demonstrates that the MVE model is successful in its GARCH formulation at predicting excess returns, but it also suggests a class of less restricted models that are more successful.

---

[9] Similar results were found, however, when money, bonds, and real estate were allowed into the portfolio (Frankel, 1985a,b, and Frankel and Dickens, 1984) and when foreign assets were allowed (Frankel, 1982, and Frankel and Engel, 1984.)

References

Amemiya, T., 1985, Advanced econometrics (Cambridge: Harvard University Press).

Baillie, R.T. and T. Bollerslev, 1989, The message in daily exchange rates: A conditional variance tale, Journal of Business and Economics Statistics 7, 297–305.

Baillie, R.T. and R.T. Myers, 1991, Bivariate GARCH estimation of the optimal commodity futures hedge, Journal of Applied Econometrics 6, 109–124.

Berndt, E. and N.E. Savin, 1977, Conflict among criteria for testing hypotheses in the multivariate linear regression model, Econometrica 45, 1263–177.

Bollerslev, T., 1986, Generalized autoregressive conditional heteroskedasticity, Journal of Econometrics 31, 307–327.

Bollerslev, T., R.F. Engle, and J.M. Wooldridge, 1988, A capital asset pricing model with time-varying covariances, Journal of Political Economy 96, 116–131.

Cox, W.M., 1985, The behavior of Treasury securities; Monthly 1942–1984, Journal of Monetary Economics 16, 227–250.

Engel, C. and A.P. Rodrigues, 1989, Tests of international CAPM with time-varying covariances, Journal of Applied Econometrics 4, 119–138.

Engel, C. and A.P. Rodrigues, 1993, Tests of mean-variance efficiency of international equity markets, Oxford Economics Papers 45, 403–421.

Fama, E.F. and K.R. French, 1992, The cross-section of expected stock returns, Journal of Finance 47, 427–465.

Ferson, W.E., S. Kandel, and R.F. Stambaugh, 1987, Tests of asset pricing with time-varying expected returns and market betas, Journal of Finance 42, 201–220.

Frankel, J.A., 1982, In search of the exchange risk premium: A six-currency test of mean-variance efficiency, Journal of International Money and Finance 2, 255–274.

Frankel, J.A., 1985a, Portfolio shares as 'beta-breakers', Journal of Portfolio Management 11, 18–23.

Frankel, J.A., 1985b, Portfolio crowding-out, empirically estimated, Quarterly Journal of Economics 100, 1041–1065.

Frankel, J.A. and W. Dickens, 1984, Are asset-demand functions determined by CAPM?, University of California, Berkeley, 1984.

Frankel, J.A. and C. Engel, 1984, Do asset demand functions optimize over the mean and variance of real returns? A six currency test, Journal of International Economics 17, 309–323.

Gibbons, M.R., S.A. Ross and J. Shanken, 1989, A test of the efficiency of a given portfolio, Econometrica 57, 1153–1169.

Giovannini, A. and P. Jorion, 1989, The time-variation of risk and return in the foreign exchange and stock markets, Journal of Finance 44, 307–326.

Hansen, L.P. and S. Richard, 1987, The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models, Econometrica 55, 587–614.

Harvey, C., 1989, Time varying conditional covariances in tests of asset pricing models, Journal of Financial Economics 24, 289–317.

Ibbotson Associates, 1986, Stocks, bonds, bills and inflation: 1986 yearbook (Ibbotson Associates: Chicago).

Markowitz, H.M., 1952, Portfolio selection, Journal of Finance 7, 77–91.

Mehra, R. and E. Prescott, 1985, The equity premium: A puzzle, Journal of Monetary Economics 15, 145–161.

Ng, L., 1991, Tests of the CAPM with time-varying covariances: A multivariate GARCH approach, Journal of Finance 46, 1507–1521.

Roll, R., 1977, A critique of the asset pricing theory's tests; Part 1: On past and potential testability of the theory, Journal of Financial Economics 4, 129–176.

Sharpe, W.F., 1964, Capital asset prices: A theory of market equilibrium under conditions of risk, Journal of Finance 19, 277–293.

Stambaugh, R., 1982, On the exclusion of assets from tests of the two parameter model: A sensitivity analysis, Journal of Financial Economics 10, 237–268.

Tobin, J., 1958, Liquidity preference as behavior toward risk, Review of Economic Studies 67, 8–28.

Tobin, J., 1969, The general equilibrium approach to monetary theory, Journal of Money, Credit and Banking 1, 15–29.

White, H., 1982, Maximum likelihood estimation of misspecified models, Econometrica 50, 1–25.