Estimating Time Varying Risk Premia in the Term Structure: The Arch-M Model

Robert F. Engle, David M. Lilien, Russell P. Robins

ESTIMATING TIME VARYING RISK PREMIA IN THE TERM STRUCTURE: THE ARCH-M MODEL[1]

By Robert F. Engle, David M. Lilien, and Russell P. Robins

The expectation of the excess holding yield on a long bond is postulated to depend upon its conditional variance. Engle's (1982a) ARCH model is extended to allow the conditional variance to be a determinant of the mean and is called ARCH-M. Estimation and inference procedures are proposed and the model is applied to three interest rate data sets. In most cases the ARCH process and the time varying risk premium are highly significant. A collection of LM diagnostic tests reveals the robustness of the model to various specification changes such as alternative volatility or ARCH measures, regime changes, and interest rate formulations. The model explains and interprets the recent econometric failures of the expectations hypothesis of the term structure.

KEYWORDS: Term structure, financial models, ARCH, risk premium, heteroskedasticity, nonlinear models.

1. INTRODUCTION

ALTHOUGH THE VALUATION of risk is the central feature of financial economics, the standard methods for measuring and predicting risk are extraordinarily simple and unsuited for time series analysis. As the degree of uncertainty in asset returns varies over time, the compensation required by risk averse economic agents for holding these assets, must also be varying. Time series models of asset prices must therefore both measure risk and its movement over time, and include it as a determinant of price. Any increase in the expected rate of return of an asset as it becomes more risky will be identified as a risk premium.

The importance of such risk premia in the term structure of interest rates has been highlighted by a series of papers which all find the traditional expectations hypothesis inadequate to explain the observed data. For some recent examples see Shiller (1979, 1981), Sargent (1979, 1972), Shiller, Campbell, and Schoenholtz (1983), Mankiw and Summers (1984), and Campbell (1984). Some of these are based upon tests which find the variance of long term rates too large to be consistent with the expectations hypothesis. Others are based on regression tests which essentially show that the implicit predictors of future interest rates, derivable from the term structure, are inefficient and biased. Information available at the time could have improved the accuracy of the forecasts. Stated another way, these tests find that the one period rate of return which should, ex ante, be unforecastable, could have been predicted using available information.

These findings are generally interpreted as implying either some form of less than fully rational expectations, or time varying premia on different term debt. Attempts by Shiller, Campbell, and Schoenholtz (1983) and Mankiw and Summers (1984) to model particular forms of irrational expectations were unsuccessful.

Consequently, the main thrust of this literature is to introduce the possibility of time varying term premia. Amsler (1984) and Pesando (1983) have extended Shiller's variance bounds to allow time varying term premia. Campbell (1984) and Mankiw and Summers (1984) estimate or derive statistics about the required properties of time varying term premia. The latter conclude: "Most of the changes in the slope of the yield curve reflect these changing liquidity premiums or expectations that do not satisfy the standard postulates of rationality. These results suggest the importance of developing models capable of explaining fluctuating liquidity premiums."

The key postulate in the current paper is that time varying premia on different term debt instruments can be well modeled as risk premia where the risk is due to unanticipated interest rate movements and is measured by the conditional variance of the one period holding yield. While this is in the spirit of Bodie, Kane, and McDonald (1983) and Fama (1976), new econometric techniques are needed to estimate and test this model and these are developed here.

The autoregressive conditional heteroscedasticity (ARCH) model introduced by Engle (1982a), explicitly models time varying conditional variances by relating them to variables known from previous periods. In its standard form the ARCH model expresses the conditional variance as a linear function of past squared innovations; in markets where price is a Martingale, price changes are innovations, and this corresponds precisely to the Mandlebroit (1963) observation: "Large changes tend to be followed by large changes—of either sign—and small changes tend to be followed by small changes..." The ARCH model is used to provide a rich class of possible parameterizations of heteroscedasticity.

This paper introduces the ARCH-M model which extends the ARCH model to allow the conditional variance to affect the mean. In this way changing conditional variances directly affect the expected return on a portfolio. This resolves many of the empirical paradoxes in the term structure. Variables which apparently were useful in forecasting excess returns are correlated with the risk premia and lose their significance when a function of the conditional variance is included as a regressor. Furthermore, the heteroscedasticity in the disturbances had biased the test statistics, leading to the false finding of significant variables.

This model is applied to six month treasury bills, to two month treasury bills, and to 20 years Aaa corporate bonds to determine whether there appear to be time varying risk premia and how large they are. Section 2 develops a theoretical model of the relationship between means and variances which is formulated as a statistical model in Section 3. Section 4 describes the ARCH-M model and Sections 5 and 6 present the applications. Section 7 is a summary.

2. A MODEL OF THE RELATION BETWEEN RISK AND RETURN

Risk averse economic agents require compensation for holding risky assets. In the simplest set-up of one risky asset with normally distributed returns and one riskless asset, the risk is measured by the variance of the returns from holding the asset, and the compensation by a rise in the expectation of the return. The

relation between the mean and the variance of the returns which will insure that the asset is fully held in equilibrium will depend upon the utility function of the agents and the supply conditions of the assets.

To investigate this relation we now suppose that in this two asset economy the variance of the payoff of the risky asset may change over time and consequently the price offered by risk averse agents will change over time. This equilibrium price determines the relation between the mean and variance of the excess returns from holding the risky asset and therefore how the risk premium is related to the variance of the returns.

Consider a world with two assets, one has price 1 and is perfectly elastically supplied at a sure total rate of return r. The other has a price p and yields a random total return q (denominated in units of the numeraire) which has mean θ and variance ϕ. Wealth W, measured in units of the riskless asset, is therefore allocated between shares of the sure asset x, and shares of the risky asset s, so that

(1) $W = ps + x.$

The excess return per dollar invested in shares of the risky asset is given by

$y = (q/p) - r,$

so that the mean and variance of the excess returns is given by

(2) $E(y) = \mu = (\theta/p) - r, \qquad V(y) = \sigma^2 = \phi/p^2.$

Agents maximize expected utility of the end-of-period wealth, which, assuming normality of the returns, means that only the first two moments of the distribution matter. Under constant absolute risk aversion, expected utility can be expressed by:

$$EU = 2E(qs + rx) - bV(qs + rx)$$

and it will be maximized by choosing

(3) $sp = \mu/(b\sigma^2).$

Now suppose ϕ has a time subscript and is known to agents although not to the econometrician. Then the equilibrium values of p, μ, σ^2, and s will also vary over time. If in equilibrium the *value* of the outstanding shares of the risky asset remains constant, then the mean return will be proportional to the variance of returns since $s_t p_t$ in (3) is a constant.

A convenient assumption is that the riskless asset is held in zero net supply so that r becomes endogenous. The value of the outstanding shares of the risky asset is simply W. The mean and variance will therefore be proportional regardless of the supply elasticity of s if both wealth and b are constant. Such a model, however, leaves no role for price in evaluating risk.

If, instead, the physical number of shares is fixed so that $s_t = s$ and r is fixed, then in equilibrium (4) can be rewritten

$\mu_t^2 + \mu_t r_t = bs\sigma_t^2 \theta$

and, suppressing time subscripts,

(4) $\mu = [-r + \sqrt{r^2 + 4bs\sigma^2\theta}]/2$

so that the mean will be zero when the variance is zero, the slope is always positive, and for large variance the mean is proportional to the standard deviation. Thus if ϕ varies over time, but r, s, and θ do not, the econometrician should expect to see a relation between observed means and variances of returns which makes them move in the same direction but not proportionally.

For more general utility functions b will itself be a function of other variables such as σ^2. Thus we can replace b in (4) with $b(\sigma^2)$. Furthermore, there may be some elasticity of supply of the risky asset so that

$$s = f(p) = f(\theta/(\mu + r))$$

can be substituted for s. With these two flexible functions it is possible to find a wide range of relationships between observed means and variances.

Thus in general, one might expect the mean to increase less than in proportion to the variance with the precise relation determined by the supply elasticity of the risky (and possibly the riskless) asset and the risk preferences of agents. This paper introduces some empirical evidence on this relationship.

3. FORMULATION OF THE MODEL

Letting μ_t be the risk premium, y_t the excess holding yield on a long bond relative to a one period treasury bill, and ε_t the difference between the ex ante and ex post rate of return which in efficient markets would be unforecastable,

(5) $y_t = \mu_t + \varepsilon_t,$ $\text{Var}(\varepsilon_t | \text{all available information}) = h_t^2.$

It is assumed that the risk in holding a long bond is not diversifiable so that only the variance matters. The initial specification takes the mean as a linear function of the standard deviation:

(6) $\theta_t = \beta + \delta h_t.$

A nonzero value of β might reflect the linearization of a nonlinear function such as that derived above, or a preferred habitat argument. The choice of the standard deviation represents the assumption that changes in variance are reflected less than proportionally in the mean. Empirically, the log of h_t is found to be even better.

A complication in the interpretation of θ_t arises from the differential tax treatment of capital gains and interest income. Under the tax laws, long term capital gains are taxed at a lower rate than ordinary interest income and short term capital gains. This feature of the tax system makes a strategy of investing in long term bonds more desirable than rolling over short term paper. Investors can, to a large extent, treat one period capital losses as ordinary income for tax purposes by selling the bond and realizing their losses. Short term capital gains can be turned into long term capital gains for tax purposes by holding the bond

for a year or longer. Because this choice can be made ex post, after Y_t is observable, risk neutral investors should be willing to hold long term bonds at a lower expected pre tax yield than is paid on treasury bills. This tax advantage may explain the fact that the average value of Y_t for many types of long term bonds, has been below the average short term treasury bill rate over the last 30 years. We might therefore expect $\beta < 0$.

To complete the specification of the model, h_t^2, the conditional variance, must be parameterized as a function of the information set available to investors. We assume that the most useful information to agents are the previous innovations or surprises ε_t. If these have been large in absolute value then, extending Mandlebroit's observation, they are likely to be large in the future. In its simplest form we postulate that

$$(7) \qquad h_t^2 = \alpha_0 + \alpha_1 \sum_{i=1}^{p} w_i \varepsilon^2{}_{t-i}.$$

The conditional variance as observed by both the economic agents and the econometrician is a weighted sum of past squared surprises. One can discount older innovations in this weighting scheme.

Other variables which are in the information set at time t could also be introduced into (7) in the fashion of more traditional heteroscedasticity corrections. One such suggestion would be to use the squared changes in price as analyzed by Mandlebroit. Such a specification misses the fact that in the bond market a portion of the price change may be anticipated and this information is unlikely to be useful in forecasting changes in variance.

In the next section, the estimation and testing of the model in (5), (6), and (7) is considered in a more general context. In the following three empirical analyses, many of the caveats discussed above are then put to test.

4. ESTIMATING AND TESTING THE ARCH-M MODEL

The economic model described in the previous section incorporates an important extension of Engle's (1982a) ARCH model or in fact any heteroscedastic model; not only are the disturbances heteroscedastic, but the standard deviation of each observation affects the mean of that observation. In this section the estimation and testing of such models, called ARCH in mean or ARCH-M models, is discussed.

The general setup is given by

$$(8) \qquad Y_t \mid X_t, \Pi_t \sim N(\beta' X_t + \delta h_t, h_t^2),$$

$$(9) \qquad h_t^2 = \alpha' W_{\eta t} + \gamma' Z_t,$$

where X_t and Z_t are $k \times 1$ and $j \times 1$ vectors of weakly exogenous and lagged dependent variables, as in Engle, Hendry, and Richard (1983). The vector Z_t includes a constant whose coefficient represents the constant variance component of h_t. The $p \times 1$ vector $\eta_t' = (\varepsilon_{t-1}^2, \ldots, \varepsilon_{t-p}^2)$ where ε_t are the disturbances given

by $Y_t - \beta' X_t - \delta h_t$. The matrix W is a $q \times p$ array of fixed constants which may be used to impose restricted parameterizations on the response of the conditional variance to past squared residuals. In the most unrestricted case, W would be the identity matrix. The variance parameter vectors α and γ are therefore $q \times 1$ and $j \times 1$ respectively while the mean parameter vectors β and δ are $k \times 1$ and 1×1. These parameters can be combined into $\phi' = (\alpha', \gamma', \beta', \delta)$, an $m \times 1$ vector where $m = q + j + k + 1$.

Conditional on the initial values of all the data, the log likelihood function can be expressed as

$$(10) \qquad L(\phi) = \sum_t L_t(\phi); \quad L_t(\phi) = -\log h_t - \varepsilon_t^2 / 2h_t^2.$$

In practice, the presample values of the disturbances are set to their expectation, zero. The first order conditions for a maximum of this likelihood are given by:

$$(11) \qquad \partial L_t / \partial \phi = \sum ([\varepsilon_t^2 - h_t^2 - h_t \delta \varepsilon_t] h_t^{-4}) \, \partial h_t^2 / \partial \phi / 2$$
$$- \sum [\varepsilon_t / h_t^2][\partial \beta' / \partial \phi$$

The derivatives of the parameters with respect to ϕ are simply matrices with zeros and ones which select which terms to include for each derivative. The second line of (11) is the term relevant for GLS estimation of the regression coefficients without ARCH complications, that is when $\alpha = 0$. The expression in (11) gives the standard ARCH model when δ is zero.

The primary complexity introduced in this model comes in evaluating $\partial h^2 / \partial \phi$. From (9) this depends upon the derivatives of previous innovations with respect to the parameters. Yet these derivatives in turn depend upon the past derivatives of h with respect to the parameters if δ is nonzero. The desired derivatives must be computed recursively from an assumption that the initial values do not depend upon the parameters.

In the early analyses presented in Engle, Lilien, and Robins (1982) summarized in Section 5, analytical derivatives were calculated recursively and used to evaluate (11). However, numerical derivatives gave similar results, were simpler to compute and gave added flexibility to changes in specification. They therefore are probably the preferred approach for the ARCH-M model.

Estimation and testing can simply be carried out in terms of these derivatives. $\partial L / \partial \phi$ can be written compactly in terms of the $T \times m$ array S with typical element

$$[S]_{ti} = \partial L_t / \partial \phi_i$$

as

$$(12) \qquad \partial L / \partial \phi = S' i$$

where i is a $T \times 1$ unit vector so the first order condition is simply

$$S' i = 0.$$

The Hessian of the log likelihood is the sum of the Hessians of the t conditional log likelihoods, L_t. Under the assumption that the likelihood function is correctly specified,

$$\mathscr{I}_t = E[\partial L_t/\partial\phi \; \partial L_t/\partial\phi'] = -E[\partial^2 L_t/\partial\phi \; \partial\phi']$$

where $\mathscr{I}_t$ is the information matrix of the tth observation. Defining the information in the sample $\mathscr{I}$ is the average of the information over each observation,

$$\mathscr{I} = E[S'S/T].$$

Under slightly stronger conditions, $S'S/T$ is also consistent for $\mathscr{I}$.

A ready solution to the maximization of this likelihood function is to adopt the Berndt, Hall, Hall, Hausman (1974) approach using the iteration

(13) $\qquad \phi^{i+1} = \phi^i + \lambda (S'S)^{-1} S'i$

with λ as a step length which is adjusted from its a priori value of unity by a simple line search, and S as the matrix of first derivatives evaluated at ϕ^i.

The likelihood is in the form analyzed by Crowder (1976). Under sufficient regularity conditions, a solution to (13) will have the property that

(14) $\qquad (S'S)^{1/2}(\phi^* - \phi^0) \overset{A}{\sim} N(0, I)$

where ϕ^* is the maximum likelihood estimator obtained from (13) and ϕ^0 is the true value of the parameters. Unlike the simple ARCH model, this information matrix is not block diagonal between the parameters of the mean and the parameters of the variance.

Pantula (1984) has carefully investigated regularity conditions sufficient to guarantee (14) in the simple first order ARCH case. His conditions are stronger than can be accepted for this study in that he requires the existence of eighth order moments of the disturbance which are only finite for very small values of the ARCH parameter. Weiss (1986) has suggested some slightly weaker conditions; however, neither has addressed the ARCH-M model. Thus the appropriateness of the asymptotic distribution theory for this analysis remains a conjecture at this point.

Subject to the above caveat, inference procedures are available directly from (14). In particular, Wald tests can be computed in standard fashion. Lagrange multiplier tests can be simpler if the model has already been estimated under the null hypothesis and are easily constructed from the matrix of scores, S. Suppose the null hypothesis specifies that $\phi \in \Phi^0$ which is a proper subset of Φ. Denote by S^0 the matrix of scores calculated assuming the more general model to be true, but evaluated at the parameter estimates under the null. The scores corresponding to the restricted parameters are the Lagrange multipliers, and their variances are given by the information matrix. The LM test can be constructed as

(15) $\qquad \Phi_{LM} = i'S^0(S^{0'}S^0)^{-1}S^{0'}i$

$\qquad\qquad = TR_0^2$

where R_0^2 is the uncentered R^2 achieved by regressing the unit vector on the matrix of scores under the null. This statistic will asymptotically be chi squared with the number of degrees of freedom of the restriction when the null is true. This is easily computed from the R^2 of the first iteration of (13) starting from the estimates found under the null. Thus the tests take a form familiar from Engle (1982b, 1984) and it is recommended to construct a battery of diagnostics to convey information on the validity of the model both to the user and the reader.

The LM tests are convenient for testing restrictions in either the mean or the variance specification since reestimation may be costly and convergence is sometimes unsure. Tests are easily constructed for variables excluded from the mean such as interest rates or other functional forms. It is just as simple to test variance restrictions such as $\alpha = 0$, α is a set of linearly declining weights, or elements of γ are equal to zero (thereby testing for variables excluded from h). Many of the variance tests, however, may be interpreted as being on the boundary of the admissible parameter space so that one-tailed tests or other adjustments may be appropriate.

For the preferred models in this study h_t depended only on the intercept and a weighted average of past squared innovations where the weights are assumed to be linearly declining. These strong restrictions are subjected to a great variety of tests which allow changes in slope, seasonal spikes, freely estimated coefficients, and a wide variety of observable variables such as interest rates, volatility, and dummy variables for policy regimes. The models generally accept the more parsimonious specification at reasonable significance levels either because they are close to the true specification or because there is little power in the data to discriminate between alternative variance formulations. If the models with less restricted parameterizations are iterated toward convergence (for example to calculate a Wald or a likelihood ratio test) we found it difficult to prevent nonnegativities in the parameters regardless of the types of penalty functions or transformations considered. In this case there were likely to be many local maxima and generally the likelihood was ill-behaved. Thus the imposition of a parsimonious specification for the variance function such as linearly declining weights appears to be statistically supportable, computationally useful, and economically sensible.

5. THE RESULTS FOR SHORT TERM T-BILLS

Using Salomon Brothers data from the Analytical Record of Yields from 1960 through 1984 II on 3 and 6 month treasury bills, the excess holding yield, y_t, was calculated as:

$$y_t = [(1 + R_t)^2 / (1 + r_{t+1})] - (1 + r_t)$$

which is approximately

$$y_t \approx 2R_t - r_{t+1} - r_t$$

where R_t is the yield on a six month bill and r_t is the three month yield, each measured at the beginning of the quarter.

Regressing the excess holding yield on a constant gives

$$(16) \quad y_t = .142 + e_t, \quad s = .351,$$
$$\quad\quad\quad (4.04)$$

$$L = 51.1.$$

Thus, the mean of the excess holding yield over the sample period is .142 per cent at quarterly rates or .568 per cent at annual rates. The standard deviation is .35 at quarterly rates. From the linearized expression for the excess holding yield above, the average yield spread was half .568 per cent or .284 per cent at annual rates. The maximum return on a three month balanced portfolio obtained by borrowing at the three month rate and lending at the six, was 8.2 per cent at annual rates. The worst return occurred in the subsequent quarter and was −3.1 per cent. The rates of return from such portfolios are quite erratic and, as expected, are not large especially if transaction costs are important in forming these portfolios.

A glance at the solid line in Figure 1 confirms the changes in variance which are hypothesized by the ARCH-M model to account for the changing risk premia. The vertical axis is measured in quarterly percentage rates of return. Clearly, the period subsequent to the 1979 change in operating procedures shows substantially more variability than earlier periods; however, there are also earlier episodes of increased variability. Regressing the squared residuals on a fourth order linearly declining weighted average of past squared residuals gives the ARCH test as $TR^2 = 10.1$ which would be X_1^2 if there were no ARCH. There is clearly strong evidence of heteroscedasticity in the errors.


FIGURE 1—Excess hold yield of 6 month Treasury Bills and estimated risk premia.

Regressing the excess holding yield on a constant and allowing ARCH disturbances of fourth order gives:

$$(17) \qquad y_t = .048 + e_t, \qquad h_t^2 = .004 + 1.90 \sum_{\tau=1,4} w_7 \varepsilon_{t-7}^2$$
$$ \qquad (3.77) \qquad\qquad (.95) \ (7.3)$$

$$L = 85.17, \qquad w_7 = (5 - \tau)/10 \qquad\qquad\qquad (\tau = 1, \dots, 4).$$

The ARCH effect is very strong, showing a t statistic of 7.3. The magnitude is also very large as values over 1 imply nonstationary variance processes. The estimate of the mean changes dramatically when the high variance periods are given less weight in the regression; the constant term premium falls to .048 per cent at a quarterly rate or .2 per cent at annual rates.

The time varying risk premium has been swept into the disturbance term in (17) and represents misspecification. The hypothesized true model, as presented in Section 2, can be formalized as:

$$(18) \qquad y_t = \beta + \delta h_t + \varepsilon_t,$$

$$\varepsilon_t / \text{past information} \sim N(0, h_t^2),$$

$$h_t^2 = \gamma + \alpha \sum_{\tau=1,4} w_\tau \varepsilon_{t-\tau}^2, \qquad w_\tau = (5 - \tau)/10 \qquad\qquad (\tau = 1, \dots, 4).$$

The maximum likelihood estimates and their t statistics are:

$$(19) \qquad y_t = -.0241 + .687 \, h_t + e_t,$$
$$ \qquad (-1.29)\,(5.15)$$

$$h_t^2 = .0023 + 1.64 \sum w_\tau \varepsilon_{t-\tau}^2,$$
$$(1.08) \qquad (6.30)$$

$$L = 96.34, \qquad w_\tau = (5 - \tau)/10 \qquad\qquad\qquad (\tau = 1, \dots, 4).$$

As can be easily seen, all the slope coefficients are highly significant, indicating that there is not only an ARCH effect ($\alpha \neq 0$), but also a time varying risk premium ($\delta \neq 0$). The expected riskless return is negative but not significantly so and the minimum possible expected return which would be achieved if all recent forecasts had been precisely correct, is very small and positive (.0009). The risk premium is two thirds of the standard deviation of the return, which is quite substantial, indicating stronger risk aversion by the borrowers than the lenders in this market.

The parameter in the ARCH equation is above one which implies that the unconditional variance of the excess holding yield is infinite with a fat tailed distribution. The conditional distribution,which for most purposes is the relevant distribution, is of course still normal with a finite variance. An arbitrarily large return could occur if a sufficiently long string of innovations were all large. Such an episode would be easily reversed by a number of innovations near their median value of zero. Simulations of this situation show rather sensibly behaved series with larger bursts of volatility than would be expected from a marginally normal random variable. It is possible that the maximum likelihood estimates will not have their standard properties, but, as in the unit root case, they may have superior

convergence rates and correctly calculated standard errors. As mentioned in the previous section, the asymptotic distribution theory for this problem remains to be solved. The infinite unconditional variance may be related to the frequent failures of the variance bounds tests for interest rates.

A series of diagnostic tests were calculated for the model in (19). Although several were significant, the tests for the functional relationship between the risk and rate of return are of particular interest. LM tests for omitted variables h_t^2, exp (h_t), and log (h_t) were computed to test the assumed linearity between the standard deviation and mean of returns. Economic theory has little to say on the nature of this trade-off as it presumably depends on the risk preferences of the traders. Only the log variable was significant with a test statistic of 4.13. Estimating the model with both h_t and log (h_t) produced t statistics of 2.0 on the log and $-.4$ on the level and a log likekihood of $L = 101.62$, thereby confirming that the model with the log of standard deviation is superior to that in the level of the standard deviation.

The final preferred model is therefore:

$$(20) \qquad y_t = .355 + .135 \log h_t + e_t,$$
$$\qquad\qquad (4.38)\,(3.36)$$

$$h_t^2 = .005 + 1.48 \sum_{\tau=1,4} w_\tau \varepsilon_{t-\tau}^2,$$
$$\quad (2.22)\ \ (5.56)$$

$$L = 101.35, \qquad w_\tau = (5 - \tau)/10.$$

In this model all the coefficients are significant and the log likelihood is substantially above that of (19). The minimum term premium occurring when all past innovations are zero is now a very small negative value of $-.008$ per cent at quarterly rates.

Several sets of diagnostic tests were performed with this model as well. These are summarized in Table I. Volatility is defined by:

$$\text{Volatility} = \sum_{\tau=1,4} w_\tau y_{t-\tau}^2, \qquad w_\tau = (5 - \tau)/10,$$

so that it differs from the ARCH variance by the time varying risk premium. One would expect that the weighted average of residuals would give a better estimate of the true residual variance than the same function of the dependent variable; however there is no guarantee. Table I shows the robustness of the model in (20) to a variety of types of misspecification. None of the tests is significant at the 5 per cent level. The tests check for nonlinearities in the risk premium, volatility, structural shifts in October 1979, and misspecifications of the ARCH process through omitted variables or inappropriately applied constraints. The ARCH model with log Volatility alone achieves only log likelihood $L = 98.4$ although the significance and size of the variables is nearly the same as in (20).

The economically most interesting test is that for the yield spread and we turn to a more careful analysis of this model. Mankiw and Summers (1984) (MS) find that the yield spread is a significant and positive determinant of the excess holding

TABLE I

DIAGNOSTIC TESTS FOR ARCH-M MODEL (20)

Variable	TR^2		Distribution
Variables Omitted from the Mean			
h_t	.31	~	χ_1^2
h_t^2	1.67	~	χ_1^2
Volatility	1.44	~	χ_3^2
Log Volatility	.50	~	χ_1^2
Post October 1979 Dummy	.38	~	χ_1^2
r_t	.60	~	χ_1^2
R_t	.83	~	χ_1^2
$R_t - r_t$	2.92	~	χ_1^2
y_{t-1}	.14	~	χ_1^2
y_{t-4}	3.38	~	χ_1^2
Variables Omitted from the Variance			
Volatility	.27	~	χ_1^2
Post October 1979 Dummy	.07	~	χ_1^2
r_t	1.64	~	χ_1^2
R_t	1.60	~	χ_1^2
$R_t - r_t$	.90	~	χ_1^2
ε_{t-1}^2	.31	~	χ_1^2
ε_{t-4}^2	.62	~	χ_1^2
$\varepsilon_{t-1}^2, \varepsilon_{t-2}^2, \varepsilon_{t-3}^2$	3.11	~	χ_3^2
$\sum w_7 \varepsilon_{t-7}^2, w_7 = (13-\tau)/78, \tau = 1,\ldots,12$	.76	~	χ_1^2

yield. This implies a failure in the expectations hypothesis and a failure of an alternative hypothesis that long rates are overly sensitive to short rates. Our data set gives the following least squares estimate for this model:

$$(21) \qquad y_t = -.50 + 2.44\,(R_t - r_t) + e_t, \qquad \sigma = .312.$$
$$ (-1.10)\,(5.46)$$

The corresponding coefficient and t statistic in MS for the yield spread are 1.72 and 3.1 respectively. Their data set is a little shorter, from a different source and embodies the Shiller linearizations.

Adding the yield spread to model (20) gives:

$$(22) \qquad y_t = .325 + .130 \log h_t + .392\,(R_t - r_t) + e_t,$$
$$ (4.28)\,(3.59) \qquad\quad (2.58)$$

$$h_t = .004 + 1.64 \sum w_\tau \varepsilon_{t-\tau}^2,$$
$$ (1.38) \quad (4.86)$$

$$L = 103.48, \qquad w_\tau = (5 - \tau)/10 \qquad\qquad\qquad (\tau = 1,\ldots,4).$$

It now can be seen that by both Wald and LR tests the yield spread is a significant determinant at the 5 per cent but not 1 per cent level and by the LM test it is significant at the 10 per cent but not 5 per cent level. By economic standards the size of the coefficient on the yield spread has fallen dramatically from the least

TABLE II

ESTIMATES OF VARIOUS ARCH-M MODELS
EXCESS HOLDING YIELD OF 6 MONTH T-BILLS

Indep	59.1–84.2	59.1–71.3	71.4–84.2	59.1–79.3	61.3–74.1
Log h_t	.135	.092	.196	.177	.093
	(3.36)	(3.88)	(2.40)	(2.96)	(2.01)
Const.	.355	.272	.455	.446	.261
	(4.38)	(4.31)	(3.36)	(3.72)	(2.52)
ARCH α	1.48	1.67	1.49	1.25	1.20
	(5.56)	(5.15)	(3.57)	(4.60)	(2.84)

squares fit. The rest of the parameter estimates are very close to those obtained before in (20). Economically, it is not surprising to find some residual effect in the yield spread. The expected value of the spread is approximately proportional to the risk premium this period. Since it is highly autocorrelated, it will be a very good predictor of the risk premium next period. If information other than past innovations is useful in forecasting risk premia, then one might expect to find a significant coefficient on the past yield spread. A useful extension would be to allow the yield spread to directly influence the variance and consequently to indirectly influence the risk premium.

As much of the variance in interest rates is concentrated at the end of the sample period, the model was reestimated using subsets of the data. Surprisingly, the results are relatively insensitive to the sample period both in magnitude and in significance. See Table II.

Figure 1 plots the excess holding yield and the estimated risk premium. The scale is in quarterly percentage rates of return. The term premium rises to its highest value (.41 per cent quarterly or 1.64 per cent annual rates) in the fourth quarter of 1980. Over the sample period there are two values which are very slightly negative. On average, the term premium is .14 per cent. Although the most interesting and noticeable rise in the term premium is 1979–1984, there are also relative increases in 1960, 1972, and 1975, each of which is accompanied by an increase in volatility of the excess holding yield.

6. MODELLING OTHER INTEREST RATES

Two additional interest rate series have been modelled using the ARCH-M model and more are in progress. The first is the monthly data set constructed by Fama (1976) on two month vs. one month treasury bills from 1953.1 to 1971.7. The data set differs from that used above in the sampling interval and in the sample period. In this case the holding period is naturally taken to be one month rather than one quarter and consequently the riskless asset is the one month treasury bill rather than two or three month treasury bills. If a quarter is the correct interval, then shorter lived assets must be rolled over at uncertain rates

and therefore, the short term asset would be the risky one. For a theoretical discussion of these issues see Woodward (1983).

The model in (18) was estimated directly although a longer lag was allowed in the ARCH process to give a comparable memory to the variance estimator. The results are:

$$(23) \qquad y_t = -.00052 + .80\, h_t, \qquad h_t^2 = c_0 + 1.13 \sum_{\tau=1,12} w_\tau \varepsilon_{t-\tau}^2,$$
$$ (-1.2) \qquad (4.7) \qquad\qquad\qquad (8.6)$$

$$w_\tau = (13 - \tau)/78 \qquad\qquad\qquad\qquad (\tau = 1, \ldots, 12).$$

These are quite similar to those in equation (19) where in both cases the ARCH parameters are in the explosive range and the coefficient of the standard deviation is highly significant with a value of .69 before and .8 here. The estimated risk premium is plotted in Figure 2.

A somewhat different result was obtained using 20 year AAA corporate bonds from 1953.1 to 1980.2. Assuming that the bonds are effectively infinitely lived, the one quarter excess holding yield can be expressed in terms of the quarterly yield to maturity, R_t, and the three month treasury bill rate, r_t:

$$y_t = R_t - r_t - 1 + R_t / R_{t+1}.$$

The average return from holding long term bonds and borrowing at the t-bill rate is $-.75$ per cent at quarterly rates or -3 per cent at annual rates. Thus bond holders have taken a loss over this sample period in spite of the fact that the average long term rate was 5.9 per cent while the short term rate was only 4.6 per cent. This is a consequence of unexpected increases in interest rates possibly due to unexpected acceleration of inflation.

Maximum likelihood estimation of (18) produced:

$$(24) \qquad y_t = -2.8 + .505\, h_t, \qquad h_t^2 = c_0 + .75 \sum_{\tau=1,4} w_\tau \varepsilon_{t-\tau}^2,$$
$$ (-2.2)(1.6) \qquad\qquad\quad (2.6)$$

$$w_\tau = (5 - \tau)/10 \qquad\qquad\qquad\qquad (\tau = 1, \ldots, 4),$$


FIGURE 2—Conditional standard errors of one month forward rates.

FIGURE 3—Conditional standard errors of quarterly holding yields for Moody's Aaa bond.

for the fourth order ARCH-M model, and

$$(25) \qquad y_t = -3.3 + .651 \, h_{t,} \qquad h_t^2 = c_0 + .86 \sum_{\tau = 1, 12} w_\tau \varepsilon_{t-\tau}^2,$$
$$\qquad\qquad (-2.4)(1.9) \qquad\qquad\qquad (3.4)$$

$$w_\tau = (13 - \tau)/78 \qquad\qquad\qquad\qquad (\tau = 1, \ldots, 12),$$

for the twelfth order model.

These estimates differ from the short end of the spectrum in that they no longer exhibit the explosive ARCH parameter, the coefficient on the risk premium is roughly the same size but has a larger standard error, and the intercept is considerably more negative. When (25) is estimated on data prior to 1978, the coefficient on h_t rises slightly to .84 but the t statistic falls to 1.7. Thus the same model appears to be appropriate; however, the significance falls due to the omission of the highly volatile period of 1979 and 1980. The estimated risk premium is plotted in Figure 3.

Further analysis of these two series is contained in Engle, Lilien, and Robins (1982).

7. CONCLUSIONS

The precision with which agents can predict the future varies significantly over time. In relatively quiet periods, like the mid-1960's, relatively accurate forecasts can be made and agents can speculate on the future without absorbing large risks. In volatile periods, like the early 1970's and early 1980's, forecasts are less certain and speculation is riskier. Risk premia therefore adjust to induce investors to absorb the greater uncertainty associated with holding the risky asset.

In this paper we have extended the simple ARCH technique of measuring conditional variances to the ARCH-M model where the conditional variance is a determinant of the current risk premium, and thus enters the forecasting equation of expected financial returns. Our results from applying this model to three different data sets of bond holding yields are quite promising. ARCH was clearly

present in the forecast errors of bond holding yields indicating substantial variation in the degree of uncertainty over time. This measure of uncertainty proved very significant in explaining the expected returns in two of the data sets, and was significant only at slightly more than the 5 per cent level for the third. We therefore conclude that risk premia are not time invariant; rather they vary systematically with agent's perceptions of underlying uncertainty.

While our initial results suggest the promise of the ARCH technique to applications that require the measurement of uncertainty, we feel that the current model is but a first step. The ARCH framework may be applied to more general models of uncertainty and risk. For example, the capital asset pricing model suggests that risk premia are not a function of simple risk, but rather of undiversifiable risk. Risk premia therefore depend on the covariance of asset returns with the returns of the market as a whole. The general ARCH framework may be extended to allow conditional covariances to vary, resulting in time varying risk betas. Such a model is beyond the scope of the current paper and is mentioned to give some indication of possible extensions of our much simpler approach.

Department of Economics, University of California—San Diego, La Jolla, CA 92093, U.S.A.

Department of Economics, University of California—Irvine, Irvine, CA 92717, U.S.A.

and

Trans World Oil, P.O. Box 1252, Hamilton 5, Bermuda

REFERENCES

AMSLER, CHRISTINE (1984): "Term Structure Variance Bounds and Time Varying Liquidity Premia," *Economics Letters*, 6, 137-144.

BERNDT, E. K., B. H. HALL, R. E. HALL, AND J. A. HAUSMAN (1974): "Estimation and Inference in Nonlinear Structural Models," *Annals of Economic and Social Measurement*, 4, 653-665.

BODIE, ZVI, ALEX KANE, AND ROBERT MCDONALD (1984): "Why Haven't Nominal Rates Declined?," *Financial Analysts Journal*, 40, 16-27.

CAMPBELL, JOHN Y. (1984): "Asset Duration and Time Varying Risk Premia," unpublished dissertation, Yale University.

CROWDER, M. J. (1976): "Maximum Likelihood Estimation for Dependent Observations," *Journal of the Royal Statistical Society*, Series B, 38, 45-53.

ENGLE, ROBERT F. (1982a): "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of U.K. Inflation," *Econometrica*, 50, 987-1008.

——— (1982b): "A General Approach to Lagrange Multiplier Model Diagnostics," *Journal of Econometrics*, 20, 83-104.

——— (1984): "Wald, Likelihood Ratio and Lagrange Multiplier Tests in Econometrics," *The Handbook of Econometrics, Vol. II*, Chapter 13, ed. by Z. Griliches and M. Intrilligator. Amsterdam: North Holland Publishing Co., pp. 776-826.

ENGLE, ROBERT F., DAVID LILIEN, AND RUSSELL ROBINS (1982): "Estimation of Time Varying Risk Premiums in the Term Structure," UCSD Discussion Paper 82-4.

ENGLE, ROBERT F., DAVID HENDRY, AND JEAN FRANCOIS RICHARD (1983): "Exogeneity," *Econometrica*, 51, 277-304.

FAMA, E. F. (1976): "Forward Rates as Predictors of Future Spot Rates," *Journal of Financial Economics*, 3, 361-367.

MANDLEBROIT, B. (1963): "The Variation of Certain Speculative Prices," *Journal of Business*, 36, 394-419.

MANKIW, N. GREGORY, AND LAWRENCE H. SUMMERS (1984): "Do Long-Term Interest Rates Overreact to Short-Term Interest Rates?" *Brookings Papers on Economic Activity*, 223-242.

PANTULA, SASTRY, G. (1984): "Autoregressive Conditionally Heteroscedastic Models," manuscript, Department of Statistics, North Carolina State University.

PESANDO, JAMES E. (1983): "On Expectations, Term Premiums, and the Volatility of Long-Term Interest Rates," *Journal of Monetary Economics*, 12, 467-474.

SARGENT, THOMAS J. (1972): "Rational Expectations and the Term Structure of Interest Rates," *Journal of Money, Credit and Banking*, 4, 74-97.

——— (1979): "A Note on Maximum Likelihood Estimation of the Rational Expectations Model of the Term Structure," *Journal of Monetary Economics*, 5, 133-143.

SHILLER, ROBERT J., JOHN Y. CAMPBELL, AND KERMIT L. SCHOENHOLTZ (1983): "Forward Rates and Future Policy: Interpreting the Term Structure of Interest Rates," *Brookings Papers on Economic Activity*, 173-217.

SHILLER, ROBERT J. (1979): "The Volatility of Long-Term Interest Rates and Expectations Models of the Term Structure," *Journal of Political Economy*, 87, 1190-1219.

——— (1981): "Alternative Tests of Rational Expectations Models: The Case of the Term Structure, *Journal of Econometrics*, 16, 71-87.

WEISS, ANDREW A. (1986): "Asymptotic Theory for ARCH Models: Estimation and Testing," *Econometric Theory*, 2, 107-131.

WOODWARD, SUSAN (1983): "The Liquidity and the Solidity Premia in the Term Structure," *American Economic Review*, 73, 348-361.