NBER WORKING PAPER SERIES

TIME-VARYING RISK PERCEPTIONS AND THE PRICING OF RISKY ASSETS

Benjamin M. Friedman

Kenneth N. Kuttner

Working Paper No. 2694

NATIONAL BUREAŬ OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 1988

This research is part of NBER's research program in Financial Markets and Monetary Economics. Any opinions expressed are those of the authors not those of the National Bureau of Economic Research.

NBER Working Paper #2694
August 1988

TIME-VARYING RISK PERCEPTIONS AND THE PRICING OF RISKY ASSETS

ABSTRACT

Empirical results based on two different statistical approaches lead to several conclusions about the role of time-varying asset risk assessments in accounting for what, on the basis of many earlier studies, appear to be time-varying differentials in ex ante asset returns. First, both methods indicate sizeable changes over time in variance-covariance structures conditional on past information. These changing conditional variance-covariance structures in turn imply sizeable changes over time in asset demand behavior, and hence in the market-clearing equilibrium structure of ex ante asset returns.

Second, at least for some values of the parameter indicating how rapidly investors discount the information contained in past observations, the implied ex ante excess returns bear non-negligible correlation to observed ex post excess returns on either debt or equity. The percentage of the variation of ex post excess returns explained by the implied time-varying ex ante excess returns is comparable to values to which previous researchers have interpreted as warranting rejection of the hypothesis that risk premia are constant over time.

Third, although for long-term debt the two statistical methods used here give sharply different answers to the question of how much relevance market participants associate with past observations in assessing future risks, for equities both methods agree in indicating extremely rapid discounting of more distant observations -- so much so that in neither case do outcomes more than a year in the past matter much at all. While the paper's other conclusions are plausible enough, the finding of such an extremely short "memory" on the part of equity investors suggests that the standard representation of equity risk by a single normally distributed disturbance is overly restrictive.

Benjamin M. Friedman
Department of Economics
Harvard University
Littauer 127
Cambridge, MA 02138

Kenneth N. Kuttner
Economic Research Department
Federal Reserve Bank of Chicago
230 South LaSalle Street
Chicago, IL 60690

The past two decades of empirical research on asset pricing, including studies of the term structure of interest rates in the monetary economics literature as well as studies of the market for corporate equities in the finance literature, have now provided a large body of evidence indicating a prominent role for time-varying risk premia in expected asset returns.[1] Part of this evidence consists of the repeated rejection of the hypothesis of equal ex ante returns on different classes of securities. Another part consists of the repeated rejection of the hypothesis that whatever differences exist between the ex ante returns on different classes of securities remain constant over time.

Despite the tenacity of many researchers' apparent dedication to these respective rejected null hypotheses, there is nothing startling a priori in the finding of risk premia that not only assume non-zero values but also vary over time. Unless the market for the assets in question is dominated by risk-neutral investors controlling large pools of capital (who are they? where does their capital come from?),[2] most standard theories of asset pricing imply that assets with different risk properties will bear different ex ante returns. Moreover, these same asset pricing theories also immediately suggest a variety of circumstances -- familiar examples include changing investor perceptions of asset risks, changing outstanding supplies of outside assets, and utility exhibiting non-constant risk aversion -- under which the ex ante return differentials that they imply will vary over time.[3] Indeed, only under specific assumptions ruling out these and other potential sources of change do standard risk-based asset pricing models imply ex ante return differentials that are constant at any non-zero level.

Each of the known major potential explanations for the variation of asset risk premia over time has attracted at least some attention. Increasingly, however, attention has focused on the possible role of changing investor risk perceptions. The most plausible reason for this emphasis is the finding, in studies reporting time-varying ex ante return differentials, that the differentials they have identified typically move about in a fairly volatile way over a short period of time. The mix of outstanding outside asset supplies changes both too slowly and too smoothly to account for this phenomenon. So do the factors governing the prevailing level of market-wide risk aversion in most models. By contrast, at least in principle the risks that investors associate with holding various assets may change quite quickly, either because the underlying probabilities have actually changed in some objective sense, or because investors' perceptions have changed independently, or both.

Because investors' risk perceptions are unobservable, for purposes of empirical research it is necessary to place some discipline on the presumed structure of risk -- either at a single point in time or as it varies through time -- used to infer ex ante asset return relationships. In the context of variations over time, the usual source of this discipline is some connection to the previously observed variation of actual asset returns. Modigliani and Shiller (1972), for example, used a moving-average standard deviation of observed short-term interest rates as an explanatory variable in an equation for the spread between long- and short-term interest rates.[4] Similarly, Friedman (1980) and Roley (1982) used moving-average variances of interest rates and equity prices as explanatory variables in equations for the demands for specific financial assets by different categories of individual or institutional investors. More recently, Friedman (1985) used a moving-sample vector autoregression to estimate the ex ante variance-covariance structure of U.S. asset returns, quarter by quarter over a period of years, and Bollerslev et al. (1988) used a generalization of Engel's (1982) autoregressive conditional heteroskedasticity method to estimate a time-varying ex ante variance-covariance structure and also use that structure in explaining the variation over time of observed U.S. asset return differentials.

The object of this paper is to explore how the variation of observed (ex post) asset returns affects investors' assessments of (ex ante) asset risks, and therefore how this variation in turn affects asset demand behavior and consequently the equilibrium structure of asset prices. To carry out this analysis, the paper relies on the respective estimates of the time-varying variance-covariance structure of ex ante asset returns delivered by a generalization of the vector autoregression (VAR) approach and by the generalized autoregressive conditional heteroskedasticity (GARCH) approach. As in both Friedman (1985) and Bollerslev et al. (1988), the analysis here focuses on three broad classes of assets traded in the U.S. financial markets -- short-term debt, long-term debt, and equity -- and applies the standard one-period capital asset pricing model as the simplest plausible framework within which to exploit and compare these two empirical approaches.[5]

From the perspective of any one point in time, both the VAR and the ARCH methods are ways of inferring investors' perceptions of these assets' risk properties from the observed variation of their ex post returns up until then. In addition, both are ways of constraining the movement of these inferred perceptions from one point in time to the next. Hence each method disciplines the inference of asset risk perceptions not only by (explicitly) tying them to observed return data but also by (implicitly) imposing certain smoothness properties. Nevertheless, the estimates of ex ante asset risks that the two methods deliver, and hence their respective implications for asset demand behavior and for equilibrium asset prices, need not be identical, or even similar. Comparing the different results given by these two approaches is therefore a further object of interest throughout the paper.

Section I briefly sets out the respective mechanics of the VAR and GARCH methods as applied to the estimation of ex ante asset risk, and makes explicit the parallels and differences between them. Section II compares the results of applying these two methods to infer the ex ante risk associated with returns on long-term debt and equity in the United States. Section III draws the implications of these respective results for investors' portfolio behavior over time. Section IV shows what these results for investors' portfolio behavior imply in turn for the market-clearing structure of equilibrium asset returns. Section V summarizes the paper's principal conclusions, and indicates a specific direction that they suggest for future research.

I. The Generalized VAR and GARCH Methods

Most theories of the pricing of financial assets assign a central role to perceptions, in the minds of the investors who hold the assets, of the risks associated with holding them.[6] For example, under the simplest version of the single-period capital asset pricing model, in which each investor takes asset returns to be joint normally distributed and maximizes a utility function characterized by constant relative risk aversion, the investor's set of asset demands at any time t is of the linear homogeneous form[7]

$$(1) \qquad \underline{\alpha}_t = B_t(\underline{r}_t^e + \underline{1}) + \underline{\pi}_t$$

where $\underline{\alpha}$ is a vector of portfolio proportions (summing to one), $\underline{r}^e$ is a vector of expected returns corresponding to the assets in $\underline{\alpha}$, $\underline{\pi}$ is a vector of proportions (also summing to one) describing the minimum-variance portfolio, and B is a matrix (with each column summing to zero) showing to what extent the investor will choose a portfolio different from $\underline{\pi}$ in response to incentives provided by non-uniform expected returns. Both B and $\underline{\pi}$ depend directly on the variance-covariance structure associated with $\underline{r}^e$. If none of the available assets is riskless -- as is plausible in most macroeconomic contexts, in which investors' utility depends on real variables, and inflation is stochastic -- then B and $\underline{\pi}$ are of the form

$$(2) \qquad \begin{aligned} B_t &= \frac{1}{\rho}\,[\Omega_t^{-1} - (\underline{1}'\Omega_t^{-1}\underline{1})^{-1}\,\Omega^{-1}{}_t\underline{1}\,\underline{1}'\,\Omega_t^{-1}] \\ \underline{\pi}_t &= (\underline{1}'\Omega_t^{-1}\underline{1})^{-1}\,\Omega_t^{-1}\,\underline{1} \end{aligned}$$

where Ω is the relevant variance-covariance matrix, and (for sufficiently small time periods to make the underlying expansion approximately correct) ρ is the coefficient of relative risk aversion. Alternatively, if one asset is riskless, the minimum-variance portfolio consists only of it, and the expression equivalent to (1) then gives the demands for the risky assets only as a function of the expected excess returns on the risky assets over whatever certain return the risk-free asset provides, with

(3) $$B_t = \frac{1}{\rho}\Omega_t^{-1}$$

and $\underline{\pi}$ equal to a vector of zeroes, where $\underline{\pi}$, B and Ω are now defined for the sub-vector of risky assets only.

The crux of the matter for empirical applications of such theories, of course, is that the ex ante variance-covariance matrix Ω is unobservable. Moreover, when stochastic asset returns vary systematically over time, the unconditional variance-covariance matrix computed over any specific sample probably overstates the degree of uncertainty perceived by investors who possess information about that systematic movement.[8] What matters for investors' asset demands, and hence for market-determined asset prices, is instead the variance-covariance structure conditional on whatever relevant information investors have.

Serial correlation is hardly the only kind of systematic variation that causes sample unconditional variance-covariance matrices to overstate the risks perceived by investors, but it is surely the most obvious kind. For an investment horizon of one calendar quarter, for example, the 91-day U.S. Treasury bill is risky only because inflation is uncertain. But the quarterly time series of price inflation in most countries is highly serially correlated. Over the 1960-85 sample used as the basis of the results presented below, the first-order serial correlation of the U.S. consumer price index was 0.75. Hence an investor who wants to forecast the real return to holding Treasury bills over the coming quarter, and who knows the most recent value of the inflation rate, can typically do much better than simply to subtract from the observed nominal interest rate the sample mean of 5.24%, and associate with this forecast the sample standard deviation of 3.92%.

<u>The VAR Approach</u>. The VAR method exploited in Friedman (1985) represents investors as taking account of generalized serial correlation in forming their asset return expectations by estimating, for each time period t, a vector autoregression of the form[9]

$$(4) \qquad \underline{r}_{\tau} = \Gamma(L|t-1)\underline{r}_{\tau-1} + \underline{\epsilon}_{\tau}$$

where $\underline{r}$ is a vector of realized asset returns corresponding to expectations $\underline{r}^e$, $\Gamma(L|t-1)$ is a matrix of polynomials in the lag operator estimated using observations on $\underline{r}$ up through (but not beyond) period t-1, and $\underline{\epsilon}$ is a vector of disturbances. After period t elapses, investors incorporate the new observation of $\underline{r}$ into the sample, re-estimate (4), and use the updated model to project $\underline{r}$ for period t+1.

For each time period t, (4) gives a conditional expectation $\underline{r}^e$. In addition, because (4) gives the estimated values of the disturbances $\underline{\epsilon}$ for each period up through t-1, it also gives the associated conditional variance-covariance structure

$$(5) \qquad \hat{\Omega}_t = E_{t-1}(\underline{\epsilon}_t\underline{\epsilon}_t').$$

Given this time-specific estimate of Ω, the minimum-variance portfolio and asset substitutability matrix describing investors' asset demands in period t follow from (1) - (3), or from any richer (perhaps intertemporal) model of portfolio behavior.

An intuitively helpful interpretation of the VAR method, therefore -- at least from the perspective of this paper's focus on time-varying asset return differentials -- is that it is a way of estimating the conditional variance-covariance structure at any time, as a function of past "surprise" (that is, nonsystematic) movements of realized returns. In particular, it allows the conditional variance-covariance structure to change over time as new observations of asset returns, and hence new estimates of the corresponding "surprises," become available.

A question that immediately arises in such a context is whether investors treat more distant observations as if they have the same information content, for this purpose, as more recent ones. Friedman (1985) experimented with an expanding-sample method in which the estimation of (4) uses all observations after a specific initial period, as well as a rolling-sample method in which each new observation takes the place of the most distant one (so that the sample size remains unchanged). A more general way to allow more distant observations to "matter" less than more recent ones -- and, to anticipate, a way that is consistent with the application of the GARCH method by Bollerslev et al. -- is to relate the conditional variance-covariance structure to the estimated past disturbances according to

(6) $$\hat{\omega}_{ijt} = \frac{1-\phi}{1-\phi^{t-1}} \sum_{k=0}^{t-2} \phi^k \hat{\epsilon}_{i,t-1-k} \hat{\epsilon}_{j,t-1-k}$$

where $\hat{\omega}_{ij}$ is the ij-th element of $\hat{\Omega}$, the $\hat{\epsilon}_i$ and $\hat{\epsilon}_j$ are elements of $\underline{\epsilon}$ as estimated in (4), and ϕ can be interpreted as an arbitrarily chosen "memory" parameter determining the relative weight placed on more distant observations for this purpose.[10] The (equal-weighted) expanding-sample method used in Friedman (1985), for example, is then just the limiting case of (6) in which $\phi \rightarrow 1$. Different values of ϕ will in general imply different values of $\hat{\Omega}$ at any given time, as well as different paths of $\hat{\Omega}$ over time, as the new information contained in the most recent forecast errors from (4) is incorporated with greater or lesser weight compared to previously available information.

The GARCH Approach. The GARCH method applied by Bollerslev et al. (1988) is an alternative way of constructing time-varying estimates of Ω. Specifically, the GARCH method directly models the respective elements of Ω at time t as a vector ARMA process, according to

$$(7) \qquad \text{vech}\,(\hat{\Omega}_t) = \theta + \sum_{k=1}^{n} \Psi_k \,\text{vech}\,(\hat{\underline{\epsilon}}_{t-k}\hat{\underline{\epsilon}}'_{t-k}) + \sum_{k=1}^{m} \Delta_k \,\text{vech}\,(\hat{\Omega}_{t-k})$$

where $\hat{\underline{\epsilon}}$ for any period is the estimated vector of unexpected components of asset returns $(\underline{r} - \underline{r}^e)$; θ, Ψ_k, k=1,...,n, and Δ_k, k=1,..., m, are matrices of coefficients to be estimated; and the vech operator stacks the lower triangular portion (that is, the unique elements) of a symmetric matrix. In the terminology of Engle (1982) and Kraft and Engle (1983), (7) is a GARCH(n,m) process.

That the θ, Ψ and Δ matrices consist of freely estimated coefficients in principle allows an extremely flexible pattern of both contemporaneous and intertemporal dependence among the elements of $\hat{\Omega}$, with each element dependent in general on all elements of $\hat{\Omega}$ and $\hat{\underline{\epsilon}}\hat{\underline{\epsilon}}'$ from prior periods. Because the model in this general form requires estimation of $(p(p+1)/2)(1 + (m+n)p(p+1)/2)$ parameters (where p is the number of equations), however, both Engel (1982) and Bollerslev et al. (1988) constrained the ψ and Δ matrices to be diagonal, so that the ij-th element of $\hat{\Omega}$ depends only on the respective ij-th elements of $\hat{\Omega}$ and $\hat{\underline{\epsilon}}\hat{\underline{\epsilon}}'$ from prior periods. This constraint reduces the number of parameters to be estimated to $(1+m+n)p(p+1)/2$.

Like the VAR method, the GARCH method can also be interpreted as a way of estimating the conditional variance-covariance structure at any time as a function of past "surprise" movements of returns. Furthermore, the GARCH(1,1) specification used both in Bollerslev et al. and in the analysis presented in this paper directly is analogous to (6) in relating each period's $\hat{\Omega}$ to past estimated "surprises" according to a pattern of geometrically declining weights, since with n=m=1 (7) reduces to

$$(8) \qquad \hat{\omega}_{ijt} = \theta_{ij} + \psi_{ij} \hat{\epsilon}_{i,t-1} \hat{\epsilon}_{j,t-1} + \Delta_{ij} \omega_{ij,t-1}$$

which is equivalent to

$$(9) \qquad \hat{\omega}_{ijt} = \frac{\theta_{ij}}{1-\Delta_{ij}} + \psi_{ij} \sum_{k=0}^{\infty} \Delta_{ij}^{k} \hat{\epsilon}_{i,t-1-k} \hat{\epsilon}_{j,t-1-k}.$$

Despite the strong correspondence, there are also several key differences between the two approaches -- as will be apparent in the results presented below. The difference that is most immediately apparent from comparing (6) and (9) is that the VAR method constrains $\hat{\Omega}$ to be homogeneous of degree one in past $\hat{\underline{\epsilon}}\hat{\underline{\epsilon}}'$ while the GARCH method does not. In other words, the VAR method as applied here not only constrains the weights in the function relating $\hat{\Omega}$ and $\hat{\underline{\epsilon}}\hat{\underline{\epsilon}}'$ to sum to unity, but also excludes a constant term, while the GARCH method does neither. An additional difference is that the GARCH method implicitly delivers an internal estimate of the memory parameter ϕ, while in the VAR method it is necessary to specify ϕ a priori.

In conjunction with some model like (1) - (3), solved out to deliver asset return expectations $\underline{r}^e$ for each time period, and hence facilitating calculation of the estimated error vector $\hat{\underline{\epsilon}}$, the GARCH method in the form (7) also generates a time series of estimates of the conditional variance-covariance structure $\hat{\Omega}$, and hence the implied asset demand behavior, for each period. The specific form of the model used by Bollerslev et al., and reproduced below, applies the single-period capital asset pricing model with a risk-free asset to express the expected returns on all risky assets in the excess return form

$$(10) \qquad \underline{r}_t^e = r_{ft} \cdot \underline{1} + \rho \hat{\Omega}_t \underline{\alpha}_t^s$$

where r_f is the return (taken to be exogenous and known with certainty ex ante) on the risk-free asset, $\underline{\alpha}^s$ is a vector stating the respective shares of the risky assets in total portfolio wealth (also taken to be exogenous and known with certainty ex ante), and the risk aversion parameter ρ is here a coefficient to be estimated.[11] An alternative procedure, which avoids the assumption of a risk-free asset, would be to use some independent process (an autoregressive model, for example) to generate expected inflation, subtract this expectation from the stated nominal return to calculate the expected short-term debt return r_S^e, and then use (1) - (3) to express the expected returns on all other assets (the vector $\underline{r}_{\cdot S}^e$) as

$$(11) \qquad \underline{r}_{\cdot St}^e = r_{St}^e \cdot \underline{1} + \rho \hat{\Omega}_{\cdot St} \underline{\alpha}_{\cdot St}^s$$

where the conditional variance-covariance matrix $\hat{\Omega}_{\cdot S}$ and supply vector $\underline{\alpha}_{\cdot S}^s$ likewise refer only to assets other than short-term debt. Either (10) or (11) generates ex ante asset returns $\underline{r}^e$ for each period and hence permits calculating for each period the estimated ex post "surprises" $\hat{\underline{\epsilon}}$ and, in turn, the conditional variance-covariance structure $\hat{\Omega}$ for the next period. Given $\hat{\Omega}$, the asset substitutability matrix and (if there is no risk-free asset) the minimum-variance portfolio again follow from (1) - (3).

II. Estimated Time-Varying Risk Structures

Table 1 summarizes the basic features of the after-tax quarterly ex post returns on long-term debt and equity, stated in each case in the form of excess returns over the corresponding return on short-term debt, based on quarterly U.S. data for 1960-85.[12] The pre-tax data are for the last day of each quarter, from Ibbotson and Sinquefield (1982, and subsequent updates). For each asset the corresponding after-tax return is calculated by applying the U.S household sector's average effective marginal tax rates in each year for interest, dividends, and capital gains to the corresponding respective components of the pre-tax returns.[13]

The sample means, shown in Table 1 in decimals at annual rates, indicate positive excess returns for both risky assets.[14] Not surprisingly, the unconditional sample variance-covariance matrix indicates a larger excess return variance for equity than for long-term debt, by a factor of somewhat less than two. Although the two respective excess returns exhibit quite different p case there is sufficient serial correlation to warrant the supposition that investors who use available information efficiently take account of it in forming their expectations.

VAR Results. Table 2 and Figures 1 and 2 show the results of applying the VAR method to estimate the quarter-by-quarter variation in the ex ante risk structure of these two excess returns. This procedure as applied here involves the estimation of a separate bivariate autoregression of form (4), in which the two stochastic variables are the two excess returns, for each of the 104 quarters in the 1960-85 sample. The procedure begins by using data spanning 1954:I - 1959:IV in the autoregression that generates $\hat{\Omega}$ for 1960:I, continues by adding one quarter's data at a time, and concludes by

Table 1

Summary Statistics for After-Tax Excess Returns 1960-85

Means

r_L	.0192
r_E	.0821

Unconditional Variance-Covariance Matrix

	r_L	r_E
r_L	.0564	
r_E	.0341	.0977

Serial Correlations

	r_L	r_E
L=1	-.061	.199
L=2	.025	-.173
L=3	.132	-.045
L=4	.007	-.075

Table 2

Average Conditional Variance-Covariance Matrix Estimated by VAR Method

ϕ	$Var(r_L)$	$Var(r_E)$	$Cov(r_L, r_E)$
1.0	.018	.072	.005
0.99	.021	.074	.008
0.9	.038	.085	.002
0.7	.042	.086	.025
0.5	.044	.086	.025
0.3	.045	.085	.025
0.2	.045	.085	.025

Figure 1: Conditional Covariance Matrix Elements, VAR (Phi=1)

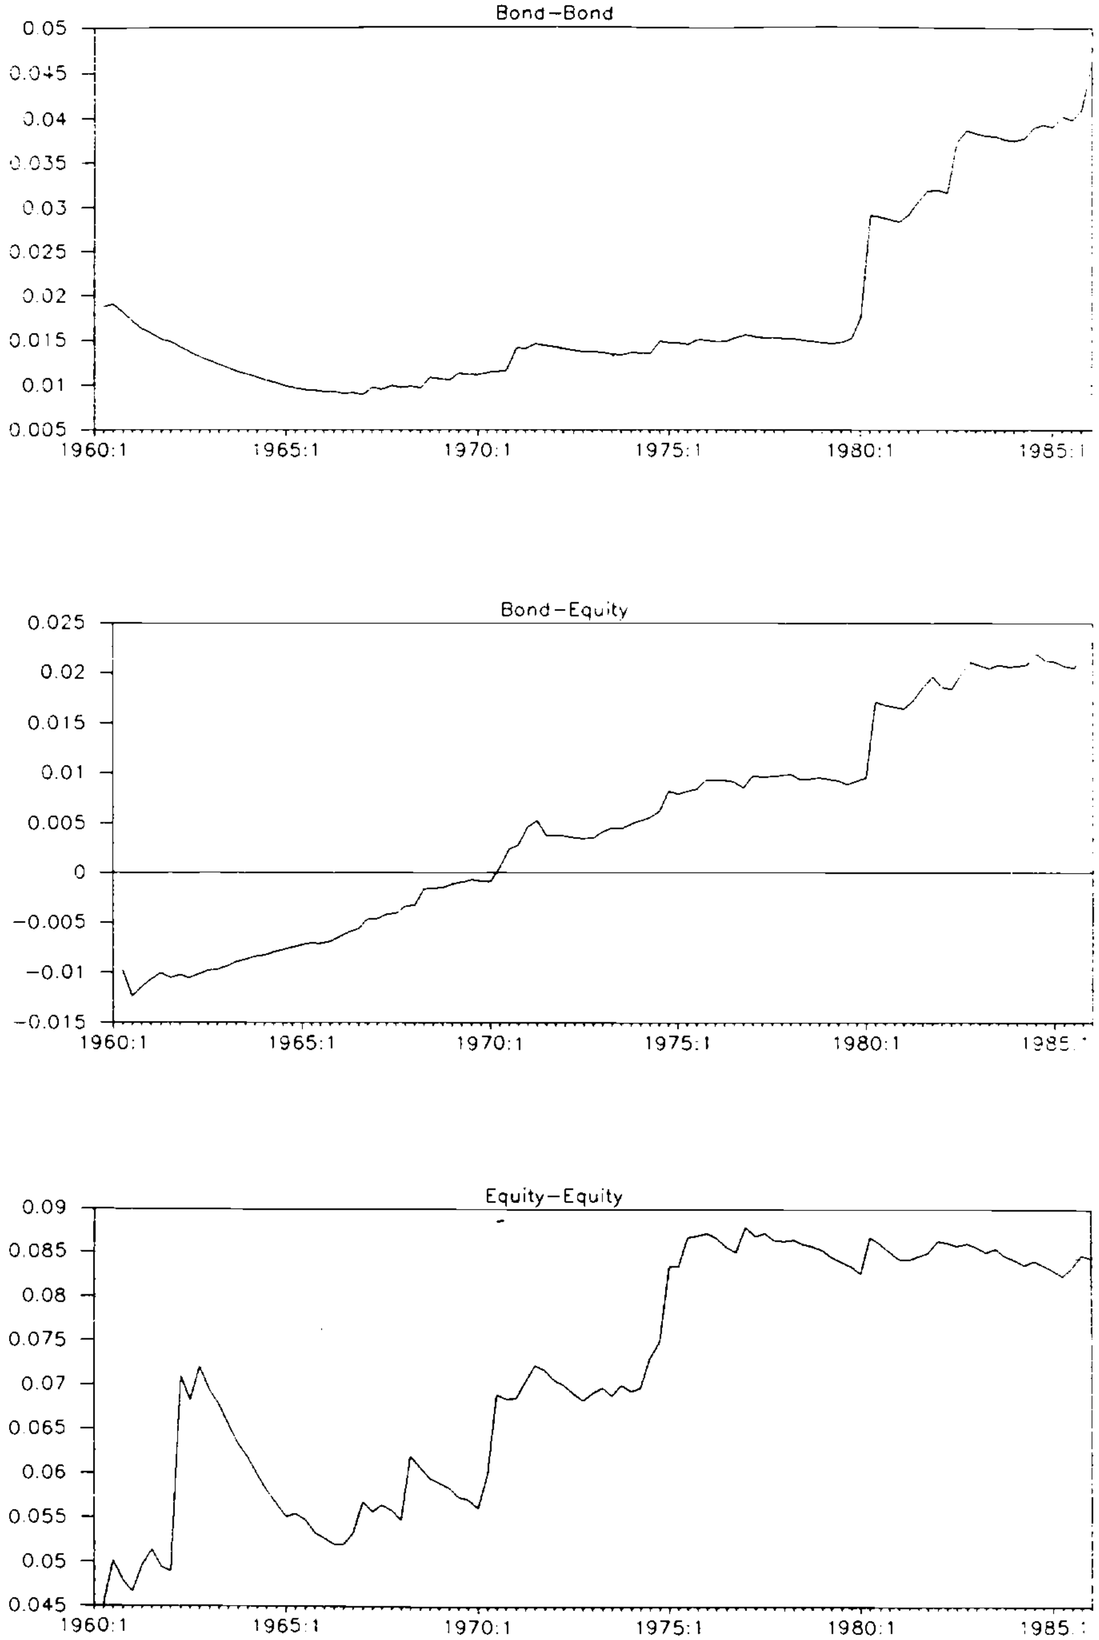

Figure 2: Conditional Covariance Matrix Elements

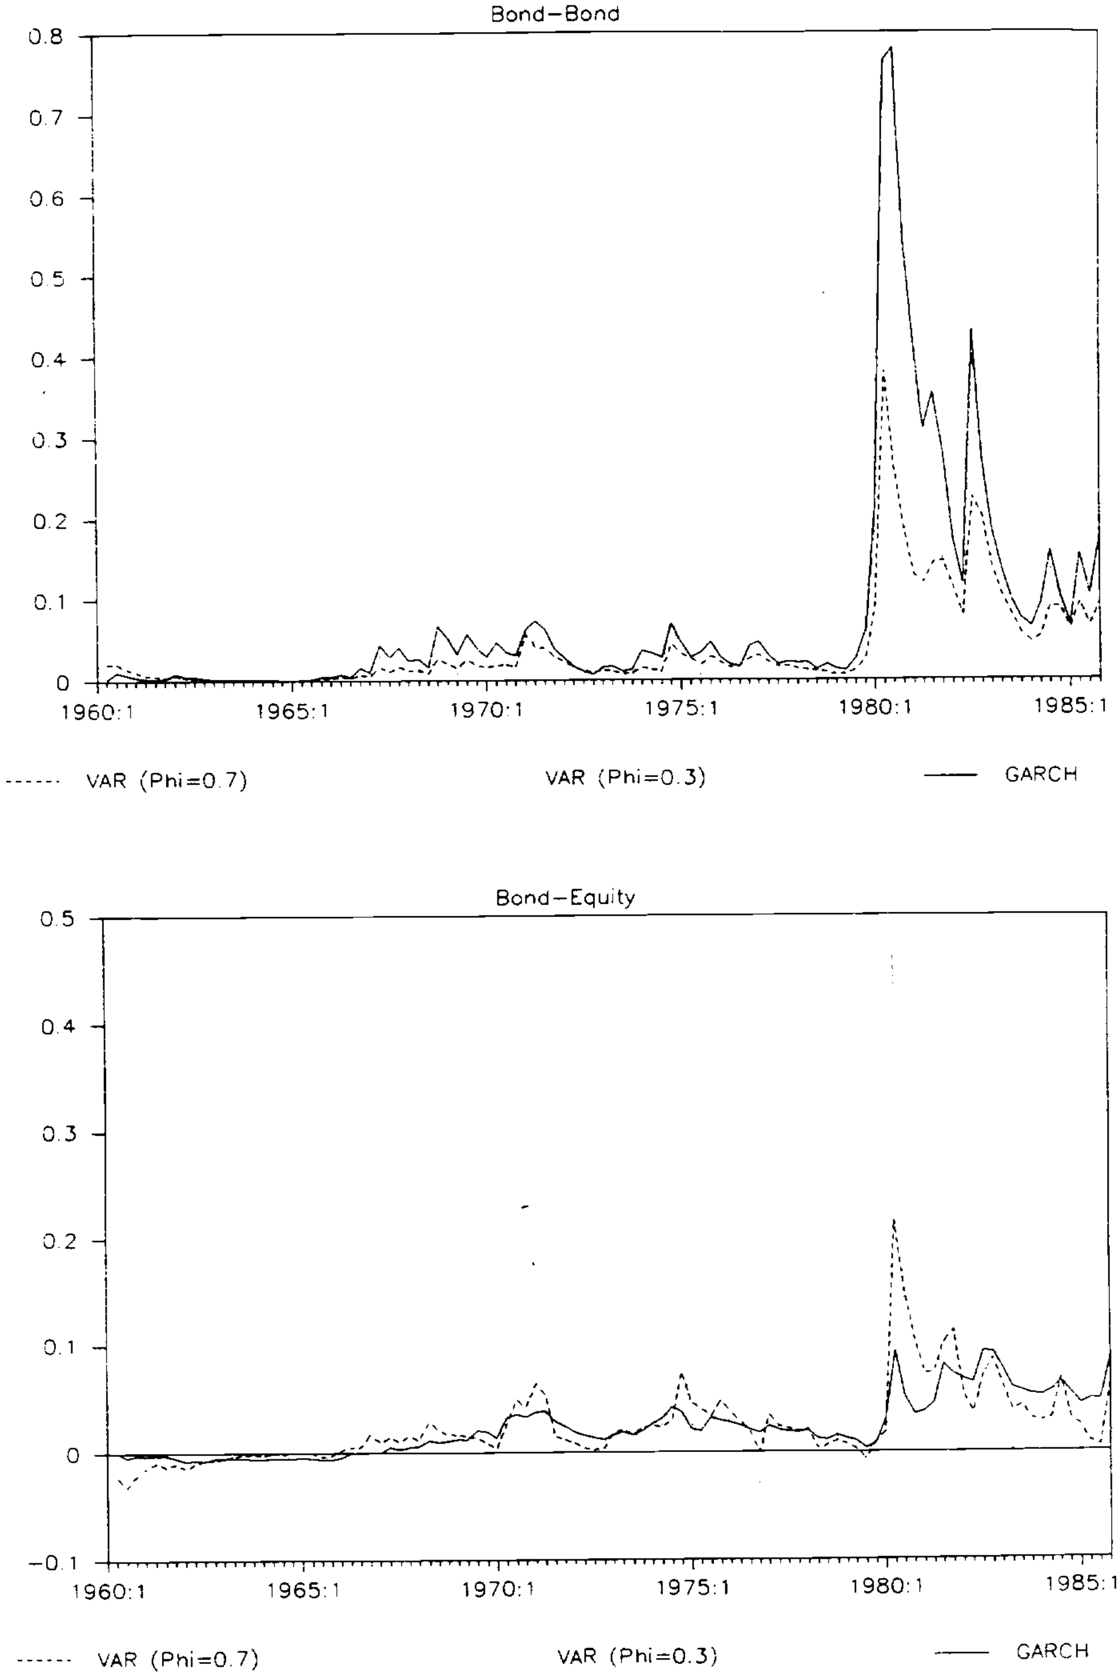

Figure 2 (continued)

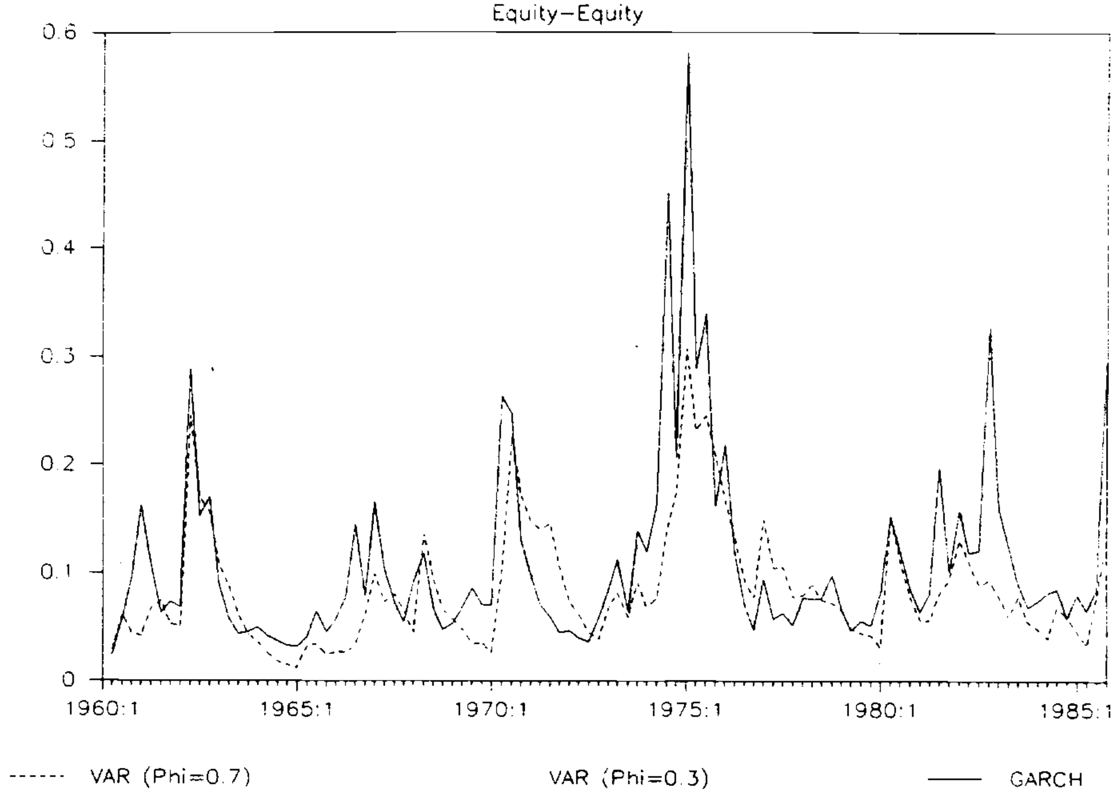

using data spanning 1954:I - 1985:III in the autoregression that generates $\hat{\Omega}$ for 1985:IV. Each autoregression includes L = 1,...4 for both variables, and also includes a constant term.

Following the discussion in Section I, a key issue in the application of any such backward-looking procedure to represent investors' perceptions about the likely distributions of future outcomes, is what part of the available past history people deem relevant and therefore use in forming these perceptions. The results summarized in the top row of Table 2 and in Figure 1 are based on a value of unity for the memory parameter ϕ, so that, as time passes, investors continue to regard as equally relevant all observations during the entire post-Korean War and post-Treasury/Federal Reserve Accord era. The top row of Table 2 summarizes the performance of the resulting bivariate vector autoregression (it makes no sense to show statistics for 104 pairs of regressions individually) by showing the mean of the estimated conditional variance-covariance matrices. Comparison with the sample unconditional variance-covariance matrix shown in Table 1 indicates that, on average, allowing for serial correlation via the VAR method removes more than two-thirds of the variance of the excess return on long-term debt, but less than one-fourth of the variance of the excess return on equity.[15]

The three panels of Figure 1 plot the quarter-by-quarter values of the three elements of the resulting conditional variance-covariance matrix. Given the assumption that the entire past history (since 1954:I) continues to be fully relevant, the variation of these $\hat{\Omega}_{ij}$ elements over time reflects a combination of investors' growing amount of information (that is, growing effective sample size) and their reaction to specific "surprise" episodes involving large ex post expectation errors. More information implies smaller variances and covariance, while the immediate effect of "surprises" is just the opposite. Prominent examples of "surprises" that have readily visible effects on the $\hat{\Omega}_{ij}$ values plotted in Figure 1 include the stock market crash in 1962, the "credit crunch" in 1966, the Penn Central default in 1970, the combination of OPEC and tight monetary policy in 1974, and the introduction of new monetary policy procedures in 1979. On balance, the effect of these and other "surprises" more than outweighs the effect due to the accumulation of additional observations, so that over time both conditional variances tend to increase (as does the conditional covariance, which also changes sign).

Alternatively, if investors do not continue to regard the entire post-1953 experience as fully relevant to future outcomes as time passes, the results shown in Figure 1 and summarized in the top row of Table 2 are based on an (ever increasing) overstatement of the information base behind the formation of investors' expectations. The remaining rows of Table 2 show analogous mean estimated conditional variance-covariance matrices based on a range of less-than-unit values for ϕ. As is to be expected, in light of the sharp increase in the volatility of long-term interest rates that has occurred during the latter part of the 1960-85 period, limiting the estimation procedure's "average memory" in this way sharply increases the mean estimated conditional variance associated with the ex ante excess return on long-term debt. At $\phi = 0.9$, for example, the mean $\hat{\Omega}_{LL}$ is already more than twice the value reported in the first row for ϕ=1. The effect on the mean $\hat{\Omega}_{EE}$ is analogous, though less pronounced. Even for ϕ values of 0.5 and smaller, however -- that is, even for memory that decays so rapidly as to place practically negligible reliance on observations more than one year in the past -- all three elements of the mean estimated conditional variance-covariance matrix are still distinctly smaller than the corresponding unconditional variances and covariance shown in Table 1.

Figure 2 -- in which the vertical scale of each panel is far larger than in Figure 1 -- further indicates the effects of discounting past observations in the VAR procedure by plotting the quarter-by-quarter values of the three elements of the conditional variance-covariance matrix estimated for $\phi = 0.7$ and $\phi = 0.3$ as dashed and dotted lines, respectively. Comparing these two sets of results to one another, and both to the results for $\phi = 1$ shown in Figure 1, indicates that discounting past observations in this manner greatly enhances the inferred impact of the major "surprise" episodes in increasing the estimated conditional variances (and covariance). As is to be expected, discounting past observations also results in a rapid decline of these estimated conditional $\hat{\Omega}_{ij}$ elements after the "surprise" has occurred.

GARCH Results. Tables 3 and 4, together with the solid lines in Figure 2, show the results of applying the GARCH (actually, "GARCH-M" for a GARCH model in which the conditional mean vector depends on the conditional variance-covariance matrix) procedure to the same problem of estimating the quarter-to-quarter variation in the ex ante risk structure of the two excess returns. This procedure as applied here involves the joint estimation of three equations of form (7) for the elements of $\hat{\Omega}$ and two equations of form (10) for the two excess returns, using data spanning 1960:I - 1985:IV.[16] Following Bollerslev et al., the specific form of (7) modeled is a GARCH(1,1) process. Also as in Bollerslev et al., each

Table 3

Estimated Values and T-Statistics for GARCH Model

Coefficient	Estimated Value	T-Statistic
θ_{LL}	.0002	1.4
θ_{LE}	-.0001	-2.3
θ_{EE}	.0018	2.1
ψ_{LL}	.600	3.1
ψ_{LE}	.152	2.4
ψ_{EE}	.540	2.3
Δ_{LL}	.600	8.1
Δ_{LE}	.877	18.3
Δ_{EE}	.417	3.9
k_L	-.0003	-0.1
k_E	.0049	0.1
ρ	2.038	1.6

Table 4

Average Conditional Variance-Covariance Matrix Estimated by GARCH Method

	r_L	r_E
r_L	.076	
r_E	.023	.109

equation of form (10) includes a constant term (although the relevant theory does not indicate any role for an ex ante return differential not related to risk).

Table 3 gives the estimated value and corresponding t-statistic for each of the model's twelve coefficients. Within the GARCH process, two of the three (diagonal) elements of the constant matrix θ in (7) are significantly different form zero at the .05 level, but all three $\hat{\theta}_{ij}$ are small in absolute value. All three elements of matrix Ψ (which relates the conditional variance-covariance matrix to the most recent observation of the error matrix) are significant at the .05 level, and those corresponding to the two variances are fairly large. All three elements of matrix Δ (which describes quarter-to-quarter persistence in the conditional variance-covariance structure and therefore corresponds to the memory parameter ϕ in the VAR method) are highly significant, again with fairly large estimated values. As the relevant theory predicts, neither of the constants arbitrarily inserted in (10), denoted by k in the table, differs from zero at any plausible level of significance. Finally, the estimated value of the coefficient of relative risk aversion is 2.0, with t-statistic 1.6.[17]

Table 4 shows the mean of the estimated conditional variance-covariance matrices which these results imply for each of the 104 quarter spanning 1960:I - 1985:IV. Comparison with Table 1 immediately indicates a sharp contrast to the mean $\hat{\Omega}$ given by the VAR procedure under any assumption about the discounting of past observations. Despite the joint estimation of (7) and (10) -- that is, despite the inclusion of the variance terms as additional explanatory variables in the excess return equations -- the GARCH procedure yields conditional variances for both excess returns that, on average, are <u>larger</u> than the corresponding sample <u>un</u>conditional variances. (The mean of the conditional covariance is smaller.) Instead of removing some of the unconditional variance, as does the VAR procedure as it allows for serial correlation, the GARCH procedure adds to it in arriving at its average estimate of the conditional variance-covariance structure.

One part of the likely explanation for this result is that the GARCH procedure, unlike the VAR procedure, is constrained to exploit the information contained in past realizations of excess returns only in so far as they affect the variance-covariance structure. Serial correlation that does not reflect changing variance-covariance structures, or even such basic information as whether a "surprise" is positive or negative, does not enter the GARCH information set in any direct way. Another part of the explanation is that both the sums $(\hat{\psi}_{LL} + \hat{\Delta}_{LL})$ and $(\hat{\psi}_{LE} + \hat{\Delta}_{LE})$ exceed unity, so that the processes estimated by GARCH for $\hat{\Omega}_{LL}$ and $\hat{\Omega}_{LE}$ are both nonstationary, while the sum $(\hat{\psi}_{EE} + \hat{\Delta}_{EE})$ is less than unity but close nonetheless.[18]

The solid lines in Figure 2 plot the quarter-by-quarter variation of the three elements of the GARCH-estimated conditional variance-covariance matrix. Once again, the effect of familiar major "surprises" -- for example, the 1962 and 1974 episodes for the equity market, and the 1979 episode for the bond market -- is readily apparent. As is to be expected from the fact that the estimated Δ_{ij} values (the persistence coefficients) are far from unity, the conditional variance-covariance elements plotted in Figure 2 display quarter-to-quarter smoothness that is more like the corresponding VAR-estimated elements shown there (with discounting of past observations) than those shown in Figure 1. The increase in conditional variance due to each major "surprise" disappears fairly rapidly. Second, again in contrast to the VAR-estimated results shown in Figure 1 but like those in Figure 2, none of the three $\hat{\Omega}_{ij}$ elements estimated by the GARCH procedure displays any noticeable tendency to grow larger over time. Finally, as is consistent with the respective means shown in Tables 2 and 4, a comparison of the respective scales used in the corresponding panels of these figures indicates that each GARCH-estimated $\hat{\Omega}_{ij}$ element tends to be larger, on average, than the corresponding VAR-estimated element, regardless of the discounting of past observations.

III. Implications for Investors' Portfolio Behavior

Given the importance of risk in theories of portfolio behavior, estimates of conditional variance-covariance structures as different as those plotted in Figures 1 and 2, and summarized in Tables 2 and 4, presumably imply widely differing asset demands. Table 5 summarizes these differences by showing the optimal asset substitutability matrix B calculated from (3) using in turn each of the average conditional variance-covariance structures estimated by the VAR and GARCH procedures, as shown in Tables 2 and 4, and in each case a value of 2.0 for the coefficient of relative risk aversion.[19]

As is to be expected from the smaller average conditional variances in Table 2, compared to those in Table 4, the VAR procedure implies much greater average sensitivity of the optimal portfolio allocation to expected excess returns than does the GARCH procedure. When the VAR procedure does not discount past observations, this greater sensitivity to expected excess returns takes the form of greater substitutability between each risky asset (individually) and short-term debt, and especially so between the two debt assets. With $\phi=1$, an increase of .01 in the expected excess return on long-term debt on average increases the optimal portfolio share invested in that asset by .28, while reducing the optimal shares invested in short-term debt and equity by .26, and .02, respectively. According to the GARCH procedure, the corresponding increase in the long-term debt share is just .07, and the reductions in the other two asset shares are .06 and .01, respectively. Similarly, the results of the VAR procedure with $\phi=1$ imply that an increase of .01 in the expected excess return on equity on average increases the optimal portfolio share invested in equity by .07, while

Table 5

Implied Average Asset Substitutability Matrix

			$(r_L - r_S)$	$(r_E - r_S)$
VAR:	ϕ=1.0	α_S	-25.91	- 5.30
		α_L	27.72	- 1.81
		α_E	- 1.81	7.11
	ϕ=0.9	α_S	-11.50	- 2.89
		α_L	15.57	- 4.07
		α_E	- 4.07	6.96
	ϕ=0.7	α_S	-10.08	- 2.93
		α_L	14.18	- 4.10
		α_E	- 4.10	7.02
	ϕ=0.5	α_S	- 9.75	- 3.01
		α_L	13.75	- 4.00
		α_E	- 4.00	7.01
	ϕ=0.3	α_S	- 9.51	- 3.07
		α_L	13.45	- 3.95
		α_E	- 3.95	7.01
GARCH:		α_S	- 5.56	-3.42
		α_L	7.02	-1.47
		α_E	-1.47	4.89

reducing the optimal shares invested in short-term debt and long-term debt by .05 and .02, respectively. According to the GARCH procedure, the corresponding increase in the equity share is .05, and the reductions in the other two asset shares are .03 and .01, respectively.

In comparison with these values, the results of the VAR procedure applied with discounting of past observations consistently indicate a greater degree of substitutability between the two risky assets. At the same time, they consistently show substitutability between long- and short-term debt that is smaller than in the VAR results based on ϕ=1, but still greater than in the GARCH results. The substitutability between equity and short-term debt that they show is consistently about the same as in the GARCH results. For values of ϕ varying from 0.9 to 0.3 (and below that, too), an increase of .01 in the expected excess return on long-term debt on average increases the optimal portfolio share invested in long-term debt by about .14, and reduces the optimal shares invested in short-term debt and equity by .10 and .04, respectively. Similarly, for values of ϕ throughout this range an increase of .01 in the expected excess return on equity increases the optimal equity share by .07, while reducing the optimal short- and long-term debt shares by .03 and .04, respectively.

The results shown in Table 5 are based on the average conditional variance-covariance structure that each procedure estimates for the entire 1960-1985 period, and therefore they describe optimal asset demands on average over this entire period. Given the quarter-to-quarter variation in the estimated conditional variance-covariances structures plotted in Figures 1 and 2, risk-based models of portfolio behavior imply that the corresponding optimal asset demands also varied substantially during this period.

Figures 3 and 4 plot the quarter-to-quarter variation in each of the six unique elements of the asset substitutability matrix corresponding to several of the estimated conditional variance-covariance structures described above, again using 2.0 as the coefficient of relative risk aversion in each case.[20] The asset demand elements plotted in Figure 3 correspond to the conditional variance-covariance structure plotted in Figure 1, estimated by the VAR method with ϕ=1. Figure 4 plots analogous asset demand elements corresponding to the VAR method with ϕ=0.5 (dashed lines) and the GARCH method (solid lines).

As is to be expected, the results shown in Figures 3 and 4 closely mirror the patterns shown in Figures 1 and 2. The asset demand parameters implied by the VAR procedure with ϕ=1 exhibit substantial smoothness, with large discontinuities in the wake of major "surprises," and in most cases a tendency to shrink (in absolute value) over time. By contrast, the asset demand parameters implied by the VAR procedure with discounting of past observations, or by the GARCH procedure, exhibit neither much smoothness nor any time trend.

The results differ in other respects as well. The VAR procedure with ϕ=1 implies that long-term debt and equity were complements from the beginning of the sample until 1970, and then substitutes thereafter. The two debt instruments were substitutes throughout the sample period, as were short-term debt and equity. With ϕ=0.5, long-term debt and equity were complements from the beginning of the sample until 1965, and then substitutes throughout the remainder of the sample except for a few

Figure 3: Asset Substitution Matrix Elements, VAR (Phi=1)

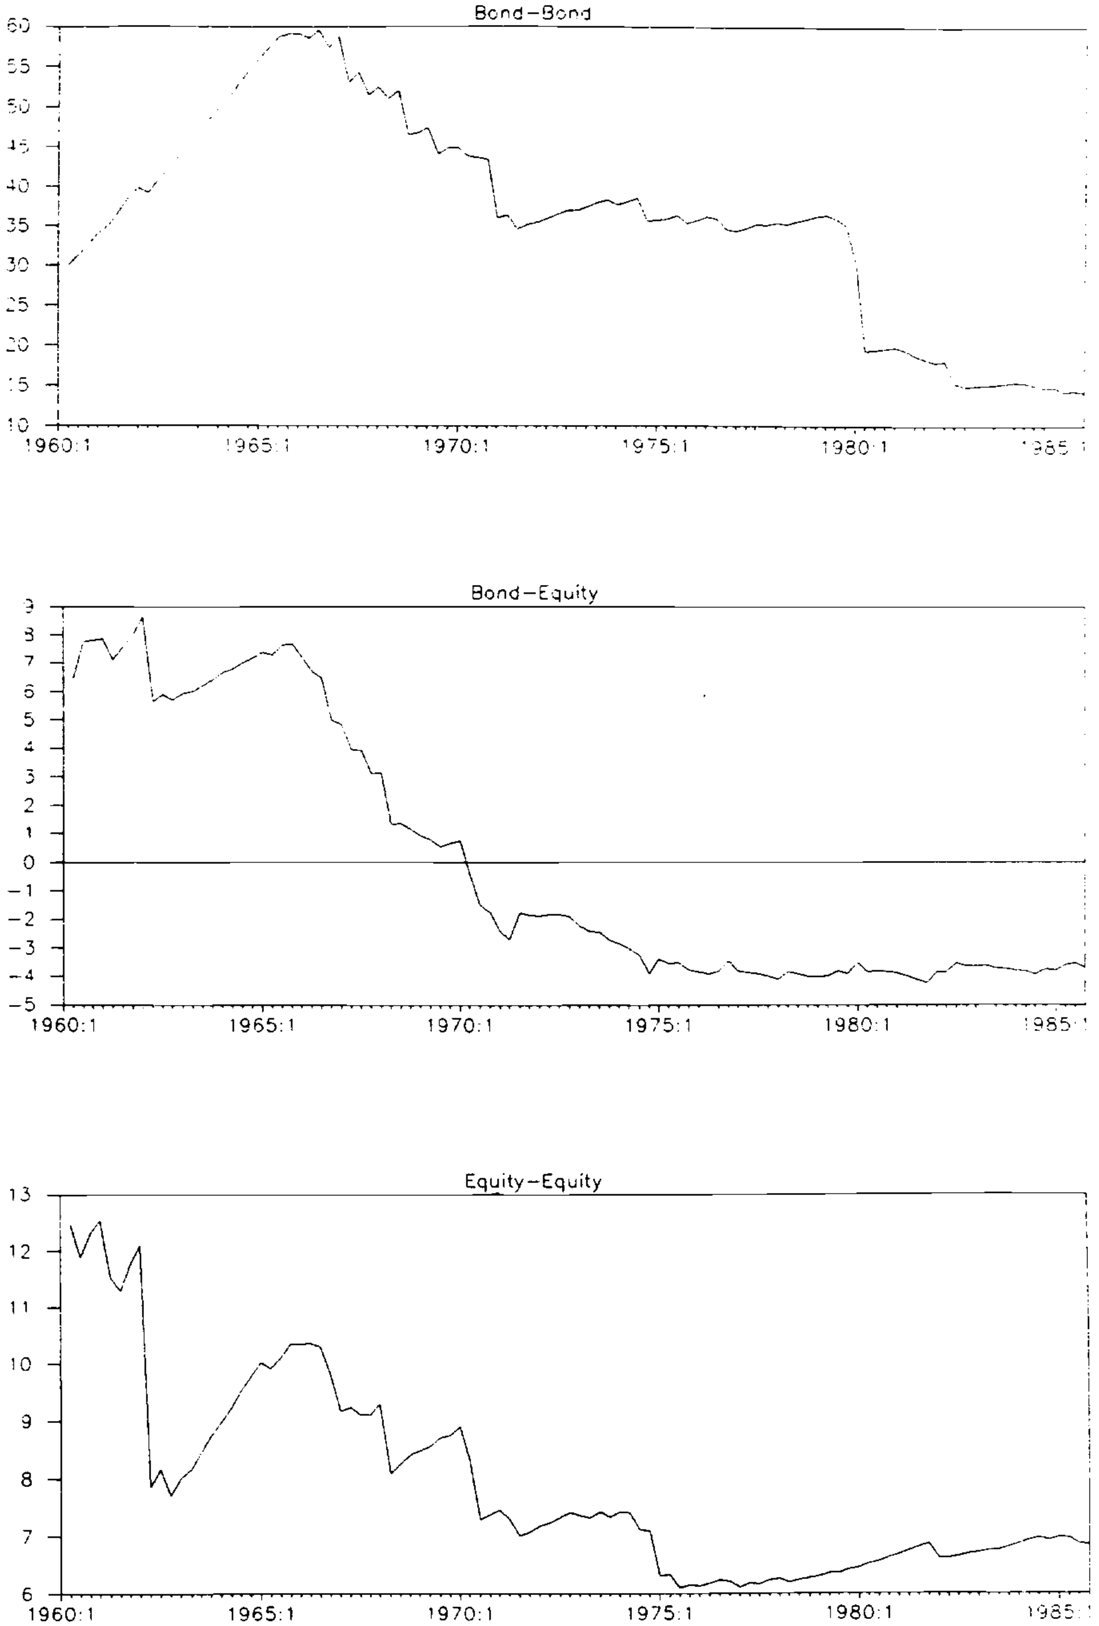

Figure 4: Asset Substitution Matrix Elements

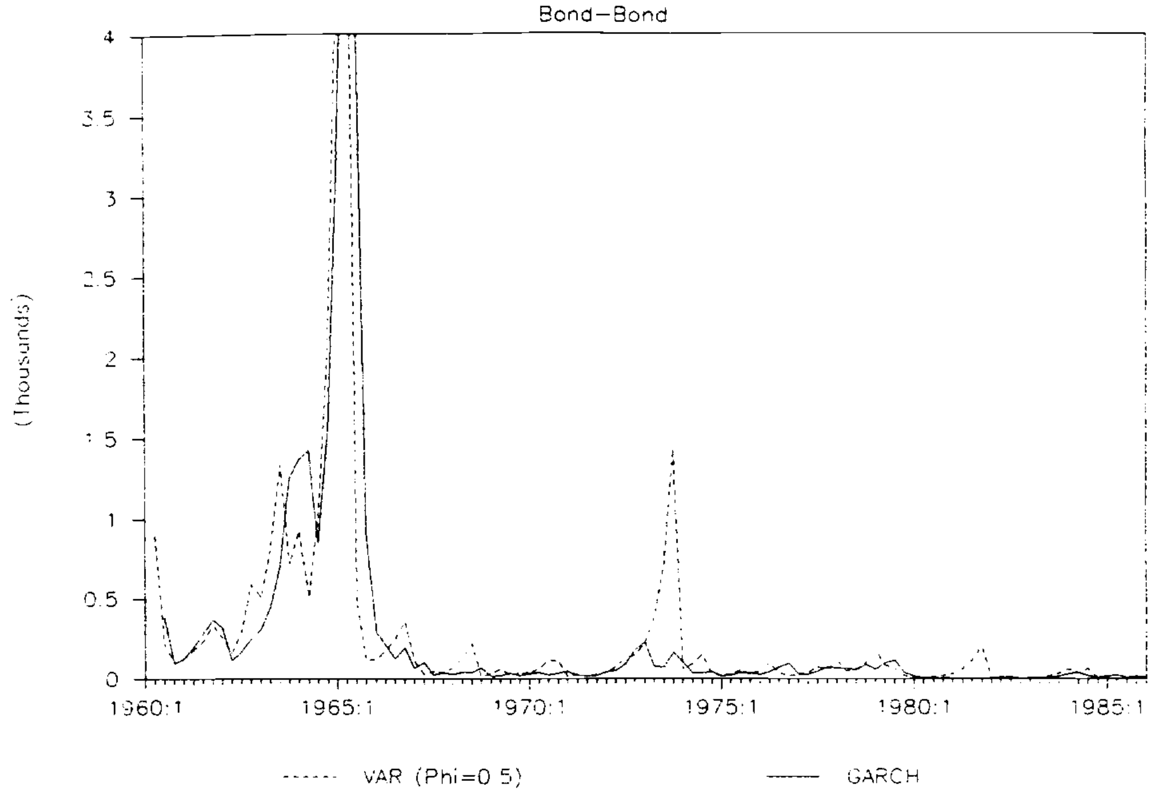

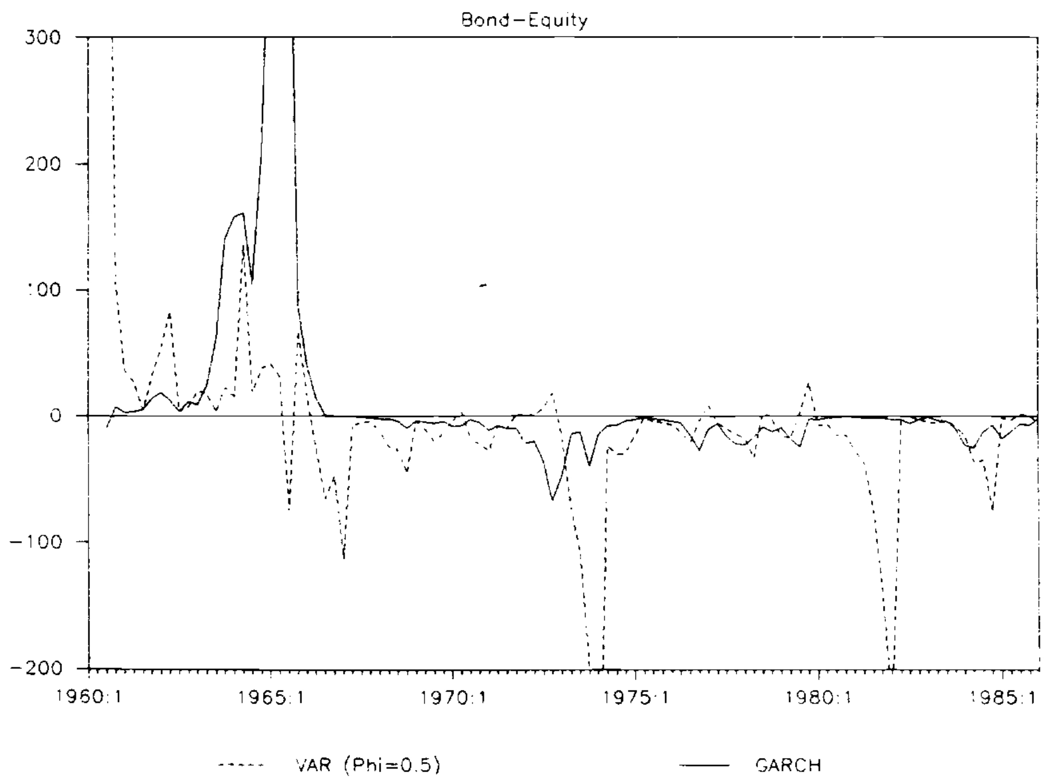

Figure 4 (continued)

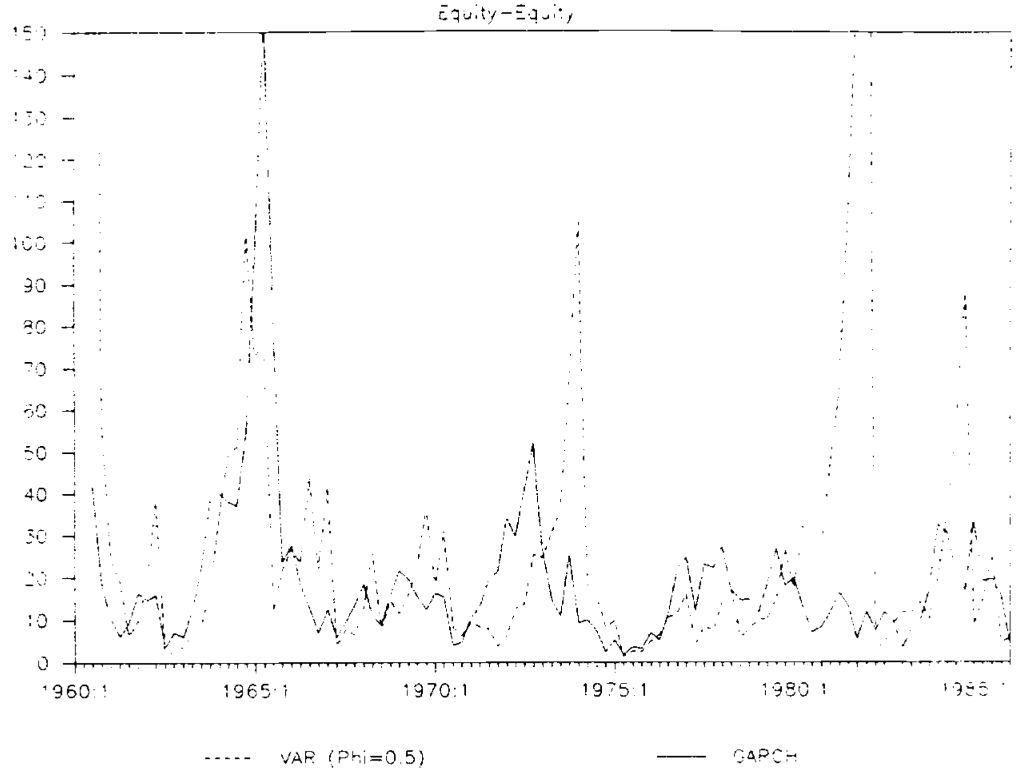

scattered quarters. Again, the other two asset pairs were substitutes throughout. The GARCH estimates indicate that long-term debt and equity were complements from the beginning of the sample until 1966, and then consistently substitutes. Here too the other two asset pairs are consistently substitutes.

Finally, as the average results shown in Table 5 suggest, the asset demand parameters implied by the VAR procedure with $\phi=1$ are each substantially smaller on average (in absolute value) than the corresponding parameters implied by the VAR procedure with $\phi=0.5$, and these in turn are mostly smaller on average than those implied by the GARCH procedure.

IV. Implications for Equilibrium Returns

Just as different variance-covariance structures imply differences in asset demand behavior, differences in asset demand behavior in turn imply different structures of market-clearing ex ante returns for any given composition of the "market portfolio" of assets to be held. Figures 5 and 6 summarize these differences by plotting the time paths of the respective sets of ex ante returns implied by the asset demand systems corresponding to the variance-covariance structures estimated by the VAR method (with different ϕ values) and the GARCH method. Following (10), the pair of ex ante excess returns (or risk premia) for each period are calculated as $\rho\hat{\Omega}_t\underline{\alpha}_t^s$ where $\hat{\Omega}_t$ is again the conditional variance-covariance matrix as estimated for that quarter by either the VAR and the GARCH procedure, $\underline{\alpha}_t^s$ is the historical value of the asset supply vector for that quarter, and ρ is again in each case the GARCH-estimated value 2.0.[21]

The upper panel of Figure 5 plots the risk premium on long-term debt, as implied by the asset demand system estimated by the VAR procedure with ϕ=1. The lower panel of the same figure -- in which the vertical scale is roughly an order of magnitude larger than in the upper panel -- plots corresponding series, again for the risk premium on long-term debt, implied by the asset demand systems estimated by the VAR procedure with ϕ=0.7 and ϕ=0.3, and by the GARCH procedure. Figure 6 plots analogous sets of results for the equity risk premium.

For long-term debt, both the GARCH procedure and the VAR procedure (for all three ϕ values) imply a systematically greater risk premium during the 1980-85 period than earlier on. This finding is a natural consequence of the larger estimated conditional variance of the excess return on long-

Figure 5: Estimated Risk Premium on Bonds

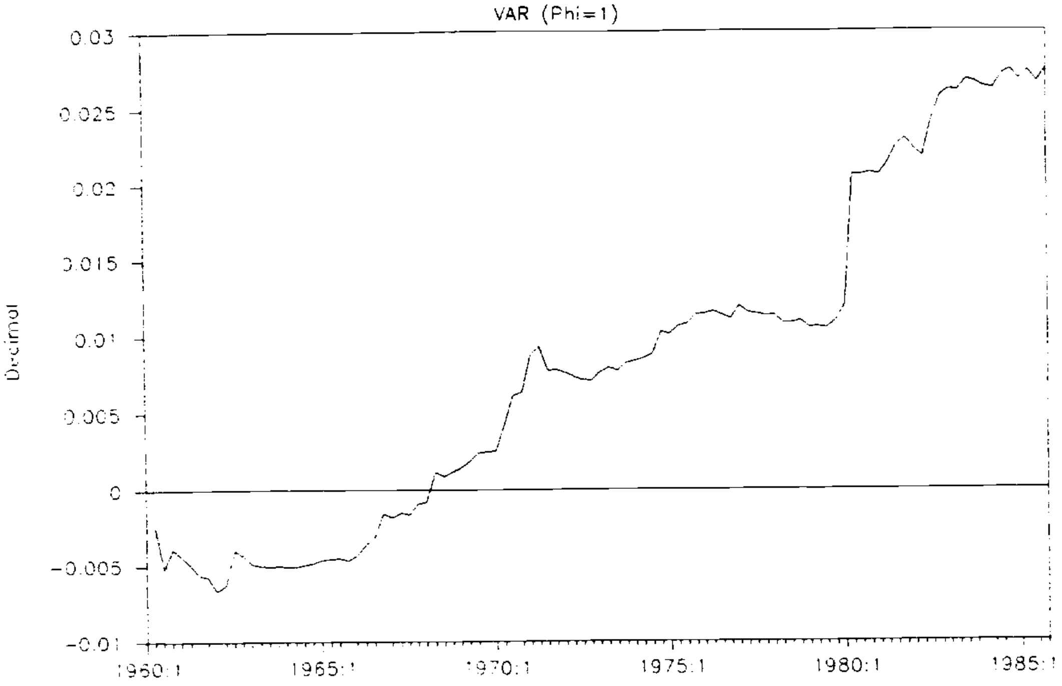

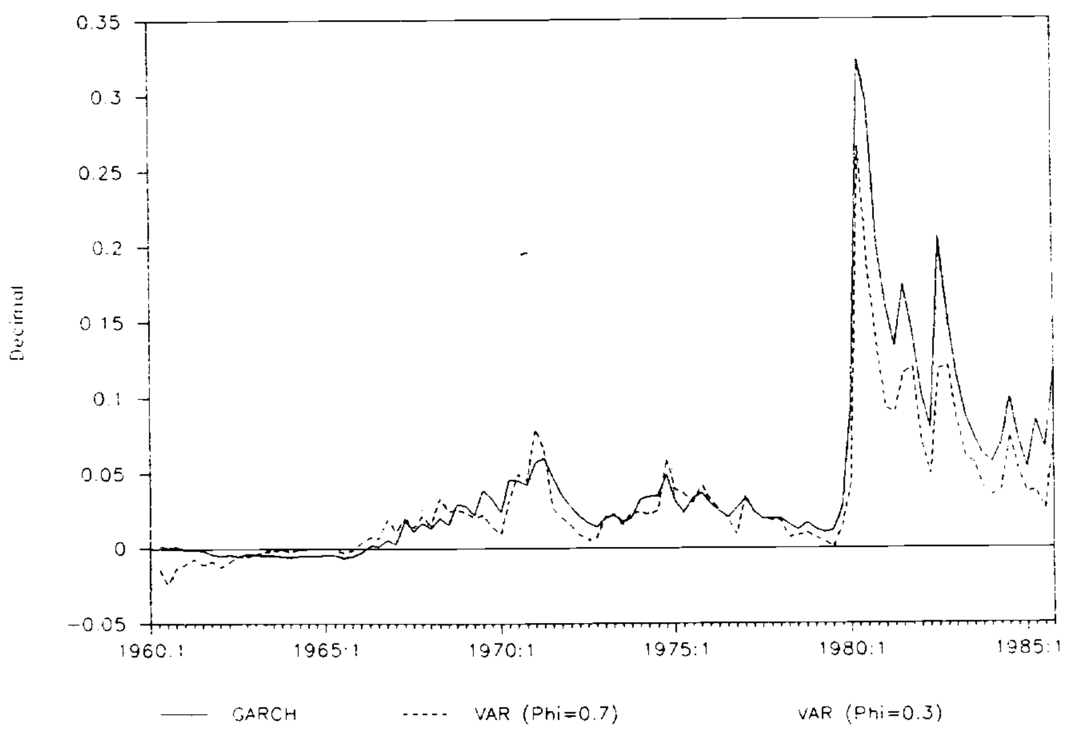

Figure 6: Estimated Risk Premium on Equity

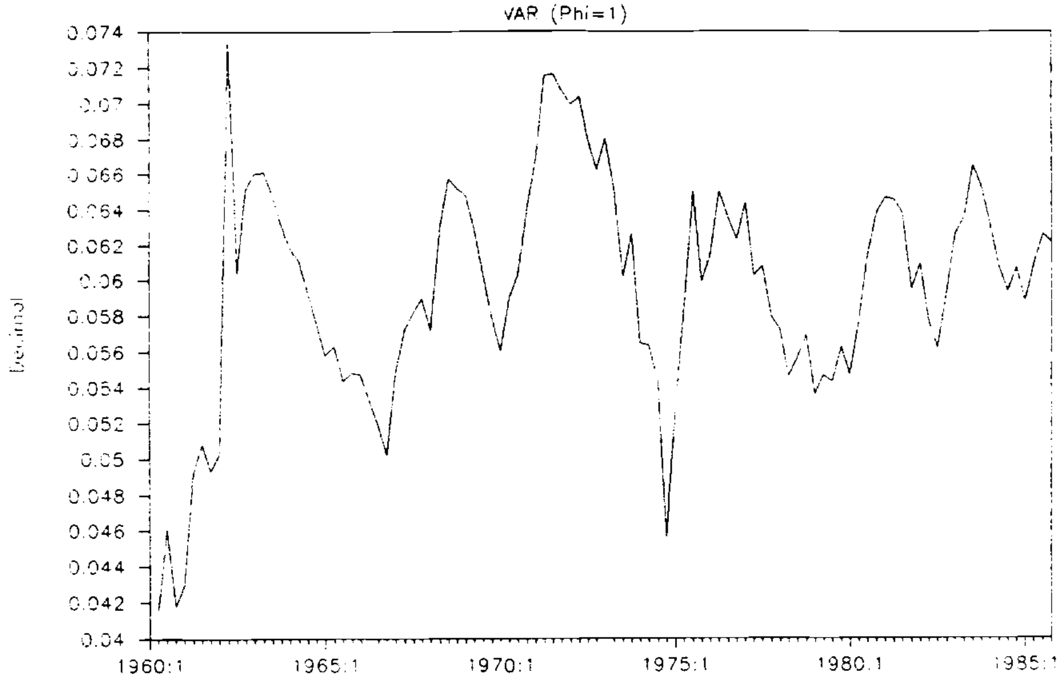

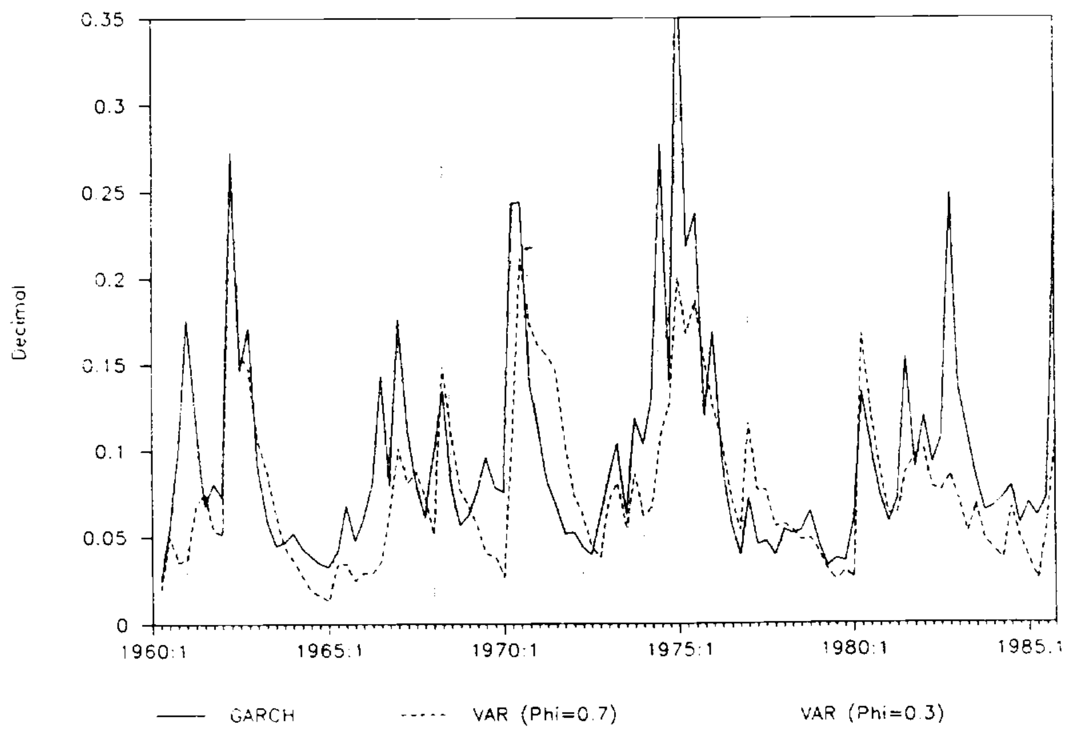

term debt, shown in Figures 1 and 2. For the early years of the sample (1960-66 under the GARCH procedure, 1960-63 under the VAR), when the estimated $\hat{\Omega}_{LL}$ was small and the estimated $\hat{\Omega}_{LE}$ was negative, the implied risk premium on long-term debt is negative. The VAR procedure with $\phi=1$ implies a risk premium that rises almost monotonically over time, reaching about .04 by the end of the sample. The other estimates exhibit substantially more volatility, in some quarters reaching what appear to be absurdly large values. The general pattern shown by the GARCH results is roughly similar to that of either of the VAR results with discounting of past observations, with less volatility than in the $\phi=0.7$ case but more than in the $\phi=0.3$ case. The mean implied risk premium on long-term debt is .0416 for the GARCH procedure, .0081 for the VAR procedure with $\phi=1$, and .0303 and .0301 for the VAR procedure with $\phi=0.7$ and $\phi=0.3$, respectively.

In contrast, none of these procedures implies any overall trend in the risk premium for equities. For both the GARCH procedure and the VAR procedure with discounting of past observations, there is substantial volatility, with the implied equity risk premium rising sharply (albeit only temporarily) in the wake of each of the major "surprises" familiar from the figures shown earlier on. The GARCH procedure again implies a degree of volatility that is intermediate between what the VAR procedure implies with $\phi=0.7$ and $\phi=0.3$. Once again, also, all three of these sets of risk premia reach what appear to be absurdly high levels in some quarters. There are some notable differences, however. For example, the largest equity risk premium implied by GARCH occurs after the 1974 market crash; three other "surprises" loom larger in the VAR results. Finally, the equity risk premium implied by the VAR procedure with $\phi=1$ is again quite smooth in comparison to the others. The mean implied risk premium on equities is .0964 for GARCH, .0596 for VAR with $\phi=1$, and .0773 and .0770 for VAR with $\phi=0.7$ and $\phi=0.3$, respectively.

How closely does any of these implied ex ante excess return series compare with the observed ex post excess returns? Table 6 shows the simple correlations between the respective ex ante risk premium series implied by the VAR and GARCH procedures and the corresponding ex post excess returns, as well as the root-mean errors of the corresponding differences, over the 104 quarters of the 1960-85 sample. For the risk premium on long-term debt, the VAR procedure with $\phi=1$ provides the best fit. The risk premium implied by the VAR procedure with any significant degree of discounting of past observations ($\phi=0.9$ or below) shows a much weaker correlation, though only marginally greater root-mean-square error. The risk premium implied by the GARCH procedure shows a poorer fit than that for the VAR procedure with little or no discounting ($\phi=0.9$ or above), but a closer fit than that for the VAR procedure with more substantial discounting. By contrast, for equities the VAR procedure with substantial discounting ($\phi=0.5$ or below) generates risk premia that are more closely related with the ex post excess returns -- for $\phi=1$ or 0.99 the correlation is negative -- and the GARCH procedure implies a risk premium with an even closer fit (albeit only marginally so).

Although the value of 2.0 for ρ estimated by the GARCH procedure is the natural one to use in calculating the asset demand behavior and consequent ex ante risk premia implied by the GARCH-estimated variance-covariance structure, its analogous use in the context of the VAR procedure is arbitrary. Simply rescaling the implied risks premia by changing the

Table 6

Measures of Closeness Between Ex Post and Implied Ex Ante Excess Returns

	Correlation Coefficients		Root-Mean-Square-Errors	
Method	r_L	r_E	r_L	r_E
VAR (ϕ=1.0)	.203	-.124	.237	.316
VAR (ϕ=0.99)	.201	-.065	.236	.315
VAR (ϕ=0.9)	.126	.131	.237	.312
VAR (ϕ=0.7)	.026	.294	.242	.304
VAR (ϕ=0.5)	.020	.336	.244	.298
VAR (ϕ=0.4)	.026	.337	.245	.297
VAR (ϕ=0.3)	.034	.331	.246	.297
VAR (ϕ=0.2)	.047	.318	.247	.298
GARCH	.103	.345	.241	.298

risk aversion parameter will not affect the correlations shown in Table 6, but in general doing so will affect the root-mean-square errors. The simplest way to explore this possibility is to project the observed ex post excess returns on the corresponding implied ex ante excess returns in an ordinary-least-squares regression of the form.

$$(12) \qquad (r_{it} - r_{St}) = \gamma + \delta\, (r_{it} - r_{St})^{*} + \xi_t$$

where $(r_i - r_S)^*$ is the ante excess return implied by either the VAR or the GARCH procedure (with ρ=2.0), ξ is a disturbance term, and γ and δ are coefficients to be estimated. A value of δ different from unity would correspond to rescaling the risk premium by substituting a value of the risk aversion parameter different from 2.0.[22]

Table 7 shows the estimated $\hat{\delta}$ values, the t-statistics testing the null hypothesis δ=1, and the regression standard error for each of the risk premium series considered in Table 6. Despite the wide range of differen· $\hat{\delta}$ estimates, including several that are significantly different from unity, the regression standard errors are not much smaller than the corresponding root-mean-square errors reported in Table 6.

Finally, in comparing the respective results given by the VAR and GARCH procedures it is important to recall that because equations (10) estimate excess returns on a weighted basis, where the weights reflect the GARCH-estimated variance-covariance structure, use of the <u>un</u>weighted correlations and root-mean-square errors shown in Table 6 potentially biases the comparison. Table 8 therefore presents correlations that are analogous to those shown in Table 6 but computed with the 104 individual quarters weighted as in the GARCH procedure.[23]

Table 7

Results of Rescaling Risk Premia by Ordinary Least Squares

	r_L		r_E	
Method	δ	SE	δ	SE
VAR (ϕ=1.0)	4.51 (1.6)	.235	-6.16 (-1.5)	.313
VAR (ϕ=.99)	3.43 (1.5)	.235	-2.71 (-0.9)	.315
VAR (ϕ=.9)	.99 (0.0)	.238	1.81 (0.6)	.313
VAR (ϕ=.7)	.15 (-1.5)	.240	1.92 (1.5)	.302
VAR (ϕ=.5)	.09 (-2.1)	.240	1.56 (1.3)	.297
VAR (ϕ=.4)	.10 (-2.3)	.240	1.38 (1.0)	.297
VAR (ϕ=.3)	.12 (-2.4)	.240	1.20 (0.6)	.298
VAR (ϕ=.2)	.13 (-2.6)	.240	1.04 (0.1)	.299
GARCH	.41 (-1.5)	.239	1.68 (1.5)	.296

Note: T-statistics are for the null hypothesis $\delta = 1$.

Table 8

Weighted Correlations Between Ex Post and Implied Ex Ante Excess Returns

Method	r_L	r_E
VAR (ϕ=1)	0.041	-0.004
VAR (ϕ=.99)	0.041	0.022
VAR (ϕ=.9)	0.025	0.113
VAR (ϕ=.7)	0.017	0.231
VAR (ϕ=.5)	0.026	0.276
VAR (ϕ=.4)	0.036	0.283
VAR (ϕ=.3)	0.048	0.139
VAR (ϕ=.2)	0.062	0.137
GARCH	0.145	0.246

These weighted correlations tell a quite different story for long-term debt, though not for equity. On this criterion, either heavy discounting of past observations or else none at all produces the strongest correlations from among the long-term debt risk premium series implied by the VAR procedure, but in any case the weighted correlations are all far smaller than their unweighted counterparts with ϕ=1 or .99. As is to be expected, the weighted correlation for the long-term debt risk premium implied by the GARCH procedure is stronger than the corresponding unweighted correlation. Even so, it remains weaker than the weighted correlations for the VAR procedure except for cases with little or no discounting.

For equities, the weighted correlations are not all that different. The weighted correlations are uniformly weaker than their unweighted counterparts, even for the GARCH model (although the pronounced negative correlations for the VAR procedure with ϕ=1 or .99 no longer arise). As with the unweighted correlations, discounting of past observations within the VAR procedure greatly strengthens the weighted correlations, at least up to a point. The maximum again occurs with ϕ=.4. Once again, the weighted GARCH correlation falls just short of this maximum.

V. Concluding Remarks About the Results

The empirical results presented in this paper lead to several conclusions about the role of time-varying asset risk assessments in accounting for what, on the basis of many earlier studies, at least appear to be time-varying differentials in ex ante asset returns. First, both the VAR method (with or without discounting of past observations) and the GARCH method indicate sizeable changes over time in variance-covariance structures conditional on past information. These changing conditional variance-covariance structures in turn imply sizeable changes over time in asset demand behavior, and hence in the market-clearing equilibrium structure of ex ante asset returns.

Second, at least for some values of the memory parameter describing the discounting of past observations, the ex ante excess returns implied by the VAR method bear non-negligible correlation to the excess returns observed ex post on either long-term debt or equity. For the GARCH model, the same is true for equities though not for long-term debt. The percentage of the variation of ex post excess returns explained by the implied time-varying ex ante excess returns is as great as .04 for long-term debt, and .11 for equity -- comparable to values to which previous researchers have interpreted as consistent with rejection of the null hypothesis of risk premia that are constant over time.[24]

Third, for long-term debt the two methods give sharply different answers to the questions of how much relevance market participants associate with more distant past observations in assessing future risks. Within the VAR method, the greatest correlation between implied ex ante excess returns and observed ex post excess returns follows from values of the memory parameter indicating little or no discounting of past observations, so that more distant outcomes are as relevant as more recent ones (or nearly so). The corresponding parameter estimated directly within the GARCH method indicates far greater discounting of past observations. The correlation between implied ex ante excess returns and observed ex post excess returns is about twice as great under the VAR method with little or no discounting as under the GARCH method, however.

Fourth, for equities both methods agree in indicating extremely rapid discounting of more distant observations -- so much so that in neither case do outcomes more than a year in the past matter much at all. Under the VAR method, the greatest correlation between implied ex ante excess returns and observed ex post excess returns follows from a memory parameter value of .4 (per quarter). The corresponding value directly estimated by the GARCH model is substantially identical, at .417.

While the first three conclusions reported here are plausible enough, the finding of such an extremely short "memory" on the part of market investors -- indeed, the almost precise agreement of both empirical methods applied here in reaching this finding -- is sufficiently startling to raise serious questions about the overall approach to risk assessment underlying the analysis both here and in the related literature. In particular, the representation of equity risk by a single normally distributed stochastic disturbance may be too restrictive.

A more general alternative, with which the authors are experimenting in further research along these lines, is to specify equity returns to include two distinct stochastic components: one that may be well described by the usual assumptions of normality and (modest) serial correlation, and the other characterized by larger, albeit only occasional nonzero values, and no serial correlation -- for example, a Poisson process.[25] An investor who correctly understood the nature of the two processes would take account of the serial correlation of the first in the manner modeled here, but would assume that a nonzero realization of the second bore no implication for the _immediate_ future. By contrast, the econometrician who failed to distinguish the two processes, and proceed in the fashion of this paper (and substantially all of the related literature), would derive a downward biased estimate of the memory parameter relevant to the first process. If the variance of the second process were sufficiently large relative to that of the first, the bias could be great enough to deliver an implausibly short "memory" like that found here for equities.

Footnotes

* The authors are grateful to Rob Engle and Tim Bollerslev for making their computer program available and for numerous helpful discussions; to John Campbell, Gary Chamberlain, Robert Engle, Robert Merton, Lawrence Summers and James Stock for helpful comments on two earlier drafts; and to the National Science Foundation, the Alfred P. Sloan Foundation and the Harvard Program for Financial Research for research support.

1. The recent surveys by Melino (1986), Shiller (forthcoming) and Singleton (forthcoming) provide numerous references.

2. This assumption was made explicit by Meiselman (1963), for example.

3. Prominent examples include the capital asset pricing model, in the original single-period form due to Sharpe (1964) and Lintner (1965), or the intertemporal form first developed by Merton (1973), or the intertemporal consumption-based form due to Breeden (1978); Stiglitz's (1970) consumption-based model of the term structure of interest rates; Ross's (1976) arbitrage pricing model; and successor models, like that laid out in Cox et al. (1985), which make the factor approach more explicit. See Merton (forthcoming) for a survey delineating the common features in all of these models.

4. Modigliani and Shiller's dependent variable was actually the long-term interest rate, but their equation included the short-term rate as a right-hand-side variable.

5. As in Friedman (1985), but unlike in Bollerslev et al. (1988), the returns here are stated in after-tax terms, using the average effective marginal tax rates applicable to individual U.S. investors in each year

during the sample.

6. Additional elements are necessary if the theory is also to encompass nonfinancial assets which bear direct service returns, or financial assets which provide transactions services in addition to any pecuniary returns.

7. See Friedman and Roley (1987) for derivations of these expressions, for references, and for qualifications about the assumptions employed.

8. If investors are aware of the possibility of events that have not occurred during the sample, however, the measured sample variance-covariance matrix may <u>under</u>state the relevant uncertainty. See, for example, Ederington's (1986) analysis of this possibility in a context closely related to this paper's focus on serial correlation.

9. The application of the VAR method in Friedman (1985) is more like a factor model of asset risk in that it first decomposes each asset's return into components known with certainty ex ante (for example, the nominal return on a Treasury bill, or the price and coupon rate on a Treasury bond) and stochastic components that the investor must forecast (for example, the price change in bonds or equities), and then specifies the vector autoregression in terms of the stochastic components directly. Our efforts to estimate the GARCH model in a parallel way proved unsuccessful, however, and so for purposes of this paper we applied the VAR method directly to return vector $\underline{r}$ so as to make the two sets of results as comparable as possible.

10. An alternative procedure, which would probably deliver about the same results, would be to estimate (4) using generalized least squares, with a block diagonal weighing matrix made up of successive powers of ϕ.

11. The assumption that the supply vector is exogenous, which is standard in much of the literature, embodies the contradiction of implicitly taking as given the prices of the assets whose expected returns the model is supposed to determine, even though for most assets it is primarily variation in price that delivers variation in expected return.

12. Our original intention was to use real rather than nominal returns and to treat all assets as risky, as in Friedman (1985). Trivariate versions of the GARCH procedure failed, however, as the parameter estimates converged to values corresponding to non-positive-definite $\hat{\Omega}_t$ for some t. Hence we followed Bollerslev et al., as well as much of the finance literature, in using nominal returns (so that short-term debt would be risk-free) and estimating the model in terms of excess returns relative to the short-term rate. (An alternative approach, as in Engle (1987), would have been to impose a factor structure onto the variance-covariance matrix and estimate the factor loadings.) We did follow Friedman (1985) in using after-tax rather than before-tax returns.

13. The marginal tax rates applied to interest and dividends for years 1960-79 are values estimated by Estrella and Fuhrer (1983), on the basis of Internal Revenue Service data, to reflect the marginal tax bracket of the average recipient of these two respective kinds of income in each year; the values applied for years 1980-85 are from an updating of the Estrella-Fuhrer analysis. The marginal tax rate applied to capital gains in each year is an analogous estimate, in principle including allowances for deferral and loss offset features, due to Feldstein and Jun (1986).

14. In samples that exclude the 1980s, the excess return on long-term debt is typically negative. See the discussion of this phenomenon in Bodie et al. (1985) and Friedman (1985).

15. The average $\bar{R}^2$s of the 104 individual regressions are -0.08 for r_L and 0.11 for r_E.

16. The maximum likelihood estimation was performed using the quadratic hill-climbing algorithm supplied in the GQOPT numerical optimization package.

17. A value near two is consistent with the findings of Friend and Blume (1975). By contrast, Grossman and Shiller's (1981) work suggested a value of four, and Friend and Hasbrouck (1982) suggested six. The $\bar{R}^2$s of the two CAPM regressions are for (r_L-r_S) and for (r_E-r_S).

18. Yet a further potential explanation is that, because equations (10) estimate excess returns on a weighted basis, where the weights reflect the estimated variance-covariance structure, the <u>un</u>weighted mean excess returns estimated by these equations in general differ from the corresponding sample-period means despite the presence of a constant term in each equation. Any difference between the estimated and actual means on an unweighted basis would, of course, add to the corresponding unweighted variances. In fact, however, differences in the unweighted means account for less than 1% of the estimated variances reported in Table 4.

19. This is the value estimated by the GARCH procedure (see again Table 3); the VAR procedure does not deliver an estimate of the risk aversion coefficient. Because the formulation of the problem here treats short-term debt as a risk-free asset, and therefore focuses on the excess returns to holding long-term debt and equity, the form of B here is 2x2

rather than 3x3. In addition, the $\underline{\pi}$ vector in each case consists simply of a unit element for short-term debt and a zero element for each of the other two assets.

20. Because B is symmetric, $\beta_{LE} = \beta_{EL}$. In addition, because the columns of B uniformly add up to zero, $\beta_{SL} = -(\beta_{LL} + \beta_{EL})$ and $\beta_{SE} = -(\beta_{LE} + \beta_{EE})$.

21. An intertemporal version of the capital asset pricing model would include an extra term to reflect the investor's ability to hedge against systematic variation over time in Ω.

22. Including the constant term also allows the regression to correct for any difference in means. The estimated $\hat{\gamma}$ was not significantly different from zero in any of the regressions run, however.

23. There is no obvious weighted analog to the root-mean-square errors reported in Table 6.

24. See, for example, Shiller (1979) and other references cited in Shiller (forthcoming).

25. See, for example, Merton (1971) for an exposition of standard portfolio theory in the presence of Poisson disturbances.

References

Bodie, Zvi, Kane, Alex and McDonald, Robert. "The Role of Bonds in Investors' Portfolios, " In Friedman (ed.), Corporate Capital Structures in the United States. Chicago: University of Chicago Press, 1985.

Bollerslev, Tim P., Engle, Robert F., and Wooldridge, Jeffrey M. "A Capital Asset Pricing Model with Time Varying Covariances." Journal of Political Economy, 96 (February, 1988), 116-131.

Breeden, D.T. "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities." Journal of Financial Economics, 7 (1979), 265-296.

Cox, John C., Ingersoll, Jonathan E., and Ross, Stepthen A. "An Intertemporal General Equilibrium Model of Asset Prices." Econometrica, 53 (March, 1985), 363-384.

Ederington, Louis H. "Mean-Variance as an Approximation to Expected Utility Maximization." Mimeo: School of Business Administration, Washington University, 1986.

Engle, Robert F. "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of U.K. Inflation." Econometrica, 50 (July, 1982), 987-1008.

Engle, Robert F. "Multivariate ARCH with Factor Structures-Cointegration in Variance." Mimeo: University of California at San Diego, 1987.

Engle, Robert F., Lilien, David M., and Robins, Russell P. "Estimating Time Varying Risk Premia in the Term Structure: The ARCH-M Model." Econometrica, 55 (March, 1987), 391-407.

Estrella, Arturo and Fuhrer, Jeffrey C. "Average Marginal Tax Rates for U.S. Household Interest and Dividend Income 1954-1980." Mimeo: National Bureau of Economic Research, 1983.

Feldstein, Martin and Jun, Joosung. "The Effects of Tax Rules on Nonresidential Fixed Investment: Some Preliminary Evidence from the 1980's." Mimeo: National Bureau of Economic Research, 1986.

Friedman, Benjamin M. "Price Inflation, Portfolio Choice, and Nominal Interest Rates." American Economic Review, 70 (March, 1980), 32-48.

Friedman, Benjamin M. "Crowding Out or Crowding In? Evidence on Debt-Equity Substitutability." Mimeo: National Bureau of Economic Research, 1985.

Friedman, Benjamin M., and Roley, V. Vance. "Aspects of Investor Behavior Under Risk." In Feiwel (ed.), Arrow and Ascent of Modern Economic Theory. London: Macmillan, 1987.

Friend, Irwin, and Blume, Marshall E. "The Demand for Risky Assets." American Economic Review, 65 (December, 1975) 900-922.

Friend, Irwin, and Hasbrouck, Joel. "Effect of Inflation on the Profitability and Valuation of U.S. Corporations." In Sarnat and Szego (eds.), Savings, Investment, and Capital Markets in an Inflationary Economy. Cambridge: Ballinger, 1982.

Grossman, Sanford J., and Shiller, Robert J. "The Determinants of the Variability of Stock Prices." American Economic Review, 71 (May, 1981).

Ibbotson, Roger G., and Sinquefield, Rex A. Stocks, Bonds, Bills and Inflation: The Past and The Future. Charlottesville: The Financial Analysts Research Foundation, 1982.

Kraft, Dennis F., and Engle, Robert F., "Autoregressive Conditional Heteroskedasticity in Multiple Time Series Models." Mimeo: University of California at San Diego, 1983.

Lintner, John. "The Valuation of Risky Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets." Review of Economic Statistics, 47 (1965), 13-37.

Meiselman, D. The Term Structure of Interest Rates. Englewood Cliffs, N.J. Prentice-Hall, 1962.

Merton, Robert C. "Optimum Consumption and Portfolio Rules in a Continuous-Time Model." Journal of Economic Theory, 3 (December, 1971), 373-413.

Merton, Robert C. "An Intertemporal Capital Asset Pricing Model." Econometrica, 41 (September, 1973), 867-888.

Merton, Robert C. "Capital Market Theory and the Pricing of Financial Securities." In Friedman and Hahn (eds.), Handbook of Monetary Economics. Amsterdam: North holland (forthcoming).

Modigliani, Franco, and Shiller, Robert J. "Inflation, Rational Expectations, and the Term Structure of Interest Rates." Economica, 40 (February, 1973), 12-43.

Melino, A. "The Term Structure of Interest Rates: Evidence and Theory." Mimeo: National Bureau of Economic Research, 1986.

Roley, V. Vance. "The Effect of Federal Debt-Management Policy on Corporate Bond and Equity Yields." Quarterly Journal of Economics, 97 (November, 1982), 645-668.

Ross, S.A. "The Arbitrage Theory of Capital Asset Pricing." Journal of Economic Theory, 13 (1976), 341-360.

Sharpe, W.F. "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk." Journal of Finance, 19 (1964, 425-442.

Shiller, Robert J., "The Volatility of Long-Term Interest Rates and Expectations Models of the Term Structure." Journal of Political Economy, 87 (December, 1979), 1190-1219.

Shiller, Robert J. "The Term Structure of Interest Rates." In Friedman and Hahn (eds.), Handbook of Monetary Economics. Amsterdam: North Holland (forthcoming).

Singleton, Kenneth J. Specification and Estimation of Intertemporal Asset Pricing Models." In Friedman and Hahn (eds.), Handbook of Monetary Economics. Amsterdam: North Holland (forthcoming).

Stiglitz, J.E. "A Consumption-Oriented Theory of the Demand for Financial Assets and the Term Structure of Interest Rates." Review of Economic Studies, 37 (July, 1970), 321-351.

</ref_section>