The Behaviour of UK Stock Prices and Returns

BY

SIMON HAYES

SUBMITTED FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY
SEPTEMBER 1995

DEPARTMENT OF ECONOMICS
UNIVERSITY OF NEWCASTLE UPON TYNE,
NEWCASTLE UPON TYNE. NE1 7RU

Abstract

In this thesis I combine VAR forecasting methods with the Campbell-Shiller log-linear approximation to the present-value formula for stock prices. Four aspects of UK stock market behaviour are studied.

The first analysis involves decomposing the variance of the unexpected stock return into components due to news about dividends, news about future returns, and the covariance between the two. I find that changing expectations about future returns accounts for around four times as much of the variance of unexpected returns as news about dividends, with a negligible covariance term. My second study is a detailed analysis of the links between macroeconomic risks and required stock returns. Using 27 industry-based stock portfolios, I attempt to determine the effect that a number of macroeconomic and financial factors have on expectations of dividends, real interest rates and future required returns. The results go some way to explaining why some risk factors appear not to be significantly priced in financial markets, whilst others (particularly inflation) appear to induce counter-theoretical reactions in stock prices.

Given an empirical proxy for equilibrium returns, the present-value model implies a set of non-linear restrictions on the parameters of a VAR, the latter being taken as a model of investors' expectations formation. In my third analysis, I test various models of equilibrium returns using aggregate UK data, and find some support for market efficiency. In particular, in accordance with the intertemporal CAPM, I find that the well-known ability of the dividend yield to forecast stock returns can be traced to the fact that the dividend yield Granger-causes the market return variance. In the final section I test two propositions: whether rejections of the CAPM at the aggregate level can be traced to rejections in specific sub-sectors of the market; and whether investors are more skilled at eliminating mis-pricing within market sub-sectors than in the market as a whole. I find mixed support for the CAPM at the disaggregated level. Furthermore, eliminating the covariance terms from the model for sector returns has little effect on the results, providing some support for the market segmentation hypothesis.

Acknowledgements

First of all, I would like to thank my supervisor, Keith Cuthbertson, whose close and enthusiastic supervision ensured that things kept ticking over. Thanks are also due to my friend and colleague Dirk Nitzsche, whose close and enthusiastic beer drinking made the whole experience much more pleasurable.

I would also like to thank everyone in the Department of Economics at Newcastle University for their guidance and support, and the Fellowhips and Studentships Committee for their financial help via the Strang Studentship.

Finally, I am greatly indebted to my wife, Linda, for her immeasureable patience and support. Without her, this thesis would have been finished ages ago; but I would have been all the sadder for it.

Contents

List of Tables

List of Figures

Chapter 1: INTRODUCTION AND OVERVIEW

1.1 Introduction

This thesis has two broad aims. The first addresses the question of which factors have, historically, been the most important influences on the behaviour of stock prices and returns in the UK. The second aim is to test the degree to which various popular asset pricing models are able to explain the time-series and cross-sectional behaviour of UK stock prices. The main difference between this study and most of the previous analyses of UK stock market behaviour is that I constrain all effects to enter only via the Rational Valuation Formula (RVF) for stock prices[1]. The RVF states that the current stock price is equal to the expected discounted present value of future dividend payments, the relevant discount rate being investors' required return on the stock. Thus all of the potential influences considered in this thesis are relevant only in so far as they affect expectations about future dividends, expectations of future discount rates, or both. This contrasts with the bulk of the literature on stock price behaviour, which does not attempt to distinguish between these separate effects.

Since the general form of the RVF is non-linear, it does not easily lend itself to econometric analysis. However, one can make the RVF tractable by imposing certain restrictive assumptions regarding investors' expectations. In particular, previous studies have assumed that dividends are expected to grow at a constant rate, or that expected returns are constant. I do not wish to impose either of these conditions on my analysis. I therefore employ several versions of the log-linearisation of the RVF (the "dividend-price ratio model") developed by Campbell and Shiller (1988, 1989). This model has two important advantages over the RVF. First, it is linear in the expectations terms, whether or not those expectations are assumed to be time-varying. Second, it can be stated in terms of the dividend-price ratio and dividend growth, rather than stock prices and dividends. Since the latter are widely believed to be non-stationary series, whereas the former are not, standard methods of statistical analysis can be applied to the Campbell-Shiller dividend-price ratio model which could not be used to study the RVF. The cost of these substantial conveniences is a degree of approximation error, but this is shown empirically to be

---

1 Miles (1993) is an exception.

relatively small, making the dividend-ratio model a powerful tool for research into stock price behaviour.

This introductory chapter is organised as follows. In Section 1.2 I briefly describe the historical development of the RVF. This is useful because it allows one to lay out the conditions required for the formula to hold, and also makes clear why certain simplifications of this formula have dominated the stock market literature until fairly recently. Section 1.3 gives a brief introduction to the Campbell-Shiller log-linearisation of the RVF. The main part of the thesis comprises four studies of different aspects of stock price behaviour. Section 1.4 takes each of these chapters in turn and discusses the motivations behind each analysis, with particular emphasis on the context in which each individual study is set.

1.2 The Historical Development of the Rational Valuation Formula

The early development of the concept of rational valuation was motivated by an attempt to reconcile classical economic theory with popular beliefs about the prices of stock market investments. I now briefly trace this development.

1.2.1 Fundamental Analysis

The popular view held by financial market practitioners and commentators is that a security has a "fundamental value" equal to the discounted present value of cash flows accruing to its holder, and the price of the security fluctuates around this value. Fundamental value changes when underlying income-generating conditions change; but these create trends rather than instantaneous jumps in prices, because most traders have imperfect knowledge of fundamentals. The future trend in prices reflects the gradual dissemination of awareness of the sign and magnitude of fundamental shifts. The prescription to market analysts is to identify assets' fundamental values, to buy assets whose current price is below fundamental value, and sell assets whose current price exceeds fundamental value. An investor who can identify and interpret changes in fundamentals before their effects are fully incorporated into asset prices will outperform a simple buy-and-hold strategy.

There are two main schools of professional analysts: the "fundamentalists" and the "technical analysts". Whilst both agree on the above description of market structure, their methods for exploiting it are quite different. Fundamental analysis consists of projecting future cash flows by studying general business conditions, demand and profit prospects for individual firms and sectors. The emphasis then is on external factors which underlie price changes. Technical analysis feeds off fundamental analysis in so far as it attempts early identification of changes in prices which are caused by trading on fundamentals. The technician studies price movements of the immediate past for signs of movement in the immediate future. For example, Charles H. Dow, founder of the Dow Jones financial news service and founder and editor of the *Wall Street Journal*, plotted daily the industrials and railroads price indices. If a move in one was confirmed by a move in the other, this was taken as an indication of a fundamental shift upon which investment advice could be based. This system became known as the Dow Theory.

The first major setback to fundamental analysis came from a study by Alfred Cowles (1933), who claimed that fundamental analysis simply did not work. Part of his study involved looking at 7,500 recommendations from 16 leading financial services about individual stocks over a four-and-a-half-year period. He found that only 6 firms on average achieved positive returns, and that the annual average effective rate of return of all the services was -1.43%. Moreover, statistically there was an evens chance of finding at least one service performing as well as the best in the sample. Thus the performance of the best firm could be ascribed as much to luck as to skill.

Cowles also studied the recommendations of William Peter Hamilton, successor of Dow as editor of the *WSJ* and principal sponsor of the Dow Theory. Over the twenty-six years of his incumbency, Hamilton's tips would have earned a return of 12% per annum, which appears to be very successful. However, over the same period, a buy-and-hold strategy would have returned 15% per annum.

1.2.2 The Random Walk Model

It has long been popular in the statistical analysis of economic time series to decompose series into a trend component and a cyclical component. The two are studied separately as the manifestations of long-term and short-term influences respectively. However, the development of autoregression schemes, in which disturbances play an integral part, led to the suggestion that the long-term was in fact a cumulation of short-term movements, and so to treat the two separately was misguided. King (1930) concluded that stock prices resembled cumulations of purely random changes. Working (1934) pointed out that such cumulations would give rise to conspicuous trends, but that these trends could not be interpreted as a base for predicting the future course of the series. He suggested ascertaining with what relative frequency patterns recognised by technicians actually occurred in purely random-difference series. "If they occur as frequently and as clearly in the random-difference series as in the actual series, it is to be supposed that they are without forecasting significance, for it is known that changes in a random-difference series are quite unpredictable" (Working (1934) p21).

Kendall (1953) searched for systematic effects in stock price changes, but found any such effects statistically undetectable; autocorrelations and cross-correlations with other stock price changes were found to be so weak as to be useless for predictive purposes. Roberts describes Kendall's results with a gambling analogy:

> "Kendall found that changes in security prices behaved nearly as if they had been generated by a suitably designed roulette wheel for which each outcome was statistically independent of past history and for which relative frequencies were reasonably stable through time. This means that, once a person accumulates enough evidence to make good estimates of the relative frequencies (probabilities) of different outcomes of the wheel, he would base his predictions only on these relative frequencies and pay no attention to the pattern of recent spins. Recent spins are relevant to prediction only insofar as they contribute to more precise estimates of relative frequencies. In a gambling expression, this roulette wheel 'has no memory'" (taken from Cootner (1964) p9).

Studies by Granger and Morgenstern (1963), who used spectral techniques, and Cootner (1962), amongst others, corroborated Kendall's findings.

Roberts labelled this the Chance Model. The simplest embodiment of the Chance Model is that prices follow a random walk:

$$P_t = P_{t-1} + \varepsilon_t \qquad \varepsilon_t \text{ white noise}$$

where P_t is the real stock price. Thus the best forecast of tomorrow's price given today's price is simply today's price: price movements are unpredictable on the basis of past prices. Moreover, the variance of the stock price increases through time (Kendall noted this phenomenon when studying wheat prices). It became customary to assume in addition that the errors were normally distributed (Working had used drawings from a normal distribution for comparison, and Bachelier (1900) had modelled asset prices in continuous time with a Brownian Motion process, the increments of which are centred normal).

If stock prices did indeed perform random walks, success by investors must be due to a one or more of the following:

i) luck;
ii) the fact that at certain times all prices rise together - a no-lose situation;
iii) having inside information to anticipate movements.

It was soon realised that random behaviour may reflect instantaneous adjustment to new information, assuming that information arrives at random intervals, and one would expect this in a market dominated by rational individuals. This insight raised an important question for the practitioners' view: if fundamental analysis works, why do new agents not enter the market and compete the gains away? Cowles's results suggested that this did in fact happen. Indeed, the condition that excess profit opportunities cannot exist in equilibrium became the foundation of the Efficient Markets Hypothesis.

Whilst the fundamentalists found themselves in a difficult position, the random walk model was not without its problems. As the literature developed, a number of empirical inconsistencies and theoretical enigmas were highlighted. On the empirical side, the focus of research shifted away from documenting the similarities between price series and random walks towards highlighting systematic discrepancies between the two. These are many and various, but for the purposes of the current exposition the most important was reported by Mandelbrot (1963). He observed that, whilst price changes appeared to be serially uncorrelated, "large changes are followed by large changes - of either sign - and small changes tend to be followed by small changes". This phenomenon, sometimes termed "volatility clustering", militates against the independence assumption embedded in the random walk formulation of the simple Chance Model.

There were two main theoretical problems with the random walk hypothesis. First, it appeared to imply that stock prices, as the cumulation of the effects of random events, had nothing to do with preferences and technology, whilst these were believed to underpin price determination in all other markets. There must exist an optimisation framework which linked these foundations to random behaviour, yet none was forthcoming. Second, the random walk hypothesis appeared to state that investors at once rationally compete all excess gains away and irrationally waste resources on fundamental analysis. Can agents be simultaneously rational and irrational in an economic equilibrium?

1.2.3 The Martingale Model for Stock Prices

From the point of view of economic theory, the assumption that excess profit opportunities cannot persist in equilibrium was paramount; but the random walk model was simply too restrictive to be generated within a reasonably broad class of optimisation models. It turned out that a necessary and sufficient condition for price processes to admit no arbitrage opportunities is that they are related to martingales. Furthermore, the martingale model can easily be derived from a bona fide economic pricing model. Samuelson's (1965) paper was the first to develop the link between financial market efficiency and martingales, and as such is viewed by many to the most important paper in the efficient markets literature.

A stochastic process x_t is a martingale with respect to some information set Φ_t if

$$E(x_{t+1}|\Phi_t)=x_t$$

i.e. the conditional first moment at time t is x_t. A stochastic process y_t is a "fair game" if

$$E(y_{t+1}|\Phi_t)=0$$

thus if x_t is a martingale, Δx_t is a fair game (and so fair games are sometimes called "martingale differences").

Notice that if x_t follows a random walk then x_t is a martingale and $\Delta x_t = \varepsilon_t$, a fair game. However, by requiring probabilistic independence between the ε_t's, the random walk is a much more restrictive stochastic process than the martingale. Volatility clustering, for instance, is not inconsistent with the martingale model, but cannot occur in a random walk.

Economic theory predicts that returns in excess of equilibrium expected returns (i.e. that return just sufficient to induce a investor to hold a security) are a fair game. Suppose that H_{t+1} is the return on a stock from t tot+1, and suppose also that the equilibrium expected return is a constant H. Using the definition of a stock return, and taking expectations, we have

$$H = E\left(\frac{P_{t+1} - P_t + D_t}{P_t} | \Phi_t\right)$$

which can be rewritten as

(1.1) $P_t = (1+H)^{-1} E(P_{t+1} + D_{t+1} | \Phi_t)$

The current stock price thus equals the discounted present value of expected dividend income next period and the expected price at time t+1, the equilibrium expected return being the relevant discount rate.

None of the variables defined so far is a martingale. However, consider a mutual fund which holds stock, the price of which follows (1.1). The fund is assumed to reinvest dividends in further share purchases. Let v_τ be the value of this mutual fund discounted back to time τ:

$$v_\tau = \gamma^{t-\tau} P_t N_t$$

where N_t is the number of shares at time t and $\gamma = (1+H)^{-1}$. The assumption that dividends are reinvested implies that

$$P_{t+1}N_{t+1}=(P_{t+1}+D_{t+1})N_t$$

Taking expectations of $v_{\tau+1}$, conditional on Φ_τ (from now on denoted E_τ),

$$\begin{aligned} E_\tau(v_{\tau+1}) &= \gamma^{t-\tau+1}E_\tau(P_{t+1}N_{t+1}) \\ &= \gamma^{t-\tau+1}E_\tau[(P_{t+1}+D_{t+1})N_t] \\ &= \gamma^{t-\tau}E_\tau[\gamma(P_{t+1}+D_{t+1})N_t] \\ &= \gamma^{t-\tau}P_tN_t = v_\tau \end{aligned}$$

Therefore v_t is a martingale[2].

Samuelson (1973) set out the conditions, in terms of preferences, which lead to the martingale model. The martingale model would arise if investors:

i) have a common and constant time preference;

ii) agree on the underlying probability distribution generating returns;

iii) are risk neutral.

In this case they will always prefer to hold the asset with the highest expected return, regardless of its risk characteristics. In equilibrium, all assets must have the same expected return - the equilibrium expected return, H - which will equal the real interest rate (a consequence of risk neutrality), and which I have assumed constant.

Risk neutrality implies the martingale model, but not the random walk model. If agents are risk neutral, they are only concerned about expected returns, and not about the higher moments of the return distribution. Therefore arbitrage will bid away serial correlation only in the mean of the series, and returns are not necessarily independently distributed across time. The fact that, in particular, variances may be conditionally forecastable means nothing since variances are of no consequence in the investment decision-making process. Fama (1970) pointed out that a random walk will arise only if the evolution of preferences and technology and the information-generating process conspire to produce equilibria in which return distributions replicate themselves through time.

2 It is important to emphasise that although the model is commonly termed "the martingale model for stock prices", stock prices themselves do not follow a martingale process. Rather, stock prices plus reinvested dividends is the martingale. Consequently, the term "price" should be understood to include reinvested dividends.

Samuelson's (1973) paper provided a result which has proved very useful in analysing the efficient markets hypothesis as described by the martingale model. He noted that the law of iterated expectations, combined with the transversality condition $\lim_{i\to\infty} E_t\gamma^i P_{t+i}=0$ (i.e. the real stock price does not explode over time) means that (1.1) can be solved forwards to yield

$$(1.2)\quad P_t = \sum_{i=0}^{\infty} \gamma^i E_t D_{t+i}$$

(1.2) is the present value formula for stock prices, and states that the current price of a stock equals the discounted present value of expected future dividends.

(1.2) is interesting because it illustrates that the efficient markets model and the fundamentalists' model are not diametrically opposed: in fact the efficient markets model is best seen as an extreme form to the fundamentalists' model. Samuelson's result implies that stock prices *always* equal the discounted present value of cash flows, whereas the fundamentalists' view is that the actual price fluctuates around this value. If we start with the fundamentalists' model and then assume that a large number of traders conducting fundamental analysis reach the same conclusion about fundamental value (common probabilities, common and constant time preference) the present value formula (1.2) will obtain. Thus the basic difference between the two models is the assumed degree of homogeneity and competition in the market, since the unpredictability or predictability of stock price changes (or stock returns) arises as a corollary of whether (1.2) holds continuously or not.

1.3 The Campbell-Shiller Log-Linearisation of the RVF

Campbell and Shiller (1988, 1989) used a first-order Taylor approximation to derive the following log-linear version of the RVF:

$$(1.3)\quad p_t = (1-\rho)\sum_{j=0}^{\infty} \rho^j E_t d_{t+1+j} - \sum_{j=0}^{\infty} \rho^j E_t \xi_{t+1+j} + k$$

where p_t is the log real stock price at the end of time t, d_t is the log real dividend payment during period t, ξ_t is the approximate log real holding period return on the stock, ρ is a linearisation constant (see Section 3.3), and k is a constant. If we define δ_t as the log dividend-price ratio, p_t-d_t (1.3) can be rewritten as follows:

$$(1.4) \quad \delta_t = \sum_{j=0}^{\infty} \rho^j E_t(\xi_{t+1+j} - \Delta d_{t+1+j}) + k$$

This states that the log dividend-price ratio is an optimal linear forecast of future real returns and real dividend growth. The dividend-ratio model (1.4) is used in Chapters 5 and 6 to test specific models of asset pricing. The idea is to replace the unobservable required return ξ_t with an empirical proxy implied by theory, and to compare the behaviour of forecasts of future dividend growth and returns with that of the dividend-price ratio. (1.4) implies that, with the correct model for expected returns, the two series should be the same.

The second important product of the dividend-ratio model is the following expression:

$$(1.5) \quad h_t - E_t h_t = (E_{t+1} - E_t)\sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} - (E_{t+1} - E_t)\sum_{j=0}^{\infty} \rho^j h_{t+1+j}$$

where h_t is the log real stock return. Thus, the unexpected return on a stock is equal to revisions in expectations about future dividend growth, minus revisions to expectations about future real returns. Given a model relating various sources of news about stock returns to expectations of future dividends and discount rates, (1.5) can be used to apportion unexpected returns between revisions in expectations of dividends and discount rates, and to calculate the effects of individual news factors on these expectations. These analyses are undertaken in Chapters 3 and 4.

1.4 Outline of Thesis

1.4.1 Chapter 3: Explaining Movements in UK Stock Prices

Rational valuation requires that investors value a stock by discounting their best forecast of the dividend stream accruing to the stockholder. If dividend payments are not expected to be constant for all time, the RVF implies a certain amount of *predictable* movement in stock prices. This occurs because, when dividends are paid, the time profile of future dividends is "shifted" forwards by one period, so that all expected future payments are discounted at a different rate. This causes stock prices to change over time even if expected returns are constant. Of course, if expected returns are not constant, but vary in a predictable fashion, even if dividends are constant there will be some predictable variation in stock prices. However, although there is a large body of literature which presents models for predicting stock price changes, the predictive power is commonly very low. Consequently, the larger portion of stock price changes results from *unexpected* price movements. From the above discussion, it follows that any unpredictable movement in stock prices must stem from either *changes* in expectations of future dividends, or *changes* in future discount rates, or both. In Chapter 3 I analyse unexpected stock price changes (returns[3]), and seek to determine whether it is changes in expected dividends or changes in expected returns which have the major impact on stock price movements.

Campbell (1991) makes the important point that the contribution of changes in expected returns to the overall variance of stock price changes depends not only on the degree to which stock returns are predictable, but also on the persistence in the expected returns process. Even if the degree cf return predictability is low, if expectations of returns are persistent (so that a current shock to expected returns has a near-permanent effect), then changes in expected returns could have a large impact on the current stock price. This is because dividends are discounted by functions of all future expected returns, so that if a change in the current expected return implies revisions in all future expected returns, all discount rates change dramatically, and there will be a correspondingly large impact on the current stock price.

The idea of persistence in return expectations leads me to consider the area of "volatility persistence". This is a separate body of research which has concentrated on modelling the time-series behaviour of the variance of unexpected stock returns.

---

3 Since the larger proportion of stock returns consists of capital gains/losses, rather than dividend income, I use the terms "price changes" and "returns" interchangeably.

The idea is similar to Campbell's, in that it is believed that changes in volatility will affect stock prices only if shocks to volatility have a persistent effect on volatility expectations, but previous analyses have never sought to combine models of return volatility with the RVF. Using a VAR system to model both expected stock returns and expected return volatility, I seek to provide a measure of the persistence in expected returns, and the impact this has on the contribution of changes in return expectations to unexpected stock price movements.

As well as studying real stock returns, unexpected excess returns (above some short interest rate) are decomposed into three factors: changes in expectations of future dividends; changes in expectations about future real interest rates; and changes in expectations about future excess returns. This three-way decomposition highlights some important offsetting covariation between these elements.

Since the decomposition requires a model for expected real and excess returns, the objectives of Chapter 3 are as follows. First, evidence is presented of the predictive power of various series known in advance. Second, the variance of both unexpected real and unexpected excess returns is decomposed as discussed above. Third, the time-series behaviour of the implicit expected return process is analysed. The analysis is performed on annual observations of an aggregate UK stock market index over the period 1918 to 1993.

1.4.2 Chapter 4: Identifying Sources of Systematic Risk in the UK Stock Market

In recent years a large amount of research effort has been directed towards multi-factor models of asset pricing. In the Linear Factor Model (LFM), the unexpected return on any asset is a linear function of so-called "factor innovations". "Factors" are sources of financial risk which impact upon investors' required asset returns. The coefficient relating the kth factor to the ith asset is known as the "factor loading", or more simply the "beta" between factor k and asset i. Although the LFM itself does not specify the nature of expected asset returns, the Arbitrage Pricing Theory (APT) of Ross (1976) uses the absence of profitable arbitrage opportunities to derive a linear relationship between the expected return on any asset and the extra return required by investors to compensate for the risk induced by a particular factor (termed the "market price" of a factor risk). Previous research has been concerned with two issues: determining the number of factors priced in various asset markets; and identifying specific macroeconomic variables which might proxy for the unidentified risk factors. In Chapter 4 the emphasis is on the factor loadings rather than the factors themselves. As mentioned in Section 1.4.1 above, the unexpected excess return on any asset depends on revisions to expectations of future dividends, future real interest rates, and future excess returns. Thus if any factor is to impact upon the stock market, innovations in that factor must lead investors to revise their expectations of one or more of these three components. I therefore decompose the betas of a set of financial and macroeconomic factors into components due to news about dividends, interest rates and excess returns, to determine through which channels each factor has its main influences. This decomposition is performed on a set of 27 industry-based portfolios, over the period 1970-1993.

Some factors are found to have a significant influence on expectations of more than one component. In this case, it is possible that the separate effects are offsetting or reinforcing. Consequently, whilst it is not always easy to predict the net effect of any factor on required returns, the decomposition of that factor's beta sheds light on the source(s) of risk associated with that factor, and may help explain the sign and significance of the factor in the linear factor model.

Including the market excess return as a factor in the model allows me to analyse several issues. By looking at the way in which dividend expectations in each sector are revised when expected dividends on the market portfolio change, I can identify certain sectors as "pro-cyclical" and "non-cyclical". This proves a useful descriptive device when cross-sectional variation in other factor betas is being analysed. Also, I am able to construct both formal and informal tests of the CAPM (which is a one-factor model, a special case of the APT), combining both the cross-sectional and the time-series information contained in the set of asset portfolios.

Of particular interest is the beta between stock returns and inflation. Although the Fisher hypothesis (which states that nominal asset returns move one-for-one with inflation) appears to imply that there should be no statistical relation between real stock returns and inflation, a significant negative relation has been consistently observed. There have been several attempts to explain this phenomenon, the most notable being that the relation is spurious and occurs because inflation, through the money market, is correlated with expected future real activity. Assuming that expectations of future dividends are revised with changes in expected future real activity, I can estimate directly the correlation between shocks to inflation and revisions to expected future activity. I can also determine the contribution of the latter to current excess returns. In this way, the decomposition of the inflation beta allows me to comment in depth on the proposed relations, and to identify the sources of cross-sectional variation in inflation betas.

1.4.3 Chapter 5: Testing the Efficiency of the UK Stock Market

The stated purpose of this chapter is to determine how well popular models of asset-price determination describe the historical behaviour of the aggregate UK stock market, using the same data as in Chapter 3. Ostensibly the same data set was used by Bulkley and Tonks (1989), who applied variance bounds tests to the real stock price series, under the assumption (supported by formal testing) that real prices and dividends are stationary series. Under weak-form rational expectations, Bulkley and Tonks could not reject the hypothesis that expected returns were constant. However, the data they used have since been revised, with the effect that the stationarity assumption is no longer supported by formal tests. I take a different approach to Bulkley and Tonks, and in Chapter 5 I apply the Campbell-Shiller dividend-price ratio model to test three different models of discount rate determination: that expected returns are constant; that expected return vary only with the safe rate of interest; and that expected returns depend on the conditional variance of the market return.

The analysis is founded on the fact that, according to the dividend-ratio model, the log dividend-price ratio is an optimal forecast of future real dividend growth and future expected real returns. Rational valuation implies that a linear combination of optimal forecasts of dividends and discount rates (which is termed the "theoretical" dividend-price ratio) should equal the observed dividend-price ratio. Using a VAR to forecast dividend growth and discount rates, a theoretical dividend-price ratio is constructed, and compared to the actual series. Many metrics are used - informal graphical comparisons, together with descriptive statistics of the behaviour of the two series, are supplemented by formal Wald tests of sets of linear and non-linear restrictions on the parameters of the VAR which are implied by the rational valuation model. The tests are performed using returns over a number of different horizons. One particularly interesting implication of the model is that predictability of one-period returns can be directly linked to the predictability of infinite-horizon returns, which in turn impacts on the success or failure of the present value relation.

1.4.4 Chapter 6: Market Segmentation and Stock Price Behaviour

The main implication of the CAPM is that covariances between asset returns are the sole source of systematic risk for the investor. The development of the ARCH model has spawned a large number of studies in which asset returns are dependent upon the conditional variance of the unexpected return on that asset. Despite the Roll Critique (which states that, since an empirical proxy for the "market" of investment opportunities is probably impossible to come by, the CAPM is virtually untestable), one would still expect a model in which covariances with other asset returns were included, in addition to the conditional variance of the return on the asset, to improve upon a model which omits these covariances. However, if investors concentrate their attention only on market sub-sectors, so that mis-pricing is eliminated within sub-sector portfolios but not across sub-sectors, one would not expect any covariance terms to ameliorate the own-variance model. If "market segmentation" exists, each sub-sector is effectively a separate market, and so required sub-sector returns should depend only on the conditional variance of returns within that sub-sector. In the case of market segmentation, conditional covariances with other asset returns should have no incremental explanatory power over returns on sub-sector portfolios.

In Chapter 6 I study whether models of equilibrium returns are improved by the inclusion of conditional covariances with returns on other assets. My study focuses on the stock market, comparing the returns on six sector portfolios - Capital Goods, Consumer Goods, Industrials Including Oil, Industrials Excluding Oil, Financial Services and Other Sectors - with the return on the aggregate stock market, over the period 1965 to 1993. This work differs from previous studies, which have used multivariate GARCH-M models (see Bollerslev, Engle and Wooldridge 1988, and Hall, Miles and Taylor 1990) to capture time-variation in conditional covariances, in two respects. On the downside, since I use a VAR model for all forecasts, my models for conditional variances and covariances are not as sophisticated as non-linear GARCH models. However, unlike previous studies, I impose the structure of the RVF, so that all effects enter solely through their impact on expected dividends and discount rates. Moreover, neither of the above studies tests the significance of the conditional covariance terms in the multivariate model.

Chapter 2: ECONOMETRIC ISSUES

2.1 Introduction

All of the subsequent analyses share two common features: VAR models are used to forecast sets of financial and macroeconomic variables; and GMM estimators are employed, including a correction for an unknown form of heteroscedasticity. Since the two are probably not yet considered to be "standard practice" in applied econometrics, this chapter details the econometric principles involved, and provides the motivation for their use. Section 2.1 deals with the issues relevant to Vector Autoregressions. In particular, it is argued that, since no underlying set of structural equations is assumed here, modelling expectations with VARs does not suffer from the criticisms of the Sims methodology which were levelled by the Cowles Comission. In Section 2.2, I discuss Generalised Method of Moments estimation, with particular emphasis on the particular applications in the subsequent chapters of this thesis. Finally, in Section 2.3, I briefly discuss the Phillips-Perron (1988) methodology for testing series for unit roots, which is used throughout this thesis.

2.2 Vector Autoregressions (VARs)

2.2.1 Introduction

Vector autoregressions (VARs) are simple multivariate time-series models in which each variable is dependent on its own past values and those of the other variables in the system. Although in the time-series literature VARs are regarded as straightforward extensions of univariate ARMA models, their use in econometrics has been the source of much controversy. Most of this stems from claims that VARs are useful and less restrictive models for hypothesis testing and policy analysis than traditional structural econometric models. Given the importance of this debate, it is useful to discuss the precise motivation behind employing VARs in the current analysis, and such a discussion is the main purpose of this section. However, I begin with a brief introduction to VAR modelling and forecasts.

2.2.2 Estimation and Forecasting with VAR Models

Suppose that one wishes to model the behaviour of two economic variables, x_{t+1} and y_{t+1}[1]. A first-order VAR (denoted VAR(1)) model is given by

---

1 All variables are defined as as deviations from their mean.

(2.1) $x_{t+1}=a_{11}x_t+a_{12}y_t+\varepsilon_{1t+1}$

(2.2) $y_{t+1}=a_{21}x_t+a_{22}y_t+\varepsilon_{2t+1}$

where the ε_{it} are random errors which may be contemporaneously correlated. Notice first that, since the RHS variables of the two equations are identical, the model parameters may be efficiently and consistently estimated by OLS. Another important point is that there are no exogenous variables in VAR models: once a vector of variables has been chosen, no restrictions are placed on feedback effects between them (apart from linearity).

The ease of forecasting provided by VAR models is best seen if we rewrite equations (2.1) and (2.2) in matrix form:

(2.3) $\mathbf{z}_{t+1}=\mathbf{A}\mathbf{z}_t+\varepsilon_{t+1}$ where $\mathbf{z}_t \equiv \begin{bmatrix} x_t \\ y_t \end{bmatrix}$ and $\mathbf{A} = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix}$

Suppose that one wishes to calculate a forecast of x_{t+2} based only on information available at time t and earlier. Defining ι_1 as a 2x1 vector with 1 as its first element and zero as its second, the forecast of x_{t+2} is given by the expression

(2.4) $E_tx_{t+2}=\iota_1'\mathbf{A}^2\mathbf{z}_t$

To see why this is true, leading equation (2.1) by one period and take expectations at time t, gives

(2.5) $E_tx_{t+2}=a_{11}E_tx_{t+1}+a_{12}E_ty_{t+1}$

Replacing the expectations on the RHS of equation (2.5) with the forecasts from equations (2.1) and (2.2), and collecting terms,

(2.6) $E_tx_{t+2}=(a_{11}^2+a_{12}a_{21})x_t+(a_{11}a_{12}+a_{12}a_{22})y_t$

The coefficients on x_t and y_t in equation (2.6) are the same as the elements in the first row of the $\mathbf{A}^2$ matrix in equation (2.4), and the latter is precisely what is picked out by the selection vector ι_1.

Fortunately, the ease with which multi-period-ahead VAR forecasts can be obtained is not restricted to first-order VAR systems, since a VAR of any order may be rewritten in first-order form. To see this, consider the general VAR(p) model:

$$(2.7)\quad x_{t+1} = \sum_{i=0}^{p-1} a_i x_{t-i} + \sum_{i=0}^{p-1} b_i y_{t-i} + \varepsilon_{1t+1}$$

$$(2.8)\quad y_{t+1} = \sum_{i=0}^{p-1} c_i x_{t-i} + \sum_{i=0}^{p-1} d_i y_{t-i} + \varepsilon_{2t+1}$$

Redefining z_t to include all past values of x and y, together with a suitable redefinition of the A matrix, we have

$$(2.9)\quad \begin{bmatrix} x_{t+1} \\ \cdot \\ \cdot \\ \cdot \\ x_{t-p+2} \\ y_{t+1} \\ \cdot \\ \cdot \\ \cdot \\ y_{t-p+2} \end{bmatrix} = \begin{bmatrix} a_1 & . & . & a_p & b_1 & . & . & b_p \\ 1 & & & & & & & \\ & \cdot & & & & & & \\ & & \cdot & & & & & \\ & & & 1 & & & & \\ c_1 & . & . & c_p & d_1 & . & . & d_p \\ & & & & 1 & & & \\ & & & & & \cdot & & \\ & & & & & & \cdot & \\ & & & & & & & 1 \end{bmatrix} \begin{bmatrix} x_t \\ \cdot \\ \cdot \\ \cdot \\ x_{t-p+1} \\ y_t \\ \cdot \\ \cdot \\ \cdot \\ y_{t-p+1} \end{bmatrix} + \begin{bmatrix} \varepsilon_{1t} \\ 0 \\ \cdot \\ \cdot \\ 0 \\ \varepsilon_{2t} \\ 0 \\ \cdot \\ \cdot \\ 0 \end{bmatrix}$$

where blank elements are zero. Expression (2.9) is a first-order VAR, just like equation (2.3), and as such j-period-ahead forecasts can be obtained as

$$(2.10)\quad E_t z_{t+1+j} = \mathbf{A}^{j+1} z_t$$

Predictions derived from expressions such as equation (2.10) will play a major role in my analysis of the UK stock market.

One of the key decisions when constructing a VAR model is the choice of maximum lag length, p. If p is large, then the number of parameters to be estimated will be large. and the VAR is more likely to pick up within-sample random variation, in addition to any systematic relationships. However, if p is too small, important lag dependencies may be omitted, and if there remains serial correlation in the VAR residuals then the estimated coefficients on the lagged dependent variables will be inconsistent. Often, it is more convenient to report results from a VAR(1) model, and then to study any qualitative effects of increasing the VAR lag length, rather than identifying the most appropriate value for p by formal testing. However, in some cases I do initially report the results using the "optimal" VAR lag length, and my initial choice of lag length is determined as follows. A likelihood-ratio test of a VAR system with p+1 lags versus a p-lag VAR is used to test the joint significance of the extra lagged variables. Extra lags are added until the next lag leads to an insignificant improvement in the likelihood function, subject to there being no detectable serial correlation in the residuals of any of the VAR equations. The latter is determined using an F-test of the hypothesis that lagged residuals have no explanatory power over the error process from each VAR equation.

2.2.3 Time-series Background

ARMA models have been popularly employed for many years to decompose economic time series into predictable and unpredictable components without reference to theory-based economic relationships. The VAR model may then be viewed as a natural extension of this tradition. In so far as the extra variables may be chosen with reference to economic theory (or at least economic intuition), the VAR potentially has a slightly more econometric flavour than the ARMA model. However, the economic content must not be over-emphasised: when constructing a VAR model the intent is to allow the data to determine the precise form of the model.

Wold's (1938) well-known result that any stationary variable can be uniquely represented by a (possibly infinite) moving average (MA) process provides the keystone for VAR modelling. As long as the MA lag polynomial is invertable, an MA process may be more parsimoniously stated as a finite autoregressive (AR) process. More generally, the methodology of Box and Jenkins (1970) relies on any time series being well approximated by a low-order mixed (ARMA) process. However, the omission of all information other than past values of the variable under consideration leads naturally to the question of whether the inclusion of other variables in the model may provide a more simple representation of its time-series evolution. This led to the introduction of one or more exogenous variables into the model, and the coining of the acronym ARMAX. However, since forecasting with such models requires forecasts of the exogenous variable(s), Tiao and Box (1981) suggested vector ARMA models, in which all variables are endogenous. For any vector of stationary variables $\mathbf{z}_{t+1}$, a vector ARMA model may be written as

$$(2.11)\quad \Phi(L)\mathbf{z}_{t+1}=\Theta(L)\varepsilon_{t+1}$$

where $\Phi(L)$ and $\Theta(L)$ are matrix-valued polynomials in the lag operator and ε is a multivariate white noise process. If $\Theta(L)$ is invertable, (2.11) can be written as a multivariate AR process:

$$(2.12)\quad \Theta^{-1}(L)\Phi(L)\mathbf{z}_{t+1}\equiv\Gamma(L)\mathbf{z}_{t+1}=\varepsilon_{t+1}$$

With the normalisation $\Gamma(1)=1$, we can rewrite (2.12) as

$$\mathbf{z}_{t+1}=-\Gamma^{*}(L)\mathbf{z}_{t}+\varepsilon_{t+1}$$

or

$$\mathbf{z}_{t+1}=\mathbf{A}\mathbf{z}_{t}+\varepsilon_{t+1}$$

which is the VAR model (2.3).

2.2.4 Econometric Background

Following Zellner and Palm (1974), any general dynamic linear simultaneous econometric model can be written as a VAR with suitable zero-restrictions on the elements of the A matrix to capture any exogeneity. More generally, the VAR model can be thought of as a linear approximation to the reduced form of any non-linear structural model (Holden, 1994). However, Sims (1980) eschewed, amongst other things, the imposition of zero restrictions and the arbitrary choice of dynamic specification typically found in econometric models. His view was that the so-called "structure" derived from economic theory and intuition served only to limit and distort the empirical analysis of economic time series. *Instead, he proposed an atheoretical* approach to modelling, in which economic paradigms are strapped firmly into the back seat, and claimed that this can produce more reliable empirical results. To this end, all the econometrician needs is a list of variables believed to summarise the relevant economic behaviour, and some notion as to what form these variables should take in the VAR. The data are allowed free rein to determine the appropriate VAR lag length.

Although in Sims's methodology the VAR is viewed as an unrestricted reduced form of an underlying structural model, there appears to be no difference between a Sims VAR and a VAR produced by a pure time-series analyst, once a particular set of variables has been alighted upon. However, it was the further claims of Sims regarding the ability of VAR modelling to supersede standard theory-based modelling techniques, in the areas of hypothesis testing and policy analysis, that caused the most controversy. Whereas VARs had been used previously with restrictions on the parameters of the **A** matrix according to some underlying structural model, Sims advocated free estimation of **A**, with restrictions being place instead on the innovations covariance matrix. Sims claimed that policy analysis and hypothesis testing could be conducted with sole reference to the unrestricted VAR parameters, and, moreover, that such a procedure was likely to give more accurate inference than traditional methods. However, as thoroughly documented by Cooley and LeRoy (1985), if the parameters of an underlying structural model ("deep parameters") were mapped into the "shallow parameters" of a reduced-form VAR system, meaningful hypothesis testing and policy analysis could be performed only if one were willing to impose restrictions on some of the VAR parameters; and these restrictions were precisely those which Sims was unwilling to entertain on prior grounds. For example, Sims argued that VAR models could be used to analyse the effects of alternative monetary or fiscal policies. If a test of Granger causality[2] indicates that the policy instrument is exogenous (i.e. Granger causality is rejected), one can view VAR model forecasts conditional on different hypothesised values of the instrument as capturing

---

2 A variable x_t Granger causes y_t if the lags of x_t in the VAR equation for y_t are jointly significant.

the effect of alternative instrument settings on the endogenous variables. However, consider the following money-income model:

$$(2.13)\ m_t = \theta y_t + \beta_{11} m_{t-1} + \beta_{12} y_{t-1} + \varepsilon_{1t}$$

$$(2.14)\ y_t = \gamma m_t + \beta_{21} m_{t-1} + \beta_{22} y_{t-1} + \varepsilon_{2t}$$

where m_t and y_t are real money and real income respectively, and the errors are contemporaneously and serially uncorrelated. Suppose we wish to take money as the policy instrument, and use the model to assess the effects of different levels of the money stock for income. Money is predetermined for income if θ=0. If (2.13) and (2.14) are solved for the reduced form, we have the following VAR(1) model:

$$(2.15)\ m_t = \pi_{11} m_{t-1} + \pi_{12} y_{t-1} + u_{1t}$$

$$(2.16)\ y_t = \pi_{21} m_{t-1} + \pi_{22} y_{t-1} + u_{2t}$$

Income fails to Granger cause money if π_{12}=0. In terms of the deep parameter of (2.13) and (2.14), π_{12} is given by

$$\pi_{12} = \frac{\theta\beta_{22} + \beta_{12}}{1 - \theta\gamma}$$

Clearly, Granger non-causality is neither necessary nor sufficient for money to be predetermined with respect to income, since θ=0 neither implies nor is implied by π_{12}=0. Since predeterminedness is the relevant form of exogeneity when analysing interventions, it can be seen that the VAR test of Granger causality is irrelevant to issues of policy analysis. For reasons such as this, VAR modelling has come to be viewed with suspicion in some areas of macroeconomics.

2.2.5 The Campbell-Shiller VAR Methodology

In a series of popular papers, Campbell and Shiller(1987, 1988, 1989) utilised VAR forecasting methods to test the implications of economic theory for the behaviour of a number of important financial time series. In this case, the VAR is viewed as a summary of the process through which financial market investors form their expectations. Friedman (1979) decomposed Muth (1961) Rational Expectations into two components:

i) the assumption that economic agents use information in an optimal manner;

ii) a specification of the precise information set with which agents are endowed.

Whilst assumption i) is behavioural, assumption ii) depends on information availability. Strong-form rational expectations assumes that agents know the true model underlying the economy. However, Friedman argued that weak-form rational expectations, in which agents use an estimate of the true model, is a more plausible assumption. In this framework, whether or not the model used by agents is accurately specified, account is taken of the fact that in finite data sets the estimated parameters will not converge fully to the true parameters of the underlying model.

Given a particular set of information, and an absence of a formal structural model, the most information-efficient linear forecasting scheme is provided by allowing all variables to be endogenous, so that the past behaviour of each variable is permitted to affect expectations of every other variable in an unrestricted way. In so far as the model parameters are estimated in a consistent way, minimising within-sample forecast errors, the forecasts are rational forecasts. However, forecasts are "weakly" rational, since no underlying structural model is assumed to be known - it is only assumed that the information at hand is used in an optimal manner. Furthermore, although in Chapters 5 and 6, restrictions on the parameters of the VAR are derived which force expectations to behave in a way which is consistent with the RVF, this is not the same as restrictions used in structural VAR modelling. The latter are used to force the VAR to conform with the deep parameters of the structural model, whereas there are no such deep parameters in my analysis.

2.3 The Generalised Method of Moments

2.3.1 Introduction

Although the parameters in all of the regression models under consideration in this thesis can be consistently and efficiently estimated using OLS, for my purposes Generalised Method of Moments (GMM) estimation is more appropriate. The reason is that the statistics of interest in each of the following four chapters are, in general, non-linear combinations of the regression model parameters. The question then arises of the best way of calculating standard errors for such statistics. My approach is to treat the estimation of the full set of model parameters, say Θ, as the indirect objective of GMM estimation. In this case, any non-linear function of the model parameters, $f(\Theta)$, has a standard error given by $\sqrt{f'(\Theta)Vf'(\Theta)'}$ where V is the regression model parameter covariance matrix.

Method of moments (MM) estimation relies on the fact that statistical models often imply "moment conditions", i.e. restrictions on expectations of certain functions of the variable(s) in question. The theoretical moment condition is replaced by a sample counterpart, and the MM estimator is the solution to the resulting equation(s). A trivial example is the MM estimator of the population mean μ of a random variable x_t. The relevant moment condition is simply

$$E(x_t)=\mu$$

or

$$E(x_t-\mu)=0$$

The sample analogue to this expectation is

$$\frac{1}{T}\sum_{t=1}^{T}(x_t-\hat{\mu})=0$$

$\hat{\mu}$ is the MM estimator of the population mean of x_t.

A more relevant example is that of a classical linear regression. For the simple model

$$y_t=\alpha x_t+\beta z_t+\varepsilon_t$$

One of the classical assumptions is that the independent variables are orthogonal to the error process:

$E(x_t\varepsilon_t)=E(z_t\varepsilon_t)=0$

The sample counterparts to these conditions are

$$\frac{1}{T}\sum_{t=1}^{T} x_t\left(y_t - \hat{\alpha}x_t - \hat{\beta}z_t\right) = 0$$

$$\frac{1}{T}\sum_{t=1}^{T} z_t\left(y_t - \hat{\alpha}x_t - \hat{\beta}z_t\right) = 0$$

The MM estimators of α and β satisfy these moment equations, and it is clear in this case that the resulting estimators are the OLS estimators of α and β. In fact, many estimators in common use can be constructed as method of moments estimators.

In the next section I outline the Generalised Method of Moments estimator, and in Section 2.3.3 I apply GMM techniques to econometric models. Finally, in Section 2.3.4 I derive the GMM estimator for the particular models used in the subsequent analyses.

2.3.2 The Generalised Method of Moments

The simple examples considered above had one thing in common: there were exactly as many moment conditions as parameters to be estimated. That is, each parameter was *exactly identified.* However, in general there will be many more moment conditions than parameters, giving rise to an *over-identified* parameter set. One might, for example, envisage a regression model in which the single independent variable is assumed to be correlated with the error term. In such circumstances, it is common to invoke an instrumental variables (IV) estimator. However, rarely, if ever, will there be a unique choice of instrument, and the more sample information used in the estimation process, in general the more efficient the estimator will be.

Suppose that the econometrician is armed with J>1 potential instrumental variables. One approach to estimation might be to take the orthogonality between each of the J possible instruments and the regression error as a set of moment conditions to be jointly satisfied. However, with only one parameter to estimate, and J moment conditions, there is no unique solution to the estimation problem. GMM estimation involves minimising a quadratic form based on the J moment conditions, giving rise to an estimator which is efficient amongst all estimators defined by the moment conditions.

To illustrate, suppose that we have a model which has a K-vector parameter set, Θ, to be estimated. Suppose also that the model implies a set of J>K moment conditions,

$$E[m_{jt}(\Theta)]=0 \qquad j=1,\dots,J$$

with J sample counterparts

$$\frac{1}{T}\sum_{t=1}^{T} m_{jt}(\hat{\Theta}) = 0$$

Method of moments estimation would require solving the J equations above for K unknowns, to which of course there is no unique solution. The GMM solution is very similar to the OLS procedure. In the latter, since the model parameters cannot be chosen to make *every individual* observed error equal to zero, the sum of squared residuals is minimised, with the result that the *mean* error is zero. In a similar vein, one way of reconciling the J moment conditions with K parameters would be to minimise the sum of squared moment conditions:

$$q_t=\mathbf{m}_t(\Theta)'\mathbf{m}_t(\Theta)$$

The resulting value for the K-vector Θ would be such that the moment conditions would be zero *on average*.

However, in the same way as GLS improves upon OLS by weighting the errors in proportions that are inversely related to their variances, if we define **W** as the moment covariance matrix (Hansen 1982), the GMM objective function is given by

$$q_t = \mathbf{m}_t(\Theta)'\mathbf{W}^{-1}\mathbf{m}_t(\Theta)$$

The GMM estimator of Θ derives from the minimisation of the quadratic form q_t.

The asymptotic covariance matrix of the GMM estimator is given by

$$\Sigma = (\mathbf{G}'\mathbf{W}^{-1}\mathbf{G})^{-1}$$

where **G** is a matrix of derivatives of the moment conditions with respect to the parameters, with a typical row,

$$\mathbf{G}_i = \frac{\delta m_i}{\delta\beta'}$$

2.3.3 GMM Estimation of Econometric Models

Suppose that economic theory, or a time-series model, suggests the following relationship:

(2.13) $y_t = \mathbf{x}_t\beta' + \varepsilon_t \qquad E(\varepsilon_t)=0,\ E(\varepsilon_t\varepsilon_t')=\Omega$

where β is a Kx1 vector of parameters that we wish to estimate. Although (2.13) is specified as a linear relationship, GMM is perfectly capable of handling non-linear functional forms. Initially we place no restrictions on the form of Ω (so the disturbances may be heteroscedastic and/or autocorrelated), and we do not assume that $\mathbf{x}_t$ and ε_t are orthogonal. However, we do assume that the econometrician has available a set of J variables $\mathbf{z}_t$ that can be used as valid instruments for $\mathbf{x}_t$, i.e. that the variables in $\mathbf{z}_t$ are orthogonal to ε_t. GMM estimation of the parameter vector β with therefore proceed on the basis of the J orthogonality/moment conditions

(2.14) $E(\mathbf{z}_t\varepsilon_t)=0$

provided that J is at least as great as K. The sample counterpart to the orthogonality conditions (2.14) is found by replacing ε_t with the regression residuals e_t, and summing over the data sample:

$$\mathbf{m}_t(\beta) = \frac{1}{T}\sum_{t=1}^{T} \mathbf{z}_t e_t = \frac{1}{T}\mathbf{Z}'\mathbf{e}$$

The objective function for GMM estimation is then

$$q_t = \mathbf{m}_t(\hat{\beta})'\mathbf{W}^{-1}\mathbf{m}_t(\hat{\beta}) = \frac{1}{T^2}(e_t'\mathbf{z}_t)\mathbf{W}^{-1}(\mathbf{z}_t'e_t)$$

for some **W**. Minimising this function is a non-linear optimisation problem.

Hansen (1982) showed that the optimal choice of weighting matrix, **W**, is the asymptotic covariance matrix of the moment conditions. For our model, this is given by

$$W = \text{Asym.Var}\left(\frac{1}{T}\mathbf{Z}'\mathbf{e}\right)$$

$$= \frac{1}{T^2}\sum_t \sum_\tau \text{Cov}(\mathbf{z}_t e_t, \mathbf{z}_\tau e_\tau)$$

$$= \frac{1}{T}\mathbf{Z}'\Omega\mathbf{Z}$$

The GMM objective function is therefore

$$\mathbf{q} = \left(\frac{1}{T}\mathbf{e}'\mathbf{Z}\right)\left(\frac{1}{T^2}\mathbf{Z}'\Omega\mathbf{Z}\right)^{-1}\left(\frac{1}{T}\mathbf{Z}'\mathbf{e}\right)$$

Estimation usually proceeds in two stages. First a consistent but inefficient estimate of β is obtained (commonly be assuming that **W=I**) for use in the construction of **W**. The latter is then used in the GMM objective function to provide efficient second-stage estimates.

For the estimation of **W** to be feasible, some restrictions must be imposed on its form. Newey and West (1987) have developed an estimator for the case of autocorrelated but homoscedastic disturbances. However, I shall in all cases be concerned only with allowing for heteroscedastic errors, as my model specification will rule out residual autocorrelation. Consequently, the cross products may be omitted from the covariance matrix, so that what is required is an estimate of

$$\mathbf{W}_{HET} = \frac{1}{T^2}\sum_t \mathbf{z}_t\mathbf{z}_t' Var(e_t(\beta))$$

White (1984) shows that this matrix may be estimated using

$$\frac{1}{T^2}\sum_t \mathbf{z}_t\mathbf{z}_t' e_t(\hat{\beta})^2$$

2.3.4 An Application of the GMM Estimatior

All of the models to be estimated in this thesis share 4 common features:

i) It is assumed that the RHS variables are independent of the error term. Thus the original set of variables can be used as their own instruments, and each parameter is consequently *exactly identified.*
ii) Contemporaneous correlation between the error terms in each equation is not ruled out.
iii) Each set of equations has identical RHS variables.
iv) Heteroscedasticity of an unknown form is allowed for in the error covariance matrix, but residual serial correlation is not.

I now examine the impact of these factors on the GMM estimator.

In general, one can write a system of linear regression equations as follows:

$$\mathbf{y} = \begin{bmatrix} \mathbf{X}_1 & & 0 \\ & \mathbf{X}_2 & \\ & . & \\ 0 & & \mathbf{X}_m \end{bmatrix} \begin{bmatrix} \beta_1 \\ \beta_2 \\ . \\ \beta_m \end{bmatrix} + \begin{bmatrix} \varepsilon_1 \\ \varepsilon_2 \\ . \\ \varepsilon_m \end{bmatrix}$$

or

(2.15) $\mathbf{y}=\mathbf{X}\beta+\varepsilon$ $\qquad E(\varepsilon)=0,\ E(\varepsilon\varepsilon')=\Omega$

Since **Z**=**X** (assumption (i)), the GMM objective function is

$$q=(\varepsilon'\mathbf{X})(\mathbf{X}'\mathbf{W}\mathbf{X})^{-1}(\mathbf{X}'\varepsilon)=[(\mathbf{y}-\mathbf{X}\beta)'\mathbf{X}](\mathbf{X}'\mathbf{W}\mathbf{X})^{-1}[\mathbf{X}'(\mathbf{y}-\mathbf{X}\beta)]$$

The objective is to choose β to minimize (2.16). In fact, in the exactly identified case there is a unique solution to this problem. Differentiating (2.16) with respect to β and setting this equal to zero we have

$$-\mathbf{X}'\mathbf{X}(\mathbf{X}'\mathbf{W}\mathbf{X})^{-1}\mathbf{X}'\mathbf{X}\mathbf{y}+\mathbf{X}'\mathbf{X}(\mathbf{X}'\mathbf{W}\mathbf{X})^{-1}\mathbf{X}'\mathbf{X}\hat{\beta}=0$$

(2.16) $\hat{\beta}=(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{y}$

However, since the regressions within all of the systems estimated in this thesis contain the same RHS variables, this can be simplified futher. If $\mathbf{X}_1=\mathbf{X}_2=\ldots=\mathbf{X}$, (2.15) becomes

$$\mathbf{y}=(\mathbf{I}\otimes\mathbf{X})\beta+\varepsilon$$

and (2.16) simplifies to

(2.17) $\hat{\beta}=(\mathbf{I}\otimes(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}')\mathbf{y}$

(2.17) is the GMM estimator of β employed throughout this thesis, and it is immediately apparent that this is identical to the application of the OLS estimator to each individual equation. However, since GMM is a systems estimation procedure, the parameter covariance matrix does not force cross-equation parameter covariances to be zero, and this is the main motivation behind its use.

2.3 Phillips-Perron Unit Root Tests

It is assumed throughout this thesis that the VAR processes used in forecasting are stationary. As an indication of the reliability of this assumption, all series are tested for stationarity using the testing procedure developed by Phillips (1987), Perron (1988) and Phillips and Perron (1988). A series is integrated (and therefore non-stationary) if it has at least one unit root in its AR representation. A sufficient condition for a series to have one unit root is that the AR coefficients sum to unity. This is the basis of the well-known Dickey-Fuller (DF) test. If a series x_t is well represented by a first-order autoregression[3], i.e.

$$(2.18)\quad x_t = \mu + \phi x_{t-1} + \varepsilon_t$$

then the null hypothesis that x_t is integrated translates into a test of the hypothesis that ϕ=1. Fuller (1976) presents critical values for two tests of this hypothesis: the standard t-statistic; and the statistic T(ϕ-1), where T is the sample size. The augmented Dickey-Fuller (ADF) test (Said and Dickey, 1984) is applied to cases where an autoregression of higher order than one is required adequately to model x_t, although the test statistics and critical values are the same as for the DF test.

The Phillips-Perron (PP) tests provide an alternative to the ADFstatistics when the error process in the DF regression is not white noise. Phillips and Perron (1988) present non-parametric corrections of the DF statistics which account for serially correlated (and heterogeneously distributed) residuals, allowing one to apply the DF critical values. Letting u_t denote the DF regression error, the PP statistics which corresponds to the DF t-statistic is

$$Z(t_{\hat{\phi}}) = \left(\frac{S_u}{S_{Tl}}\right)t_{\hat{\phi}} - \frac{1}{2}\left(S_{Tl}^2 - S_{\hat{u}}^2\right)\left[S_{Tl}\left(T^{-2}\sum_{t=2}^{T}(y_{t-1} - \bar{y}_{-1})^2\right)^{\frac{1}{2}}\right]^{-1}$$

whilst that for the T(ϕ-1) statistic is

$$Z\left(\hat{\phi}\right) = T\left(\hat{\phi} - 1\right) - \frac{1}{2}\left(S_{Tl}^2 - S_{\hat{u}}^2\right)\left[T^{-2}\sum_{t=2}^{T}(y_{t-1} - \bar{y}_{-1})^2\right]^{-1}$$

where $S_{\hat{u}}^2 = T^{-1}\sum_{t=1}^{T}\hat{u}_t^2$ and $S_{Tl}^2 = T^{-1}\sum_{t=1}^{T}\hat{u}_t^2 + 2T^{-1}\sum_{j=1}^{l}\sum_{t=j+1}^{T}\hat{u}_t\hat{u}_{t-j}$. Clearly, if u_t is iid, S_u^2=S_{Tl}^2 and the PP and DF statistics are identical. Since little is known of the relative

[3] "Well represented" means that ε_t in (2.18) is white noise.

power of these two PP statistics, I report both so that one might be better able to judge the appropriateness of the stationarity assumption.

Chapter 3: EXPLAINING MOVEMENTS IN UK STOCK PRICES

3.1 Introduction

According to the rational valuation formula (RVF), stock prices change because of time-variation in expected dividends and expected returns (discount rates). A large number of empirical studies have found that stock returns are, in part, predictable, and it is common to ascribe this to systematic movement in expected returns. However, the low degree of predictability typically found in return regressions means that a large proportion of stock price movements remains unexplained. Unexpected movements in stock prices must be due to revisions in investors' expectations about future dividends and revisions to expected returns. In this chapter I seek to determine which of these factors has, historically, been the more important.

The literature on stock return predictability is vast, but can be broadly split into two categories. First, univariate studies are often based on some measure of the autocorrelations of returns over different horizons, and some consistent patterns have been uncovered[1]. However, in such studies any variability in ex-ante returns that is linked to information other than past returns will not be captured. Second, multivariate *single-equation* regressions have confirmed the importance of variables such as the dividend-price ratio, the yield spread, a default spread, the gilt-equity yield ratio and some measure of volatility as having incremental explanatory power[2]. These variables are often interpreted (somewhat loosely) as determinants of the *expected* return.

Other studies have attempted to model *unexpected* returns as innovations in those variables assumed to determine expected returns. For example, Fama (1990) took the residuals from AR(1) models fitted to a term premium and a default premium as proxies for shocks to returns. Cutler et al (1989) performed a similar analysis but using a multivariate rather than a univariate forecasting model to obtain innovations in a number of financial and macroeconomic variables, which were then used to explain stock returns. However, in these studies, the measured news variables are ad hoc

[1] For example, Fama and French (1988a), Poterba and Summers (1988), Cutler et al (1991) have found that stock returns typically display positive autocorrelation at short horizons and negative autocorrelations over long horizons.

[2] See, *inter alia* Clare, Thomas and Wickens (1994), Shah and Wadhwani (1993), Fama and French (1988a, 1988b), Campbell (1987).

and may be due to changes in dividend expectations or changes in future discount rates; there is no attempt to discriminate between these two sources.

Campbell (1991) noted that the importance of revisions to discount rates in explaining stock price movements depends not only on the degree of return predictability, but also on the time-series properties of expected returns. In particular, even if return predictability is low, nevertheless news about expected returns can have a powerful effect on stock prices provided that expected returns are "persistent" (i.e. a shock to the current expected return has a near-permanent effect on expectations of all future returns).

Following Campbell (1991) and Campbell and Ammer (1993), I avoid the problems associated with univariate and single-equation studies by using a multivariate (VAR) model to forecast expected returns. In contrast to many earlier regression-based tests which examine the impact of news on stock returns, I impose the theoretical structure of the RVF, so that any news variable must influence unexpected returns either by influencing expectations of dividends, or returns, or both. I also examine in tandem the importance of "predictability" and "persistence" in forecasts of dividends and future discount rates on unexpected movements in stock prices.

A complementary strand to the literature outlined above has attempted to introduce time-varying conditional variances as determinants of one-period expected returns (e.g. Poterba and Summers 1986, Hall, Miles and Taylor 1989, French, Schwert and Stambaugh 1987, Chou 1988), as suggested by the intertemporal CAPM (Merton 1973, 1980). Poterba and Summers (1986) assumed that the discount rate in the RVF depended on a measure of risk, which they took to be the volatility of stock prices. A key result was that if "news" about volatility has a protracted effect on expected future volatility (i.e. expected volatility is persistent) then this can have an important effect on stock prices. In their model, Poterba and Summers assumed that dividend growth was constant and "volatility" was modelled as a univariate ARIMA process for the return variance. They found that incorporating this risk premium into the RVF did not provide a good explanation of the movement in stock prices. French et al (1987) measured monthly volatility as the average squared *daily* return within the month (corrected for the AR(1) element of daily returns). They then estimated expected volatility as predictions from an ARIMA(0,1,3) model and used the residuals as a measure of unexpected volatility. The one-period excess return was regressed on expected and unexpected volatility and only the latter was found to be statistically significant, with a negative impact on the one-period return. They interpreted the negative effect of unexpected volatility as being consistent with a positive effect of *expected* volatility on returns. This interpretation is possible because the persistence in volatility implies that an unexpected increase in volatility causes all future values of expected volatility to increase, hence the discount rate in the RVF increases and prices fall immediately, so that the current one-period return is negative. The French et al (1987) interpretation can be directly examined in the current framework since I explicitly incorporate the impact of forecasts of volatility on future returns (discount rates) and hence on unexpected changes in current returns. This analysis therefore extends the above studies by including a time-varying measure of risk (i.e. conditional variance) into the determination of stock returns, but in a multivariate (VAR) framework. Chou (1988) and French et al examined an ARCH-M model for expected returns, but such complexities are beyond the scope of the multivariate VAR framework. In any case, as noted by French et al (1987), non-stochastic ARCH models do not allow for unanticipated volatility. In addition, the adopted framework does not constrain dividends to grow at a constant rate, and so I do not force all of the variability in stock prices to be "explained" by the variability in future discount rates.

To summarise. This study uses a multivariate VAR framework to predict UK stock returns. The multiperiod predictions from the VAR are used in the RVF to apportion *unexpected* stock returns into news about dividends, news about future returns (discount rates) and the covariance between the two. I assess the sensitivity of my results to alternative explanatory variables for the returns equation, alternative lag lengths in the VAR and for alternative sample periods, using a long UK data set on an aggregate stock market index from 1918 to 1993.

The rest of this Chapter is organised as follows. In Section 3.2 I review the literature on stock return predictability. I discuss results for both univariate and single-equation multivariate models. In Section 3.3 I introduce the Campbell-Shiller log-linear approximation to the rational valuation formula. This facilitates the analysis of stock price movements without assuming either that dividend growth or expected returns are constant. Section 3.4 shows how the approximate RVF can be manipulated to obtain a variance decomposition of unexpected stock returns. In Section 3.5 the data are discussed and some sample statistics presented. The empirical results from the first stage of the decomposition are presented in Section 3.6.

The notions of volatility and volatility persistence are introduced in Section 3.7, beginning with a review of the volatility literature followed by a discussion of the effects of volatility persistence on the variance decomposition. Whilst the bulk of the analysis is couched in terms of real stock returns, Section 3.8 presents results for the variance decomposition of unexpected *excess* returns. Section 3.9 considers briefly the possibility that return expectations are well modelled by a first-order autoregressive process, and 3.10 discusses the implications of the findings. In Section 3.11, the robustness of the results to changes in the VAR lag length and the data period is considered, whilst in Section 3.12 alternative measures of cash flow news are compared. Results from Campbell's variance decomposition using US stock market data are discussed in Section 3.13. Section 3.14 concludes.

3.2 Literature on Return Predictability

3.2.1 Univariate Studies

Early studies (for example Fama 1965) focused on autocorrelations of returns of individual stocks. Although some significant findings were reported, the general conclusion was that since predictable variation in daily and weekly returns appeared to be such a small portion of the overall return variance, such time-variation was not economically important. However, Fisher's (1966) results suggested that the variance reduction obtained in portfolio diversification would lead to statistically stronger findings for portfolio returns than for individual stocks. Also, Shiller (1984) and Summers (1986) argued that studies of such short return horizons might miss degrees of predictability that were in fact economically important.

Fama and French (1988a) provided a useful framework within which the relationship between expected returns and return autocorrelations may be analysed.

They began by positing the natural log of a stock price at time t, p_t, as the sum of a random walk, q_t, and a stationary component, u_t:

$$p_t=q_t+u_t \tag{3.1}$$

$$q_t=q_{t-1}+\mu+\varepsilon_t \tag{3.2}$$

where μ is expected drift, and ε_t is white noise. In addition, u_t was specified as following an AR(1) process:

$$u_t=\phi u_{t-1}+\eta_t \qquad 0<\phi<1 \tag{3.3}$$

where η_t is white noise.

In the model (3.1)-(3.3), shocks to stock prices come from two sources. A positive innovation to q_t through ε_t has a permanent effect on the level of stock prices. However, a positive innovation to u_t through η_t is gradually eliminated from prices as time passes. It is the transitory component of stock prices induced by the latter effect that implies predictability (negative autocorrelation) of returns.

The continuously compounded return from time t to t+T is given by

$$\begin{aligned} r_{t,T}&=p_{t+T}-p_t \\ &=[q_{t+T}-q_t]+[u_{t+T}-u_t] \end{aligned} \tag{3.4}$$

The first term on the RHS of (3.4) is the unpredictable component of returns generated by the random walk component in stock prices. Fama and French showed that the mean reversion of the stationary price component leads to negative autocorrelation in returns.

The first-order autocorrelation coefficient of T-period changes in u_t is

$$\rho(T)=\frac{Cov(u_{t+T}-u_t,\,u_t-u_{t-T})}{Var(u_{t+T}-u_t)} \tag{3.5}$$

The numerator covariance is

$$(3.6) \qquad Cov(u_{t+T}-u_t,\ u_t-u_{t-T})=-Var(u)+2Cov(u_t,u_{t+T})-Cov(u_t,\ u_{t+2T})$$

Since u_t is stationary, the covariances on the RHS of (3.6) approach zero as T increases, so that the whole expression tends to -Var(u). The variance in the denominator of (3.5) is

$$(3.7) \qquad Var(u_{t+T}-u_t)=2Var(u)-2Cov(u_{t+T},u_t)$$

which approaches 2Var(u). Given these results, the first-order autocorrelation coefficient of T-period changes in u_t tends to -0.5 as T increases.

Two more expressions are useful in interpreting the autocorrelation $\rho(T)$. First, if u_t is an AR(1) process, the expected change in u_t from t to T is

$$(3.8) \qquad E_t(u_{t+T}-u_t)=(\phi^T-1)u_t$$

and so we may rewrite expression (3.6) as

$$(3.9) \qquad Cov(u_{t+T}-u_t,\ u_t-u_{t-T})=-(1-\phi^T)Var(u)$$

Now suppose that ϕ is close to 1[3], and we are looking at a one-period return horizon (T=1). Expression (3.9) indicates that in this scenario, $\rho(T)$ will be close to zero. However, as we increase T, the autocorrelation coefficient slowly approaches -0.5. The implication is that if the transitory component is highly persistent, so that shocks decay only very slowly, it may not be detected by tests of autocorrelations over short return horizons. This point was first made by Summers (1986). Conversely, one might expect to find evidence of mean-reversion in stock prices only by studying autocorrelations of long-horizon returns. Fama and French therefore claimed that

---

[3] If $\phi=1$, u_t is a random walk. In this case we would be unable to distinguish between the shocks ε_t and η_t and stock returns would be unpredictable. ϕ close to one is a slight loosening of this condition which, as we shall see, can have significant effects.

returns are likely to be more highly autocorrelated (more predictable) over long horizons than short horizons.

u_t is unobservable; but its behaviour will be reflected in the behaviour of returns, which we do observe. However, the relationship between return autocorrelations and the transitory price component is not straightforward. The first-order autocorrelation coefficient for the return series is given by

$$\text{(3.10)} \qquad \theta(T) = \frac{Cov(r_{t,t+T}, r_{t,t-T})}{Var(r_{t-T,t})}$$

$$\text{(3.10a)} \qquad = \frac{\rho(T)Var(u_{t+T} - u_t)}{Var(u_{t+T} - u_t) + Var(q_{t+T} - q_t)}$$

Equation (3.10a) helps us to interpret the behaviour of return autocorrelations. If stock prices do not have a transitory component then $\theta(T)$ is zero for all T. On the other hand, if the random walk component is absent from prices then $\theta(T)=\rho(T)$, and these both approach -0.5 as T increases. However, the interesting result presented by Fama and French is that if the stock price has both a random walk and a mean-reverting component then *first-order return autocorrelations display a U-shaped pattern with respect to the return horizon.* Thus for small T, $\theta(T)$ takes on values close to zero, as T increases $\theta(T)$ becomes more negative, but beyond some point moves back towards zero.

To see this, notice that, from expression (3.10a), the stationary component of the stock price tends to push the return autocorrelation towards -0.5 as the return horizon increases. However, at the same time, the variance of the random walk innovation depresses $\theta(T)$ towards zero. Since $Var(u_{t+T}\text{-}u_t)$ approaches a constant 2Var(u) as T increases, whilst the variance of the white noise component increases linearly with the return horizon, the latter component will dominate at long return horizons.

Fama and French estimated return autocorrelations using New York Stock Exchange data from 1926 to 1985. They found that the predicted U-shaped pattern was observed for all size deciles, industry portfolios and two market portfolios. That is, estimated autocorrelations were around zero for short horizons (up to 1 year), became negative over intermediate horizons (3-5 years), and approached zero again after 10 years or so. Although there was no obvious pattern in the variation of estimated autocorrelation coefficients across industries, the lower deciles (smaller firms) tended to produce lower minimum autocorrelations than the higher deciles. Moreover, around 40% of the variation of 3-5 year returns on small-firm stocks was picked up by this model.

One important caveat was that the analysis of sub-periods suggested that the negative autocorrelation was weaker after 1940, implying a less important, or maybe absent, stationary price component after this period. However, return variances dropped substantially after 1940, making inference less precise. In fact, the statistical imprecision of these tests due to the small number of observations of long-horizon returns greatly reduced the force of any conclusions about the predictability of long-horizon returns.

Lo and MacKinlay (1988) also studied return autocorrelations using the NYSE data. The most striking conclusion from their analysis was that returns were significantly *positively* autocorrelated at weekly and monthly horizons, which contradicted the Fama-French model.

The findings of Lo and Mackinlay were supported by Poterba and Summers (1988), whose point estimates identified positive autocorrelation in return horizons shorter than one year, and negative autocorrelation thereafter. However, Poterba and Summers questioned the statistical power of the methods employed, and expressed doubts about the usefulness of this approach, as even the broad characteristics of the data could not be estimated precisely.

Finally, Cutler, Poterba and Summers (1991) studied a large number of asset markets in several countries. Although some inconsistencies existed, their broad conclusion was that return autocorrelations are positive at short horizons and negative over longer horizons.

A different approach was taken by Conrad and Kaul (1988). Rather than looking directly at return autocorrelations, they tracked time-variation in expected returns using a Kalman filter. Conrad and Kaul noted that the use of daily observations on portfolio returns may give spurious evidence of return autocorrelations. The argument is that for stocks which do not trade very often, news which occurs on day 1 may not be reflected in prices until, say, day 2, whereas for heavily-traded stocks such information will be reflected in prices almost instantaneously. The upshot is that following a piece of news on day 1, prices of heavily-traded stocks react on day 1, whilst those of occasionally-traded stocks react on day 2. Returns on portfolios that include both types of stocks will appear to be serially correlated, but this may be due wholly to this "non-synchronous trading" effect. Fisher (1966) emphasised that this effect may be more important for portfolios containing a significant portion of stocks with low market capitalisations.

Conrad and Kaul therefore used weekly returns of ten size-based portfolios over the period 1962-85, they concluded that expected returns were well-characterised by an AR(1) model[4]. Variation in expected returns was found to be a significant portion of the variance of ex post returns. In addition, there was a strong monotonic relation between the size rankings of the portfolios and the proportion of the returns variance accounted for by expected returns. For the smallest portfolio, variation in expected returns accounted for 26% of the return variance, whilst for the largest portfolio the proportion dropped to only 1%.

A number of conclusions can be drawn from univariate studies of return predictability. First, estimated return autocorrelations do follow a U-shaped pattern with respect to the return horizon, although autocorrelations are actually positive (rather than zero) at the shortest horizons. Second, because of the reduction in the return variance achieved through diversification, the predictable component of returns is likely to be a greater proportion of portfolio returns that for returns on individual stocks. Finally, evidence of return autocorrelation appears to be strongest for small-firm portfolios and weakest for large-firm portfolios.

All of these conclusions are clouded by two effects. First, non-synchronous trading may induce spurious autocorrelation in portfolio returns, and this is likely to be greater for small-firm portfolios. Second, the small number of observations available for long-horizon returns significantly reduces the statistical power of the techniques employed, and so the strength of any conclusions drawn.

---

[4] I test this proposition on UK data in Section 3.9.

3.2.2 Multivariate Studies

The focus here has been on discovering variables known in advance which have predictive power over stock returns (or more usually excess returns over and above some risk-free rate of interest). If risk premia are related to perceived economic conditions then variables which proxy for such expectations should forecast ex post returns.

Keim and Stambaugh (1986) asked whether variables which proxy for asset prices have predictive power over excess returns. Their argument was as follows. The RVF implies that changes in asset prices reflect changes in expected cash flows (dividends) and changes in discount rates (equilibrium expected returns). To the extent that any changes are due to movements in discount rates, prices should have predictive power over equilibrium returns. They constructed three variables to model asset prices:

i) the difference between yields on long-term low-grade (under BAA-rated) corporate bonds and one-month US Treasury Bills, which reflects the market's perception of bankruptcy risk through the level of low-grade bond prices;

ii) minus the log of the ratio of the real S&P Composite Index to its previous long-run level, which serves as an ex ante detrended price series;

iii) a small-firm price variable. Beginning with the proposition that small-firms' stocks' risk premiums are the most volatile, Keim and Stambaugh argued that small-firm prices may provide a sensitive measure of expected risk.

Regressing monthly excess returns on these three variables over the period 1928-1978, they found a significant role for all of them when entered separately. (Collinearity was such that when entered jointly no single variable had a significant effect.) However, the practical interest in the results was somewhat diminished by the fact that the regression R^2's rarely climbed above 1%. The conclusion was therefore that there was some systematic element to stock returns, but here it was identified as an extremely small portion.

Campbell (1987) also looked at excess returns, relating them to risk premia on bills ("maturity premia", which catch the effect of an uncertain end-of-period price for a bill with more than one period to maturity). He regressed excess stock returns on the spread between the two-month and one-month rate, and the spread between the six-month and one-month rate. He also included the one-month bill rate to test whether the coefficient on this variable was significantly different from unity (the construction of excess returns imposes a unit coefficient). The sample was split between 1959(5)-1979(8) and 1979(9)-1983(11) to account for a change in operating procedures of the Ferderal Reserve. Over the first sample period, all but the six-month over one-month bill rate had predictive power over stocks, with an R^2 of 11.2%. In the second sample period, only the one-month bill rate was significant, but R^2 rose to 22.8%. No great correlation between this variable and the others was reported, and so Campbell concluded that a high one-month bill rate predicts a low stock return. This was in line with Fama (1981) who rationalised this as being due to a negative influence of inflation proxied by the bill rate (see Chapter 4, especially Section 4.10).

Fama and French (1988b) reported that dividend yields forecast returns on the value- and equal-weighted NYSE index for return horizons ranging from one month to four years. Explanatory power increased with the return horizon, with the dividend yield accounting for 25% of the variance of two- to four-year returns. As in their 1988a paper, Fama and French argued that if expected returns are highly autocorrelated, a large proportion of the variance of long-horizon returns will constitute variation in expected returns (the variance of expected returns rises faster than the variance of unexpected returns as the returns horizon increases). Thus one should expect predictability to increase with the return horizon. They also argued (as did Keim and Stambaugh) that shocks to returns will be reflected in price changes, and this was captured by the dividend yield.

In their 1989 paper, Fama and French considered variables reflecting a maturity premium and a default premium on stocks, in addition to the dividend yield. The former was captured by the difference between the yields on a AAA bond portfolio and the one-month bill rate. The latter was modelled as the difference in yields between a portfolio of corporate bonds and the AAA portfolio.

The majority of parameter estimates were positive. Having compared the behaviour of the three forecasting variables to long- and short-term business cycles, they concluded that the dividend yield and the default spread tracked that portion of expected returns which is high during long downward swings in business conditions (such as the Great Depression) and high when business is persistently strong. The maturity premium captured that component of expected returns that varies with shorter-term business cycles. The general conclusion was that expected returns vary counter-cyclically.

In line with the conclusions of Fama and French, Pesaran and Timmerman (1992) found that return predictability increased with the return horizon. Their regression R^2's were 9%, 21.6% and 63.4% for monthly, quarterly and annual excess returns respectively (although they had only 36 observations for the annual regression, posing the same problems with statistical power as highlighted by Fama and French). The variables employed comprised various lags and transformations of those mentioned already plus the rate of inflation and the change in industrial production. Pesaran and Timmerman noted the weak theoretical rationale for the inclusion of some of these variables; their objective was to show that there existed some information, known ex ante, which helped to predict returns.

Campbell and Shiller (1988) presented evidence that a long moving average of accounting earnings had predictive power over returns through forecasting dividend growth. They argued that, since earnings are constructed by accountants with the objective of providing investors with an idea of the value of a company, and that the latter is directly related to the firm's expected future income, earnings data might properly hold a place in the econometrician's bank of relevant information. This holds regardless of the fact that earnings data are often viewed as inconsistent series because of continual changes in accounting definitions: as long as they provide a reasonable proxy for expected future income, one would expect them to be useful as predictors of stock returns.

Using annual observation on the real Standard and Poor Composite Index from 1871 to 1987, they regressed log returns on the log dividend-price ratio, real dividend growth and a thirty-year moving average of earnings. For real returns, the regression R^2's were 7.6% for 1-year returns, 20.4% for 3-year returns and 63.7% for a 10-year return horizon. The R^2's for excess returns were 8.6%, 19.5% and 49.3% respectively.

Clare, Thomas and Wickens (1994) found a significant explanatory role for the ratio of a long government bond yield to the equity market dividend yield (the gilt-equity yield ratio, geyr). They began by noting the practice amongst investors of setting buy-sell thresholds for equities on the assumption that the geyr has a long-run level reflecting an arbitrage relation between bonds and shares. When the geyr is above this level, equities are thought to be overpriced relative to bonds, and without a fall in bond yields, equity yields are bound to rise. The implication is usually taken to be that equity prices will fall to restore "equilibrium". Conversely, a lower-than-normal geyr is taken to suggest an imminent rise in equity prices.

They suggested that the use of the geyr by UK analysts reflects a preoccupation with income flows brought about by the preponderance of pension fund managers in the UK market. Short-term price fluctuations may be given less weight by such investors since they are primarily concerned with regular income flows to meet their liabilities[5].

Clare et. al. tested the predictive power of the geyr using quarterly observations of the FT All-Share Index from 1968(1) to 1992(2). Their "simple geyr model", in which real returns were regressed only on the lagged level and changes in the geyr, produced an R^2 of 51.3%. Combining the geyr with the term spread between consols and Treasury Bills, the reverse yield gap (the spread between the consol and dividend yields), the lagged real return and the 3-month TB rate raised the return regression R^2 to 63.27%.

Although the above does not by any means constitute an exhaustive catalogue of the literature, there is sufficient information to draw two broad conclusions on the current state of knowledge vis a vis the predictability of stock returns. First, there appears to be a substantial improvement in explanatory power when variables other than lagged returns are included. In particular, the dividend yield appears extremely powerful for US data, whilst the geyr performs well for the UK. Second, the claim by Fama and French (1988a) that explanatory power increases with the return horizon is unanimously supported in both univariate and multivariate analyses.

However, Shah and Wadhwani (1990) provided a note of caution. They studied the predictive power of the term spread and the dividend yield on excess

---

[5] Since, in a rational market, the prediction of capital gains requires no more information than the prediction of future dividends, the logic behind this statement is unclear.

returns from fifteen countries on monthly data over the period 1970(5) to 1988(10). Although their results for the US were consistent with previous work, in no other country was the forecasting ability of these two variables confirmed. Shah and Wadhwani noted that one interpretation which has been given to the ability of the term spread to forecast US stock returns was that the term spread helps predict future real activity. However, whilst in other countries the links between the term premium and future real activity, and between future real activity and stock returns, are apparent, they devised no systematic link between the term premium and excess stock returns. The striking lack of uniformity in their results moved them to interject a note of caution into attempts to develop an all-encompassing model of expected return determination on the basis of previous findings.

Fama (1991) also emphasised the precarious balance to be struck between the discovery of reliable empirical proxies for return expectations and the implications of mass data mining: "With many clever researchers ... rummaging for forecasting variables, we are sure to find instances of 'reliable' return predictability that are in fact spurious" (Fama 1991, p1585).

3.3 *The Campbell-Shiller Dividend-Ratio Model*

With non-constant expected dividend growth, the RVF is econometrically tractable in a linear setting only if discount rates are constant, and I do not wish to make this assumption. Campbell and Shiller (1988, 1989) introduced a log-linearisation of the rational valuation formula which has subsequently formed the basis of a sizeable body of research into stock price movements. This approximation is used throughout this thesis, and so I now discuss its derivation and properties.

The ex post one-period real holding-period return on a stock is given by

$$(3.11) \qquad H_{t+1} \equiv \frac{P_{t+1} + D_{t+1}}{P_t}$$

where P_t is the real stock price at the end of period t and D_{t+1} is the real dividend paid during period t+1[6]. (3.11) can be rewritten in terms of dividend-price ratios:

[6] In the empirical analyses undertaken below I have end-of-period observations of stock prices, whereas Campbell and Shiller used opening prices. Consequently my time scripts differ from those in the original papers.

$$H_{t+1} = \frac{D_{t+1}}{D_t}\left(\frac{D_t}{D_{t+1}}\frac{P_{t+1}}{P_t} + \frac{D_t}{P_t}\right)$$

(3.12) $$H_{t+1} = \frac{D_{t+1}}{D_t}\left(\frac{D_t/P_t}{D_{t+1}/P_{t+1}} + \frac{D_t}{P_t}\right)$$

Using lower-case characters to denote the natural logarithm of the corresponding upper-case variables, and letting δ_t denote the log dividend-price ratio (d_t-p_t), taking logs of (3.12) gives

(3.13) $$h_{t+1} = \Delta d_{t+1} + \ln(e^{(\delta_t - \delta_{t+1})} + e^{\delta_t})$$

Campbell and Shiller linearised (3.13) around the mean log dividend-price ratio δ and the mean real dividend growth rate g, and obtained the following expression for the approximate real one-period return:

(3.14) $$h_{t+1} \approx \xi_{t+1} \equiv \delta_t - \rho\delta_{t+1} + \Delta d_{t+1} + k \qquad \text{where } k \equiv \ln(1+e^{\delta}) - \frac{\delta e^{\delta}}{1+e^{\delta}}$$

and ρ is a number a little smaller than unity (this is discussed more fully below).

An intuitive feel for the approximation (3.14) may be procured by first noting that taking logs of (3.11) we have

(3.15) $$h_{t+1} = \ln(P_{t+1} + D_{t+1}) - p_t$$

Also, unpacking the dividend-price ratio terms in expression (3.14), and rearranging, results in the following approximation to equation (3.15):

(3.16) $$\xi_{t+1} = (1-\rho)d_{t+1} + \rho p_{t+1} + k - p_t$$

(3.15) and (3.16) differ because the log of the sum of the real price and dividend in (3.15) have been replaced in (3.16) by a constant k plus a weighted average of the log real price and log real dividend. That these two expressions are approximately equal can be seen by first demonstrating that the *difference* $(1-\rho)\Delta d_{t+1}+\rho\Delta p_{t+1}$ approximates the *difference* $\Delta\ln(P_{t+1}+D_{t+1})$. It is a commonly-applied result that the change in the log of any variable is approximately equal to the proportionate change in the level, so that we have

$$(3.17) \qquad \Delta\ln(P_{t+1}+D_{t+1}) \approx \frac{P_{t+1}+D_{t+1}-P_t-D_t}{P_t+D_t} = \frac{P_{t+1}-P_t}{P_t+D_t} + \frac{D_{t+1}-D_t}{P_t+D_t}$$

If we assume that the ratio of the price to the sum of the price plus the dividend is roughly constant through time, and denoted by ρ, then equation (3.17) simplifies to the required approximate equality:

$$\Delta\ln(P_{t+1}+D_{t+1}) \approx \rho\frac{(P_{t+1}-P_t)}{P_t} + (1-\rho)\frac{D_{t+1}-D_t}{D_t} \approx \rho\Delta p_{t+1} + (1-\rho)\Delta d_{t+1}$$

Now that I have demonstrated the approximate equivalence of the difference in the two expressions (3.15) and (3.16), it remains to show that the constant k is such that the approximation holds exactly in levels in a static world in which dividends grow at a constant rate and stock returns are constant.

In a static world, if we add and subtract $(1-\rho)p_{t+1}$ to the RHS of equation (3.16), we find that the (constant) approximate log return is equal to

$$\xi = k + (1-\rho)(d_{t+1}-p_{t+1}) + (p_{t+1}-p_t)$$

$$(3.18) \qquad \xi = k + (1-\rho)\delta + g$$

Two more pieces of information are needed to obtain the required result. First, given the definition of ρ we can write the mean log dividend-price ratio as

$$\delta = \ln\left(\frac{1}{\rho}-1\right)$$

Substituting this into the expression for k gives

$$(3.19) \qquad k = -\ln\rho - (1-\rho)\delta$$

(3.19) in (3.18) gives

(3.20) $\xi = -\ln\rho + g$

Equation (3.20) for the static-world approximate log return is in fact exactly equal to the static-world actual log return h. To see this, we start from the definition of ρ:

$$\rho \equiv \frac{P_{t+1}}{P_{t+1}+D_{t+1}} = \frac{P_{t+1}}{P_t}\frac{P_t}{P_{t+1}+D_{t+1}} = e^{\ln\left(\frac{P_{t+1}}{P_t}\right)-\ln\left(\frac{P_{t+1}+D_{t+1}}{P_t}\right)} = e^{(g-h)}$$

Taking logs,

$$\ln\rho = g - h$$

$$h = -\ln\rho + g$$

Thus k ensures that the approximation (3.14) holds in levels.

Equation (3.14) can be thought of as a difference equation in terms of the log dividend-price ratio. If we solve this forward, imposing the terminal condition $\lim_{i\to\infty}\rho^i\delta_{t+i} = 0$, we obtain

(3.21) $\delta_t \approx \sum_{j=0}^{\infty}\rho^j(\xi_{t+1+j} - \Delta d_{t+1+j}) - \frac{k}{1-\rho}$

This expression states that the log dividend-price ratio is approximately equal to the discounted present value of all future returns in excess of real dividend growth, plus a constant. Equation (3.21) is an ex post relationship, completely devoid of economic content: it has been obtained only by the Taylor approximation to the log real return and the imposition of the dividend-price ratio terminal condition. However, we can also derive a useful ex ante relationship. Taking expectations of (3.21) using information available at the end of time t, we have

(3.22) $\delta_t \approx E_t\sum_{j=0}^{\infty}\rho^j(\xi_{t+1+j} - \Delta d_{t+1+j}) - \frac{k}{1-\rho}$

which states that the log dividend-price ratio equals the expected discounted present value of returns in excess of real dividend growth. Equation (3.22) will form the basis of all of the empirical analyses to come.

3.4 *A Variance Decomposition for Stock Returns*

3.4.1 *Decomposing the Variance of Unexpected Stock Returns*

Campbell (1991) used the log-linearisation (3.14) to decompose the variance of the unexpected one-period stock return into components due to news about future dividend growth, news about future returns, and the covariance between the two. Using equation (3.22) to substitute out δ_t and δ_{t+1} in (3.14), and collecting terms, we obtain the following expression for the unexpected one-period return:

(3.23) $$h_{t+1} - E_t h_{t+1} = (E_{t+1} - E_t)\sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} - (E_{t+1} - E_t)\sum_{j=1}^{\infty} \rho^j h_{t+1+j}$$

where h_{t+1} is the natural log of $H_{t+1.}$ Equation (3.23) has the following interpretation. Other things being equal, a positive shock to dividends, which the RVF states is coupled with an unexpected capital gain, is associated with a higher-than-expected one-period return. On the other hand, a rise in expectations of future returns leads to an immediate capital loss, and so a lower-than-expected current real return. For notational simplicity, one may re-write equation (3.23) as follows:

(3.24) $$\nu_{ht+1} = \eta_{dt+1} - \eta_{ht+1}$$

where

(3.24a) $$\nu_{ht+1} \equiv h_{t+1} - E_t h_{t+1}$$

(3.24b) $$\eta_{dt+1} \equiv (E_{t+1} - E_t)\sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j}$$

(3.24c) $$\eta_{ht+1} \equiv (E_{t+1} - E_t)\sum_{j=1}^{\infty} \rho^j h_{t+1+j}$$

Hence,

$$(3.25) \qquad Var(\nu_{ht+1})=Var(\eta_{dt+1})+Var(\eta_{ht+1})-2Cov(\eta_{dt+1},\eta_{ht+1})$$

Thus ν_{ht+1} is the unexpected stock return, η_{dt+1} represents news about future dividends whilst η_{ht+1} denotes news about future returns. It is clear from equation (3.25) that the impact of either news about dividends or future returns on unexpected one-period returns may be amplified or attenuated depending on the covariance between the two[7]. In my empirical work I divide through by $Var(\nu_{ht+1})$ in equation (3.25) and so the figures reported are the proportionate contributions of each news element.

In order to analyse unexpected returns, I first need a forecasting model for returns. To this end I construct a VAR model for returns which includes other variables which are believed to affect expected returns. If $\mathbf{z}_{t+1}$ is a k-vector, the first element of which is the real log stock return, h_{t+1}, I assume that $\mathbf{z}_{t+1}$ follows a VAR(1) process:

$$\mathbf{z}_{t+1}=\mathbf{A}\mathbf{z}_t+\mathbf{w}_{t+1}$$

The first-order representation[8] is useful because forecasts of future $\mathbf{z}_t$'s are easily obtained as

$$(3.26) \qquad E_t(\mathbf{z}_{t+1+i})=\mathbf{A}^{i+1}\mathbf{z}_t$$

Moreover, if we define the k-vector ι_1 as having unity as its first element and zeros elsewhere, multiperiod return forecasts are generated by the equation

$$(3.27) \qquad E_t h_{t+1+i}=\iota_1'\mathbf{A}^{i+1}\mathbf{z}_t$$

---

[7] The effect of the covariance term is excluded in the Poterba-Summers (1986) and Chou (1988) approaches.

[8] A VAR of any order can be rewritten as a first-order VAR (see Section 2.2.2).

Hence ν_{ht+1}, the unexpected stock return, is given by

(3.28) $\nu_{ht+1}=\iota_1'\mathbf{w}_{t+1}$

It follows, using (3.27), that the component of unexpected returns due to news about future returns can be written as

$$\eta_{ht+1} \equiv E_{t+1}\sum_{j=1}^{\infty}\rho^j h_{t+1+j} - E_t\sum_{j=1}^{\infty}\rho^j h_{t+1+j} = \sum_{j=1}^{\infty}\rho^j \iota_1{}'\mathbf{A}^j(\mathbf{z}_{t+1} - \mathbf{A}\mathbf{z}_t)$$

(3.29) $\eta_{ht+1}=\iota_1{}'\rho\mathbf{A}(\mathbf{I}-\rho\mathbf{A})^{-1}\mathbf{w}_{t+1} \equiv \lambda'\mathbf{w}_{t+1}$ where $\lambda'\equiv\iota_1'\rho\mathbf{A}(\mathbf{I}-\rho\mathbf{A})^{-1}$

From (3.24) and (3.29) I obtain the discounted sum of revisions to expected dividends as a residual:

(3.30) $\eta_{dt+1}=(\iota_1'+\lambda')\mathbf{w}_{t+1}$

Having estimated the VAR coefficient matrix **A**, I can calculate ν_{ht+1}, η_{dt+1} and η_{ht+1} from the above equations, and hence apportion the variability in one-period unexpected returns $Var(\nu_{ht+1})$ between news about dividends, news about future returns and any covariance between the two.

Excess Returns

So far I have considered only real stock returns; but much applied work focuses on the equity risk premium, defined as the excess return over some short-term interest rate. If the log of (one plus) the real interest rate is r_{t+1}, the excess return is defined as

(3.31) $e_{t+1}\equiv h_{t+1}-r_{t+1}$

One advantage of working with excess returns is that the price deflators in h_{t+1} and r_{t+1} cancel. This means that e_{t+1} is not dependent on measured price indices, which have a high likelihood of mis-measurement. Combining (3.31) with (3.23) we obtain an expression for the unexpected excess return:

$$\text{(3.32)} \qquad e_{t+1} - E_t e_{t+1} = (E_{t+1} - E_t)\sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} - (E_{t+1} - E_t)\sum_{j=0}^{\infty} \rho^j r_{t+1+j} - (E_{t+1} - E_t)\sum_{j=1}^{\infty} \rho^j e_{t+1+j}$$

or

$$\text{(3.33)} \qquad \nu_{et+1} = \eta_{dt+1} - \eta_{rt+1} - \eta_{et+1}$$

In the empirical analysis of this version of the model, the vector $\mathbf{z}_{t+1}$ now has e_{t+1} as its first element and r_{t+1} as its second. If we now define ι_2 as a k-vector with unity as its second element and zeros elsewhere, then from (3.26), (3.32) and (3.33) we obtain

$$\text{(3.34a)} \qquad \eta_{et+1} = \lambda' \mathbf{w}_{t+1}$$

$$\text{(3.34b)} \qquad \eta_{rt+1} = \iota_2'(\mathbf{I} - \rho\mathbf{A})^{-1}\mathbf{w}_{t+1} = \mu' \mathbf{w}_{t+1}$$

$$\text{(3.34c)} \qquad \eta_{dt+1} = (\iota_1' + \lambda' + \mu')\mathbf{w}_{t+1}$$

In a similar fashion as for real returns, I can use the VAR estimates to apportion the variance of unexpected excess returns between the constituent news elements on the RHS of equation (3.33).

3.4.2 *The Importance of Persistence in Return Expectations*

Campbell (1991) argued that two factors are important in determining the proportion of the variance of unexpected returns which is due to changes in expectations about future returns. These are the *forecastability* of stock returns, and the degree of *persistence in return expectations*. He demonstrated that if expected returns follow an AR(1) process, (3.35) $E_{t+1}h_{t+2}=\phi E_t h_{t+1}+u_{t+1}$

then the ratio of the variance of news about future returns to the overall variance of unexpected returns is approximately

$$\frac{Var(\eta_{ht+1})}{Var(\nu_{ht+1})} \approx \left(\frac{1+\phi}{1-\phi}\right)\left(\frac{R^2}{1-R^2}\right)$$

where R^2 is the proportion of the variance of returns which is forecastable. Campbell's focus was on monthly returns, for which R^2 is typically low. His point, therefore, was that news about future returns, η_{ht+1}, could still be important, providing ϕ were large. I analyse instead annual returns, which have a moderately high degree of predictability. However, the more general point regarding the joint influence of predictability and persistence still applies: for any given R^2, the contribution of η_{ht+1} depends on the persistence of return expectations.

Since, in the VAR context, there is no single natural measure of expectations persistence, Campbell proposed looking at the variability of the change in expectations of *all* future returns relative to the variability of the innovation in the *one-period-ahead* expected return. The VAR persistence measure for expected returns is then

$$P_h \equiv \frac{\sigma(\eta_{ht+1})}{\sigma(u_{t+1})} = \frac{\sigma(\lambda'\mathbf{w}_{t+1})}{\sigma(\mathbf{i}_1'\mathbf{A}\mathbf{w}_{t+1})} \equiv P_e$$

where $\sigma(x)$ denotes the standard deviation of x. P_h can also be viewed as (minus) the elasticity of the unexpected stock price change with respect to innovations in expectations about all future returns. A 1% positive innovation in expected returns will be associated with an unexpected P_h% capital loss on the stock[9].

[9] In the excess returns case I report a similar persistence measure for expected real interest rates, P_r. A 1% positive innovation in the expected real interest rate is associated with a P_r% capital loss.

3.5 *Data and Sample Statistics*

3.5.1 *Data*

The returns data are calculated using the BZW (value-weighted) equity stock price index and related dividends, and other variables used in the VAR are also taken from the *BZW Equity-Gilt Study.* Observations are taken annually at the end of each year from 1918 to 1993. Five variables are considered for inclusion in the investors' information set: the dividend-price ratio (D/P), the gilt-equity yield ratio (geyr), the real 1-month TB rate (rr), the spread between the gilt yield and the TB rate (term) and the real gilt yield (rgy). The Phillips-Perron (1988) tests shown in Table 3.1 indicate that all of the variables used in this VAR analysis are stationary, with the exception of the gilt-equity yield ratio[10]. Consequently, although I continue to report results from using this variable, one should bear in mind that their statistical validity is questionable.

3.5.2 *Approximation Error and Estimation of Rho*

Table 3.2 contains evidence on the accuracy of the Campbell-Shiller approximation when applied to the BZW data, presenting summary statistics of h_t, ξ_t and the difference between the two. The correlation between the actual and approximate log returns is very high (0.997), and although their means are a little different, this does not matter since I define all variables as deviations from their mean. The approximation (3.14) therefore seems empirically reliable for these data.

To implement the empirical study, I also need to estimate the linearisation constant ρ. I use the expression $\rho=1/(1+e^{\delta})$, where δ is the sample mean of the δ_t series, and the resultant estimate is 0.951, corresponding to a mean dividend-price ratio of 5.15%.

[10] Over the full sample period, there is a noticeable trend in the geyr series. However, after the late 1950's this trend is less apparent. There are then two possibilities with regards the robustness of the findings of Clare et al (1994) reported in Section 3.2.2. On the one hand, the nature of the gilt-equity yield ratio might have changed in recent decades, so that modelling its recent behaviour as a stationary process could be perfectly valid. However, if the apparent stationarity of the series is a small-sample problem, which is only apparent from my longer data set, the validity of Clare et al's conclusions are called into question.

Table 3.1: Phillips-Perron Unit Root Tests

Variable	PP Statistic[1]	
	t-statistic	$T(\phi-1)$
h	-8.915	-74.541
D/P	-5.537	-44.081
geyr	-1.712	-5.122
rr	-6.242	-49.990
term	-3.874	-23.339
rgy	-7.640	-67.949

[1] 10% critical values are -2.09 for the t-statistic and -13.4 for the $T(\phi-1)$ statistic. All reported statistics are calculated using a truncation lag of 1. Increasing the lag made no qualitative difference to the results.

Table 3.2: Comparison of Actual and Approximate Log Real Returns

Variable	Mean	Standard Deviation	Minimum	Maximum	Correlation with h_t
h_t	0.072	0.227	-0.856	0.692	1.000
ξ_t	0.064	0.229	-0.886	0.684	0.997
$h_t-\xi_t$	0.008	0.018	-0.126	0.038	-0.091

Table 3.3: BZW Return Autocorrelations

Lag	Coeff	Q-stat	p-value
1	-0.043	0.142	0.706
2	-0.156	2.069	0.356
3	-0.069	2.455	0.484
4	0.045	2.618	0.624
5	0.029	2.686	0.748

Table 3.4: Single Equation Return Regressions

X	$\hat{\beta}$	p-value	$\bar{R}^2$
D/P	7.79	0.00	23.75
geyr	0.05	0.09	1.98
rr	0.78	0.01	4.14
term	-1.72	0.31	0.59
rgy	0.82	0.04	2.93

3.5.3 *Return Autocorrelations and Single Equation Regressions*

Before proceeding with the VAR analysis, I first present a brief summary of the BZW return autocorrelations and the predictive power of several economic and financial variables. Table 3.3 shows the return correlogram for 1-5 years, together with Q-statistics and associated marginal probabilities. The first-order autocorrelation is negative, which coincides with the findings of Fama and French (1988a) and Poterba and Summers (1988). However, none of the estimated autocorrelations is statistically significant: there is no evidence from this source that the UK aggregate stock price is mean-reverting.

Next, I regressed the log stock return h_t on the first lag of each of the five variables considered here as candidates for the VAR. The estimated coefficients, p-values and adjusted R^2's are reported in Table 3.4. (All standard errors are heteroscedasticity-consistent.) It is clear from Table 3.4 that the dividend-price ratio is by far the most powerful predictor of returns, accounting for over 20% of annual return variation. The real interest rate, real gilt yield and the gilt-equity yield ratio also possess significant predictive power, although whether this would be true in the presence of the dividend-price ratio remains to be seen. In contrast to findings with US data, the term premium does not have any significant effect in the return equation.

3.6 *Estimation and Empirical Results*

Although I analyse the robustness of my results to changes in the VAR lag length, my initial choice of lag is determined by the data, on both informational and statistical grounds. Specifically, I add extra lags until the next lag is insignificant, subject to there being no detectable residual serial correlation in any of the equations. In addition, I correct the covariance matrix for possible heteroscedasticity using White's (1984) estimator.

3.6.1 *Results for Real Returns*

Before looking at the results of the variance decomposition, it is useful to view first a direct comparison of the current unexpected real stock return, v_{ht}, and the discounted value of revisions to expectations of future real returns, η_{ht}. This graphical comparison shows how the relative importance of revisions in expectations of future returns to current unexpected returns is related to the forecastability of stock returns. If expected returns are significantly time-varying, as measured by my five empirical proxies for expected returns, the RVF implies that one should see a negative relation between the two series. This is because, if future required returns are revised upwards, because, say, of a perceived increase in the riskiness of the portfolio, ceteris paribus the current stock price will fall. The latter leads to an immediate unexpected capital loss, and a negative current unexpected real return. However, note that if expected returns do not vary much over time (so that η_{ht} is practically zero for all t), it is revisions in expectations of future dividends which will be the major source of unexpected variation in stock prices. There will therefore be little or no relationship between ν_{ht} and η_{ht}.

Figures 3.1-3.5 show this comparison when each of the five return-predictor variables is included separately in the VAR. Figure 3.1 is for a VAR which includes only the log real stock return and the dividend-price ratio, and shows clearly a negative correlation between the current unexpected real stock return and revisions to expectations of all future real returns. Although this relationship is also apparent when the gilt-equity yield ratio is used to forecast returns (Figure 3.2), it is much less marked for the remaining three variables (the real interest rate, the real gilt yield and the term premium). This is consistent with the above findings on the predictive ability of these variables. Since these last three variables have little predictive power over returns, the VAR which uses them to forecast returns produces forecasts that are roughly constant, the corollary being that unexpected returns are *more* variable than when D/P or geyr are used. As expected returns are measured to be roughly constant, revisions to future expected returns are minimal, so that the model ascribes nearly all of the movement in current stock prices to changes in expected future dividends. The observed relationships between ν_{ht} and η_{ht} in Figures 3.3, 3.4 and 3.5 are therefore extremely weak.

Table 3.5 reports the results of the decomposition of the variance of unexpected real returns. These statistics are scaled by $\mathrm{Var}(\nu_h)$ and so show the relative importance of news about dividends, future returns and the covariance between the two. Results in the first row are for a VAR containing only returns and

Figure 3.1: Unexpected Current and Discounted Future Log Real Returns
z=[h, D/P]

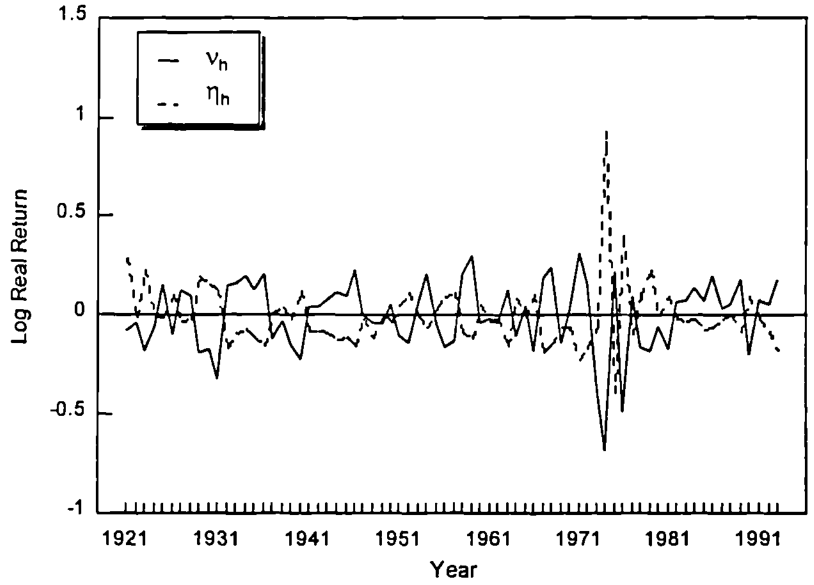

Figure 3.2: Unexpected Current and Discounted Future Log Real Returns
z=[h, geyr]

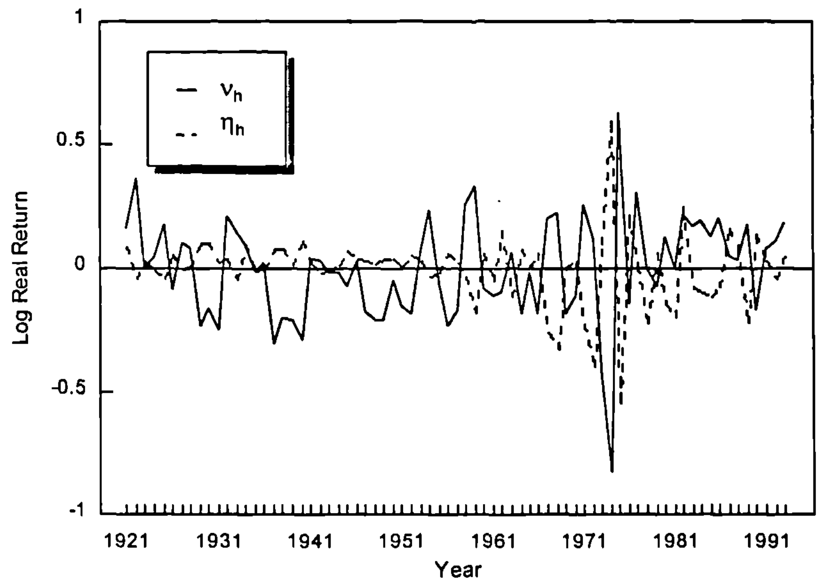

Figure 3.3: Unexpected Current and Discounted Future Log Real Returns
z=[h, rr]

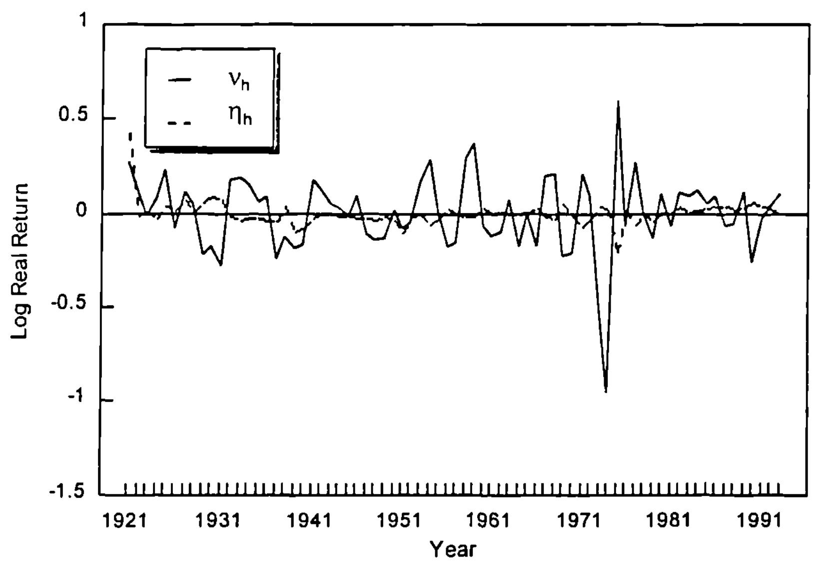

Figure 3.4: Unexpected Current and Discounted Future Log Real Returns
z=[h, rgy]

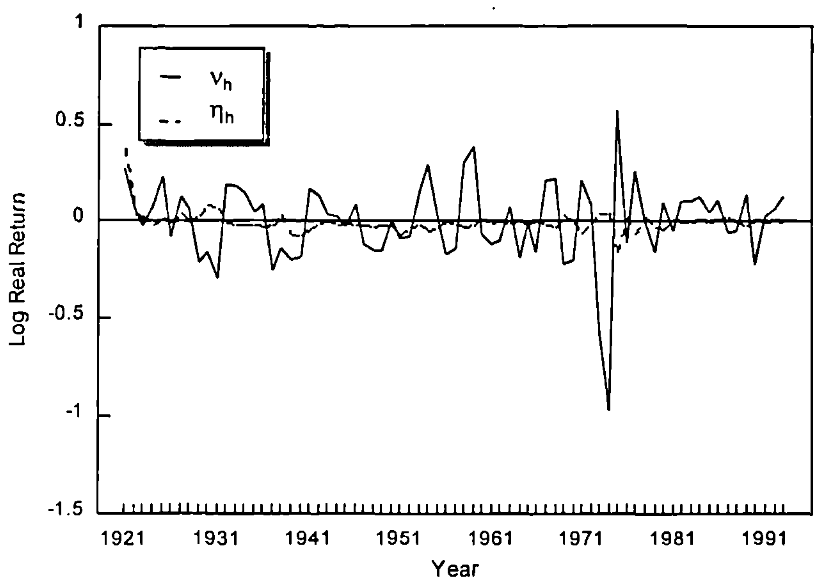

Figure 3.5: Unexpected Current and Discounted Future Log Real Returns
z=[h, term]

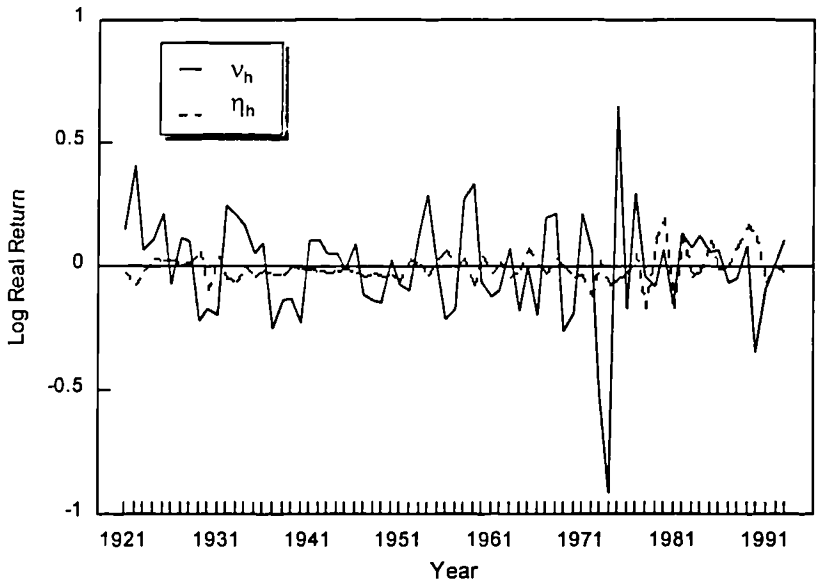

Table 3.5: Variance Decomposition of Unexpected Real Returns

VAR Variables	Lags	$\bar{R}_1^2$	$\bar{R}_2^2$	P_h	$Var(\eta_h)^1$	$Var(\eta_d)$	$-2Cov(\eta_h,\eta_d)$	$Corr(\eta_h,\eta_d)$
h, D/P	2	39.23	34.72	1.39	0.74	0.28	-0.02	0.02
				(0.36)	(0.17)	(0.09)	(0.14)	(0.15)
h, geyr	2	24.53	85.65	1.21	0.46	1.07	-0.53	0.38
				(0.69)	(0.55)	(0.37)	(0.58)	(0.23)
h, rr	2	7.20	47.61	1.78	0.07	0.92	0.01	-0.01
				(0.74)	(0.06)	(0.43)	(0.42)	(0.83)
h, rgy	2	6.13	36.73	1.52	0.04	0.84	0.12	-0.34
				(0.57)	(0.05)	(0.39)	(0.35)	(0.89)
h, term	2	4.39	48.15	35.64	0.18	1.10	-0.29	0.32
				(611.29)	(0.26)	(0.76)	(0.99)	(0.79)

[1] The statistics reported for the variance decomposition (Columns 7-10) are standardised by division by $Var(\nu_h)$, with the result that they sum to unity (see equation (3.25)). Thus, for example, what is reported as $Var(\eta_h)$ is strictly $Var(\eta_h)/Var(\nu_h)$. The statistics give the *proportional* contribution of the relevant news varable to the variance of the unexpected current return.

Table 3.6: Variance Decomposition with Multiple Predictors

VAR Variables	Lags	$\bar{R}_1^2$	$\bar{R}_2^2$	$\bar{R}_3^2$	P_h	$Var(\eta_h)^1$	$Var(\eta_d)$	$-2Cov(\eta_h,\eta_d)$	$Corr(\eta_h,\eta_d)$
h, D/P,geyr	2	43.34	37.98	89.10	1.51	0.73	0.25	0.02	-0.02
					(0.45)	(0.26)	(0.10)	(0.28)	(0.15)
h, D/P, rr	3	38.89	36.94	51.04	1.75	0.80	0.27	-0.07	0.07
					(0.66)	(0.24)	(0.11)	(0.26)	(0.26)
h, D/P, rgy	3	37.29	36.02	33.96	1.70	0.78	0.24	-0.02	0.02
					(0.66)	(0.22)	(0.09)	(0.23)	(0.26)
h, D/P, term	3	42.18	38.70	49.59	1.96	1.02	0.35	-0.37	0.31
					(0.76)	(0.25)	(0.20)	(0.38)	(0.22)

[1] See notes to Table 3.5.

the dividend-price ratio, and the adjusted R^2 for both equations is quite high at 39% and 35% respectively[11]. The degree of persistence in returns is not particularly high with P_h = 1.39 (s.e. = 0.36) indicating that a 1% positive innovation in the discounted present value of all future expected returns is associated with a 1.39% current capital loss, *ceteris paribus*. However, the relatively high degree of predictability of returns, coupled with this modest degree of persistence of expected returns, results in a high proportion (0.74) of unexpected returns being attributable to news about future returns. The proportion of the unexpected return variance attributed to revisions in expected future dividends is correspondingly quite low (0.28). Both contributions are statistically different from zero, with virtually no contribution from the covariance term. This finding suggests that news about future dividends and news about future returns are largely independent.

In fact, this last conclusion is supported by all of the other models in the table: the covariance between news about dividends and news about returns is never statistically significant. However, estimates of the other statistics do differ according to model specification. In row 2 of Table 3.5, almost 25% of the variation in the real return is accounted for by the gilt-equity yield ratio. The persistence measure is modest (1.21), and not very different from that in row 1. One striking result here is that a large part of the return innovation variance is apportioned to news about dividends (1.07), and that ascribed to news about returns (0.46) is not statistically significant. A similar pattern emerges with the models that include the real interest rate (rr), the real gilt yield (rgy) and the term premium (term) (table 3, rows 3-5). Much of the unexpected return variance accrues to shocks to dividend growth, with shocks to returns playing no significant role. This is perhaps not surprising given the relation between R^2 and η_h discussed above in Section 3.4.2, and the previous discussion of Figures 3.1-3.5.

Notice finally that in the term premium model none of the effects is statistically significant. We saw above that the term premium does not appear to be a significant predictor of returns. (Indeed the marginal probability of the F-test that the coefficients in the VAR return equation are jointly zero is 0.34 for this model.) Expected returns are practically constant, giving rise to an implausibly high (and imprecisely estimated)

11 All reported R^2's are adjusted for degrees of freedom.

return expectations persistence measure. It is plainly the case that if time variation in expected returns is minimal, the bulk of the variability in stock prices must come through dividend surprises. However, if we believe that expected returns do vary significantly through time, and the VAR is an attempt to model this time-variation through the use of expectation proxies, the results for the term premium model are not very informative.

I next look at the results from VAR models which all include the dividend-price ratio and one other variable, in addition to the real return. These are reported in Table 3.6. There is little change in the estimates from the addition of the extra variables: the results display a remarkable degree of uniformity. Only with the addition of the gilt-equity yield ratio does the adjusted R^2 rise (from 39.23% to 43.34%). The return expectations persistence statistics are all higher, although not significantly so, ranging from 1.51 (s.e.=0.45) to 1.75 (s.e.=0.66), as compared to 1.39 (s.e.=0.36) in row 1 of Table 3.5. In all three cases the contribution of news about real returns to the unexpected return variance is around three times the contribution of dividend news. None of the covariances is significant.

The dividend price ratio is clearly the most important variable in the above analysis. When it is allowed to influence expected returns, the majority of variation in unexpected returns derives from news about future discount rates, with news about cash flows playing a smaller and statistically marginal role. However, there is a substantial literature indicating that expected market volatility is the key to stock price movements, and so I now review this literature prior to taking explicit account of the effects of expected volatility on the variance decomposition.

3.7 *Volatility, Volatility Persistence and Expected Returns*

3.7.1 *Review of Volatility Literature*

Between 1965 and 1982, the New York Stock Exchange Index fell by nearly 70% in real terms. Initially, this dramatic decline was attributed to the concurrent increase in the average rate of inflation, with several causal links being identified. These included the suggestion of confusion between real and nominal returns (Modigliani and Cohn 1979), and the interaction of inflation with structural factors such as the tax system (Feldstein 1980, Summers 1981).

Pindyck (1984) dismissed empirical evidence in support of the second of these factors as being due to misspecification of the model relating expected returns to inflation. Pindyck took up Malkiel's (1979) suggestion that a marked increase in the "riskiness" of capital investments over the period in question was empirically a more powerful explanation of the stock market decline. The argument is that the variance of firms' real gross marginal return on capital rose significantly, which in turn increased the relative riskiness of investors' real returns from holding stocks, reducing the attractiveness of stock market investment. Pindyck's empirical analysis showed that the mean and variance of the inflation rate, and the variance of the real gross marginal return of capital, increased significantly over the twenty years to 1985. Combining these findings with a simple model of asset demand, Pindyck concluded that much of the decline in share values was due to greater uncertainty about the gross marginal return on capital.

The issue of *volatility persistence* was raised by Poterba and Summers (1986). They argued that if shocks to the volatility of stock returns are to have a significant impact on stock prices then these shocks must die away slowly. Only if the market becomes more volatile for a long period after an innovation will investors re-calculate their risk premia on stocks, and hence the discount rate applied to expected future dividends (assuming rational valuation). Through this channel, volatility may affect the level of the market. Conversely, if shocks to volatility are only transitory, volatility will not affect mean returns.

Poterba and Summers based their analysis on four assumptions:

i) firms are not levered. This avoids Black's (1976) observation that the level of share prices, by affecting the degree of leverage, can have a direct impact on volatility;
ii) dividends grow at a constant rate g, and the real rate of return, h, is constant;
iii) there is a linear relationship between the risk premium α^* and the variance of equity returns, σ^2;
iv) the conditional variance of returns σ^2_t follows an AR(1) process, with AR(1) coefficient ϕ measuring persistence in the conditional variance.

Poterba and Summers derived the following expression for the elasticity of the stock price with respect to the conditional return variance:

$$\frac{d \ln P_t}{d \ln \sigma_t^2} = -\frac{\alpha}{1 + h + \alpha - \phi(1 + g)}$$

where α is the mean risk premium. The absolute value of this elasticity is increasing in ϕ: the larger is ϕ, the greater is the responsiveness of prices to changes in conditional volatility (as measured by the conditional variance of equity returns).

Poterba and Summers evaluated the above elasticity using "plausible parameter values" for all but ϕ, which they estimated directly. Their estimates produced a maximum value of the elasticity of -0.225. The latter implies that a 50% increase in volatility would depress share prices by only 11%. They also noted that this elasticity estimate included the Great Depression years, during which the volatility of market returns reached an unprecedented level for several months, and thus had a significant positive effect on the estimate of ϕ. They concluded that their results "cast serious doubt on the view that changes in volatility...have a substantial effect on stock market values".

Pindyck (1986) developed a simple portfolio choice model to quantify the relative effects of firms' profitability (a "fundamental"), interest rates, inflation and risk on the stock price. In support of his earlier paper, only the inflation rate was found to have no significant effect on stock prices. Whilst conceding, on the basis of the Poterba and Summers paper, that shocks to volatility are transitory, and so the effect of these surprises was less marked than he had originally claimed, Pindyck concluded that the variance of stock returns nevertheless is the most powerful explanatory variable of those considered, with an elasticity of -2.9 for the 1949-83 period (rising to -6.5 for 1962-83).

French, Schwert and Stambaugh (1987) addressed the question of whether the expected market risk premium is positively related to risk, as measured by the volatility of the stock market. They used daily values of the Standard and Poor Composite Portfolio to estimate the monthly standard deviation of stock market returns from January 1928 to December 1984. Regressions of the form

$$R_{mt} - r_t = \alpha + \beta\hat{\sigma}_{mt} + \lambda\hat{\sigma}^u_{mt} + \varepsilon_t$$

were performed, where R_{mt} is the return on the market portfolio, r_t is a risk-free rate, and stock market volatility was decomposed into its predicted and unpredicted elements using an ARIMA(0,1,3) model. Although λ was found to be reliably and significantly negative, so that excess returns rise (fall) when unexpected volatility falls (rises), the estimates of β provided little evidence of a direct effect of predicted volatility on expected risk premia. Their findings pointed them towards the following argument for the relationship between volatility and stock prices. If the discount rate applied to dividends is positively related to expected volatility, a positive shock to the latter (i.e. an *unexpected* rise in volatility) will increase the dividend discount rate. Assuming that cash flows themselves are unaffected by this innovation[12], the higher discount rate reduces the current stock price. Thus a positive relation between expected returns and expected volatility is consistent with a negative relation between stock prices and unexpected volatility.

French et al also modelled volatility by fitting a restricted ARCH model and a GARCH(1,2) model to their excess return series. For their whole sample period they estimated the sum of the GARCH coefficients to be 0.996, and calculated a comparable ARCH persistence parameter of 0.938. Although they did not make this point, these results are in direct conflict with the Poterba and Summers conclusion that shocks to volatility die away quickly.

Chou (1988) re-examined the issue of volatility persistence using the GARCH-M model developed by Engle, Lilien and Robbins (1987), in which the conditional variance of the return regression errors is allowed to affect the conditional mean return. Chou's complete model was

$$R_t = r + \gamma V_{t-1} + \varepsilon_t$$

$$V_t = \alpha_0 + \alpha \varepsilon_{t-1}^2 + \beta V_{t-1}$$

The motivation behind this specification was twofold. First, it is interesting to see if conditional volatility has a direct effect on returns, regardless of the question of persistence. Second, Chou derived an expression for the elasticity of the stock price

[12] In the current study, this independence of shocks to volatility and shock to dividends is evidenced by insignificant correlations between η_{ht} and η_{dt}, reported in Section 3.2.7.

with respect to conditional volatility which was very similar to that of Poterba and Summers. The main difference was that the AR(1) persistence parameter ϕ was replaced by the GARCH(1,1) persistence measure $(\alpha+\beta)$. That $(\alpha+\beta)$ measures volatility persistence is most easily seen by noting that the GARCH(1,1) model implies the following relationship:

$$E_t(V_{t+s}-\sigma^2)=(\alpha+\beta)^s(V_t-\sigma^2)$$

where σ^2 is the unconditional error variance. If the current conditional variance V_t differs from the unconditional variance, then current expectations of future divergences between the two will tend to zero in proportion to the forecast horizon s at a rate determined by $(\alpha+\beta)$. The closer is $(\alpha+\beta)$ to unity, the longer any difference between V_t and σ is expected to persist. The estimates of α and β from the above model would therefore provide more evidence on the volatility persistence question.

Using weekly returns on the NYSE from July 1962 to December 1985, Chou estimated γ to be around 4.5 and statistically significant, indicating that excess returns were positively related to expected volatility. The estimated GARCH persistence measure was 0.986, which implies that even a year after a shock occurs, 42% of the initial impact still affected volatility.

Chou evaluated the elasticity of the real stock price with respect to conditional volatility using Poterba and Summers's value for γ and his own estimates converted into monthly measures. He calculated that, on average over the whole period, a doubling of stock volatility reduced the stock value by 11%. This does not seem a very potent effect. However, the implication is that the doubling of stock volatility in 1974 would cause a 26% drop in the market index, which is very close to the actual drop of 27%. Thus Chou's model appeared to confirm that increased volatility was an important determinant of the plunge in the US market.

Because of the proximity of his estimate of $(\alpha+\beta)$ to unity, Chou tested his GARCH model against an Integrated-GARCH specification[13]. He found that the data could not discriminate between the two, and concluded that the variance process was in fact non-stationary. However, contradictory evidence was presented by Akgiray

[13] The Integrated GARCH (I-GARCH) model has $(\alpha+\beta)=1$.

(1989) who analysed daily returns on the CRSP value- and equal-weighted indices for the period 1963-1986. On the basis of a series of likelihood ratio tests, Akgiray also settled upon a GARCH(1,1) model. His estimates of (α+β) were always close to, but less than, one, and when a Dickey-Fuller test was applied to the conditional variance series, the null of a unit root was rejected for all but one sub-period. However, the conclusion that conditional volatility possesses a high degree of persistence remained.

Lamoureaux and Lastrapes (1990a) also fitted a GARCH(1,1) model, but with the intention of explaining, and thus removing the GARCH effect. The excess kurtosis typically exhibited by stock returns could be due to them being generated by a mixture of distributions. Lamoureaux and Lastrapes proposed the rate of daily information arrival as the stochastic mixing variable, as Diebold (1986), amongst others, has suggested that ARCH might capture the time-series properties of this mixing variable. Taking daily trading volume as a proxy for the rated of information arrival, Lamoureaux and Lastrapes compared the estimates of (α+β) with and without the trading volume effect appearing in the conditional variance equation. Without the trading volume effect, the persistence measure was estimated to be 0.728 (which was markedly lower than those of French et al, Chou or Akgiray). However, when volume was included, the mean estimate of (α+β) fell to only 0.073. They concluded that lagged squared residuals contributed little additional information about the variance of the stock returns process once the rate of information flow had been accounted for, and suggested that for asset return series for which no appropriate measures of information arrival are available, an ARCH model would suffice.

Lamoureaux and Lastrapes (1990b) investigated the extent to which estimates of the persistence of volatility may be overstated due to structural breaks in the model. They expressed discomfort with the apparent near-integratedness of the conditional variance of stock returns since it lacked theoretical underpinning, and pointed out that findings on persistence seemed to be related to the sample size/return horizon (LL's previous estimate was relatively low, and over a shorter horizon). One possibility was that the parameters of the GARCH model were unstable over long horizons. They therefore compared a standard GARCH(1,1) model with one augmented by thirteen dummy variables allowing shifts in the intercept. They found that the inclusion of the dummies resulted in a significant change in the estimated GARCH persistence measure, and concluded that the possibility of time-varying GARCH parameters accounted for the difference in results over different sample horizons. However, it remained to identify exactly when the shifts occur, and whether they were one-off events or contain relevant information regarding the evolution of the variance process.

A couple of conclusions can be drawn from the above discussion. First, there appears to be consensus on the view that if market volatility is to have an important effect on stock prices then expected volatility must be persistent, so that shocks to volatility have a near-permanent effect. However, there is by no means consensus on the degree of persistence in volatility expectations implicit in US data.

There are three more points worth making at this juncture. First, none of the above studies considers time-variation in expected dividends: all movements in stock prices are assumed to be the consequence of changes in expected returns. The VAR methodology which forms the basis of the current study not only allows dividend expectations to vary through time, but allows for covariation between changes in expected returns and changes in expected dividends. Second, although the time-series modelling of returns effected by the VAR is less sophisticated than ARCH-type models, non-stochastic formulations of the latter do not allow for unexpected volatility, whereas both expected and unexpected components of volatility can be retrieved from the VAR. Finally, just as VAR modelling of returns may be preferred to univariate time-series models, to the extent that volatility is not well modelled by its own past values, the VAR model provides a more parsimonious and more powerful decomposition of measured volatility than the more traditional ARIMA models.

3.7.2 *The Effects of Volatility on the Variance Decomposition*

I examine the importance of the market's perception of the impact of risk by including a measure of stock price volatility V (the squared ex post real stock return[14]) in the VAR. Before discussing the results of the variance decomposition, it is

[14] The qualitative results are invariant to the use of squared returns or mean-adjusted squared returns.

instructive to look first at a couple of the VAR models for returns and volatility. First of all, the estimated VAR equations for the two-variable model z_t=[h_t,V_t] are

$$h_t = \underset{(0.331)}{0.125}\, h_{t-1} - \underset{(0.520)}{0.508}\, V_{t-1} - \underset{(0.239)}{0.380}\, h_{t-2} + \underset{(0.333)}{0.442}\, V_{t-2} - \underset{(0.137)}{0.072}\, h_{t-3} - \underset{(0.171)}{0.011}\, V_{t-3}$$

$$V_t = -\underset{(0.154)}{0.411}\, h_{t-1} + \underset{(0.221)}{0.449}\, V_{t-1} - \underset{(0.080)}{0.049}\, h_{t-2} + \underset{(0.115)}{0.053}\, V_{t-2} + \underset{(0.059)}{0.013}\, h_{t-3} - \underset{(0.080)}{0.050}\, V_{t-3}$$

with standard errors in parentheses. The adjusted R^2's are 1.37% and 44.68% respectively. It is noticeable that the first lag of the log real return has significant forecasting power for volatility, although lagged volatility appears to have no marginal explanatory power over returns. However, there is a market improvement in the overall performance of the VAR when the dividend-price ratio is also included. Table 3.7 reports the VAR estimates, their marginal probability values, and the individual equation adjusted R^2's for the model which includes both the dividend-price ratio and the volatility measure. This VAR model seems particularly successful at modelling volatility, with an adjusted R^2 of 65.32%. The most striking feature of these results is that the dividend-price ratio plays a major role in the volatility model, with highly significant coefficients at lags 1 to 3. This raises the possibility that the significance of the dividend-price ratio in the return equation is due to it being closely related to perceived risk. The question of whether the dividend-price ratio's success in predicting returns is due *solely* to its ability to forecast risk (as would be predicted by the CAPM) is explicitly tested in Chapter 5.

The R^2 of the equations for returns and the dividend-price ratio increase quite substantially when V is included (compare the R^2's in row 1 of Table 3.8 with those in row 1, table 3.5) with that for returns being 49% and for the dividend-price ratio being 47%. Hence, in conformity with the CAPM, in this VAR model the variance of the market portfolio does have incremental explanatory power for returns. (Although none of the lags of V itself is significant in the returns equation, the apparent collinearity between V and the dividend-price ratio makes inference based on individual coefficients difficult. In a VAR, where none of the variables is exogenous

Table 3.7: Volatility Model VAR Estimates

	h_t	D/P_t	V_t
h_{t-1}	0.862 (0.022)	-0.073 (0.031)	-0.039 (0.812)
D/P_{t-1}	16.401 (0.004)	-0.497 (0.286)	6.497 (0.023)
V_{t-1}	-0.407 (0.207)	0.046 (0.078)	0.649 (0.000)
h_{t-2}	-0.943 (0.001)	0.056 (0.001)	-0.469 (0.001)
D/P_{t-2}	-20.631 (0.000)	1.463 (0.000)	-13.980 (0.000)
V_{t-2}	0.106 (0.676)	-0.024 (0.163)	-0.149 (0.294)
h_{t-3}	-0.235 (0.331)	0.028 (0.086)	-0.067 (0.555)
D/P_{t-3}	11.901 (0.051)	-0.220 (0.586)	7.677 (0.014)
V_{t-3}	0.256 (0.081)	-0.014 (0.120)	0.105 (0.147)
h_{t-4}	0.539 (0.000)	0.024 (0.001)	0.158 (0.065)
D/P_{t-4}	7.038 (0.0.67)	-0.478 (0.034)	2.458 (0.220)
V_{t-4}	-0.669 (0.000)	0.032 (0.812)	-0.185 (0.020)
$\bar{R}^2$	48.890	47.190	65.320

Table 3.8: Variance Decompostion of Unexpected Real Returns: Volatility Model

VAR Variables	Lags	$\bar{R}_1^2$	$\bar{R}_2^2$	$\bar{R}_3^2$	P_h	$Var(\eta_h)^1$	$Var(\eta_d)$	$-2Cov(\eta_h,\eta_d)$	$Corr(\eta_h,\eta_d)$
h, D/P,V	4	48.89	47.19	65.32	1.95	1.05	0.27	-0.32	0.30
					(0.73)	(0.24)	(0.14)	(0.33)	(0.22)
h, V,geyr	2	27.61	45.53	85.74	1.10	0.41	0.81	-0.22	0.19
					(0.63)	(0.48)	(0.27)	(0.46)	(0.30)
h, V, term	2	12.21	41.02	46.79	2.18	0.24	0.91	-0.15	0.16
					(1.88)	(0.23)	(0.61)	(0.76)	(0.71)

[1] See notes to Table 3.5.

Table 3.9: Variance Decomposition of Unexpected Excess Returns

VAR Variables	Lags	$\bar{R}_1^2$	$\bar{R}_2^2$	$\bar{R}_3^2$	$\bar{R}_4^2$	P_e	P_r	$Var(\eta_e)^1$	$Var(\eta_r)$	$Var(\eta_d)$	$-2Cov(\eta_e,\eta_r)$	$-2Cov(\eta_e,\eta_d)$	$-2Cov(\eta_r,\eta_d)$
e, r, D/P	3	42.12	50.70	37.11	-	1.55	3.77	0.62	0.91	0.85	0.03	0.15	-1.57
						(0.59)	(0.75)	(0.17)	(0.38)	(0.37)	(0.35)	(0.37)	(0.71)
e, r, D/P, V	2	41.06	49.38	44.08	63.42	1.20	4.19	0.62	0.51	0.50	0.18	-0.14	-0.68
						(0.30)	(0.82)	(0.18)	(0.17)	(0.16)	(0.26)	(0.25)	(0.27)
e, r, D/P, geyr	3	46.20	49.07	36.35	88.64	1.59	3.68	0.83	0.92	0.91	-0.41	0.37	-1.62
						(0.51)	(0.73)	(0.30)	(0.36)	(0.41)	(0.63)	(0.61)	(0.70)
e, r, V	2	13.46	47.34	52.12	-	1.09	4.68	0.11	0.65	1.50	-0.03	0.46	-1.68
						(0.81)	(0.78)	(0.10)	(0.28)	(0.71)	(0.30)	(0.42)	(0.79)
e, r, geyr	3	27.18	45.04	84.44	-	1.16	4.74	0.56	1.30	2.34	-0.30	0.30	-3.20
						(0.43)	(1.24)	(0.45)	(0.72)	(1.18)	(0.57)	(0.81)	(1.82)
e, r, term	2	1.27	62.50	46.92	-	1.70	6.05	0.04	0.98	2.21	-0.04	0.35	-2.54
						(1.59)	(2.15)	(0.08)	(0.61)	(1.42)	(0.53)	(0.80)	(1.66)

[1] See footnote to Table 3.5.

and the potential for collinearity - and so estimation bias - is high, the *overall rise* in explanatory power from the inclusion of an extra variable is perhaps a better measure of that variable's contribution to forecasting than tests of the significance of individual coefficients.) As already mentioned, the equation for V also has a high R^2 of 65% and identifies some persistence in volatility, with the sum of the (own) lagged dependent variables being 0.43[15]. V also depends on lagged returns, a feature noted by French et al (1987, footnote 2), but not explained in their single-equation return regressions. Once again, in so far as lagged returns predict future volatility, the significance of the lagged dependent variables in the returns equation may yet be accounted for.

When z=[h, D/P, V], the variance decomposition (row 1, Table 3.8) indicates an increase in the point estimate of the contribution of news about future returns (from 0.74) to 1.05. The contribution to unexpected movements in current returns of news about future dividends is virtually unchanged at 0.27, but is now statistically insignificant (s.e.=0.14), whilst the covariance term remains insignificantly different from zero. Thus the independence of news about returns and news about dividends remains a key feature of the data.

The return expectations persistence statistic, P_h, is actually a little higher than before, rising from 1.39 (s.e.=0.36) without V in the VAR (see row 1, Table 3.5), to 1.95 (s.e.=0.73). The observed persistence in expected volatility noted above appears to induce marked extra persistence in return expectations. A 1% positive shock to expected returns will, according to this model, cause an immediate capital loss of around 2%, which is substantially larger than that predicted by the Poterba and Summers (1986) model for the US. With regards, then, to Pindyck's point about volatility persistence and the likely effects on stock price movements, for the UK I find both substantial persistence in volatility expectations, and a large effect of shocks to expected future returns on current stock prices.

If one excludes the dividend-price ratio from the VAR (see rows 2 and 3, Table 3.8) then the explanatory power of the returns equation falls dramatically and all the statistics of interest are ill-determined[16]. None of the other models has much ability to

---

[15] This provides a lower bound on the degree of expectations persistence, since the equation for V also includes the dividend-price ratio, which also exhibits persistence.

[16] Campbell also obtains this result using an aggregate stock price index for the US (see Section 3.13).

forecast either volatility or returns. Just as we have come to expect, the lower the return equation R^2, the more weight is given to dividend news and less to cash flow news.

Overall, the results using real returns indicate that the dividend-price ratio is a crucial determinant of one-period returns, whilst volatility also provides incremental explanatory power. Also, these variables imply that the contribution of news about future discount rates is three to four times that of news about future dividends in determining unexpected capital gains and losses on stock prices. Finally, we see that the covariance between news about future dividend growth and news about future returns, whilst usually positive, is never statistically significant.

3.8 *Results Using Excess Returns*

Initially I examine the variance decomposition using excess returns (e), the real interest rate (r) and the dividend-price ratio in the VAR (row 1, Table 3.9). The persistence measure for excess returns P_e is 1.55 (s.e. = 0.59) and the proportionate contribution of news about future excess returns to changes in current unexpected returns is 0.62 (s.e. = 0.17). Both of these statistics are of a similar order of magnitude as in the real returns model. The excess returns model allows one to examine the importance of the real interest rate and its interaction with news about future dividends and excess returns. From Table 3.9 (row 1) we see that news about dividends and the real rate have a "direct" effect on unexpected excess returns of 0.85 (s.e. = 0.37) and 0.91 (s.e. = 0.38) respectively. However, the covariance between these two sources of news is positive (-2 $\text{cov}(\eta_d, \eta_r)$ = -1.57, s.e. = 0.71) so that their *net contribution* to unexpected excess returns is close to zero (i.e. 0.85 + 0.91 - 1.56). As the remaining covariances of 0.03 and 0.15 are relatively small (and statistically insignificant), it is news about future excess returns that ultimately drives unexpected capital gains and losses. This result is therefore broadly consistent with that found for the real returns model.

The results in rows 2 and 3 in Table 3.4 show that variations in the variables in the VAR do not appreciably affect the qualitative conclusions outlined above. When the dividend-price ratio is excluded from the VAR the explanatory power of the (excess) returns equation falls dramatically (see Table 3.9, rows 4 and 5) so that the statistics of interest are very poorly determined. This was also the case for the real returns equation. I therefore do not discuss these results further.

Note that although real interest rates are highly persistent (P_r is in excess of 3.5 in Table 3.9[17]) and therefore would have a powerful impact on stock prices, *ceteris paribus*, nevertheless in this data set their net impact is small because increases in real interest rates (which tend to depress stock prices) are accompanied by a rise in real dividends (which tends to increase stock prices).

3.9 *Time-Series Behaviour of Expected Returns*

Conrad and Kaul (1988) claimed that expected returns are well modelled by a first-order autoregressive process. It would be convenient if this were found to be the case, as the "persistence" measure ϕ is more intuitive than the P_h measure. I therefore test the appropriateness of the AR(1) assumption. Taking expectations of (3.35) at time t, we have

$$E_t h_{t+2} = \phi E_t h_{t+1}$$

This expression can be rewritten in terms of the VAR forecasts as follows:

$$\mathbf{\iota_1' A^2 z_t = \phi \iota_1' A z_t}$$

which gives us the following set of restrictions on the VAR parameters:

$$\mathbf{\iota_1' A(A - \phi I) = 0}$$

I computed Wald tests of these restrictions to the models which include the dividend-price ratio. When only h and D/P were included in the VAR, the marginal probability value of the test was 0.02, and so the AR(1) restriction was rejected at the 5% level, and for the other VAR specifications, the p-values were much lower. I therefore conclude that whilst the AR(1) model is a useful heuristic device, its implications are not consonant with the data.

[17] My estimates of P_r are around double those obtained by Campbell (1991) with US data.

3.10 *Implications*

The implications for studies of stock prices and returns are quite clear. When examining the behaviour of stock prices using the RVF, it does not seem an unreasonable approximation to assume that dividend growth is constant (as in Poterba and Summers 1986, French et al 1987) since news about future dividends has relatively little impact on unexpected changes in stock prices. However, it is very important to model the time variation in returns in order to capture any persistence in return expectations. The latter applies even though one-period returns are not found to be highly predictable.

Single-equation studies that regress *excess* stock returns on variables which purport to measure expected returns (e.g. dividend-price ratio) together with "news variables" will fail to provide an adequate statistical explanation (e.g. Fama 1990, Roll 1988) unless they capture the covariance between news about dividends and news about real interest rates. Usually such studies do not explicitly model this interaction and it may therefore only be captured by serendipity in the ad hoc surprise variables of the single-equation regression.

The advantage in using a VAR framework is that the impact of the persistence in expected returns on movement in stock prices can be explicitly analysed. Through this, I find clear support for the French et al (1987) conjecture that a negative relationship between one-period returns and *unexpected* volatility is consistent with a positive relationship between the ex ante return and *expected* volatility. The implication of persistence in volatility expectations is that a current shock to volatility translates into news about returns in all future periods. This in turn has a powerful effect on current prices and hence on current one-period unexpected returns.

3.11 *Some Variants*

In this section I give a breviloquent outline of the effects on the variance decomposition of changes in the VAR lag length and a change in the sample period.

All of the VAR models were initially estimated with VAR lag lengths of 1, 2, 3 and 4 (not reported). No clear pattern is discernible between the lag length and the statistics of interest, but the latter are never very far from those reported. For both real and excess returns (when D/P is included in the VAR), there is some tendency for the estimate of the persistence of expected returns to rise above 2 when the VAR lag length is 1; but this is never statistically different from the estimates reported in tables 3, 4, 6 and 7. Moreover, the concomitant fall in the return-equation R^2 leaves the contribution of news about future returns virtually unchanged. In any case, as already mentioned, the existence of significant residual serial correlation in some of these models should make one very wary of drawing inference from them.

Unfortunately, the low frequency of the data makes rigorous sensitivity analysis with regard to sample period very difficult. However, I did re-estimate the reported VAR models over the post-war period (end-1945 to 1993) to study the effects on the variance decomposition. I found no major changes in the results. Concentrating on the results for the models which include both D/P and V in the VAR, the return-regression adjusted R^2's increased substantially to 57.20% in the case of real returns, and to 59.82% for excess returns[18]. The expected return persistence estimates fell slightly from 1.95 (Table 3.8) to 1.52 for real returns, and from 1.2 (Table 3.9) to 0.99 for excess returns. However, the contribution of news about future real returns of 1.05 (Table 3.8, row 1) changed very little, increasing slightly to 1.18 (s.e.=0.40) in the post-war period. For excess returns, Table 3.8 row 2 gives 0.62 for the contribution of future excess returns, which again increases slightly to 0.78 (s.e.=0.26) in the post-war period. Consequently, the results appear to be robust to the exclusion of the earliest points in the data set.

3.12 *Direct Measurement of Dividend Surprises*

In the preceding analysis, news about future dividends has been treated as the residual component of the unexpected stock return, after having calculated the component due to shocks to expected future returns. However, it is possible to obtain a more "direct" cash flow component by including real dividend growth as a variable in the VAR. In such a case, lagged real dividend growth is allowed to affect expected stock returns, whilst the residual from the dividend growth equation provides a measure of dividend surprises. The latter may be entered directly into the discounting formula for the cash flow component.

---

[18] Campbell (1991) also found a marked improvement in return predictability after excluding the earliest data points.

More formally, define $\mathbf{z}_{t+1}$ to include the real return h_{t+1} as the first element, and Δd_{t+1} (the change in the natural logarithm of real dividends) as the second. The following expression for the portion of the unexpected real stock return due to news about dividends is then easily derived:

$$\eta_{t+1} \equiv (E_{t+1} - E_t)\sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} = \iota_2'(\mathbf{I} - \rho\mathbf{A})^{-1}\mathbf{w}_{t+1}$$

When $\mathbf{z}_{t+1}$ is defined to include the dividend-price ratio and the volatility measure V, the return regression R^2 is 45.47%, the persistence measure estimate is 1.59 (s.e.=0.35), and the contribution of news about future returns is 0.82 (s.e.=0.21). All of these are comparable to the results reported in row 1 of table 6, whereΔd_{t+1} was omitted. The residual and direct measures of cash flow surprises have a correlation of 0.93, and neither measure gives a significant covariance between return news and dividend news. The contribution of news about future dividends is 0.35 (s.e.=0.17) for residual dividends and that for direct dividend is 0.48 (s.e.=0.64), which are also broadly comparable.

3.13 *Previous Findings from Variance Decompositions*

Campbell (1991) applied the above analysis to monthly observations of the NYSE index from 1926 to 1988. He focused mainly on one VAR specification, which included the real return, the dividend-price ratio and the "relative" interest rate (the short-term interest rate minus a one year moving average of short rates, which effects a stochastic detrending of the possibly non-stationary short-term interest rate). He obtained a very low return equation R^2 (2.4%), and no cross effects were found between the dividend-price ratio and the relative interest rate, so that they could have been well modelled by univariate processes. This is in sharp contrast to my findings using the dividend-price ratio and volatility, where interdependencies were complex and significant.

Campbell found that just over one third of the variance of unexpected returns was attributable to news about future dividends, and just under one third was attributable to news about future returns. The correlation between the two types of news was around -0.5 and statistically significant, implying that good news about cash flows was associated with a decline in expected future returns. Also, Campbell's estimate of the returns expectations persistence statistic implied that a 1% positive innovation to expected returns led to a 5% current capital loss. Finally, for excess returns P_r was relatively small (1.64) compared to P_h (4.38), and the contribution of news about future interest rates was very small, with small insignificant covariances with the two other news sources. All of these findings contrast sharply with those reported above using annual UK data.

In common with the current study, Campbell found that the dividend-price ratio was the critical variable: without it, the variance of news about future returns is less important, cash flow news is more important, and the covariance term becomes insignificant. Also, the AR(1) process for expected returns is not supported by the data.

3.14 *Discussion and Conclusions*

I have modelled the behaviour of returns in a multivariate VAR framework. Using a log-linear version of the RVF, I apportioned unexpected changes in real returns between news about future returns (discount rates), news about future dividends and any covariance between the two. This approach enables one to avoid some of the limitations of univariate models and single-equation return regressions. In particular, one can include many variables that might affect returns and work out the implications for stock prices of both predictability and expectations persistence. I find that most of the variance in unexpected real stock returns is due to news about future expected returns, with no significant covariance between the two. Hence models of stock price movements must be able to take account of the short-term predictability and persistence in returns.

My results are invariant to the inclusion of alternative variables in the equation to explain expected returns, as long as the dividend-price ratio is included. Exclusion of the latter variable results in an equation for expected returns with a very low R^2, with the result that the statistics of interest are poorly determined. The results are also robust to alternative VAR lag lengths and for the post-1945 sample period.

When examining *excess* returns in a VAR framework I again find that news about future returns is important in explaining the volatility in one-period excess returns. However, in addition there is a strong positive covariance between news about dividends and news about real interest rates. The strength of this effect is such that, in my data set, the *net* effect on the volatility of returns of news about real interest rates and dividends is close to zero. This is because positive news about dividends is accompanied by positive news about real interest rates and, via the RVF, these have opposite effects on unexpected returns.

My VAR models (although not the theoretical framework of the RVF) are empirically based, yet my results clearly imply that it is the persistence (as well as some predictability) in expected returns that is important in explaining movements in stock prices. Whether such persistence can be explained in a coherent theoretical model is now being actively pursued in the literature (e.g. Cochrane 1991, Campbell and Cochrane 1994).

On the issue of volatility and volatility persistence, the addition of the squared ex post real return to the VAR does significantly increase explanatory power of the returns equation when the dividend-price ratio is present. The dividend-price ratio also appears to be a significant predictor of volatility, opening up the possibility of a plausible explanation for the significance of the dividend-price ratio in this data set. Although I do not have a direct measure of volatility persistence, evidence from the estimated coefficients on the lagged values of volatility, together with the observed persistence in the dividend-price ratio imply that expected volatility is indeed persistent. The results from the volatility model imply that a 1% shock to expected returns is associated with a current capital loss of around 2%.

The key to all of the above results is the quality of the time-series decomposition afforded by the VAR. Different specifications can lead to different inferences, and as long as the chosen variables are devoid of a theoretical base, the economic meaning of my results is open to challenge. In particular, if the performance of the predictive variables cannot be rationalised in a formal asset pricing framework, return predictability could be an indication that the stock market is in fact inefficient, and my concentration on the implications of the rational valuation formula may be misguided. However, if one takes the view that the market is in fact efficient, in the absence of theoretical guidance as to which variables are actually proxies for expected returns, the multivariate time-series decomposition combined with the log-linear RVF provides a useful and tractable framework within which to study stock price movements.

The question of whether the predictability found in the VAR systems is consistent with popular asset pricing models is the central issue of the empirical analysis in Chapter 5.

Chapter 4: IDENTIFYING SOURCES OF SYSTEMATIC RISK IN THE UK STOCK MARKET

4.1 Introduction

According to the Linear Factor Model (Burmeister and McElroy 1988)[1], the unexpected excess return on any asset depends on a set of factor innovations and their factor loadings (betas), plus an idiosyncratic innovation. The factor betas measure the extent to which investors adjust their required risk premium on an asset in response to news about non-diversifiable (systematic) risk. Since the theory gives no indication of the likely identity of the factors, the latter have been determined empirically, and previous researchers have used variables such as real output growth, the real exchange rate, interest rates and the dividend-price ratio (see, for example, Clare and Thomas 1994, Chen 1991, Chen, Roll and Ross 1986).

Using a log-linear version of the rational valuation formula (RVF), Campbell and Mei (1993) demonstrated that factor innovations can impact upon required excess returns in three ways: by affecting expectations of future dividend payments; by affecting expectations about future real interest rates; and by affecting future risk premia. In the APT literature, the focus has been on testing which factors are priced, rather than the channels through which the factors impact upon an asset's systematic risk. In this chapter, I combine the cross-sectional analysis of multifactor models with the fundamental analysis of the RVF. I take a number of macroeconomic and financial factors and attempt to determine how each factor impinges on expectations of future dividends, future real interest rates and future excess returns. I am then able to ascertain how an asset's factor betas are determined by covariances between fundamentals and macroeconomic risks.

My methodology is quite straightforward. Excess returns on a set of UK stock portfolios are assumed to depend linearly on a set of state variables. The latter are modelled as a VAR process. I am then able to decompose the excess stock returns and the state variables into expected and unexpected components. The unexpected portions of the state variables are taken to be factor innovations, and the factor betas are estimated as scaled covariances between the unexpected excess asset returns and these factor innovations. In order to study the precise source of factor risk, vis-à-vis any particular portfolio, the Campbell-Shiller (1988, 1989) log-linearisation of the rational valuation formula is used to decompose estimated betas into components

---

1 The Arbitrage Pricing Theory (Ross 1976) is a special case of the Linear Factor Model.

due to news about future dividends, news about future real interest rates, and news about future excess returns (risk premia).

The rest of this chapter is organised as follows. In Section 4.2 I discuss the restrictions placed on the LFM by the APT in order to obtain a unique expression for the expected return on any asset. In Section 4.3 I review a series of papers that has attempted to identify empirically which factors are actually priced in financial markets. The particular form of the Campbell-Shiller log-linearisation used in this analysis is introduced in Section 4.4, and Section 4.5 shows how this effects the decomposition of asset-i's beta with factor k. Section 4.6 outlines the econometric methodology, and 4.7 introduces the data and reports some sample statistics. Sections 4.8 to 4.13 present the empirical results of the beta decomposition for the various factors used. In particular, Section 4.8 considers the roll of the market excess return, and presents a formal and informal test of the CAPM, whilst 4.10 addresses the issue of the relationship between stock returns and inflation. Section 4.14 considers the robustness of the findings to changes in the VAR lag length, and 4.15 concludes.

4.2 Multi-Factor Models: The Linear Factor Model and the APT

The CAPM states that if a portfolio of assets is mean-variance efficient, the expected return on the portfolio depends on one factor only: the expected covariance between the portfolio return and the return on the market portfolio. The APT, on the other hand, allows expected asset returns to depend on more than one factor. Moreover, given the factor generating model, the absence of riskless arbitrage profits leads immediately to the APT prediction that the expected return on any asset is a linear combination of that asset's sensitivities to the various factor risks. The CAPM turns out to be a special case of the APT.

Burmeister and McElroy (1988) assumed that the return on asset i, R_{it}, is generated by a Linear Factor Model:

$$(4.1)\quad R_{it} = E(R_{it}) + \sum_{k=1}^{K} \beta_{ik}\tilde{f}_{kt} + \varepsilon_{it}$$

where $E(\varepsilon_{it}|\tilde{f}_{kt}) = E(\tilde{f}_{kt}) = E(\varepsilon_{it}) = 0$. Thus, the unexpected return depends on K factor innovations, $\tilde{f}_{kt}$, times their factor loadings (more usually termed "betas"), plus an idiosyncratic innovation, ε_{it}. Apart from its linearity, the LFM represents a very general statement of asset returns, giving no insight into the nature of the relevant factors, nor the form of the expected return. The Arbitrage Pricing Theory (APT) of Ross (1976) places restrictions on the LFM which give rise to a unique expression for the expected return on asset i. For ease of exposition, suppose that we have two factors, so that the LFM gives the following expression for R_{it}:

$$(4.2) \quad R_{it} = E(R_{it}) + \beta_{i1}\tilde{f}_{1t} + \beta_{i2}\tilde{f}_{2t} + \varepsilon_{it}$$

The most important maintained assumption in the APT is the absence of opportunities for profitable arbitrage. However, in order to derive mathematically an expression for the expected return on an asset, it is useful to construct a portfolio of assets as follows. The APT requires that there are enough assets in the market for an investor to form a portfolio with the following characteristics:

$$(4.3a) \quad \sum_{i=1}^{N} w_i = 0$$

$$(4.3b) \quad \sum_{i=1}^{N} w_i\beta_{i1} = 0, \quad \sum_{i=1}^{N} w_i\beta_{i2} = 0$$

$$(4.3c) \quad \sum_{i=1}^{N} w_i\varepsilon_{it} \approx 0$$

where w_i is the proportion of wealth invested in asset i. Condition (4.3a) states that this portfolio of N assets involves zero investment. Equations (4.3b) and (4.3c) imply that the portfolio is riskless. If profitable arbitrage opportunities are absent, an investment which involves no expenditure and no risk must have an expected return of zero. That is, the above conditions imply that

$$(4.4) \quad \sum_{i=1}^{N} w_i E(R_{it}) = 0$$

In the terminology of linear algebra, (4.3a) states that the vector of N asset proportions is orthogonal to a vector of ones; (4.3b) states that the vector of asset proportions is orthogonal to the vector of betas; and these imply that the vector of asset proportions is orthogonal to the vector of expected returns, which is (4.4). A well-known theorem of linear algebra states that, if the fact that a vector is orthogonal to N-1 vectors implies that it is orthogonal to the Nth vector, then the Nth vector may be written as a linear combination of the N-1 vectors. In the current case, the implication is that the vector of expected returns may be written as a linear combination of a vector of ones and the vector of asset betas. Thus, given the above conditions, one may write the expected return on any asset i as a constant times one, plus the sum of a set of constants times the betas. In our two-factor model, the expected return on asset i may be written as

$$(4.5) \quad E(R_i) = \lambda_0 + \lambda_1\beta_{i1} + \lambda_2\beta_{i2}$$

This expression holds for all assets and all portfolios of assets. Thus the expected return for any asset i differs to that for asset j only because the β_{ik} differ from β_{jk}: the λ_k's are identical across assets. It is straightforward to show that λ_0 is equal to the risk-free rate of interest, whilst each remaining λ_k is the excess return (risk premium) on a portfolio which has a beta of unity with factor k.

The APT gives no indication of the number or identity of risk factors relevant to asset pricing. Consequently, two branches of research have developed. The first, "factor analysis", attempts to determine the *number* of factors relevant to the pricing of a particular set of securities, without putting precise interpretations on the factors so identified[2]. The second branch involves the pre-specification of a set of variables which *a priori* might be expected to influence asset returns. The analysis then involves testing whether, in fact, each of these factors has a significant risk premium (λ) across a set of assets. It is this latter body of research that is relevant to the current study, and which I now review.

4.3 Identification of Observable Factors

Chen, Roll and Ross (1986) (CRR) were the first to address directly the question of the likely identity of the factors in a multi-factor asset pricing model.

---

2 See, *inter alia*, Roll and Ross (1980) and Beenstock and Chan (1986).

Within the confines of the APT, they took a number of macroeconomic time series and applied single-equation models to each to obtain series of unanticipated movements, which they took to measure the factor innovations, $\tilde{f}_k$. They estimated the risk premia associated with each factor (the λ_k's) as follows. The betas were estimated by a multiple time-series regression of the portfolio returns on the factor innovations over a five-year period. These estimates were then used as the *independent* variables in 12 cross-sectional regressions, one regression for each of the next 12 months. The coefficients from the latter regressions provided estimates of the risk premium associated with each factor. This two-step procedure was repeated for each year in the sample, yielding a time series for each factor risk premium, λ_{kt}. The question of whether each of the factors was priced by the stock market was translated into a test of whether the means of the λ_{kt} were significantly different from zero.

Using monthly observations on 20 stock portfolios grouped by market capitalisation, over the period 1958-84, CRR found that industrial production, unexpected inflation and the spread between government bonds and BAA-grade corporate bonds were significantly priced. The risk premium for consumption growth was found to be insignificant and had the "wrong sign".

Chen (1991) studied the relationships between various macroeconomic variables and asset returns in a more informal setting. His main hypothesis was as follows. The expected market premium on any asset is negatively related to the recent growth of economic activity (which he used as a proxy for the current state of the economy), and positively related to the expected future growth of economic activity and its conditional variance. He therefore took a set of variables commonly found to forecast stock returns and related them to various measures of past and future real activity. Using quarterly data from 1954(1) to 1986(4), Chen found that the market dividend yield and a measure of the default premium were indicators of the current and near-future economic state. In particular, an above average dividend yield and default spread observed at the end of quarter t-1 indicate that the growth of GNP over the preceding four quarters has been lower than average, and growth over the next two quarters was expected to be lower than average. The current short-term interest rate, the current term structure and the lagged industrial production growth rate forecast changes in the future growth of GNP beyond the power of the dividend yield and the default spread. Finally, Chen presented some evidence that the above state variables (except for the term spread) were related to the conditional variance of the GNP growth rate, although the results were not statistically strong.

There have been several studies of risk factors in the UK market. Beenstock and Chan (1988), using a similar methodology to CRR, found that four factors were significantly priced: a short-term interest rate, fuel and materials costs, the money supply and inflation. Poon and Taylor (1991) tested whether the factors found by CRR to affect the US market, were also significant in the UK. Using univariate ARIMA models to obtain factor innovations, they found that none of the CRR factors appeared significant in the UK. Clare and Thomas (1994), again using univariate time-series models to obtain factor innovations, identified six factors for the UK: oil prices, default risk, the RPI, private sector bank lending, the current account balance and the redemption yield on a corporate debenture index.

4.4 An Approximation to the Rational Valuation Formula

In this chapter I use the Campbell (1991) expression for the unexpected excess return presented in Section 3.4.1 above, i.e.

$$e_{it+1} - E_t e_{it+1} \approx (E_{t+1} - E_t)\left[\sum_{j=0}^{\infty} \rho_i^j \Delta d_{t+1+j} - \sum_{j=0}^{\infty} \rho_i^j r_{t+1+j} - \sum_{j=1}^{\infty} \rho_i^j e_{it+1+j}\right]$$

or, more simply,

$$(4.6) \quad \tilde{e}_i \approx \tilde{e}_{di} - \tilde{e}_{ri} - \tilde{e}_{ei}$$

where ρ_i is a number a little smaller than unity (see Section 4.7 below). The unexpected excess return on asset i, $\tilde{e}_i$, is approximately equal to the revision in expectations of the discounted present value of three components: future real dividend growth, $\tilde{e}_{di}$, future real interest rates, $\tilde{e}_{ri}$, and future excess returns on asset i, $\tilde{e}_{ei}$. Since expectations are revised in response to relevant news, we say that the unexpected one-period excess return depends on news about future dividends, news about future real interest rates and news about future excess returns.

4.5 A Decomposition of Asset Betas

The k-th factor loading of asset i is given by

$$(4.7a)\quad \beta_{i,k} \equiv \frac{Cov(\tilde{e}_i, \tilde{f}_k)}{Var(\tilde{f}_k)}$$

where $\tilde{f}_k$ is the innovation in the kth factor. $\beta_{i,k}$ is therefore the covariance between the unexpected excess return on asset i and the innovation in the kth factor, scaled by the variance of the factor innovation. Given equation (4.6), the factor beta (4.7a) can be decomposed into three elements:

$$(4.7b)\quad \beta_{i,k} = \frac{Cov(\tilde{e}_{di}, \tilde{f}_k)}{Var(\tilde{f}_k)} - \frac{Cov(\tilde{e}_{ir}, \tilde{f}_k)}{Var(\tilde{f}_k)} - \frac{Cov(\tilde{e}_{ei}, \tilde{f}_k)}{Var(\tilde{f}_k)}$$

$$(4.8)\quad \beta_{i,k} = \beta_{di,k} - \beta_{ri,k} - \beta_{ei,k}$$

The three terms on the RHS of equation (4.8) denote the betas between the kth-factor innovation and news about future dividends, news about future real interest rates and news about future asset-i excess returns, respectively. A factor beta $\beta_{i,k}$ will tend to be larger, the greater is the covariance between factor surprises and revisions to future expected dividends, $\beta_{di,k}$, and smaller, the greater is the covariance between the factor innovation and revisions in expectations of future real interest rates, $\beta_{ri,k}$, and asset-i excess returns, $\beta_{ei,k}$. For example, suppose that a factor, such as real output growth, has a large positive dividend beta. Then a negative shock to output growth implies large revisions to expected dividends in the same direction, which will lead, ceteris paribus, to a sizeable fall in the current price of the asset. Investors will therefore require a large current excess return to compensate for this factor risk. However, if negative surprises in output growth also lead to downward revisions to expectations of future real interest rates and future excess returns (i.e. $\beta_{ri,k}$ and $\beta_{ei,k}$ are positive), then this will attenuate the dividend effect. The current stock price will therefore be less vulnerable to innovations in output growth, and the required premium to compensate for this factor news will be smaller. Thus the overall contribution of a factor to the required return on any asset depends in an intuitive way on the relative sizes of the three "fundamental" beta components on the RHS of equation (4.8).

4.6 Econometric Implementation

I postulate that the expected excess return on each of the stock portfolios considered is linear in a set of variables, $\mathbf{x}_t$, which are known to the market at the end of period t, and which summarise the known current economic state[3]. Given a vector of returns on s sector portfolios, $\mathbf{e}_{t+1}$, this model can be written as

(4.9) $$\mathbf{e}_{t+1} = \mathbf{A}\mathbf{x}_t + \tilde{\mathbf{e}}_{t+1}$$

The state-variable vector is assumed to follow a first-order VAR[4]:

(4.10) $$\mathbf{x}_{t+1} = \Pi\mathbf{x}_t + \tilde{\mathbf{x}}_{t+1}$$

so that

$$\mathrm{E}_t\mathbf{x}_{t+j} = \Pi^j\mathbf{x}_t$$

and

$$(\mathrm{E}_{t+1} - \mathrm{E}_t)\mathbf{x}_{t+j} = \Pi^j\tilde{\mathbf{x}}_t$$

Given these relationships, together with equation (4.9), and substituting in equation (4.6), the news components can be written as follows:

(4.11a) $$\tilde{e}_{di} = \tilde{e}_{it+1} + (\iota_r' + \rho_i a_i')(I - \rho_i\Pi)^{-1}\tilde{\mathbf{x}}_{t+1}$$

(4.11b) $$\tilde{e}_{ei} = \rho_i a_i'(I - \rho_i\Pi)^{-1}\tilde{\mathbf{x}}_{t+1}$$

(4.11c) $$\tilde{e}_r = \iota_r'(I - \rho_i\Pi)^{-1}\tilde{\mathbf{x}}_{t+1}$$

---

[3] Although the choice of state variables is largely arbitrary, the real interest rate must be included in $\mathbf{x}_t$ in order to estimate the real interest rate beta $\beta_{ri,k}$.

[4] The first-order VAR form here is not restrictive since a VAR of any order can be written as a first-order VAR (see Section 2.2.2).

where $\tilde{e}_i$ is the ith row of **e** and a_i it the ith row of **A**. ι_r is a selection vector which picks out the real interest rate equation from the VAR, i.e. $\iota_r' \mathbf{x}_{t+1} \equiv r_{t+1}$. The factor innovations are the residuals from the k individual VAR equations, i.e.

$$\text{(4.11d)} \qquad \tilde{f}_k = \tilde{x}_{kt+1}$$

where $\tilde{x}_{kt+1}$ is the kth row of $\tilde{\mathbf{x}}_{t+1}$. Having estimated (4.9) and (4.10), and obtained the variables in (4.11a)-(4.11d), it is straightforward to calculate the relevant variances and covariances, and hence the betas in (4.7a) and (4.8).

4.7 Sample Statistics

The data used in this analysis comprise end-of-month observations from January 1970 to January 1993 of returns on 27 UK industry-based portfolios from the FT Actuaries database. I have three very broad portfolios - Capital Goods, Consumer Goods and Financial Services - and 24 more specific portfolios.

I study six state variables: the aggregate market (FT500) excess return, the market dividend yield, the real 1-month TB rate, the inflation rate, the industrial production growth rate and the percentage change in the real Sterling Effective Exchange Rate[5]. All variables are defined as deviations from their mean. Results are reported initially using a VAR lag length of one. In Section 4.14 I discuss results when the VAR has a lag of three.

The linearisation constant, ρ_i, is estimated as $1/(1+\exp(\delta_i))$, where δ_i is the mean log dividend-price ratio on portfolio i (see Campbell 1991). Since the range of estimated ρ_i's across the portfolios is not very large (0.9946-0.9971), and the results are not sensitive to variations of ρ within a plausible range, I set ρ equal to 0.9958 for all portfolios. This corresponds to a mean market dividend-price ratio of 4.95%. From equation (6c), it is clear that the use of the same value of ρ_i for all portfolios restricts the impact of each factor innovation on revisions to expectations of future real interest rates ($\beta_{ri,k}$) to be the same across all portfolios.

Table 4.1 presents the contemporaneous correlations between the six state variables (upper triangular) and between the six corresponding factor innovations

---

[5] All monthly rates are expressed as percent per annum, except for the dividend-price ratio which is in basis points per annum.

Table 4.1: Correlations of Factors and Factor Innovations

	e_m[1]	D/P	r	ipg	infl	Δs
e_m		-0.1260	-0.1520	-0.0080	-0.1280	0.0700
D/P	-0.9410		-0.2580	-0.1090	0.3850	0.0440
r	-0.1280	0.1030		0.0030	-0.9430	-0.3380
ipg	-0.0080	-0.0030	-0.0370		-0.0230	-0.0880
infl	0.1000	-0.0790	-0.9950	0.0410		0.3710
Δs	0.0890	-0.0890	-0.4440	-0.0290	0.4210	

[1] The upper triangular contains the contemporaneous correlations between the factor variables, whilst the lower triangular contains contemporaneous correlations between factor *innovations*. The variables are, respectively: the market excess return, e_m; the market dividend-price ratio, D/P; the real interest rate, r; industrial production growth, ipg; *inflation, infl*; *and the percentage* change in the real exchange rate, Δs.

(lower triangular). Four observations are worth making. First, although there is only a small correlation between the market dividend-price ratio and the market excess return (-0.126), the correlation between *innovations* in these two variables is quite substantial (-0.941). Second, one observes a large negative correlation between innovations in inflation and innovations in the real interest rate (-0.995). Third, there is some correlation between innovations in the percentage change in the real exchange rate and innovations in inflation and real interest rates (0.421 and -0.444 respectively). Finally, neither the expected nor the unexpected industrial production growth series is highly correlated with any of the other variables.

4.8 Market Betas

The popularity of the single-factor CAPM naturally generates particular interest in the impact of the market excess return on the required return on asset i. The market beta $\beta_{i,m}$ measures the change in the required return on asset i associated with a 1% expected change in the market excess return.

The first column of numbers in Table 4.2 gives the market beta point estimates with estimated standard errors in parentheses[6,7]. Not surprisingly, all of the market betas are significantly positive, so that positive news about the market return is associated with an upward revision in investors' required return on all of the industry portfolios. Although only seven of the stocks analysed have an estimated market beta beyond two standard errors from unity, the point estimates imply potentially large deviations in risk premia. Considering two extreme cases, the point estimate of the market beta for Shipping and Transport is 0.8387, whilst that for Contracting and Construction is 1.2457. Thus every 1% of the market price of market risk places a 0.4% wedge between the risk premia of these two assets. If one takes the unconditional mean excess return on the market of 6% as a proxy for the conditional expected market excess return, the implication is that the required return on the Contracting and Construction portfolio is around 2.5% higher than that for Shipping and Transport.

---

6 All estimated standard errors are corrected for an unknown form of heteroscedasticity using White's (1984) estimator.

7 In all of the tables that present results of the beta decompositions, estimated betas that are beyond two standard errors from zero are emboldened.

Table 4.2: Decomposition of Market Beta

	$\beta_{l,m}$	$\beta_{dl,m}$	$\beta_{el,m}$	$\beta_{dl,dm}$	R^2
Capital Goods	**1.0789** (0.0207)	0.0864 (0.0891)	**-1.1580** (0.1005)	**1.0324** (0.0717)	10.02
Consumer Goods	**0.9937** (0.0149)	**0.1996** (0.0660)	**-0.9597** (0.0833)	**1.0255** (0.0399)	8.40
Financial Services	**1.0382** (0.0239)	**0.2152** (0.0732)	**-0.9885** (0.0955)	**1.1807** (0.0800)	8.12
Building Materials	**1.2125** (0.0377)	**0.1901** (0.0815)	**-1.1880** (0.1060)	**1.1266** (0.1086)	7.90
Contracting and Construction	**1.2457** (0.0429)	-0.0937 (0.1012)	**-1.5050** (0.1373)	**0.8439** (0.1589)	11.52
Electricals	**1.0453** (0.0354)	-0.0161 (0.0818)	**-1.2269** (0.0937)	**0.8754** (0.1097)	7.49
Aerospace	**0.9987** (0.0342)	0.1077 (0.1011)	**-1.0566** (0.1001)	**0.9848** (0.1034)	9.06
Engineering (General)	**1.0273** (0.0347)	0.1489 (0.1042)	**-1.0440** (0.0994)	**1.1418** (0.0858)	9.95
Metals and Metal Forming	**1.0307** (0.0465)	**0.4460** (0.1012)	**-0.7502** (0.0854)	**1.4841** (0.1029)	6.43
Motors	**0.9710** (0.0927)	**0.4502** (0.1476)	**-0.6864** (0.0825)	**1.8186** (0.1123)	10.71
Brewers and Distillers	**0.9642** (0.0346)	**0.2478** (0.0640)	**-0.8819** (0.0820)	**0.9477** (0.0889)	5.36
Food Manufacturing	**0.9423** (0.0344)	0.1051 (0.0621)	**-1.0027** (0.0970)	**0.8356** (0.0925)	9.18
Food Retailing	**0.9907** (0.0403)	-0.0957 (0.0668)	**-1.2520** (0.0953)	**0.7187** (0.1253)	10.33
Hotels and Leisure	**1.1141** (0.0305)	0.0446 (0.0977)	**-1.2350** (0.1173)	**0.9659** (0.1309)	9.42
Packaging, Paper and Printing	**1.0306** (0.0339)	**0.2405** (0.0778)	**-0.9557** (0.0752)	**1.1732** (0.0745)	7.25
Stores	**1.1024** (0.0420)	0.1175 (0.0850)	**-1.1504** (0.0906)	**0.9412** (0.1237)	8.28
Textiles	**1.0470** (0.0368)	**0.4659** (0.0773)	**-0.7466** (0.0778)	**1.3505** (0.1236)	4.35
Chemicals	**0.9556** (0.0323)	**0.1896** (0.0755)	**-0.9315** (0.0850)	**0.9396** (0.0970)	5.77
Shipping and Transport	**0.8387** (0.0414)	**0.3363** (0.0664)	**-0.6680** (0.0569)	**1.1852** (0.1293)	5.81
Miscellaneous	**1.0420** (0.0368)	**0.3796** (0.0844)	**-0.8280** (0.0753)	**1.3347** (0.0983)	6.05
Industrials	**1.0128** (0.0099)	**0.1902** (0.0694)	**-0.9882** (0.0845)	**1.0575** (0.0223)	8.65
Oil and Gas	**0.9196** (0.0539)	0.0846 (0.0549)	**-1.0005** (0.0754)	**0.5789** (0.1245)	5.12
Banks	**1.0026** (0.0431)	**0.3466** (0.0792)	**-0.8216** (0.1050)	**1.2898** (0.1459)	6.47
Insurance	**1.0100** (0.0398)	**0.1561** (0.0775)	**-1.0194** (0.0794)	**1.0825** (0.1090)	5.42
Merchant Banks	**1.1901** (0.0622)	**0.4068** (0.0553)	**-0.9489** (0.0894)	**1.4301** (0.1624)	9.02
Property	**1.0385** (0.0532)	0.1422 (0.1195)	**-1.0598** (0.1146)	**1.1194** (0.1566)	7.33
Investment Trusts	**0.9906** (0.0362)	0.0179 (0.0690)	**-1.1383** (0.0965)	**0.8568** (0.0944)	8.71

The estimated interest rate beta, $\beta_{r,m}$, is **0.1656** (s.e.=0.0273).

For each portfolio i, in columns 2 and 3 I decompose the market beta $\beta_{i,m}$ into the components due to news about dividends $\beta_{di,m}$ and news about excess returns $\beta_{ei,m}$. The absolute value of the excess return component is always much larger than that of dividend news. This confirms the findings of Campbell (1991) and Campbell and Mei (1993) for the US, and those presented in Chapter 3 for the UK, that the major source of variation in stock prices arises from changing expectations about future returns, rather than from news about expected future dividends. This finding is consistent with results from the variance bounds literature (Shiller 1981, and see Chapter 5) that stock prices are found to be excessively volatile when only dividend expectations (and not discount rates) are allowed to vary through time.

With only two (statistically insignificant) exceptions, all of the dividend components are positive, so that an unexpected rise in the market excess return coincides with an upward revision in expected dividends on each asset, whilst all of the excess return effects are negative. Thus, a positive shock to the market return co-impacts with higher expected dividends and lower discount rates, so that there is a significant rise in the current asset-i stock price. This "explains" the origins of the significant upward revision in the current required returns indicated by the positive total market betas.

Cash flow (dividend) news is taken as a residual after the explicit modelling of expected excess returns and real interest rates (see equations (4.11a)-(4.11c)). However, this does not appear to influence the results in any systematic fashion since there is no obvious relation between the size of the cash-flow betas and the degree of return predictability. For example, Contracting and Construction has a relatively high degree of predictability (R^2=11.52%), contrasting with that for Oil and Gas (R^2=5.12%), but their cash-flow betas are both small and insignificant (-0.0937, s.e.=0.1012 and 0.0846, s.e.=0.0549 respectively). Also, the Motor industry returns are relatively forecastable (R^2=10.71%) and have a relatively large cash-flow component of 0.4502 (s.e.=0.1476), whilst Textiles have a similar estimated cash-flow beta of 0.4659 (s.e.=0.0773) but low return predictability (R^2=4.35%).

The fourth column of numbers in Table 4.2 presents estimates of the asset-i dividend expectations reaction to changes in expected *market* dividends. This provides a measure of cyclicality of the various industries. The closer is $\beta_{di,dm}$ to unity, the more closely expectations of dividends paid on asset i covary with aggregate market dividends. A number greater than unity indicates a high degree of cyclicality, with revisions to dividend expectations being greater than the market in times of optimism, and lower than the market when the overall outlook is poor. Conversely, the further is the number below unity, the more stable are dividend expectations compared with the market.

All of the portfolios show a significant positive movement of dividend expectations with expectations of market dividends, but there are marked differences between the extent of the covariation. For example, the Motor industry has a dividend beta of 1.8186, indicating that a 1% improvement in expectations of aggregate dividends is associated with an upward-revision of motor-industry dividends of nearly 2%, whilst in an economic downturn expectations for the motor industry are proportionately much worse than the overall outlook. This is precisely what one would expect from an industry for which demand is highly pro-cyclical. Financial Services (1.1807), Banks (1.2898), and Merchant Banks (1.4301) also significantly overreact to the market. On the other hand, Brewers and Distillers (0.9477), and Food Manufacturing (0.8356) and Retailing (0.7187) show significant "under-reaction", reflecting the extent to which demand for the products of these industries is relatively inelastic with respect to general economic conditions. Not surprisingly perhaps, the smallest reaction is in the Oil and Gas industry, with an estimated beta of 0.5789, so that movements in expected dividends in this sector are only a little more than half those expected for the market as a whole.

The CAPM

The above analysis can be used to test the one-factor CAPM representation of asset returns. The CAPM is a special case of the APT, which in turn can be shown to be nested within my (unrestricted) multifactor framework. The APT states that the expected excess return on any asset is given by the sum of K factor betas times the prices of K factor risks:

$$(4.12)\quad E_t e_{it+1} = \sum_{k=1}^{K} \beta_{i,k} \lambda_{kt}$$

where λ_{kt} is the *conditional* risk premium on factor k. If one assumes that the conditional factor risk premia, λ_{kt}, are linear in the vector of N (which may or may not equal K) state variable, $\mathbf{x}_t$, we have the following expression:

$$(4.13)\quad \lambda_{kt} = \sum_{n=1}^{N} \theta_{kn} x_{nt}$$

Combining equations (4.12) and (4.13), one can see that the coefficients of the **A** matrix in equation (4.9) are restricted as follows:

$$(4.14)\quad a_{in} = \sum_{k=1}^{K} \beta_{i,k}\theta_{kn}$$

That is, each coefficient on the N state variables is equal to the beta of asset i with factor k (which is different in each equation) times θ_{kn}, which differs across factors but not across equations (assets).

The CAPM is a one-factor model (K=1), and I have six state variables (N=6). With 27 portfolio returns, the unrestricted **A** matrix has 162 elements. There are several ways in which the number of free parameters can be reduced. If one assumes that asset returns follow a one-factor model, but that the factor is unobserved, (4.14) implies a set of restrictions on the parameters of the **A** matrix. For example, if there were only two assets and two state variables, the unrestricted equations are

$$e_{1t} = a_{11}x_{1t-1} + a_{12}x_{2t-1} + \tilde{e}_{1t}$$

$$e_{2t} = a_{21}x_{1t-1} + a_{22}x_{2t-1} + \tilde{e}_{2t}$$

whereas the equations when there is a single unobservable factor are

$$e_{1t} = \beta_1(\theta_1 x_{1t-1} + \theta_2 x_{2t-1}) + \tilde{e}_{1t}$$

$$e_{2t} = \beta_2(\theta_1 x_{1t-1} + \theta_2 x_{2t-1}) + \tilde{e}_{2t}$$

The unobservable single factor model therefore implies that

$$\frac{a_{11}}{a_{12}} = \frac{a_{21}}{a_{22}}$$

In my full system of 27 assets and six state variables, the fact that the θ_{1n}'s are constant across equations means that the system now has 33 parameters to be estimated. The likelihood ratio statistic which compares the unrestricted system with the unobserved one-factor model is found to be 144.85, which has a marginal probability value of 0.16. However, if one is willing to use the return on the FT500 index as the "market" return, then the number of parameters can be reduced further. Taking the expected market return as the fitted values from a regression of the market excess return on the lagged state variables, one can proceed in either of two ways. On the one hand, imposing the expected market return on the 27 portfolio return regressions (which effectively fixes the θ's), leaves only the 27 β's to be estimated. This produces an LR statistic of 159.22, which has a p-value of 0.08. Thus, the CAPM is rejected at the 10% level but not at the 5% level. One might, therefore, be willing tentatively to accept the CAPM simplification of the multi-factor model. However, some doubt may be cast on this conclusion by instead looking directly at the implications of the CAPM for the cross-section of portfolio returns. If the CAPM is correct then there is a linear relationship between asset i's market beta, $\beta_{i,m}$, and the beta of asset i's news about future cash flows and the market return, denoted $\beta_{di,m}$. Given my cross-section estimates of $\beta_{i,m}$ and $\beta_{di,m}$, I can compare these with the "theoretical" linear relationship implied by the CAPM.

The linear relationship between $\beta_{i,m}$ and $\beta_{di,m}$ can be derived as follows. The expected excess return on asset i is equal to beta times the expected market excess return:

$$(4.15)\quad E_t e_{it+j} = \beta_{im} E_t e_{mt+j}$$

where e_m is the market excess return. Combining equations (4.15) and (4.6), we have

$$\begin{aligned} \tilde{e}_{ei} &= (E_{t+1} - E_t)\sum_{j=1}^{\infty} \rho^j e_{it+1+j} \\ &= (E_{t+1} - E_t)\sum_{j=1}^{\infty} \rho^j E_{t+j}\beta_{im} e_{mt+1+j} \\ (4.16) \qquad &= \beta_{i,m}\tilde{e}_{em} \end{aligned}$$

where $\tilde{e}_{em}$ is news about future excess returns on the market. Substituting (4.16) in (4.7b), with factor k being the market return, gives

$$\beta_{i,m} = \frac{Cov(\tilde{e}_{di} - \tilde{e}_r, \tilde{e}_m)}{Var(\tilde{e}_m)} - \beta_{i,m}\frac{Cov(\tilde{e}_{em}, \tilde{e}_m)}{Var(\tilde{e}_m)}$$

which can be rewritten as

$$(4.17) \quad \beta_{i,m} = \alpha_0 + \alpha_1 \beta_{di,m} \qquad \text{where} \quad \alpha_0 \equiv -\frac{\beta_{r,m}}{1+\beta_{em,m}} \quad \alpha_1 \equiv \frac{1}{1+\beta_{em,m}}$$

where $\beta_{em,m}$ is the market beta of news about future excess returns on the market. From (4.17) it is clear that the CAPM implies a linear relationship between the dividend component of asset i's market beta and asset i's total market beta.

I estimate $\beta_{r,m}$ to be 0.1656, and $\beta_{em,m}$ to be -0.9869, giving the theoretical CAPM line an intercept and slope of -12.54 and 75.76 respectively. In Figure 4.1 the theoretical linear relationship (4.17) is plotted as a solid line. The CAPM predicts that the larger is an asset's dividend beta, the larger will be its total market beta. The question now is how well does this prediction fare when confronted with my estimates of ($\beta_{i,m}$, $\beta_{di,m}$) pairs.

The triangles in Figure 4.1 plot the ($\beta_{i,m}$, $\beta_{di,m}$) coordinates for each of the 27 assets in my data set, and the dotted line is an OLS line of best fit through these points. Quite clearly the predictions of the CAPM do not match my findings. The estimated line is in fact nearly horizontal, with a small positive intercept. Thus, although there is a considerable range of dividend betas$\beta_{di,m}$ (see Table 4.2, column 2 and Figure 4.1), the $\beta_{em,m}$ component appears to attenuate rather than amplify the relationship between $\beta_{i,m}$ and $\beta_{di,m}$. Thus there appear little support for the CAPM within the maintained hypothesis of the multifactor model.

Figure 4.1: Theoretical CAPM Line Versus Estimated Relationship

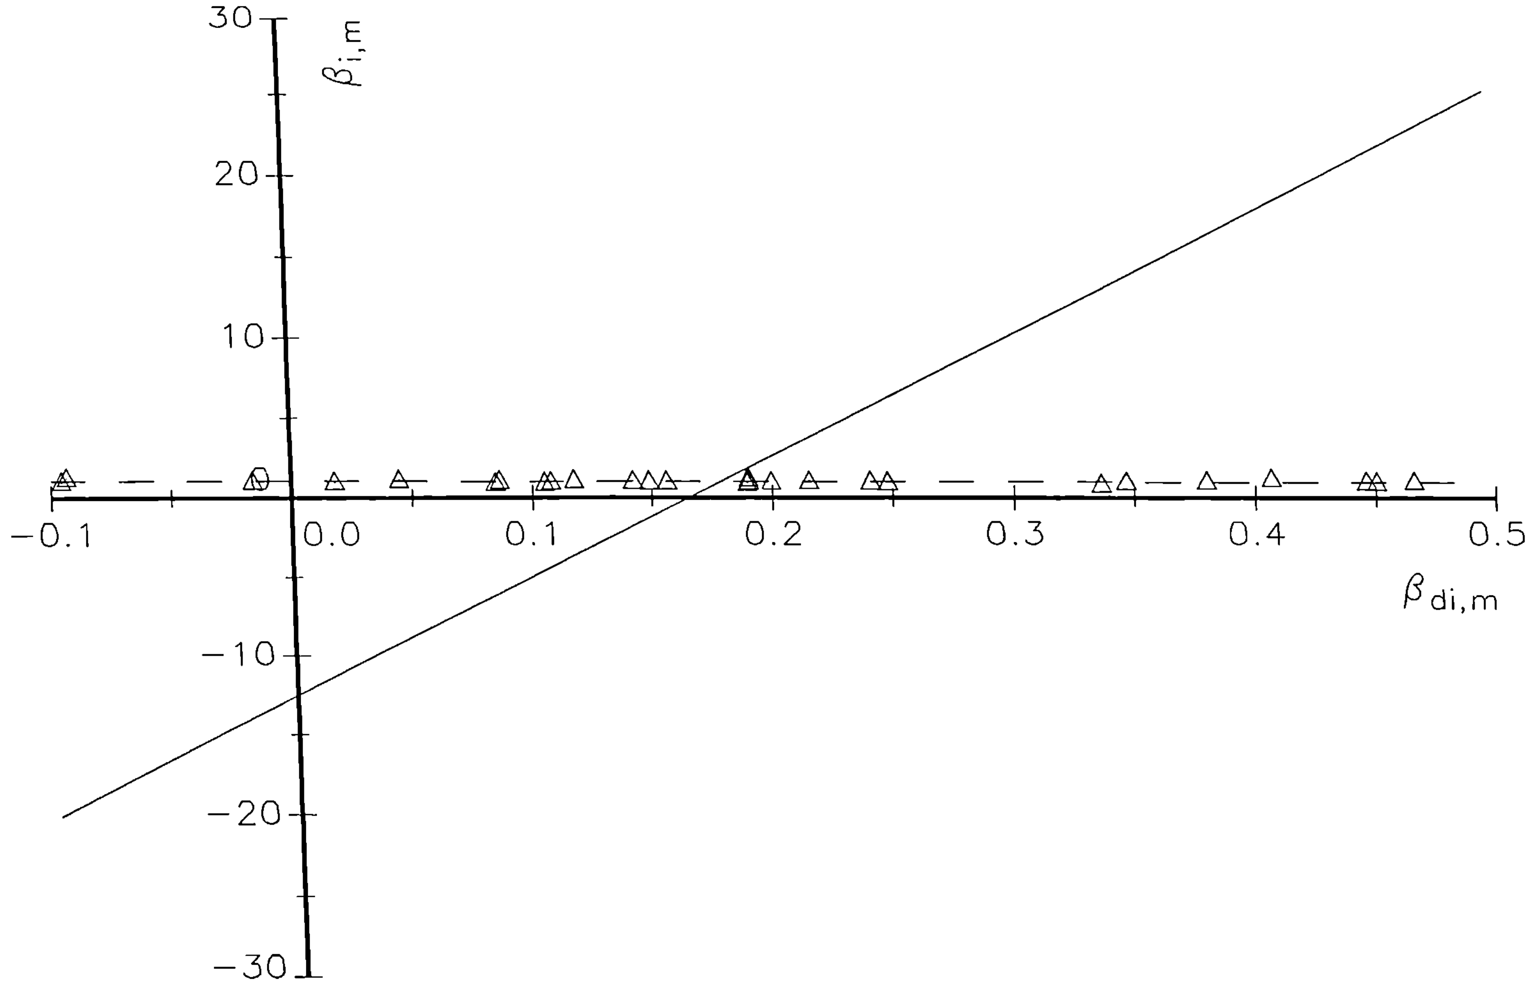

4.9 Dividend-Price Ratio Betas

The dividend-price ratio has been widely used in predicting stock returns. In this section I consider the covariance between the unexpected return on asset i and the innovation in the dividend-price ratio, $\beta_{i,D/P}$, and the relative contribution to the latter from the impact of D/P on future cash flows, $\beta_{di,D/P}$, future real interest rates, $\beta_{r,D/P}$, and future returns, $\beta_{ei,D/P}$.

Since the findings of Fama and French (1988b) (see Section 3.2.2), the dividend-price ratio has become pervasive as a proxy for required returns. Also, recall that Chen, Roll and Ross (1986) presented evidence that an above average dividend-price ratio forecasts lower-than-average GNP growth in the near future. Fama and French (1989) suggested that the dividend-price ratio might track the counter-cyclical component of expected returns: expected returns rise when the economic outlook is poor, and the latter is indicated by a higher dividend-price ratio. Although no comparable study has been done using UK data, I believe that there is sufficient interest in this variable to include it in the set of factors, and one might bear the US findings in mind when interpreting our results. In addition, since innovations in the dividend-price ratio appear not to be highly correlated with any of the other macroeconomic variables considered below (see Table 4.1), it may provide useful incremental information about return predictability.

The estimates in Table 4.3 do not fit in well with the US findings. All of the estimated dividend-price ratio total betas are significantly negative: a positive shock to the dividend-price ratio corresponds with a reduction in the current required return. The effect arises mainly from the future expected excess returns beta, $\beta_{ei,D/P}$, which is significantly positive in all cases. Given the present value relation, this is to be expected, since higher future excess returns require a fall in the current price level, and the latter will, ceteris paribus, cause the dividend-price ratio to rise. If this consistency effect is not to dominate, $\beta_{ei,D/P}$ needs to be outweighed by a large positive dividend beta and/or a large negative real interest rate beta. However, whilst the interest rate beta is indeed significantly negative (-0.3515, s.e.=0.0333), there is no consistent pattern in the signs of the dividend betas. Thus, no general statement about a link between shocks to the dividend-price ratio and revisions in expectations of future dividends/output can be made.

Table 4.3: Decomposition of Dividend Yield Beta

	$\beta_{l,D/P}$	$\beta_{d,D/P}$	$\beta_{ei,D/P}$
Capital Goods	**-1.8927** (0.1582)	0.0883 (0.1576)	**2.3324** (0.0240)
Consumer Goods	**-1.7577** (0.1410)	-0.1802 (0.1400)	**1.9290** (0.0250)
Financial Services	**-1.8420** (0.1351)	-0.1886 (0.1517)	**2.0049** (0.0497)
Building Materials	**-2.1462** (0.1446)	-0.0999 (0.1511)	**2.3977** (0.8602)
Contracting and Construction	**-2.2358** (0.1330)	**0.4605** (0.1480)	**3.0477** (0.0548)
Electricals	**-1.8473** (0.1516)	0.2418 (0.1308)	**2.4407** (0.0128)
Aerospace	**-1.7407** (0.1752)	0.0550 (0.1878)	**2.1472** (0.0415)
Engineering (General)	**-1.7829** (0.1867)	-0.0096 (0.1994)	**2.1248** (0.0385)
Metals and Metal Forming	**-1.7512** (0.2047)	**-0.5487** (0.2332)	**1.5540** (0.0593)
Motors	**-1.5817** (0.3015)	-0.4918 (0.3309)	**1.4414** (0.0549)
Brewers and Distillers	**-1.7137** (0.1299)	**-0.2840** (0.1387)	**1.7812** (0.0384)
Food Manufacturing	**-1.6945** (0.1124)	-0.0103 (0.1264)	**2.0357** (0.0487)
Food Retailing	**-1.7884** (0.1132)	**0.3483** (0.0961)	**2.4883** (0.0094)
Hotels and Leisure	**-1.9740** (0.1610)	0.1793 (0.1799)	**2.5048** (0.0540)
Packaging, Paper and Printing	-1.7821 (0.1745)	-0.2286 (0.1569)	**1.9050** (0.0135)
Stores	**-1.9648** (0.1865)	-0.0255 (0.1716)	**2.2908** (0.0145)
Textiles	**-1.8129** (0.1627)	**-0.6341** (0.1799)	**1.5303** (0.0457)
Chemicals	**-1.6810** (0.1420)	-0.1493 (0.1445)	**1.8833** (0.0330)
Shipping and Transport	**-1.4754** (0.1475)	**-0.4841** (0.1417)	**1.3428** (0.0132)
Miscellaneous	**-1.8103** (0.1910)	**-0.4867** (0.1942)	**1.6751** (0.0241)
Industrials	**-1.7843** (0.1452)	-0.1497 (0.1410)	**1.9861** (0.0192)
Oil and Gas	**-1.6719** (0.0997)	-0.0361 (0.0979)	**1.9873** (0.0194)
Banks	**-1.8027** (0.1254)	**-0.4417** (0.1747)	**1.7125** (0.0954)
Insurance	**-1.7605** (0.1676)	-0.0875 (0.1522)	**2.0245** (0.0169)
Merchant Banks	**-2.1563** (0.1185)	**-0.5755** (0.1287)	**1.9322** (0.0283)
Property	**-1.8133** (0.1949)	0.0043 (0.2350)	**2.1691** (0.0830)
Investment Trusts	**-1.7785** (0.1188)	0.1531 (0.1158)	**2.2831** (0.0235)

The estimated interest rate beta, $\beta_{r,D/P}$, is **-0.3515** (s.e.=0.0333).

4.10 Stock Returns and Inflation

Investigating the relation between stock returns and inflation has become a subject of interest in its own right. The inclusion of inflation in the multifactor model allows me to analyse several issues relating to this field of study. Central to this analysis is the so-called Fisher hypothesis.

4.10.1 The Fisher Hypothesis

The Fisher hypothesis states that expected nominal rates of return on assets should move one-for-one with expected inflation. This stems from the view that the real sector of the economy is causally independent of the monetary sector. Consequently, real returns depend only on real factors, such as productivity, risk aversion and time preference. The Fisher hypothesis is usually taken to imply that $\beta=1$ in the following regression

$$R_{it+1}=\alpha+\beta E_t\omega_{t+1}+\varepsilon_{t+1}$$

where R_{it+1} is the nominal rate of return on asset i and ω_{t+1} is the rate of inflation. As far as the stock market is concerned, for a long time the consensus opinion was that since stocks are claims on real assets, a stock portfolio was an effective hedge against inflation. However, in the mid-1970's a series of empirical studies cast substantial doubt on the Fisher hypothesis and the degree to which stocks actually were inflation-proof.

4.10.2 Early Empirical Evidence

Tests of the Fisher hypothesis initially involved regressing the nominal stock return on inflation. Using monthly NYSE data spanning 1953(1)-1971(2), Jaffe and Mandelker (1976) found a significant negative relationship between the two, although using the Standard and Poor annual data (1876-1970), the relationship was insignificantly different from zero. Nelson (1976) presented similar results using the Scholes Index and the S&P 500. However, the Fisher hypothesis relates asset returns to expected, rather than ex post, inflation, and Nelson pointed out that even if the Fisher hypothesis were true, using actual rather than expected inflation could lead to an erroneous rejection of the model. To see this, if R_t is the nominal stock return, and ω is the rate of inflation, one may write the two as follows:

$$R_t=E_t(R_t)+\varepsilon_t$$

$$\omega_t=E_t(\omega_t)+u_t$$

Separating the expected real return H_t into its unconditional mean, H, and deviations from the mean, H_t', the relationship between nominal returns and inflation can be written as

$$R_t=H+\beta\omega_t+(H_t'+\varepsilon_t-\beta u_t)$$

The Fisher hypothesis implies that H_t' and ω_t are uncorrelated, and so the probability limit of the OLS estimator of β is

$$(4.18)\quad \text{p}\lim\hat{\beta}=1-\frac{V(u)}{V(\omega)}+\frac{Cov(\varepsilon,u)}{V(\omega)}$$

The second term on the RHS of (4.18) is the effect of measurement error induced by the use of actual rather than expected inflation, and is larger the larger is the variance of inflation shocks. If actual inflation is not a good proxy for expected inflation then $V(u)$ will be large relative the $V(\omega)$, and the estimate of β will be substantially less than unity. Moreover, the third term on the RHS of (4.18) indicates that if the market reacts negatively to unexpected changes in the rate of inflation ($Cov(\varepsilon,u)<0$) then the estimate of β will be further reduced, and could be negative, even though the Fisher hypothesis holds.

Jaffe and Mandelker modelled expected inflation using both lagged ex post inflation rates, and, following Fama (1975), the real rate of interest. Nelson on the other hand used an ARMA(1,1) model to decompose inflation into expected and unexpected components. Both studies found a negative relationship between stock returns and expected inflation. However, Jaffe and Mandelker also discovered a negative relationship between returns and unexpected inflation, whilst Nelson presented results which suggested a negative relationship between returns and changes in expected inflation.

Bodie's (1976) approach was somewhat different. He asked to what extent can investors reduce the risk (variance) of the real return on a nominal bond by combining it with a stock portfolio. If the Fisher hypothesis is true, so that stocks are an effective hedge against inflation, it must be optimal for any holder of a nominal asset to hold a long position in stocks. Bodie's analysis is quite simple. The real return on a risk-free nominal bond, r_n, can be written in terms of changes in the purchasing power of money. If R_n is the nominal risk-free rate of interest, and S_t denotes the aggregate price level, we have the following definition:

(4.19) $1 + r_n \equiv (1 + R_n)Q + (1 + R_n)q$

where $Q \equiv E_t\left(\frac{S_t}{S_{t+1}}\right)$, and $q \equiv \frac{S_t}{S_{t+1}} - E_t\left(\frac{S_t}{S_{t+1}}\right)$

or, more simply,

(4.20) $1 + r_n \equiv E(1 + r_n) + (1 + R_n)q$

q is the unexpected change in the purchasing power of money, and so may be thought of as "inflation risk". The real return on equity may be similarly written as

(4.21) $1 + H \equiv E(1 + H) + \varepsilon$

where ε is again the unexpected stock return. The latter may be decomposed into inflation risk, αq, and "non-inflation risk", μ:

(4.22) $\varepsilon = \alpha q + \mu$

Combining (4.21) and (4.22), we have

(4.23) $1 + H = E(1 + H) + \alpha q + \mu$

Finally, if the investor holds a portfolio comprising the nominal bond and a quantity of stocks in proportions (1-w) and w respectively, the real return on this portfolio, r_p is given by

(4.24) $1 + r_p \equiv 1 + (1 - w)r_n + wH = (1 + r_n) + w(H - r_n)$

The problem is to choose w to minimise the variance of this portfolio. The variance-minimising value of w may be written as

(4.25) $$w = \frac{(1 + R)(1 + R - \alpha)}{(1 + R - \alpha)^2 + (\sigma_\mu^2/\sigma_q^2)}$$

where σ^2_f is the variance of q and σ^2_μ is the variance of μ. The sign and magnitude of w depends critically on the term (1+R-α), with the following possibilities. If α<0 (the unexpected stock return is positively related to unexpected inflation) then (1+R-α)>0 and the investor holds a long position in stocks. If 0<α<1+R, then the unexpected stock return is negatively related to unexpected inflation and the investor takes a long position in stocks. However, if 0<1+R<α then w is negative: the investor must sell stocks short in order to minimise the portfolio risk.

Using monthly data, over the period 1953(1)-1972(12), Bodie estimated α to be positive and large, certainly larger than 1+R. The implication is that the unexpected real return on stocks is strongly negatively related to unexpected inflation, and so the investor should sell stocks short to hedge against inflation. Bodie concluded that the data did not support the Fisher hypothesis.

In response to the above three studies, Firth (1979) analysed whether the apparent relation between stock returns and inflation was also a feature of the UK market. Using monthly (1935-1976) and annual (1919-1976) data, and employing an autoregressive model to obtain expected rates of inflation, Firth discovered a positive relation between nominal stock returns and expected inflation. Moreover, for the monthly data the coefficients on expected inflation were greater than unity, indicating that investors were more than compensated for expected inflation when holding UK stocks. The coefficient for the annual data, although positive, was not significantly so. Firth concluded that there was broad support for the Fisher hypothesis.

Gultekin (1983) studied the stock return-inflation relation using data from 26 countries. Over the period 1947-1979, nominal stock returns were found to be negatively related to actual inflation, and negatively related to expected inflation as measured by an ARIMA model. When the Treasury Bill rate was employed as an alternative proxy for expected inflation, the negative relations between both expected and unexpected inflation were particularly strong. However, the notable exception was the UK. Only when the TB rate was used did the UK data provide a negative relation between returns and expected inflation, and even in this case the coefficient on unexpected inflation was positive.

4.10.3 The Proxy Hypothesis

Fama (1981) argued that the observed negative relation between inflation and expected stock returns was spurious. Changes in inflation did not*cause* investors to revise their required returns: however, the correlation between the two could be explained by the fact that they were both related to a third variable, namely expected future real activity. Fama began by positing the following expression for the demand for real money:

$$\Delta \ln M_t - \omega_t = b_0 + b_1 \Delta \ln A_t + b_2 \Delta \ln(1 + R_{nt}) + v_t$$

where M_t is nominal money, A_t is a measure of anticipated real activity and R_{nt} is the nominal interest rate. Economic theory postulates that $b_1 > 0$ (because a larger amount of real money is required to accommodate the higher level of transactions which result from increased real activity), and $b_2 < 0$. Assuming that, within the monetary sector, anticipated future real activity and the interest rate are exogenous, and that money is also exogenous ("money causes prices"), the money demand equation can be inverted to produce a model for inflation:

$$\omega_t = -b_0 - b_1 \Delta \ln A_t - b_2 \Delta \ln(1 + R_{nt}) + \Delta \ln M_t - v_t$$

We now have the following chain of events. If anticipated real activity rises, money demand theory states that the demand for real money will also rise. Given the nominal money stock and the rate of interest, an increase in real money is achieved by a fall in the price level. Thus a positive relation between real money and anticipated real activity implies a negative relation between inflation and anticipated real activity. Fama therefore argued as follows. The present value formula for stock prices has prices depending on expected future dividend payments. To the extent that expectations of dividend payments follow expectations of future real activity, stock price changes (returns) will be positively related to changes in anticipated real activity. This effect, coupled with the above negative relation between inflation and anticipated real activity, produces a spurious negative relation between stock returns and inflation. Fama tested this argument by looking to see if measures of expected and unexpected inflation remained significant in the stock returns equation when measures of future real activity were also included. He found that when the growth of base money and future real activity growth rates were included as explanatory variables, expected inflation rates never had marginal explanatory power over stock returns[8]. He concluded that the negative stock return-inflation relation was therefore proxying for the positive relation between returns and expected future real activity.

Geske and Roll (1983) pointed out that there was no justification given for the inclusion of the money base growth rate in Fama's model for returns[9]. They argued that the significance of the money base implicates the money supply process, in addition to money demand, as a channel through which stock returns an inflation might be spuriously related. However, they went further and suggested that a "reverse causality" effect was at work: changes in stock prices signalled changes in the rate of money growth. They argued that stock prices change according to anticipated economic conditions, and that government revenue is highly pro-cyclical. If government expenditures are roughly constant, deficit cycles emerge, and if the

[8] However, with monthly and quarterly (but not annual) data, the negative relation between stock returns and unexpected inflation remained. Fama suggested that this might reflect shortcomings in the measures of anticipated real activity in these regressions.

[9] Fama simply noted that the relative weights on money base growth and anticipated real activity changed significantly between the equations for inflation and for stock returns, with much more weight in the stock return regression on anticipated real activity. His point therefore was that the inclusion of expected inflation in the returns regression, which forced the relative weights on the two explanatory variables to remain constant, was an imperfect proxy for the true relationship.

debt is monetised, deficits will presage inflation. Thus, when the economic outlook worsens, stock prices fall, a government deficit is expected, and rational agents will anticipate higher future inflation. Geske and Roll suggested that the money demand effect postulated by Fama would be reinforced by this reverse causality effect, so that both the negative relation between stock returns and inflation, and the significance of the monetary base in Fama's stock return regressions, were consistent.

4.10.4 Decomposition of the Inflation Beta

Two relations follow from the proxy hypothesis. First, assuming that expectations of future dividends follow expectations of future real activity, the apodosis is that *a positive shock to inflation will be associated with a downward-revision in expected future dividends,* so that there is a fall in the current stock price. Second, according to the consumption CAPM (C-CAPM)[10], a negative shock to expected future real activity reduces the value of current consumption relative to (lower) expected future consumption. This has two implications for required returns on investments. The required return on investments will be lower *now,* as agents are more willing to transform low-value current consumption into high-value future consumption. However, required returns will be higher in the future, as agents require higher returns if they are to divert income from high-value consumption to investment. Thus, if a positive shock to inflation portends a fall in expected future real output, *future required returns will rise,* which, through the RVF, puts further downward pressure on stock prices. The unexpected fall in the current stock price effects the reduction in *current* real returns required by the C-CAPM.

The current framework is particularly apposite for analysing these issues. This is because, unlike previous work, I am able to identify separately the covariances between shocks to inflation and changes in expectations of future dividends ($\beta_{di,infl}$), and changes in future required returns ($\beta_{ei,infl}$). Since the studies referred to above concentrate on aggregate stock indices, I shall first concentrate on the Total Market return betas with inflation. I discuss the sectoral results in a separate section below. I estimate the following market inflation betas:

---

[10] The C-CAPM is discussed in more detail in Section 4.5.2.

$\beta_{m,infl} = \underset{(0.6317)}{0.1982}$ $\beta_{dm,infl} = \underset{(0.0864)}{-2.2447}$ $\beta_{em,infl} = \underset{(0.6957)}{-1.1461}$

First, notice that the market dividend beta, $\beta_{dm,infl}$, is significantly negative. Thus, in conformity with the proxy hypothesis, a positive shock to inflation coincides with a downward-revision in expected future dividends. However, contrary to the predictions of the C-CAPM, I estimate a negative future excess returns beta. The fact that my empirical findings are opposite to the prognoses C-CAPM is not surprising, given the findings of Campbell and Shiller (1989) (who use US data) and Lund and Engsted (1993) (who use data for the UK, Denmark, Germany and Sweden), who find that real consumption growth has the "wrong sign" when forecasting discount rates in the RVF (see Chapter 5)[11].

My estimated total market inflation beta is positive, meaning that a positive shock to inflation coincides with a rise in the current excess return. The question is, can a *positive* relation between *excess* returns and inflation be reconciled with the *negative* relation between *real* returns and inflation reported in previous studies. The answer is yes. Define h_{mt} to be the real stock return on the market, e_{mt} to be the excess return on the market and ω_t to be the rate of inflation. My empirical results focus on the relationship between shocks to inflation and shocks to excess market returns:

$$(4.26)\ \tilde{e}_{mt} \equiv \tilde{h}_{mt} - \tilde{r}_t = \beta_{m,infl}\,\tilde{\omega}_t$$

In modelling excess rather than real returns, I do not directly consider the relation between shocks to the real interest rate and shocks to inflation:

$$(4.27)\ \tilde{r}_t = \gamma\tilde{\omega}_t$$

However, combining (4.26) and (4.27), the relationship between the real market return and inflation is

$$(4.28)\ \tilde{h}_{mt} \equiv \tilde{r}_t + \tilde{e}_{mt} = \left(\gamma + \beta_{m,infl}\right)\tilde{\omega}_t$$

---

11 Also, as reported in Section 4.3 above, Chen, Roll and Ross (1986) found the consumption growth risk premium to be of the "wrong sign".

Given the beta decomposition (4.8), we can rewrite (4.28) as

$$(4.29)\ \tilde{h}_{mt} = \left(\gamma + \beta_{dm,infl} - \beta_{r,infl} - \beta_{em,infl}\right)\tilde{\omega}_t$$

$$= \beta^h_{m,infl}\,\tilde{\omega}_t \qquad \text{where } \beta^h_{m,infl} \equiv \left(\gamma + \beta_{dm,infl} - \beta_{r,infl} - \beta_{em,infl}\right)$$

$\beta^h_{m,infl}$ is the inflation beta of the *real* market return. Equation (4.29) illustrates all of the different effects being picked up when other researchers have directly estimated the coefficient $\beta^h_{m,infl}$. The sign of the estimated relation will depend both on the signs and the relative magnitudes of its constituent parts. The strong negative correlations reported in Table 4.1 between inflation and the real interest rate (both their levels and innovations) imply that γ is negative: regressing the real return innovation from the VAR on the innovation in inflation, I estimate γ to be -0.9840. The dividend effect of the proxy hypothesis predicts a negative $\beta_{dm,infl}$, and this is what I find. The effect of a shock to inflation on expectations of future real interest rates, $\beta_{r,infl}$, is negative (see the foot of Table 4.4), and I have already discussed the negative estimated $\beta_{em,infl}$. The first two elements in parentheses in (4.29) therefore imply a negative value for $\beta^h_{m,infl}$, whilst the final two terms tend to increase its value. Because, from the above estimates, the first two elements outweigh the second two, I find the implied value for $\beta^h_{m,infl}$ to be negative (-0.7858). An estimated positive relation between excess returns and inflation therefore coexists with a negative relation between real returns and inflation, as found by other researchers. However, in my analysis all shocks must influence stock prices working via the consistency condition of the RVF, and I am therefore able to provide additional insights into the source of the negative real return-inflation nexus documented by other researchers.

4.10.5 Inflation Betas of Industrial Portfolio Returns

Boudoukh, Richardson and Whitelaw (1994) pointed out that the stock return-inflation relation depends on the degree to which the income from an asset varies with the economic cycle. Since the degree of cyclicality varies across industries, one should expect to find that different industry portfolios have different return-inflation relations. In particular, stock returns of highly cyclical industries will tend to covary negatively with inflation, whilst those of non-cyclical industries may have a positive covariance with inflation.

The results of the decomposition of the inflation beta using my industrial portfolios are presented in Table 4.4. All of the dividend betas, except that for Oil and Gas, are significantly negative, which again provides support for the proxy hypothesis. However, there is a large cross-sectional variation in the size of these effects, which is suggestive of cross-sectional variation in the proxy effect. For example, the dividend beta for Merchant Banks is -3.1729, whilst that for Chemicals is -0.4465. However, the cross-sectional pattern predicted by Boudoukh et al is not replicated in my findings. To see this, recall that in Section 4.2 I identified four sectors as highly cyclical - Motors, Banks, Merchant Banks and Financial Services - and four as non-cyclical - Brewers and Distillers, Food Manufacturing, Food Retailing and Oil and Gas. In conformity with Boudoukh et al's hypothesis, the pro-cyclical sectors do have relatively large dividend betas (ranging from -2.0414 for Motors to -3.2783 for Merchant Banks). However, although the dividend betas for the non-cyclical industries are smaller, the difference is often not marked (e.g. the dividend beta for Food Retailing is -2.0275).

Combining the estimated betas reported in Table 4.4 with the estimate of γ reported above, provides us with the estimated betas between the industry *real* returns and inflation (see equation (4.28)) as shown in Table 4.5. The theory predicts that the Motor industry real returns should be negatively related to inflation, whereas my estimate is positive. The beta decomposition allows one to identify the reason for this contradiction. It is because the excess return beta, $\beta_{ei,infl}$ (-1.7584), and the real interest rate beta, $\beta_{r,infl}$ (-1.8057), are so large as to outweigh the contemporaneous real interest rate effect and the dividend beta. Similarly, three out of the four non-cyclical industry returns are negatively related to inflation because their excess return betas are relatively small. Thus the net effects appear to depend in large part on the size of the excess return beta, on which, as I have discussed, existing asset pricing models are unable to shed much light.

Table 4.4: Decomposition of Inflation Beta

	$\beta_{i,infl}$	$\beta_{di,infl}$	$\beta_{ei,infl}$
Capital Goods	1.1908 (0.6672)	**-1.3182** (0.3791)	-0.7033 (0.8317)
Consumer Goods	0.9490 (0.6231)	**-1.7892** (0.3249)	-0.9325 (0.6965)
Financial Services	0.1892 (0.7621)	**-2.7626** (0.5543)	-1.1461 (0.7581)
Building Materials	1.3122 (0.7915)	**-1.3736** (0.5145)	-0.8801 (0.8671)
Contracting and Construction	1.4497 (0.8902)	-0.9731 (0.7027)	-0.6170 (1.1120)
Electricals	0.3626 (0.5631)	**-1.4814** (0.4982)	-0.0382 (0.8537)
Aerospace	**1.7056** (0.7197)	-0.7004 (0.3890)	-0.6002 (0.7794)
Engineering (General)	1.2602 (0.7196)	**-1.2756** (0.3650)	-0.7301 (0.7560)
Metals and Metal Forming	**1.7246** (0.8046)	**-1.3601** (0.5463)	**-1.2790** (0.5924)
Motors	1.5228 (0.9332)	**-2.0414** (0.7430)	**-1.7584** (0.5065)
Brewers and Distillers	0.5863 (0.5720)	**-1.3126** (0.4710)	-0.0932 (0.6638)
Food Manufacturing	0.7577 (0.7562)	**-1.6096** (0.3422)	-0.5616 (0.7609)
Food Retailing	0.7010 (0.5961)	**-2.0275** (0.5365)	-0.9228 (0.8658)
Hotels and Leisure	1.2017 (0.9131)	**-1.3268** (0.5149)	-0.7228 (0.9097)
Packaging, Paper and Printing	1.4557 (0.8328)	**-1.6289** (0.4595)	-1.2788 (0.6809)
Stores	1.3822 (0.7440)	**-1.7716** (0.4757)	-1.3481 (0.8226)
Textiles	1.0881 (0.8072)	**-1.2418** (0.5116)	-0.5242 (0.5678)
Chemicals	0.9089 (0.7436)	-0.4465 (0.4171)	0.4503 (0.6524)
Shipping and Transport	0.5366 (0.6287)	**-2.7391** (0.5023)	**-1.4700** (0.4537)
Miscellaneous	0.7154 (0.7402)	**-1.8956** **(0.4373)**	-0.8052 (0.5855)
Industrials	0.8947 (0.6296)	**-1.6544** **(0.3048)**	-0.7434 (0.7060)
Oil and Gas	1.8140 (0.8657)	0.2241 (0.2008)	0.2158 (0.6597)
Banks	-0.2243 (0.8224)	**-3.1792** (0.6076)	-1.1491 (0.7305)
Insurance	0.2948 (0.8327)	**-2.2291** (0.5853)	-0.7181 (0.7258)
Merchant Banks	0.4794 (0.8516)	**-3.2783** (0.7783)	-1.9519 (0.6905)
Property	0.8306 (0.9328)	**-2.2540** (0.9194)	-1.2789 (0.8711)
Investment Trusts	0.5415 (0.7508)	**-1.8612** (0.4851)	-0.5970 (0.8028)

The estimated interest rate beta, $\beta_{r,infl}$, is **-1.8057** (s.e.=0.1601).

Table 4.5: Betas Between Real Returns and Inflation

	$\beta^{h}_{m,inf}$
Pro-cyclical	
Motors	0.5338
Banks	-1.2083
Merchant Banks	-0.5046
Financial Services	-0.7948
Non-cyclical	
Brewers & Distillers	-0.3977
Food Manufacturing	-0.2263
Food Retailing	-0.2830
Oil & Gas	0.8300

4.11 Real Interest Rate Betas

Turning now to the direct real interest rate effects, the first column of numbers in Table 4.6 shows that real interest rate betas are negative: a positive shock to the level of the current real interest rate is associated with a downward revision in required stock returns. This result arises because the positive effect of real interest rate shocks on expected dividends is more than outweighed by the negative influence on excess returns through higher expected future real interest rates ($\beta_{r,r}$=1.8692), whilst for most portfolios, $\beta_{ei,r}$>0 (last column, Table 4.6). These two effects reduce the current stock price when current real rates rise. Because of the high negative contemporaneous correlation between shocks to real interest rates and inflation shocks (lower triangular, Table 4.1), the estimates in Table 4.6 are very similar (but of opposite sign) to those in Table 4.4.

4.12 Industrial Production Growth Betas

Innovations in the growth rate of industrial production do not have a discernible effect on current required returns. Table 4.7 reports that none of the industrial production growth betas is significantly different from zero. In fact only one statistic - the dividend beta for the Aerospace portfolio - is statistically significant. There does not even appear to be a consistent pattern in the signs of the estimated betas. I can only conclude that the shocks to industrial production do not contain any information relevant to the pricing of UK shares. This maybe suggests that shocks to monthly industrial production are not perceived by investors to be useful indications of changes in longer-term output trends, the latter being relevant to assessing the likely pattern of future dividend payments.

4.13 Real Exchange Rate Betas

Shocks to the real Sterling exchange rate provide a portmanteau measure of unexpected changes in the international competitiveness of British firms, and might therefore be expected to have some influence on investors' required returns. From Table 4.8 it appears that, on the whole, an unexpected appreciation of Sterling coincides with a rise in current expected returns. This effect is particularly important in the export-intensive Capital Goods sector and sub-sectors, where the total betas

Table 4.6: Decomposition of Real Interest Rate Beta

	$\beta_{i,r}$	$\beta_{di,r}$	$\beta_{ei,r}$
Capital Goods	**-1.5436** (0.6631)	**1.2029** (0.3651)	0.8773 (0.8283)
Consumer Goods	-1.2597 (0.6152)	**1.6338** (0.3070)	1.0243 (0.6911)
Financial Services	-0.4965 (0.7469)	**2.4759** (0.5228)	1.1032 (0.7467)
Building Materials	**-1.7418** (0.7823)	**1.1196** (0.4903)	0.9922 (0.8602)
Contracting and Construction	-1.9279 (0.8792)	0.7346 (0.6746)	0.7933 (1.1049)
Electricals	-0.7051 (0.5626)	**1.5429** (0.4918)	0.3788 (0.8594)
Aerospace	**-1.9780** (0.7150)	0.5445 (0.3738)	0.6533 (0.7714)
Engineering (General)	**-1.5716** (0.7168)	**1.1064** (0.3546)	**0.8087** (0.0994)
Metals and Metal Forming	**-2.0213** (0.7971)	**1.0061** (0.1012)	**1.1582** (0.5810)
Motors	-1.8197 (0.9251)	**1.7692** (0.7355)	**1.7197** (0.4983)
Brewers and Distillers	-0.9010 (0.5598)	**0.2478** (0.0640)	0.1063 (0.6562)
Food Manufacturing	-1.0545 (0.7559)	**1.3625** (0.3388)	0.5478 (0.7508)
Food Retailing	-1.0548 (0.5781)	**2.1031** (0.5344)	**1.2887** (0.0953)
Hotels and Leisure	-1.5497 (0.9048)	**1.0443** (0.4996)	0.7248 (0.8986)
Packaging, Paper and Printing	**-1.7630** (0.8257)	**1.5566** (0.0778)	**1.4504** (0.6785)
Stores	-1.7729 (0.7286)	**1.6492** (0.4559)	1.5529 (0.8203)
Textiles	-1.4154 (0.7986)	0.9194 (0.5020)	0.4655 (0.5588)
Chemicals	-1.0856 (0.7386)	0.4185 (0.4139)	-0.3652 (0.6485)
Shipping and Transport	**-0.7562** (0.6923)	**2.6898** (0.5030)	**1.5768** (0.4516)
Miscellaneous	-1.0508 (0.7327)	**1.7036** (0.4266)	0.8852 (0.5813)
Industrials	-1.2075 (0.6238)	**1.5483** (0.2921)	0.8866 (0.7029)
Oil and Gas	**-1.9292** (0.8594)	0.0527 (0.5625)	0.1127 (0.6678)
Banks	0.0128 (0.8093)	**2.7021** (0.5783)	0.8200 (0.7141)
Insurance	-0.6165 (0.8141)	**2.1141** (0.5659)	0.8614 (0.7224)
Merchant Banks	-0.8458 (0.8399)	**3.1455** (0.7571)	2.1221 (0.6883)
Property	-1.2411 (0.9103)	1.6798 (0.8754)	1.0516 (0.8538)
Investment Trusts	-0.7689 (0.7381)	**1.8480** (0.4745)	0.7477 (0.7992)

The estimated interest rate beta, $\beta_{r,r}$, is **1.8692**(s.e.=0.1562).

Table 4.7: Decomposition of Industrial Production Growth Beta

	$\beta_{i,ipg}$	$\beta_{di,ipg}$	$\beta_{ei,ipg}$
Capital Goods	0.0253 (0.2329)	0.0531 (0.1141)	0.0518 (0.2472)
Consumer Goods	-0.0584 (0.2451)	-0.1644 (0.1192)	-0.0819 (0.2075)
Financial Services	-0.1066 (0.2650)	-0.2477 (0.1666)	-0.1140 (0.2236)
Building Materials	-0.0284 (0.2742)	-0.0028 (0.1525)	0.0498 (0.2569)
Contracting and Construction	0.1670 (0.3213)	0.2217 (0.2026)	0.0788 (0.3240)
Electricals	-0.1408 (0.2207)	-0.0645 (0.1261)	0.1004 (0.2571)
Aerospace	0.2457 (0.2309)	0.3264 (0.1328)	0.1047 (0.2316)
Engineering (General)	0.1397 (0.2335)	0.2038 (0.1377)	0.0882 (0.2274)
Metals and Metal Forming	0.2007 (0.2451)	0.1270 (0.1782)	-0.0496 (0.1780)
Motors	0.1512 (0.3097)	0.0583 (0.2662)	-0.0687 (0.1604)
Brewers and Distillers	-0.0818 (0.2818)	-0.3187 (0.1879)	-0.2129 (0.1961)
Food Manufacturing	0.0927 (0.2336)	-0.1038 (0.1136)	-0.1725 (0.2243)
Food Retailing	-0.1065 (0.2833)	-0.2366 (0.1661)	-0.1061 (0.2622)
Hotels and Leisure	-0.1124 (0.3033)	-0.1752 (0.1578)	-0.0387 (0.2746)
Packaging, Paper and Printing	0.0667 (0.2494)	0.2714 (0.1448)	0.2287 (0.2038)
Stores	-0.1518 (0.2953)	-0.1979 (0.1563)	-0.0221 (0.2447)
Textiles	0.1071 (0.2428)	0.0791 (0.1664)	-0.0039 (0.1708)
Chemicals	-0.0723 (0.2178)	-0.1294 (0.1344)	-0.0330 (0.2054)
Shipping and Transport	-0.0438 (0.2203)	-0.0317 (0.1621)	0.0362 (0.1465)
Miscellaneous	0.0320 (0.2606)	-0.0540 (0.1654)	-0.0619 (0.1795)
Industrials	-0.0102 (0.2306)	-0.0535 (0.1051)	-0.0193 (0.2112)
Oil and Gas	-0.1874 (0.2603)	-0.1857 (0.2008)	0.0257 (0.2129)
Banks	-0.1258 (0.3014)	-0.3909 (0.2472)	-0.2410 (0.2188)
Insurance	-0.0963 (0.2496)	-0.2499 (0.1659)	-0.1295 (0.2177)
Merchant Banks	0.1538 (0.3281)	0.1417 (0.2179)	0.0121 (0.2061)
Property	-0.1933 (0.3301)	-0.2670 (0.2457)	-0.0496 (0.2580)
Investment Trusts	-0.0308 (0.2292)	-0.2319 (0.1318)	-0.1770 (0.2441)

The estimated interest rate beta, $\beta_{r,ipg}$, is -0.0241 (s.e.=0.0742).

Table 4.8: Decomposition of Real Exchange Rate Beta

	$\beta_{l,\Delta s}$	$\beta_{dl,\Delta s}$	$\beta_{el,\Delta s}$
Capital Goods	**0.5011** (0.2137)	0.0231 (0.0996)	-0.0479 (0.2235)
Consumer Goods	0.2738 (0.1919)	-0.1282 (0.0865)	0.0280 (0.1791)
Financial Services	0.2254 (0.2273)	0.0991 (0.1522)	0.3037 (0.1830)
Building Materials	**0.5860** (0.2538)	0.2198 (0.1523)	0.0639 (0.2245)
Contracting and Construction	**0.6783** (0.2741)	**0.4360** (0.1696)	0.1878 (0.2910)
Electricals	**0.5144** (0.2225)	-0.1068 (0.1635)	-0.1912 (0.2550)
Aerospace	**0.4727** (0.2251)	0.1369 (0.1414)	0.0942 (0.2011)
Engineering (General)	0.4159 (0.2218)	-0.0836 (0.1323)	-0.0694 (0.1994)
Metals and Metal Forming	0.4586 (0.2714)	0.2056 (0.2398)	0.1771 (0.1447)
Motors	0.2235 (0.2521)	**-0.5219** (0.2030)	**-0.3153** (0.1298)
Brewers and Distillers	**0.4192** (0.2088)	**0.2798** (0.1347)	0.2906 (0.1687)
Food Manufacturing	0.0722 (0.2038)	-0.0621 (0.1233)	0.2958 (0.1890)
Food Retailing	0.4381 (0.2235)	-0.3618 (0.1412)	-0.3699 (0.2550)
Hotels and Leisure	0.3598 (0.2319)	0.0700 (0.1439)	0.1402 (0.2334)
Packaging, Paper and Printing	0.2724 (0.2431)	-0.2666 (0.1537)	-0.1089 (0.1811)
Stores	**0.7058** (0.2383)	0.2180 (0.1242)	-0.0577 (0.2180)
Textiles	0.1723 (0.2486)	-0.1012 (0.1940)	0.1566 (0.1431)
Chemicals	0.0725 (0.2197)	**-0.4508** (0.1456)	-0.0933 (0.1902)
Shipping and Transport	0.2207 (0.2054)	**-0.5929** (0.1693)	**-0.3836** (0.1232)
Miscellaneous	0.2036 (0.2240)	**-0.3904** (0.1392)	-0.1639 (0.1561)
Industrials	0.3096 (0.1955)	**-0.1869** (0.0821)	-0.0665 (0.1887)
Oil and Gas	0.2848 (0.2405)	**-0.6960** (0.1847)	**-0.5508** (0.2182)
Banks	0.0975 (0.2334)	0.4205 (0.2346)	**0.7531** (0.1813)
Insurance	0.0535 (0.2795)	-0.3547 (0.1856)	0.0219 (0.1921)
Merchant Banks	0.4352 (0.2689)	-0.2503 (0.1938)	-0.2554 (0.1815)
Property	**0.6671** (0.2713)	**0.9206** (0.2731)	**0.6836** (0.2072)
Investment Trusts	-0.0091 (0.2097)	**-0.5586** (0.1188)	-0.1195 (0.2183)

The estimated interest rate beta, $\beta_{r,\Delta s}$, is **-0.4301** (s.e.=0.0713).

Table 4.9: Results Using a VAR(3) Model

	Capital Goods	Consumer Goods	Financial Services
		Market Betas	
$\beta_{i,m}$	1.0822 (0.0176)	0.9964 (0.0155)	1.0324 (0.0254)
$\beta_{di,m}$	0.0685 (0.0650)	0.2432 (0.0600)	0.1800 (0.0626)
$\beta_{ei,m}$	-1.2216 (0.0791)	-0.9610 (0.0731)	-1.0602 (0.0821)
		Dividend-price Ratio Betas	
$\beta_{i,D/P}$	-1.9002 (0.1445)	-1.7629 (0.1416)	-1.8349 (0.1324)
$\beta_{di,D/P}$	0.0784 (0.1180)	-0.2656 (0.1377)	-0.1525 (0.1332)
$\beta_{ei,D/P}$	2.4080 (0.0330)	1.9267 (0.0233)	2.1118 (0.0545)
		Inflation Betas	
$\beta_{i,infl}$	1.1770 (0.6943)	0.8744 (0.6189)	0.1129 (0.6685)
$\beta_{di,infl}$	-1.3368 (0.3098)	-2.3728 (0.3301)	-2.5580 (0.4638)
$\beta_{ei,infl}$	-0.5014 (0.8949)	-1.2349 (0.7492)	-0.6586 (0.8352)
		Real Interest Rate Betas	
$\beta_{i,r}$	-1.5549 (0.7046)	-1.2123 (0.6249)	-0.4472 (0.6702)
$\beta_{di,r}$	1.1255 (0.2914)	2.1498 (0.3109)	2.1335 (0.4309)
$\beta_{ei,r}$	0.6625 (0.8984)	1.3442 (0.7492)	0.5627 (0.8208)
		Industrial Production Growth Betas	
$\beta_{i,ipg}$	-0.0192 (0.2359)	-0.1002 (0.2359)	-0.1625 (0.2541)
$\beta_{di,ipg}$	-0.0804 (0.1041)	-0.1941 (0.1278)	-0.4406 (0.1629)
$\beta_{ei,ipg}$	0.0082 (0.2415)	-0.0245 (0.1898)	-0.2087 (0.2260)
		Real Exchange Rate Betas	
$\beta_{i,\Delta s}$	0.4644 (0.2232)	0.2494 (0.2027)	0.1701 (0.2318)
$\beta_{ei,\Delta s}$	-0.0332 (0.0945)	-0.1195 (0.1025)	0.1376 (0.1692)
$\beta_{di,\Delta s}$	-0.0134 (0.2399)	0.1152 (0.1855)	0.4516 (0.2028)

are generally significantly positive. The major source of this effect appears to come through adjustments in expectations about future real interest rates. Whilst most of the dividend and future excess return betas are insignificant (columns 2 and 3, Table 4.8), I estimate the real interest rate beta to be -0.4301 (s.e.=0.0713), so that an unexpected appreciation of Sterling coincides with a fall in expected future real interest rates. Since for much of the data period the Bank of England has actively used interest rates to influence the exchange rate, this is a highly plausible result. Thus, the direct effect of an unexpected appreciation of Sterling on future dividends is fairly minimal compared with the reduction in future real interest rates, and so stock prices increase.

4.14 Effects of Increasing the VAR Lag Length

My results appear to be robust to changes in the VAR lag length. Table 4.9 presents the beta estimates from a 3-lag VAR for three portfolios: Capital Goods, Consumer Goods and Financial Services. Comparison with previous tables indicates that, not only are the signs and statistical significance of the beta estimates largely unchanged, but the point estimates are all very close to those obtained from a VAR(1) model.

4.15 Conclusions

In this chapter, I have used a liearised RVF to apportion unexpected changes in asset returns into news about fundamentals, namely future dividends, real interest rates and future returns. Macroeconomic factors which might influence expected asset returns can do so only if they contain news about these three fundamentals. The analytical framework therefore combines the time-series VAR method for estimating "news" components with the cross-section approach more familiar in the APT framework.

The basic metric used is the asset's beta with a risk factor, $\beta_{i,k}$. Any $\beta_{i,k}$ can be decomposed into betas between the three fundamentals and the chosen factor. The CAPM is a special case of the (unrestricted) multifactor model. My main findings are as follows. When I consider the market return as a factor, I find that for most sectors the positive market beta is primarily due to the influence of the market return on future expected asset returns. Cash flow betas are relatively small, and therefore play only a minor role in influencing returns. The one-factor CAPM is nested in my general model, and whilst I find some statistical support for the restrictions placed on the LFM by the CAPM, informal analysis of the cross-sectional implications of these restrictions casts doubt on this finding.

I examine the channels through which other macroeconomic factors influence fundamentals, and hence asset prices. Here, in contrast to studies that look at multivariate determinants of returns without imposing the consistency requirement of the RVF, I find that simple "causal" relationships cannot be made. This is because the chosen factor can have offsetting effects on the fundamentals. In particular, the relation between real stock returns and inflation depends on four separate effects. For the market index, the net result is a negative impact of inflation on real returns, which is consistent with previous single-equation studies. However, a negative relation is not always found using industry portfolio returns, and existing explanations for cross-sectional variation in this relationship are not supported by the data.

My other findings are as follows. A positive surprise to the dividend-price ratio has a strong positive effect on expected future returns (consistent with the RVF), which outweighs any effect from future cash flows, so that the net effect is a fall in the current stock price. Similarly, higher real interest rates have a direct negative effect on stock prices which is reinforced by higher future excess returns; and these effects swamp any effect via future cash flows. Again, a higher real exchange rate has its major impact via lower expected future real interest rates, and so leads to higher current stock prices. Finally, and perhaps surprisingly, I find that shocks to real output growth have little or no effect on most asset returns.

Chapter 5: TESTING THE EFFICIENCY OF THE UK STOCK MARKET

5.1 Introduction

In this chapter I employ the VAR methodology pioneered by Campbell and Shiller (1987, 1988, 1989) to test the efficiency of the UK stock market. The evidence presented here complements that obtained using variance ratio tests (see Section 5.2 below) and direct tests of the rational valuation formula (i.e. that stock prices equal the discounted present value of expected future dividends) using cross-section data (e.g. Miles 1993).

Bulkley and Tonks (1989) applied Shiller-type variance bounds tests to annual UK time-series data. Assuming constant equilibrium returns, they found that the variance bound was strongly violated, and proffered an explanation based on the strong-form/weak-form rational expectations distinction. In this chapter, I consider an alternative possibility; that equilibrium returns are non-constant, and that existing models of time-varying equilibrium returns provide sufficient variation in discount rates so that prices are not excessively volatile.

The VAR approach has several potential advantages over alternative test procedures. Within a single framework I can test for the predictability of one-period returns and multi-period returns, and also perform tests based on stock prices using the rational valuation formula (RVF). Tests of the predictability of one-period returns (e.g. Clare, Thomas and Wickens 1994, MacDonald and Power 1991, Keim and Stambaugh 1986) may be less powerful in detecting deviations from market efficiency than tests using either multi-period returns (Fama and French 1988a) or stock prices.

Bulkley and Tonks conducted their analysis assuming real (detrended) stock prices and dividends are stationary processes (an assumption borne out by their DF statistics). However, I find that the extra data since 1985[1], together with revisions to the original series, mean that this assumption is no longer empirically tenable. The VAR methodology takes explicit account of the non-stationarity in the data, as well as allowing tests of the efficient markets hypothesis (EMH) under a wide variety of assumptions about the determinants of equilibrium returns. In addition, the analysis provides several metrics which allow one to judge the degree to which the data conform to the hypothesis.

---

[1] 1985 is the end of Bulkley and Tonks's data period.

The rest of this chapter is as follows. In Section 5.2 I review the early literature which sought to apply direct tests of the RVF to stock price data. Section 5.3 details a number of criticisms of the original tests, and Section 5.4 looks at the second generation of tests developed in response to these criticisms. Much of the early analyses concentrated on models of constant expected real returns. In Section 5.5 I discuss the CAPM and the consumption CAPM, both of which predict time variation in expected returns, and tests of these models are reviewed. Section 5.6 outlines the models of expected returns considered in the current study. Sections 5.7 and 5.8 describe the dividend-price ratio model and the way in which the implications of this model may be translated into restrictions on the parameters of a VAR. In Section 5.9, the results of the VAR methodology applied to aggregate UK data are presented. Section 5.10 concludes.

5.2 A Review of the Early Efficient Markets Literature

Until around 1980, researchers concentrated their efforts on estimating return autocorrelations, with a view to testing the martingale/random walk hypothesis for stock prices (see Section 1.2.3). The broad consensus (see Fama 1970) was that stock price changes were more or less random, a finding consistent with the view that prices changed only in response to new information about dividends. Samuelson's (1973) derivation of stock prices as the present value of expected future dividends, discounted by a constant real rate, was taken as a useful descriptive model of stock price behaviour. However, there was also a popular view that movements in stock price indices could not realistically be attributed to dividend news. In particular, prices seemed excessively volatile, given the relatively smooth pattern of dividends. This conflict led Shiller (1981) and LeRoy and Porter (1981) to develop direct tests of the present value relation. These so-called "variance bounds tests" are discussed in some detail below. The statistical validity of the early tests has been seriously questioned, and I do not apply such tests in this study. However, they are worth studying in detail, since the debate over the robustness of the early findings serves to clarify some of the finer statistical points relevant to the issue of excess volatility. As such, the Campbell-Shiller VAR methodology employed here emerges as a likely candidate for "best practice" in the area of tests of stock market efficiency. However, the directness and simplicity of Shiller's variance bounds tests, together with the failure of many of the more recent studies to overthrow his conclusions, has meant that Shiller's paper remains the most important (and most commonly cited) study in this area.

The excess volatility literature raises broad questions about the descriptive ability of neoclassical economics. Indeed, evidence against market efficiency has been used as ammunition against the utilitarian paradigm, and in the development of so-called "quasi-rational" alternatives (see Thaler 1994, esp. Part 5). However, for the most part the concentration has been on honing the statistical tools used to guide inference on the issue of volatility bounds. The debate has been characterised by extremely rigorous analysis of the power of various statistical procedures, and I now review the most important contributions to this process. The focus is on the development that has taken place in attitudes towards the discriminatory power of the econometric techniques employed.

5.2.1 The Variance Bounds Theorems

To begin with, assume that expected returns areconstant, so that the emphasis is wholly on the relationship between prices and dividends. Much of the early literature was concerned with the appropriate treatment of prices and dividends vis a vis the removal of trends, and with the robustness of the findings to the assumptions regarding the dividend process.

With a constant discount rate, the present value formula is

$$(5.1) \qquad P_t = \sum_{i=1}^{\infty} \gamma^i E_t D_{t+i}$$

The early variance bounds tests were based on a theoretical construct termed the "perfect foresight", or "ex post rational" stock price, P_t^*. Given the present value formula, this is the stock price that would obtain were investors able to predict dividends with perfect certainty, and is defined as the discounted present value of *actual* future dividends:

(5.2) $$P_t^* = \sum_{i=1}^{\infty} \gamma^i D_{t+i}$$

The actual stock price is then the conditional expectation of the perfect foresight price:

(5.3) $$P_t = E_t P_t^*$$

(5.1) can be rewritten in terms of the discounted conditional expectation of next-period's dividend plus next period's closing price:

(5.4) $$P_t = \gamma E_t(D_{t+1} + P_{t+1})$$

Using the following definition

(5.5) $$R_t \equiv D_t + P_t - E_{t-1}(D_t + P_t)$$

(5.4) can be written tautologously as

(5.6) $$P_t = \gamma(D_{t+1} + P_{t+1} - R_{t+1})$$

R_{t+1} is the "surprise return/payoff" (which should not be confused with the unexpected *rate* of return analysed in Chapter 3). If we treat (5.6) as a difference equation in $P_{,t}$ and solve it forwards subject to the terminal condition $\lim_{i\to\infty} \gamma^i D_{t+i} = 0$, we have

(5.7) $$P_t = P_t^* - \sum_{j=1}^{\infty} \gamma^j R_{t+j}$$

Assuming that R_t is covariance stationary, from (5.7) the variance of the perfect foresight price (noting that the R_t's are serially independent) is

(5.8) $$V(P_t^*) = V(P_t) + 2\mathrm{Cov}(P_t, \sum_{h=1}^{\infty} \gamma^j R_{t+j}) + \frac{\gamma^2 V(R_t)}{1-\gamma^2}$$

Rational forecasting implies that the covariance term is zero, so that (5.8) becomes[2]

$$(5.9) \qquad V(P^*) = V(P) + \frac{\gamma^2 V(R)}{1-\gamma^2}$$

Thus the variance of the perfect foresight price is equal to the variance of the actual price plus a proportion of the variance of the unexpected payoff. The apportionment of the variance of the perfect foresight price between the two RHS components depends crucially on the amount of information investors have about future dividend payments. If agents have a large amount of relevant information, so that they are fairly successful at dividend forecasting, the variance of expected dividends will be almost as large as that of actual dividends. This will translate into an actual stock price that is almost as variable as the perfect foresight price, and the surprise return variance will be low. On the other hand, if dividend forecasts are not very accurate, because of a lack of relevant information[3], expected dividends will have low variability relative to ex post dividends, whilst the variance of dividend surprises will be large. In this case, the second term on the RHS of (5.9) will account for the larger portion of the ex post rational price variance, and the variance of the actual price will be relatively small. The basic principle at work is that within rational forecasting schemes, more information cannot worsen predictive performance, or, alternatively, the variance of the forecast tends to the variance of the actual series from below, increasing with the degree of relevant information.

Since the econometrician will most likely be endowed with a smaller information set than actual investors, who in turn have less-than-complete information for forecasting purposes, the variability of the actual stock price should be greater than that of a price forecast by an econometrician. If we define P'_t to be a price series constructed by an econometrician using dividend forecasts based on a relatively small information set, we expect to find the following variance hierarchy:

$$(5.10) \qquad V(P') \leq V(P) < V(P^*)$$

2 Since all series are assumed to be covariance stationary, the time subscripts in the variance terms are unnecessary, and so are dropped.

3 The maintained assumption is that all information is used optimally. Therefore the only source of poor forecasts is poor information.

The equality (5.9) and the inequalities (5.10) formed the basis of the first generation of direct tests of the present value model.

5.2.2 Shiller (1981) and LeRoy and Porter (1981)

In a seminal paper, Shiller (1981) claimed that stock prices were too volatile to accord with the Efficient Markets Hypothesis (EMH). His main focus was on the relationship between the variance of the observed stock price and the variance of the ex post rational stock price (the right-hand inequality in (5.10)). In order to determine this relationship, two hurdles had to be overcome: the real stock price and dividend series needed detrending to ensure the existence of their second moments, and the unobservable ex post rational stock price needed to be constructed.

Shiller detrended the price and dividend series by discounting them by a long-run growth factor. This results in a re-statement of the present value formula in which detrended prices are equal to discounted expected detrended dividends, the (constant) discount rate now being a function of the expected return and the long-run growth rate.

If the price series in expression (5.4) is replaced by the ex post rational stock price, and the expectations operator is removed, forward recursion over T periods produces the following expression:

$$(5.11) \qquad P_t^* = \sum_{t=1}^{T-1} \gamma^t D_t + \gamma^T P_T^*$$

that is, the current perfect foresight price is equal to the T discounted cumulative dividends plus the discounted terminal perfect foresight price. Shiller replaced P_T^* with the average actual stock price over the T periods, i.e.

$$(5.12) \qquad P_T^* = \frac{1}{T} \sum_{t=1}^{T} P_t$$

Using expressions (5.11) and (5.12), the detrended price and dividend series were used to construct an empirical proxy for the perfect foresight stock price, and the variances of the two price series were calculated directly.

Shiller performed these tests on two data sets. The first comprised annual observations of Standard and Poor's Monthly Composite Stock Price index (currently 500 stocks) for January from 1871 to 1979, divided by the Bureau of Labor Statistics wholesale price index. The dividend series was total dividends for the calendar year accruing to the portfolio. His second data set was annual observations of the real price and dividends of the portfolio of 30 stocks which form the Dow Jones Industrial Average for 1928 to 1979. The series was modified to ensure that it reflected the performance of a single unchanging portfolio.

Shiller found that the variance of the actual stock price was at least five times as great as that of the perfect foresight price series. Although no confidence intervals were presented, Shiller claimed that the "gross nature of the violation" (Shiller 1989, p85) suggested that the apparent rejection of the EMH is unlikely to be due to data errors or other technical problems. He therefore concluded that stock prices were too volatile to accord with the present value model.

In a coincident study, LeRoy and Porter (1981) also tested the empirical content of the variance bounds theorem. Unlike Shiller's more direct approach, LeRoy and Porter calculated the relevant variances from a bivariate AR model for dividends and prices. The benefit of this model-based testing strategy is that, unlike Shiller, LeRoy and Porter were able to calculate confidence intervals for their test statistics.

The variance inequalities in (5.10), and the equality (5.9) were both tested using quarterly earnings and price data for Standard and Poor's Composite Index for 1955(1) to 1973(4) together with data for three large corporations - American Telephone and Telegraph, General Electric and General Motors. The point estimates indicated rejection of (5.10) for all data sets; but the asymptotic variances of the tests were so high that the variance inequalities could be rejected confidently only for General Electric. Regarding the variance equality (5.9), the individual firm data clearly rejected (5.9), whilst the aggregate index did not.

Apart from inference being clouded by the large confidence intervals, it is now accepted that LeRoy and Porter's method of trend correction was flawed. They attempted to remove trends from their series by reversing two effects: inflation and corporate earnings retention. However, they did note that the adjusted series showed some evidence of a downward trend. This led them to comment that "the dependence of our results on the assumption of stationarity is probably their single most severe limitation" (LeRoy and Porter, 1981, p569).

LeRoy and Porter noted that rejection of the Theorems could be due to any one or more of four factors:

i) the present value relation (5.9) does not hold;

ii) expected returns are not constant;

iii) expectations may not be rational;

iv) the tests may be subject to statistical and/or measurement problems.

Consequently, whilst their point estimates rejected dramatically the EMH, their conclusion that the stock market was not efficient was much more tentative than Shiller's.

5.3 Methodological Criticisms

Flavin (1983) and Kleidon (1986) raised serious questions regarding the robustness of the conclusions drawn from the above two studies. Flavin focused on the bias involved in the estimation of the variance of the perfect foresight price. She argued that the degree of serial correlation in the series, together with the need for a terminal price, made the estimated variances biased downwards, with the result that the variance bounds tests were prepossessed towards the finding of excess volatility. Kleidon studied two issues: the use of time-series plots of P_t and P_t^* as a replacement for formal statistical analysis of the two series; and the stationarity assumptions implicit in the original implementations. Both articles presented evidence that the rejection of the variance bounds theorems was far from being clear cut.

5.3.1 Flavin (1983)

Flavin's first point resulted from some elementary statistical analysis. Given a set of observations drawn from a common distribution and with a known mean, the mean squared deviation from the mean of these sample observations provides an unbiased estimator of the population variance, regardless of whether the observations are mutually uncorrelated. In the more usual case where the population mean has to be estimated, this variance estimator would be biased downwards, since the sample variance of any series around some constant is minimised when the constant is chosen to be the sample mean. The remedy is to reduce the degrees of freedom in the denominator by one. However, in the latter case it does matter if the drawings are correlated. In particular, if the observations are positively autocorrelated then this degrees of freedom correction will not be sufficient to remove the downward bias induced by using the sample mean. The reason is that a series which is highly positively autocorrelated is characterised by protracted deviations from the mean of its underlying distribution, so that the mean of any given sample could well be quite different from the population mean. Consequently, the use of the (variance minimising) sample mean will always understate the true variance.

Given that both the P_t and P_t^* series were highly positively autocorrelated, the variance of both will be under-estimated. However, the problem is more serious than this, since P_t^* is more highly autocorrelated ("smoother") than the actual price series. The estimated variance of P_t^* was therefore more heavily biased downwards than that of P_t, making rejection of the variance bound more likely.

There are two points to make about the terminal condition used to construct the perfect foresight stock price. First, Shiller's method of setting P_T^* equal to the mean sample stock price (see (5.12)) produces a biased estimate of P_t^*, since only if P_T^* is set equal to the terminal price P_T is the estimated perfect foresight series a rational forecast of the theoretical series i.e. $P_t=E_tP_t^*$(see Gilles and LeRoy, 1991). Second, using the actual stock price in the construction of the perfect foresight price will bias the estimated variance of P_t^* downwards. The reason is that using the actual stock price at the terminal point T forces all dividend surprises beyond this date to be zero, so that the variance of the estimated perfect foresight series will always be lower than the true value. The question of how important this bias is likely to be is discussed in Section 5.4.2 below.

5.3.2 Kleidon (1986)

Kleidon took exception to the practice adopted by Shiller (1981) and Grossman and Shiller (1981) of comparing the behaviour of P_t and P_t^* using time-series plots rather than formal statistical analysis. In Kleidon's view, there was no reason to expect the two series to appear to behave similarly. The reason is that the P_t^* series arises ex post from a single state of the world, whereas P_t estimates P_t^* across all possible states of the world from an ex ante perspective. Uncertainty regarding the future state of the world will inevitably make any estimate of a state-contingent series more variable than the series itself. For example, since, empirically, any changes in current dividends tend to imply changes in future dividends, a shock to current dividends can have a large impact on the current stock price; but since P_t^* is constructed using only realised dividends, there are no dividend surprises with implications for the variability of P_t^* as there are for actual prices. The question to be addressed is not whether P_t is less variable than any given P_t^*, but rather whether P_t is less variable than all possible P_t^*'s.

The second consideration of Kleidon's study was the effects of incorporating non-stationary price and dividend processes into the variance bounds tests. If prices and dividends are non-stationary series, their unconditional variances do not exist. One remedy is to transform them into stationary series and test unconditional variance bounds. Shiller (1981) divided the series by a growth trend, and so is implicitly assuming that the series were *trend stationary*. However, if the series were in fact *difference stationary*[4], extracting a time trend would not transform them into stationary series. Thus the unconditional variance inequalities (5.10) could not be tested.

The long-held convention in the finance literature is that log stock prices follow a geometric random walk with drift (a difference-stationary process). Given the present value formula, dividends which follow a similar process are sufficient to produce such a series for prices. Kleidon first pointed out that only one of Shiller's variance inequalities[5] is consistent with random walk prices and dividends, and this is the only inequality not violated by Shiller's calculations. Kleidon then generated artificial data for the price and dividend processes to correspond with Standard and Poor's (deflated) annual series 1926-79, and also an artificial perfect foresight price series. He applied Shiller's first variance bounds test to the generated series for various sample sizes and discount rates. He found that although his price and

---

4 Nelson and Plosser (1982,p139) could not reject that stock prices were "non-stationary stochastic processes with no tendency to return to a trend line".

5 Shiller (1981) presented several alternatives to the inequality (5.10), but for brevity I do not discuss them in any detail.

dividend series were consistent with the present value formula *by construction*, the variance bound was violated over 70% of the time. Moreover, when Shiller's method of detrending was applied to the series, the tendency to reject the inequality was exacerbated (the lowest rejection rate being 89%). He discovered further that for the detrended series, 397 replications out of 1,000 gave violation ratios greater than 5.0 (the level that Shiller had termed "gross violations") for a 5% expected return, and 148 replications for an expected return of 6.5%. He concluded that the "gross violations" of the variance bound reported by Shiller were consistent with the application to non-stationary series of estimation techniques that assume stationarity.

Whilst non-stationary series do not have unconditional variances, conditional variances can be calculated. Kleidon therefore tested the conditional inequality

$$(5.13) \qquad Var(P_t^* | \Phi_{t-k}) \geq Var(P_t | \Phi_{t-k}) \qquad k = 0, 1,$$

where Φ_{t-k} is information available at time t-k. For the Standard and Poor series 1926-79 for k=1,2,5 none of the point estimates violated the inequality. Also, whilst the point estimates for k=10 violated (5.13) - as Shiller had found - the violation was not significant at the 10% level. He concluded that the results were consistent with the hypothesis that changes in expectations about future cash flows cause changes in stock prices.

Marsh and Merton (1986) also argued that it is the assumption of stationarity rather than the failure of the EMH that is the cause of Shiller's rejections. They assumed that managers set dividends to grow at a rate g, but deviate from this long-run growth path in response to changes in permanent earnings, ε, that deviate from the long-run growth path[6]:

$$(5.14) \qquad \Delta D_t = gD_t + \sum_{k=0}^{N} \vartheta_k[\Delta\varepsilon_{t-k} - g\varepsilon_{t-k}]$$

Thus the dividend process contains a unit root

[6] This is consistent with Litner's (1956) model based on stylised facts established in a set of interviews with managers about dividend policies.

Marsh and Merton's Theorems 1 and 2 stated that if the present value formula holds, expected returns are constant, and dividends are generated by (5.14), then in any finite sample, and in the limit almost surely,

$$(5.15) \qquad \mathrm{Var}(P^*) \leq \mathrm{Var}(P)$$

This is the exact opposite to Shiller's inequality. The fact that any P* series which fails Shiller's test passes Marsh and Merton's test suggests that the reliability of variance bounds tests as tests of the EMH is questionable. Furthermore, in the absence of any strong theoretical reason for assuming trend-stationary real dividends, Shiller's "gross violations" militate more against his stationarity assumptions than against the EMH.

5.4 Volatility Tests: The New Generation

These criticisms of Shiller's inequalities motivated a second generation of tests which addressed the statistical problems that plagued the original formulations.

5.4.1 West (1988)

West (1988) developed another variance inequality which is valid when P_t and D_t are non-stationary, and does not require the construction of a perfect foresight price series. Defining prices as cum dividend, and once again denoting the econometrician's (limited information) stock price as P_t', West derived the following analogues to equation (5.7):

$$P_t^* - P_t = \sum_{j=1}^{\infty} \gamma^j \tilde{P}_{t+j} \qquad \text{where } \tilde{P}_t \equiv P_t - E_{t-1}P_t$$

$$P_t^* - P_t' = \sum_{j=1}^{\infty} \gamma^j \tilde{P}'_{t+j} \qquad \text{where } \tilde{P}'_t \equiv P'_t - E_{t-1}P'_t$$

It follows that

$$\mathrm{Var}(P_t^* - P_t') = \mathrm{Var}(P_t - P_t') + \mathrm{Var}(P_t^* - P_t)$$

$$\mathrm{Var}(P_t^* - P_t') \geq \mathrm{Var}(P_t^* - P_t)$$

or

$$\mathrm{Var}(\tilde{P}_t') \geq \mathrm{Var}(\tilde{P}_t)$$

This is the West inequality. The intuition is straightforward. The EMH implies that the stock price adjusts unexpectedly only in response to news about dividends. If the market uses all information optimally, there should be no subset of information on which forecasts can be based which "outperforms" the market price in terms of having a lower average dividend shock/lower return variance. The West inequality states that the forecast error variance with a limited information set should be at least as great as the full-information error variance.

West used the same data as Shiller to evaluate this inequality He allowed only lagged dividends to enter the information subset, and modelled the series with an ARIMA(p,1,0) process. For both data sets, the inequality was rejected. The implication was that the ARIMA model forecast dividends (and hence P*) so well that the variance in price innovations could not be due only to forecast errors.

5.4.2 Mankiw, Romer and Shapiro (1985, 1991)

Mankiw, Romer and Shapiro (1985, hereafter MRS) defined a "naive forecast" stock price, P_t^0, as

$$P_t^0 = \sum_{i=1}^{\infty} \gamma^i F_t D_{t+i}$$

where $F_t X_{t+i}$ is a naive forecast of X_{t+i} given information at time t. This naive forecast is assumed to be available to rational agents at time t so that

$$E_t[(P_t^* - P_t)(P_t - P_t^0)] = 0 \tag{5.15}$$

i.e. the naive forecast error is uncorrelated with the rational forecast error.

MRS begin with the identity

(5.16) $$(P_t^* - P_t^0) = (P_t - P_t^0) + (P_t - P_t^0)$$

Squaring (5.16), taking expectations, and using the result (5.15), we have

(5.17) $$E_t(P_t^* - P_t^0)^2 = E_t(P_t^* - P_t)^2 + (P_t - P_t^0)^2$$

which in turn implies that

(5.18) $$E_t(P_t^* - P_t^0)^2 \geq E_t(P_t^* - P_t)^2$$

and

(5.19) $$E_t(P_t^* - P_t^0)^2 \geq E_t(P_t^* - P_t^0)^2$$

Employing the law of iterated expectations, we can replace E_t with expectations conditional on information available prior to the beginning of the sample period, E':

(5.17') $$E'(P_t^* - P_t^0)^2 = E'(P_t^* - P_t)^2 + E'(P_t - P_t^0)^2$$

(5.18') $$E'(P_t^* - P_t^0)^2 \geq E'(P_t^* - P_t)^2$$

(5.19') $$E'(P_t^* - P_t^0)^2 \geq E'(P_t - P_t^0)^2$$

(5.18') states that the actual stock price is a better forecast of the perfect foresight price than is the naive forecast in so far as it has a smaller MSE. (5.19') says that P* is more volatile around the naive forecast than is the market price. MRS claimed three advantages of their tests over Shiller's. First, since the naive forecast can grow as dividends grow, $(P_t^*-P_t^0)$ and $(P_t-P_t^0)$ do not need detrending, thus avoiding problems of non-stationarity. Second, because these tests are not constructed using sample means, they do not suffer from the small sample bias documented by Flavin (1983). Finally, as discussed above, Flavin attributed some of the bias to Shiller's method of obtaining the terminal perfect foresight price (which was to use the average of detrended P's). MRS therefore used P_T as the terminal price, which ensures that rationality implies that $P_t = E_t P_t^*$.

The naive forecast of dividends was assumed simply to be

$$F_t D_{t+i} = D_{t-1}$$

i.e. dividends are perceived to follow a random walk. The corresponding naive stock price is then

$$P_t^0 = \frac{\gamma}{1-\gamma} D_{t-1}$$

Using the same data as Shiller (extended to 1983), and after a simple correction for heteroscedasticity, MRS found that inequality (5.18') was violated for all assumed (constant) values of the expected market return, and (5.19') was violated only for low values of the discount rate. Subject to the unresolved question of statistical significance, and under the assumption of constant expected returns, they concluded that stock prices do not accurately reflect fundamentals.

Shea (1989) presented two criticisms of MRS's tests, two of which centred on their choice of terminal price for the construction of the ex post rational price series. First, the test results were sensitive to the terminal date. Second, because the use of P_T as an approximation to P_T^* introduces some non-stationarity into the analysis, it is not possible to calculate confidence intervals for the test statistics, only point estimates.

Merton (1987) also took issue with MRS's tests. He argued that if P_T is "important" (i.e. the bulk of the dividends are paid out-of-sample), statistical inference using MRS's tests is not possible. In their 1991 paper, MRS argued that, since, empirically, end-of-sample prices are not a particularly important part of current prices as compared with within-sample dividends, Merton's criticism was of little practical importance. They also dealt with Shea's main criticisms by employing a rolling terminal date in the construction of P_t^*, so that the perfect foresight price was calculated at each point under the assumption of a fixed holding period, rather than using the end-of-sample price for all observations. This allowed them to calculate confidence intervals for their test statistics. Using Shiller's data extended to 1988, they found that although efficiency was rejected, the rejections were not very strong. They concluded that the data did not allow one to distinguish confidently between competing views of the stock price DGP.

5.4.3 Scott (1985)

Scott (1985) proposed a simple regression test for the present value model, based on the relation (5.7). Having constructed a series for P* (with $P_T^*=P_T$), if P_t is a rational forecast of P_t^* then we expect a=0 and b=1 in the regression

$$(5.20) \qquad P_t^* = a + bP_t + \varepsilon_t$$

Of course, if P_t and D_t are non-stationary, ε_t may be non-stationary and so standard tests of the hypotheses a=0 and b=1 may not be statistically valid. Scott therefore considered two alternative regression tests. The first uses variables detrended as in Shiller (1981). The second detrends the series by dividing through by dividends:

$$(5.21) \qquad \frac{P_t^*}{D_t} = a + b\frac{P_t}{D_t} + \eta_t$$

As long as the price-dividend ratio is stationary, η_t will be stationary and standard tests are valid. As (5.7) makes clear, the error terms in these regressions will be autocorrelated, and Scott used a spectral method when constructing the test statistics.

Scott tested the model using quarterly observations of the Standard and Poor Index from 1947(1) to 1983(2). He presented t-statistics of the individual hypotheses a=0 and b=1, and a χ^2 statistic for the hypothesis that these hold simultaneously. For model (5.20), he found $\hat{a} = 35.47$ and $\hat{b} = -0.025$, with t-statistics respectively 26.33 and -42.82. These indicated rejection of the present value model at very low marginal significance levels. The $\chi^2(2)$ statistic was 5014, which was an equally resounding rejection. For the model (5.21), the t-statistics were 6.53 and -10.02, which were much lower than for (5.20) but still very large. The $\chi^2(2)$ statistic was 366. Finally, Scott used Monte Carlo simulations to evaluate the finite-sample properties of his estimators and tests. He found that, although there was some small-sample bias in the estimates of a and b, the t-statistics were not biased against the present value model. The joint test was, however, biased against the efficient markets null. Scott concluded on the basis of the t-statistics that the present value model was not consistent with the Standard and Poor data.

Durlauf and Hall (1989) demonstrated a relation between variance bounds tests and Scott's regression tests. To see this, we first write both the actual stock price and the ex post rational stock price as the discounted sum of expected dividends, Q_t, plus a forecast error:

(5.22) $$P_t = Q_t + \tilde{P}_t$$

(5.23) $$P_t^* = Q_t + \tilde{P}_t^*$$

The present value model implies that $\tilde{P}_t = 0$. In a regression of P_t on $(P_t - P_t^*)$, the slope coefficient ω will be

$$\omega = \frac{Cov(P_t, P_t - P_t^*)}{V(P_t - P_t^*)}$$

which, using (5.22) and (5.33), can be written as

(5.24) $$\omega = \frac{Cov(Q_t, \tilde{P}_t) + V(\tilde{P}_t) - Cov(\tilde{P}_t, \tilde{P}_t^*)}{V(\tilde{P}_t) - 2Cov(\tilde{P}_t, \tilde{P}_t^*) + V(\tilde{P}_t)}$$

Again using (5.22) and (5.23), the variance inequality $V(P_t) < V(P_t^*)$ can be written as

(5.25) $$V(Q_t) + 2Cov(Q_t, \tilde{P}_t) + V(\tilde{P}_t) < V(Q_t) + V(\tilde{P}_t^*)$$

Combining (5.24) and (5.25), the null hypothesis of the variance bound can be expressed as $\omega < 0.5$. However, in Scott's framework, the present value model should be rejected for any value of ω that differs significantly from zero. On the basis of this, Durlauf and Hall concluded that Scott's regression tests were likely to be superior to volatility tests.

5.4.4 Campbell and Shiller (1987)

Campbell and Shiller (1987) developed several tests of the present value relation which are valid when prices and dividends are difference stationary. Define the spread, S_t, as[7]

$$(5.26) \qquad S_t \equiv P_t - \frac{\gamma}{1-\gamma}D_{t-1}$$

Expanding this using the present value formula we have

$$S_t = \gamma E_t D_t + \gamma^2 E_t D_{t+1} + \gamma^3 E_t D_{t+2} + \cdots - \frac{\gamma}{1-\gamma}D_{t-1}$$

$$= \frac{\gamma}{1-\gamma}[-D_{t-1} + (1-\gamma)E_t D_t + \gamma(1-\gamma)E_t D_{t+1} + \gamma^2(1-\gamma)E_t D_{t+2} + \cdots]$$

Since $E_t(E_{t+1}D_{t+1}) = E_t D_{t+1}$, we can rewrite this as

$$S_t = \frac{\gamma}{1-\gamma}[\Delta D_t + E_t(\gamma E_{t+1} D_{t+1} + \gamma^2 E_{t+1} D_{t+2} + \cdots - \gamma E_t D_t - \gamma^2 E_t D_{t+1} - \cdots)]$$

$$(5.27) \qquad S_t = \frac{\gamma}{1-\gamma}(\Delta D_t + E_t \Delta P_t)$$

Now if D_t and P_t are both I(1), then (5.27) indicates that S_t will be stationary. (5.26) thus implies that P_t and D_t cointegrate, with cointegrating parameter $\gamma(1-\gamma)^{-1}$.

Campbell and Shiller entered S_t and ΔD_t in a vector autoregression (VAR). They demonstrated that the present value model (5.1) implies a set of non-linear restrictions on the VAR parameters, on which they performed a Wald test. The restricted VAR was also used to obtain the "theoretical spread", S_t', which should be closely related to the actual spread. As well as being perfectly positively correlated, D_t and S_t' should satisfy the variance ratio

$$\frac{Var(S_t)}{Var(S_t')} \leq 1$$

[7] Campbell and Shiller used ex dividend prices. However, the dating of their variables differs from that previously used because they had start-of-period (opening) prices rather than closing prices.

Finally, a weak implication of the model was that S_t must linearly Granger-cause ΔD_t

Campbell and Shiller used Shiller's (1981) data extended to 1986. On implementing this approach, they did not find convincing evidence that P_t and D_t cointegrated. Despite this, the VAR model was estimated and the following results obtained:

i) the Wald test rejected the present value model, although not very strongly;

ii) S_t did Granger-cause ΔD_t;

iii) graphically, S_t and S_t' seemed to be unrelated, and the correlation coefficient between the two was not particularly large;

iv) although the estimated variance ratio was dramatically greater than unity (it was in fact 67.22), the standard errors of the ratio were so large that the null hypothesis could not be rejected.

They concluded that the data were not consistent with the present value model.

5.4.5 Bulkley and Tonks (1989)

In the only major application of variance bounds tests to UK data, Bulkley and Tonks (1989) argued that insufficient attention had been paid to the precise interpretation of the tests. In particular, Shiller's tests made the strongest assumptions regarding the information at the disposal of investors. Bulkley and Tonks pointed out that in Shiller's tests, forecast errors were due solely to the randomness in the true model, corresponding to strong-form rational expectations. Uncertainty about the model itself can only be taken account of in a weak-form setting. They therefore derived an alternative variance bound that stated that the variability of the actual stock price around a weak-form rational expectation of the perfect foresight price, must be strictly less that the variability of the ex post rational stock price around the same measure.

Both Shiller's strong-form test and the weak-form bound were applied to the annual BZW Equity Price Index and associated dividends over the period 1918-85 (ostensibly the same data as used in Chapter 3, although see Section 5.9 below). The real series were detrended using an estimate of the exponential growth rate and, in contrast to findings with US data, the resultant series were found to be stationary on the basis of Dickey-Fuller unit root tests. Bulkley and Tonks were therefore able to side-step the non-stationarity issues which had dominated the US literature. Using Shiller's (biased) method for calculating the terminal perfect foresight price, they found that the actual price series had a standard deviation of 7.028, whereas that for the ex post rational stock price was 1.383. The strong-from variance bound test therefore produced a gross violation with UK data comparable with those found for the US. However, when the weak-form test, which involved recursive estimation of the growth parameter, was applied, the bound was not violated. Subject, again, to the issue of statistical significance, more careful analysis of information availability appeared to reverse the original test findings.

5.5 Time-Varying Real Discount Factors

So far the assumption has been maintained that expected real discount rates are constant. The popularity of the constant real return model was motivated by three main factors. First, it can be derived from models in which investors are assumed to be risk neutral, and this was seen as a useful starting point for analysis. Second, a common practice in capital budgeting was to estimate discount rates using averages of historical time series of ex post returns. Finally, the fact that a constant discount rate makes the present value formula linear in expected dividends, making it more mathematically tractable, probably prolonged the life of the constant expected real return model. However, it is quite possible that the findings documented above which purport to reject the EMH are, in fact, indicating simply that the constant expected returns model is invalid, rather than implying some irrational alternative to market efficiency.

If expected returns vary through time then the present value model is

$$P_t = E_t\left[\sum_{i=1}^{\infty} D_{t+i}\prod_{j=1}^{i}(1+H_{t+j})^{-1}\right] \tag{5.28}$$

(5.28) is not directly testable, first because we do not observe expected discount rates, and second because for any given price and dividend series there will always be a corresponding discount rate series which makes (5.28) hold exactly (see the discussion of Shiller (1981) in Section 5.5.3 below). However, I now outline and discuss the two most popular models which give rise to expressions for time-varying required returns.

5.5.1 The Intertemporal CAPM

The CAPM requires little introduction. Assuming that investors' utility depends not only (positively) on investment returns, but also (negatively) on the risk attached to those returns, the CAPM predicts that investors' required return on a risky asset will be proportional to non-diversifiable (or systematic) risk. Measuring risk as the return variance, the Sharpe-Litner one-period CAPM predicts that in equilibrium all investors will hold the market portfolio (i.e. all risky assets will be held in the same proportions in each investor's portfolio), and the required return on this portfolio depends on the market return variance. Merton (1973) developed this idea in an intertemporal framework and showed that the excess return (that is the return over and above some risk-free rate, r_t) on the market portfolio is proportional to the expected variance of the market return:

(5.29) $$E_t(H_{t+1} - r_{t+1}) = \alpha E_t V_{mt+1}$$

where α is the harmonic mean of investors' coefficients of relative risk aversion and V_{mt+1} denotes the instantaneous market return variance in period t+1. Consequently, the more volatile the market is expected to be, and the more risk averse are investors, the greater the return required on the market portfolio to compensate investors for the perceived risk.

Of course, if the market return variance is constant then (5.29) reduces to a model of constant expected excess returns. Also, if investors are risk neutral (so that α=0), (5.9) implies that expected excess returns are zero (expected rates of return on all safe and risky assets are equated). In both of these cases, the only variation in expected returns derives from expected variation in the risk-free interest rate r_t.

5.5.2 The Consumption CAPM

In the CAPM, the individual investor's objective function is assumed to be fully determined by the standard deviation and return on the portfolio. The investor's problem is to maximise the expected return for any given level of portfolio risk. An alternative view of the determination of equilibrium returns in a well-diversified portfolio is provided by the consumption CAPM (C-CAPM). In this model, the investor maximises expected utility which depends only on current and future consumption (see Lucas 1978). Financial assets play a role in this model in that they help to smooth consumption over time. Securities are held to transfer purchasing power from one period to another. An asset is therefore more desirable if its return is expected to be high when consumption is expected to be low. Thus the systematic risk of an asset is determined by the covariance of the asset's return with consumption (rather than its covariance with the market return, as in the standard CAPM).

The individual investor is assumed to maximise the following objective function:

$$(5.30) \qquad U = E_t \sum_{j=0}^{\infty} \theta^j U(c_{t+j})$$

where $U(c)$ is the investor's utility function and θ is the utility discount factor which depends on the investor's subjective rate of time preference for consumption today versus consumption tomorrow (and is assumed to be constant). The budget constraint takes the form

$$(5.31) \qquad c_t = D_t X_{t-1} + P_t(X_{t-1} - X_t)$$

where D_t are dividends received, X_t denotes holdings of securities at time t that will yield dividend income at t+1, and P_t is the price of risky assets. The first term on the RHS of (5.31) is dividend income and the second term represents receipts from the sale of risky asset. The individual must divide current wealth between current consumption and holdings of risky assets to provide consumption for the future.

The first-order condition for this problem is

$$(5.32) \qquad U'(c_t) = E_t[(1 + H_{t+1})\theta U'(c_{t+1})]$$

or

$$(5.33) \qquad E_t(1+H_{t+1})S_{t+1}=1 \qquad\qquad \text{where } S_{t+1}\equiv\theta\frac{U'(c_{t+1})}{U'(c_t)}$$

S_{t+1} is the marginal rate of substitution of current for future (discounted) consumption (termed the intertemporal marginal rate of substitution, IMRS). (5.32) states that in equilibrium the consumer equates the expected discounted marginal utility of future consumption arising from the return from investment (RHS), with the marginal utility of current consumption. Using the definition of the covariance between any two series, (5.32) can be rewritten as follows:

$$E_t(1+H_{t+1})=\frac{[1-\mathrm{Cov}(H_{t+1},S_{t+1})]}{E_tS_{t+1}}$$

The C-CAPM therefore has two implications for the expected return on risky assets. First, the expected return on any asset depends negatively on the covariance of the asset return with the IMRS of consumption. Thus, investors place more value on assets which provide higher returns when the marginal utility of extra consumption is high relative to the future (i.e. S_{t+1} is low). The second implication is that the expected return varies inversely with the expected IMRS. This means that when investors expect that the value of future consumption will be high relative to current consumption (i.e. S_{t+1} is high), the required return is lower: agents are more willing to invest in risky assets in order to transform low-value current consumption into high-value future consumption.

From equation (5.33) we can see that the C-CAPM has the expected return equal to the MRS between current and future consumption. The present value formula then becomes

$$(5.34) \qquad P_t=E_t\sum_{j=1}^{\infty}D_{t+j}\prod_{i=1}^{j-1}S_{t+i}$$

5.5.3 Tests of Time-Varying Discount Rate Models

Shiller (1981) did not directly test a model of time-varying expected returns, but rather asked whether the discount rate series needed to satisfy (5.28) given P_t and D_t (denoted $\bar{H}_t$) was plausible. The perfect foresight price is, in this case,

$$P_t^* = \sum_{i=1}^{\infty} D_{t+i} \prod_{j=1}^{\infty} (1+\bar{H}_{t+j})^{-1}$$

Shiller used this to derive a lower bound on the variability of $\bar{H}_t$ given the variability of ΔP_t. He found that the standard deviation of $\bar{H}_t$ would have to be 4.36% for the Standard and Poor data and 7.36% for the Dow Jones data, which are very large numbers. These are the lowest possible standard deviation consistent with the model, and so Shiller concluded that time-varying real interest rates did not account for the observed excess volatility in stock prices.

Shiller and Grossman (1981) studied time-varying expected returns within the C-CAPM framework. The stochastic Euler equation (5.32) may be written in terms of prices and dividends rather than returns:

$$(5.35) \qquad P_t U'(c_t) = \theta E_t[(P_{t+1} + D_{t+1})U'(c_{t+1})]$$

Iterating on (5.35), the price of the risky asset at time t is the expected present value of dividends and a terminal price, discounted by the intertemporal marginal rate of substitution:

$$P_t = E_t\left[\sum_{j=1}^{T-t} \theta \frac{U'(c_{t+j})}{U'(c_t)} D_{t+j} + \theta^{T-t} \frac{U'(c_T)}{U'(c_t)} P_T\right]$$

The corresponding perfect foresight price is

$$P_t^* = \sum_{j=1}^{T-t} \theta \frac{U'(c_{t+j})}{U'(c_t)} D_{t+j} + \theta^{T-t} \frac{U'(c_T)}{U'(c_t)} P_T$$

Shiller and Grossman assumed a constant relative risk aversion utility function of the form

$$U(c) = \frac{1}{1-\alpha} c^{1-\alpha}$$

where α is the coefficient of relative risk aversion. They constructed P_t^* using annual observations of the dividend series relevant to the Standard and Poor Composite

5.6 Expected Returns and Discount Rates

Tests of the EMH are conditional on a model of equilibrium returns. I consider four models of the determinants of equilibrium expected returns. The first is that the expected real stock return is equal to a constant r,

$$(5.36) \qquad E_t h_{1t+1} = r$$

where h_{1t+1} is the one-period log real holding period return on a stock. In the second formulation, the safe rate varies through time, but the risk premium, rp, is assumed constant:

$$(5.37) \qquad E_t h_{1t+1} = E_t r_{t+1} + rp$$

Alternatively, we can allow for a time-varying risk premium with a constant safe rate:

$$E_t h_{1t+1} = r + rp_{t+1}$$

In particular, Merton's (1973, 1980) intertemporal CAPM has rp_{t+1}determined by the instantaneous market return variance V_t times the coefficient of relative risk aversion (see Model #1, Merton 1980), α, so that

$$(5.38a) \qquad E_t h_{1t+1} = r + \alpha E_t V_{t+1}$$

Finally, we may allow the safe rate to vary in the CAPM specification:

$$(5.38b) \qquad E_t h_{1t+1} = r_{t+1} + \alpha E_t V_{t+1}$$

5.7 The Dividend-Price Ratio Model

We can move from a model of equilibrium expected returns to testable hypotheses of the EMH as follows. The ex post one-period log real holding-period return on a stock is

$$(5.39) \qquad h_{1t+1} \equiv \log(P_{t+1} + D_{t+1}) - \log P_t$$

where P_t is the real stock price at the end of period t and D_{t+1} is the real dividend paid during period t+1[9]. A first-order Taylor expansion of equation (5.39) (Campbell and Shiller 1989) gives us the *approximate* one-period log real return ξ_{1t+1}:

$$(5.40) \qquad h_{1t+1} \approx \xi_{1t+1} \equiv \delta_t - \rho\delta_{t+1} + \Delta d_{t+1} + k$$

where k is a constant, ρ is a number a little smaller than unity (see Section 3.3), δ_t is the log dividend price ratio d_t-p_t and Δd_{t+1} is real dividend growth. Now define ξ_{it} as the discounted i-period log real return:

$$(5.41) \qquad \xi_{it} \equiv \sum_{j=1}^{i-1} \rho^j \xi_{1t+j}$$

ξ_{it} is the discounted sum of approximate one-period log real returns from tto t+i-1[10]. Combining equations (5.40) and (5.41) we can write the discounted i-period return as a linear function of δ_t, δ_{t+i} and Δd_{t+1+j}:

$$(5.42) \qquad \xi_{it+1} = \delta_t - \rho^i\delta_{t+i} + \sum_{j=0}^{i-1} \rho^j \Delta d_{t+1+j} + \frac{k(1-\rho^i)}{1-\rho}$$

This equation allows us to calculate the implications for the behaviour of the dividend-price ratio of a particular model of equilibrium returns. As Campbell and Shiller (1988) demonstrated, it is interesting to take the limit of (5.42) as i tends to infinity (with a stationary log dividend-price ratio):

$$\lim_{i\to\infty} \xi_{it} = (1-\rho)\sum_{j=0}^{\infty} \rho^j d_{t+j} - p_t + \frac{k}{1-\rho} = p_t^* - p_t + \frac{k}{1-\rho}$$

The first term in the rightmost expression is the discounted present value of actual dividends, which is more familiarly a log-linearisation of P_t^*, Shiller's (1981) perfect foresight stock price (with a constant discount rate). Thus we may think ofP_t^*-P_t as a

---

9 I have end-of-period observations, whereas Campbell and Shiller (1988, 1989) used opening prices. Consequently my time subscripts differ slightly from those in the original papers.

10 It is discounted so that the limit as i tends to infinity exists (assuming that δ_t and Δd_t are stationary), so that the concept of an infinite-horizon return is meaningful.

kind of infinite-horizon return. If infinite-horizon returns are predictable, Shiller's variance bound will be violated, and vice versa. Thus "excess volatility and predictability of multiperiod returns are not two phenomena, but one" (Campbell and Shiller 1988, reproduced in Shiller 1989, p155).

Rearranging (5.42), replacing ξ_{it+1} with the discounted sum of one-period returns, and taking expectations at the end of time t, we have the dividend-price ratio model

$$(5.43) \qquad \delta_t = \sum_{j=0}^{i-1} \rho^j E_t(\xi_{1t+1+j} - \Delta d_{t+1+j}) + \rho^i E_t \delta_{t+i} - \frac{k(1-\rho^i)}{1-\rho}$$

Equation (5.43) states that the dividend-price ratio depends on the discounted present value of expected one-period returns in excess of real dividend growth, and the terminal dividend-price ratio. Taking the limit of (5.43), assuming that $\lim_{i \to \infty} \rho^i \delta_{t+i} = 0$, we obtain

$$(5.44) \qquad \delta_t = \sum_{j=0}^{\infty} \rho^j E_t(\xi_{1t+1+j} - \Delta d_{t+1+j}) - \frac{k}{1-\rho}$$

Thus the log dividend-price ratio may be seen as an optimal forecast of all future required returns and real dividend growth. It will be useful to think of (5.44) as a log-linear approximation to the RVF which has been transformed to include only stationary variables. Given the definition of δ_t, (5.44) can be rewritten as

$$(5.45) \qquad p_t = (1-\rho)\sum_{j=0}^{\infty} \rho^j E_t d_{t+1+j} - \sum_{j=0}^{\infty} \rho^j E_t \xi_{1t+1+j} + \frac{k}{1-\rho}$$

which states that the log stock price is equal to the discounted present value of expected future log dividends minus the present value of future log required returns, plus a constant. The tests I employ to judge the degree to which the RHS and LHS of equation (5.44) are equal are equivalent to testing the equality of the two sides of (5.45). Only the non-stationarity of prices and dividends prevents one from analysing (5.45) directly.

The expected return is unobservable, but we can replace $E\xi_{1t+1}$ with the observable variable(s) implied by one of the four models of equilibrium returns outlined in Section 5.6 above. For example, in the case of constant expected excess returns (equation (5.37)), equation (5.43) becomes

$$(5.46) \qquad \delta_t = \sum_{j=0}^{i-1} \rho^j E_t(r_{t+1+j} - \Delta d_{t+1+j}) + \rho^i E_t \delta_{t+i} + \frac{(rp-k)(1-\rho^i)}{1-\rho}$$

or, in the limit,

$$(5.47) \qquad \delta_t = \sum_{j=0}^{\infty} \rho^j E_t(r_{t+1+j} - \Delta d_{t+1+j}) + \frac{(rp-k)}{1-\rho}$$

If we replace the expectations in (5.46) and (5.47) with rational forecasts, we can create a new series, the *theoretical* dividend-price ratio δ_t'. δ_t' is the dividend-price ratio that would obtain if the EMH, as embodied in a particular model of equilibrium returns, were true.

5.8 The VAR Methodology

The VAR methodology, which was developed by Campbell and Shiller (1988, 1989) to test various discount rate models, involves combining the dividend-price ratio model with a log-linear vector autoregression. The unrestricted VAR is used to forecast real dividend growth and discount rates. We then face two possibilities. First, it can be demonstrated that the present value model implies a set of non-linear restrictions on the parameters of the VAR. These can be tested directly to produce marginal rejection probabilities. Second, the VAR forecasts can be combined using the log-linear present value formula (5.43), and the resulting series compared informally with the observed log dividend-price ratio.

There are two major benefits of this approach as compared with previous studies. First, the equality of the actual and theoretical dividend-price ratios is equivalent to a test of the predictability of long-horizon returns, and the latter may be more powerful in detecting deviations from fundamentals than tests of the predictability of one-period returns (see Section 3.2 of Chapter 3). Second, the large number of metrics afforded to us by the VAR methodology allows one to make a judgement on the degree to which the model is in conformity with the data in a more intuitive and descriptive way than the rigorous study of marginal probability values. Close analysis of the full set of statistics allows one to pinpoint the major influences on the formal test results, and to assess the economic importance of any deviations from fundamentals which may show up in the statistical analysis.

5.8.1 VAR Forecasts and Present Value Restrictions

If we have a vector of variables $\mathbf{z}_t$, such that

$$(5.48) \qquad \mathbf{z}_t = \mathbf{A}\mathbf{z}_{t-1} + \varepsilon_t$$

then forecasts of future $\mathbf{z}_t$'s are easily obtained as[11]

$$(5.49) \qquad E_t(\mathbf{z}_{t+j}) = \mathbf{A}^j\mathbf{z}_t$$

For example, for the constant expected excess returns model, we define $\mathbf{z}_t \equiv [\delta_t, (r_t - \Delta d_t)]'$[12]. Also define $\iota_1' \equiv [1\ 0]$ and $\iota_2' \equiv [0\ 1]$, so that $\iota_1'\mathbf{z}_t \equiv \delta_t$ and $\iota_2'\mathbf{z}_t \equiv r_t - \Delta d_t$. I can now discuss more formally the implications of the present value model for the VAR estimates.

The first point to make is that the log dividend-price ratio must Granger-cause the measure of discount rates and dividend growth, in the case of the example, $r_t - \Delta d_t$. This is because, from equation (5.44), if the present value relation is true, the dividend-price ratio is an optimal forecast of future discount rates in excess of real dividend growth. Alternatively, given the fact that the dividend-price ratio is a predictor of stock returns (Section 3.2), if the market is efficient then this predictive power can be manifested only through an ability to predict discount rates and/or dividend growth.

The full set of cross-equation restrictions imposed on the VAR by the present value formula is quite complex. Replacing the expectations in equations (5.46) and (5.47) with the VAR forecasts, we have

[11] A VAR of any order can be rewritten in the form of (5.48), the 'companion form" (see Section 2.2.2). For ease of exposition I shall outline the VAR methodology assuming a VAR lag length p=1.

[12] All variables are entered as deviations from their mean.

$$\delta_t = \iota_1'\mathbf{z}_t = \sum_{j=0}^{\infty} \rho^j \iota_2'\mathbf{A}^{j+1}\mathbf{z}_t + \rho^i \iota_1'\mathbf{A}^i\mathbf{z}_t \equiv \delta_t'$$

where for the finite-horizon case (i=1,2,...,n),

(5.50a) $\delta_t' = [\iota_2'\mathbf{A}(\mathbf{I}-\rho\mathbf{A})^{-1}(\mathbf{I}-\rho^i\mathbf{A}^i)+\rho^i\iota_1'\mathbf{A}^i]\mathbf{z}_t$

and for the infinite-horizon case (i=∞),

(5.50b) $\delta_t' = \iota_2'\mathbf{A}(\mathbf{I}-\rho\mathbf{A})^{-1}\mathbf{z}_t$

Having obtained an estimate of **A** from the VAR, the construction of δ_t' is straightforward. One test of the hypothesis $\delta_t=\delta_t'$ is then a Wald test of the *non-linear* restrictions (for the finite- and infinite-horizon cases):

(5.51a) $\iota_1'(\mathbf{I}-\rho^i\mathbf{A}^i)-\iota_2'\mathbf{A}(\mathbf{I}-\rho\mathbf{A})^{-1}(\mathbf{I}-\rho^i\mathbf{A}^i)=0$

(5.51b) $\iota_1'-\iota_2'\mathbf{A}(\mathbf{I}-\rho\mathbf{A})^{-1}=0$

Intuitively, note that since δ_t is included in the VAR, we can rewrite (5.50a) or (5.50b) as

(5.52) $\delta_t' = f_1(\mathbf{A})\delta_t + f_2(\mathbf{A})(r_t-\Delta d_t)$

where $f_1(\mathbf{A})$ and $f_2(\mathbf{A})$ are non-linear functions of the elements of **A**. The joint hypothesis $f_1(\mathbf{A})=1$ and $f_2(\mathbf{A})=0$ constitutes the null in the non-linear Wald test. If we denote the actual stock price as P_t, and the expected discounted present value of future dividend as PV_t, then if the two are not equal there must exist a costless arbitrage opportunity. For example, if on any asset $P_t<PV_t$ then buying and holding this asset will provide the investor with dividend income of a greater value than the price of the stock. The hypothesis that $\delta_t=\delta_t'$ is the log-linear analogue to $P_t=PV_t$.

Notice that setting i=1 in (5.51a), we obtain a set of *linear* restrictions

(5.53) $\iota_1'(\mathbf{I}-\rho\mathbf{A})-\iota_2'\mathbf{A}=0$

Some further intuition about the interpretation of the VAR restrictions can be gained by noting that these linear restrictions imply that one-period returns are unpredictable. To see this, post-multiplying (5.53) by $\mathbf{z}_t$ and expanding the term in parentheses, we have

$$(5.54) \qquad \iota_1'\mathbf{z}_t - \rho\iota_1'\mathbf{A}\mathbf{z}_t - \iota_2'\mathbf{A}\mathbf{z}_t = 0$$

Noting that $\mathbf{Az}_t = E_t\mathbf{z}_{t+1}$, and given the definitions of ι_1 and ι_2, (5.54) becomes

$$(5.55) \qquad E_t(\delta_t - \rho\delta_{t+1} + \Delta d_{t+1} - r_{t+1}) = 0$$

or, combining (5.55) with (5.40),

$$(5.56) \qquad E_t(\xi_{1t+1} - r_{t+1}) = 0$$

Thus the linear restrictions (5.53) ensure that expected returns, in excess of a particular model of expected returns (here, constant expected excess returns), are zero.

The *non-linear* Wald statistics in equation (5.51a) test the (linearised) RVF for the finite-horizon cases (i=2,3..,n), i.e. with different finite terminal dividend-price ratios (see Mankiw, Romer and Shapiro 1991 for a similar analysis using volatility bounds statistics). These non-linear Wald tests are comparable to the implementation of standard variance bounds tests where, for practical purposes, the RVF is tested over a finite horizon (i.e. with the additional assumption that after the terminal date all forecast errors are zero and the EMH holds). However, because I use explicit forecasts of future dividend growth and discount rates, I can also test the RVF over an infinite horizon ($i=\infty$), invoking the transversality condition noted above. For i=1 the Wald test is linear in the parameters[13].

Expanding the VAR to include 3 variables allows one to determine whether, historically, changes in expected dividends or changes in expected discount rates have been more important in determining the dividend-price ratio. If we redefine $\mathbf{z}_t=[\delta_t, \Delta d_t, r_t]'$, the dividend-price ratio is

[13] However, note that Wald tests are not invariant to non-linear transformations, so that (5.51b) and (5.53) could yield different inferences (Gregory and Veall 1985).

$$\delta_t \equiv \iota_1' \mathbf{z}_t = \sum_{j=0}^{\infty} \rho^j (\iota_3' - \iota_2') \mathbf{A}^{j+1} \mathbf{z}_t \equiv \delta_t'$$

(5.57) $$\delta_t' = \iota_3' \mathbf{A}(\mathbf{I} - \rho\mathbf{A})^{-1} \mathbf{z}_t - \iota_2' \mathbf{A}(\mathbf{I} - \rho\mathbf{A})^{-1} \mathbf{z}_t$$

The first term on the RHS of (5.57) is the expected discounted present value of future discount rates, and the second term is that for future real dividend growth. One may write (5.57) more concisely as

(5.58) $$\delta_t' = \delta_{\xi t}' + \delta_{dt}'$$

where $$\delta_{\xi t}' = \iota_3' \mathbf{A}(\mathbf{I} - \rho\mathbf{A})^{-1} \mathbf{z}_t$$
$$\delta_{dt}' = -\iota_2' \mathbf{A}(\mathbf{I} - \rho\mathbf{A})^{-1} \mathbf{z}_t$$

$\delta_{\xi t}'$ is the component of δ_t' reflecting expected future discount rates, δ_{dt}' forecasts (minus) dividend growth and ι_3'=[0 0 1] "picks out" r_t from $\mathbf{z}_t$. This decomposition allows a further test of the EMH. If $\delta_t=\delta_t'$, then from (5.58) $\delta_{\xi t}'=(\delta_t-\delta_{dt}')$, and therefore one should find that the ratio $\sigma(\delta_t-\delta_{dt}')/\sigma(\delta_{rt}')$ and the correlation between $(\delta_t-\delta_{dt}')$ and $\delta_{\xi t}'$ both equal unity.

5.8.2 The Coefficient of Relative Risk Aversion

Both the CAPM and the C-CAPM models for equilibrium returns involve the coefficient of relative risk aversion, α. The Wald restrictions for these models are, in matrix notation,

$$\iota_1'(\mathbf{I} - \rho\mathbf{A}) - (\alpha\iota_3' - \iota_2')\mathbf{A} = 0$$

α is overidentified in this set of equations, and so they do not afford us a unique estimate of this coefficient. Also, if the VAR lag length p is greater than one, the companion form VAR matrix **A** is singular and so non-invertible.

Campbell and Shiller (1989) estimated α by choosing the value that minimises the Wald statistic of the above restrictions. Their choice of α therefore maximised the chances of passing the Wald test or, equivalently, made forecast returns correspond as closely as possible to forecast consumption growth of market volatility. However, the precise interpretation of any estimate of α using this technique is sometimes unclear. For example, for the C-CAPM they estimated α to be negative (see Section 5.8.5 below). Since the corresponding Wald statistic rejected the model, the negative α was taken as evidence against the model, rather than evidence that α is truly negative. There is no consensus as to the likely value of α for the UK economy, and so, following West (1988), I prefer to study the robustness of my findings to variations of α within a plausible range, which I take to be 1 to 5. Thus I impose a value of α for each round of estimation, and all tests are conditional on that value.

5.8.3 Actual and Theoretical Log Real Returns

The log dividend-price ratio is not the only variable that can be studied in this framework. From equation (5.40), we can see that the actual approximate log one-period return has the following definition (ignoring the constant):

$$(5.59) \qquad \xi_{1t} = \delta_t - \rho\delta_{t+1} + \Delta d_{t+1}$$

If we replace δ_t and δ_{t+1} in (5.59) with their theoretical counterparts, we obtain a theoretical approximate log one-period return ξ_{1t}'

$$(5.60) \qquad \xi_{1t}' = \delta_t' - \rho\delta_{t+1}' + \Delta d_{t+1}$$

ξ_{1t}' is what the returns on the stock would be if the expected returns model chosen from the alternatives in equations (5.36) to (5.38) holds. One may then compare the behaviour of ξ_{1t+1} and ξ_{1t+1}', the null hypothesis being that the two are equal. Since the means of the series are unconstrained, I report the ratio of their standard deviations, $\sigma(\xi_{1t}')/\sigma(\xi_{1t})$, and the correlation between the two.

5.8.4 VAR Methodology Tests

To summarise. A 2-variable VAR is used for the constant real and constant excess returns models. For each of these cases, I present the following tests:

(i) the Wald test of the RVF that is, $\delta_t=\delta_t'$ over different return horizons i (when i=1 this is a test of the unpredictability of one-period returns);

(ii) the ratio of the standard deviations of δ_t' and δ_t;

(iii) the correlation between δ_t and δ_t';

(iv) the ratio of the standard deviations of ξ_{1t} and ξ_{1t}';

(v) the correlation between ξ_{1t} and ξ_{1t}'.

Statistics (ii) to (v) equal unity under the null. For the 3-variable VAR, when the return horizon is infinite, I compute in addition the standard deviation ratio $\sigma(\delta_t-\delta_{dt}')/\sigma(\delta_{\xi t}')$ and the correlation between $(\delta_t-\delta_{dt}')$ and δ_{ξ}', which should both be unity.

In implementing the VAR approach, there are a couple of points to note. In estimation I use the dividend-price ratio and not the stock price because I expect the former to be stationary. However using the identity $p_t'= d_T\delta_t'$ I can easily calculate the theoretical stock price p_t' and compare it with the series for the actual stock price. Second, note that one should expect $\delta_t=\delta_t'$ and $\xi_{1t}=\xi_{1t}'$ even if the market has superior information to the econometrician. This is because, from equation (5.44), the dividend-price ratio is an optimal predictor of future dividend growth and discount rates. Therefore, the inclusion of δ_t in the VAR means that the VAR includes all relevant information for forecasting dividends and discount rates - the dividend-price ratio is a sufficient summary of the market's true information set. Put somewhat differently, if the null of market efficiency is rejected with a limited information set which includes the dividend-price ratio, then it would also be rejected using a larger information set.

5.8.5 Evidence from Previous Studies

Campbell and Shiller (1988, 1989) employed the VAR methodology to study the performance of various discount rate models on US data. In their 1988 paper, they confined their attention to the constant real returns and constant excess returns models. As documented in Section 3.2.2, they included a long moving average of accounting earnings in the VAR as a predictor of stock returns.

Using annual observations of the Standard and Poor Composite Index over the period 1871 to 1987, and for the constant real returns model, they estimated Corr(δ_t, δ_t')=0.175 and $\sigma(\delta_t')/\sigma(\delta_t)$=0.672. The two series did not therefore appear to be closely related. However, on the basis of time-series plots of the two variables, Campbell and Shiller reported some short-run correspondence between them. The correlation between actual and theoretical returns was quite high (0.915), but actual returns were much more volatile than their theoretical counterparts. Finally, the Wald tests rejected the efficient markets null more strongly as the return horizon increased. The hypothesis that $\delta_t=\delta_t'$ was rejected at the 0.1% level.

In their 1989 paper, Campbell and Shiller tested in addition the C-CAPM and the CAPM. For the C-CAPM, their estimate of the coefficient of relative risk aversion always had the wrong sign, and the null was always strongly rejected by the Wald tests. The reason appeared to be that the relationship between the dividend-price ratio and consumption growth was opposite to that required for the model to be successful. In the VAR, a high dividend-price ratio forecast low consumption growth over the year. Since the dividend-price ratio and the expected stock return were positively related, there was a resultant negative relation between the expected return and expected consumption growth, which required α to be negative for consistency with the C-CAPM. For the CAPM model, the estimates of α were positive but insignificant. The problem here was that the log dividend-price ratio did not Granger-cause the market return variance, so that the former's ability to predict returns was not due to it proxying for expected market volatility. The Wald tests strongly rejected this model. Campbell and Shiller concluded that, whilst there was some evidence that actual and theoretical dividend-price ratios moved together in the short run, the discount rate models were unhelpful in explaining observed stock price movements.

Lund and Engsted (1993) concentrated on analysing the performance of the C-CAPM using data from the Danish, German, Swedish and UK stock markets. Overall, their findings were similar to those presented by Campbell and Shiller. In particular, covariation between consumption growth and the dividend-price ratio had the opposite sign to that predicted by the model, with the result that estimates of the coefficient of relative risk aversion were generally negative. Consequently, whilst the likelihood ratio tests of the present value relation were not found to reject the consumption-based asset pricing model, Lund and Engsted concluded that there was little evidence in favour of the C-CAPM.

5.9 Data and Empirical Results

The tests are applied to the BZW (value-weighted) equity stock price index and related dividends (these are the same data as used in Chapter 3). The interest rate used is the return on four consecutive investments in 3-month Treasury Bills. Observations are taken at the end of each year from 1918 to 1993.

5.9.1 Stationarity Issues

Bulkley and Tonks (1989) noted the problems with first-generation variance bounds tests when the real stock price is non-stationary, and so applied a Dickey-Fuller test to the log real BZW price series, suitably detrended. Their conclusion was that the UK data were in fact stationary, and proceeded on this assumption. However, BZW have recently revised their cost of living index, with the effect of reducing its rate of increase. Consequently there is now a more prominent trend in their real stock price series.

Bulkley and Tonks reported a regression of the log real stock price on a constant and its own lag from 1918 to 1985:

$$\ln P_t = \underset{(0.216)}{0.632} + \underset{(0.071)}{0.795} \ln P_{t-1}$$

The DF t-statistic is -2.894 which, when compared to the 10% critical value of -2.60, implies that the log real stock price is a stationary process. However, when I perform the more usual DF regression on the revised data from 1918 to 1993, the estimated equation is

$$\Delta \ln P_t = \underset{(0.378)}{0.931} - \underset{(0.065)}{0.158} \ln P_t$$

which gives rise to a DF statistic of -2.418 with a marginal probability value of 0.260, indicating non-rejection of the null of non-stationarity at the 10% level. However, since the constant term is significant, it is preferable to base inference on a DF regression that includes a time trend (Perron 1988). This regression produces the following results

$$\Delta \ln P_t = \underset{(0.394)}{1.281} + \underset{(0.001)}{0.002}\, t - \underset{0.072}{0.238} \ln P_{t-1}$$

The resultant DF statistic of -3.278 has a marginal p-value of 0.079. The non-stationarity null is therefore not rejected at the 5% level but is rejected at the 10% level.

Since Bulkley and Tonks used the detrended real stock price, rather than the log real stock price, it seems more relevant to their study to look at the stationarity properties of this series. A regression of the log real stock price on a time trend gives us an estimated long-run growth coefficient of 0.013. Using this to detrend the real stock price series, the DF regression is

$$\Delta P_t^d = \underset{(16.617)}{48.143} - \underset{(0.192)}{0.056}\, t - \underset{(0.073)}{0.218}\, P_t^d$$

where P^d_t denotes the detrended real stock price. The DF statistic is -2.978, which has a marginal p-value of 0.130. The conclusion is therefore that the hypothesis that the detrended real stock price series is non-stationary cannot be rejected at the 10% level. Omitting the time trend from the above regression has little effect on the results.

However, the VAR methodology does not use stock prices; instead I use the variables δ_t, Δd_t, r_t and V_t, and as long as these are found to be stationary the application of the VAR methodology is statistically valid. Table 5.1 reports the results of applying Phillips-Perron (1988) unit root tests to the four series. These statistics are preferable to the DF statistics reported above because they allow for heterogeneity in the distributions of the underlying processes[14]. The results reported were obtained using a truncation lag of 1, although increasing this value was found to have no qualitative impact on the findings. For all four series, the null of non-stationarity is easily rejected.

---

14 DF tests were used above only to aid comparison with the findings of Bulkley and Tonks.

Table 5.1: Phillips-Perron Unit Root Tests

Variable	PP Statistic[1]	
	t-statistic	$T(\phi-1)$
δ_t	-5.043	-34.249
Δd_t	-6.159	-48.094
r_t	-6.242	-49.990
V_t	-6.791	-55.911

[1] 10% critical values are -2.09 for the t-statistic and -13.4 for the $T(\phi-1)$ statistic. All reported statistics are calculated using a truncation lag of 1. Increasing the lag length made no qualitative difference to the results.

Table 5.2: Constant Expected Real Returns

		Return Horizon (years)					
		1	2	3	5	10	infinity
(i)	Wald statistic (p-value)	22.553 (0.000)	38.433 (0.000)	52.193 (0.000)	64.466 (0.000)	67.997 (0.000)	68.152 (0.000)
(ii)	$\sigma(\delta_t')/\sigma(\delta_t)$ (s.e.)	0.574 (0.089)	0.395 (0.114)	0.310 (0.117)	0.246 (0.110)	0.223 (0.105)	0.223 (0.105)
(iii)	$Corr(\delta_t,\delta_t')$ (s.e.)	0.953 (0.067)	0.891 (0.128)	0.833 (0.153)	0.761 (0.205)	0.724 (0.247)	0.722 (0.249)
(iv)	$\sigma(\xi_{1t}')/\sigma(\xi_{1t})$ (s.e.)	0.567 (0.093)	0.405 (0.080)	0.340 (0.070)	0.300 (0.058)	0.289 (0.052)	0.288 (0.052)
(v)	$Corr(\xi_{1t},\xi_{1t}')$ (s.e.)	0.959 (0.052)	0.891 (0.097)	0.819 (0.129)	0.730 (0.192)	0.687 (0.225)	0.685 (0.227)

5.9.2 Sample Statistics

As reported in Section 3.5.2, I estimate ρ to be 0.951, and the approximation error between the real log return h_t and the approximate real log return ξ_t is tolerably small.

5.9.3 Empirical Results

Initially I report results assuming a VAR lag length of one. This makes the exposition a little easier, and the sensitivity of the findings to variations in the VAR lag length is analysed in Section 5.9.7 below. A GMM estimator was employed to correct the covariance matrix for possible heteroscedasticity (see Section 2.3).

The first three rows of Tables 5.2-5.6 show the basic statistics: the non-linear Wald tests, the standard deviation ratio $\sigma(\delta'_t)/\sigma(\delta_t)$ and the correlation $Corr(\delta'_t, \delta_t)$ for return horizons i=1,2,3,5,10 and infinity. In the case of expected real returns depending on the variance of returns I report results for the coefficient of relative risk aversion equal to 2.

5.9.4 Constant Expected Real Returns

The estimated equations for the constant real returns model are[15]

$$\delta_t = \underset{(0.103)}{0.522}\,\delta_{t-1} - \underset{(0.523)}{0.162}\,\Delta d_{t-1} + e_{1t}$$

$$\Delta d_t = -\underset{(0.041)}{0.055}\,\delta_{t-1} + \underset{(0.095)}{0.315}\,\Delta d_{t-1} + e_{2t} \tag{5.61}$$

The first thing to note is that the log dividend-price ratio does not Granger-cause real dividend growth. This is an inauspicious start for the constant real returns model, since the fact that the dividend-price ratio is a significant predictor of returns (price movements) can only be reconciled with the constant real returns model if the dividend-price ratio forecasts changes in dividends. Given the insignificant coefficient on δ_{t-1} in equation (5.61), the constant real returns model cannot explain why returns are forecast by the dividend-price ratio. All of the more formal analysis bears out this conclusion.

---

15 Heteroscedasticity-consistent GMM standard errors in parentheses.

Table 5.2 contains the test statistics for the constant real returns model. The statistic for one-year returns is 22.553 which is significant at the 0.001% level. Thus, as expected, when one-period returns are regressed on the dividend-price ratio and real dividend growth, the null hypothesis of unpredictability is easily rejected. The marginal rejection probabilities are extremely small for all return horizons, providing no support whatsoever for the log-linear version of the present value formula. The fact that this model is easily rejected on the basis of the Wald statistics over all horizons, with the rejection getting worse as the horizon increases, is indicative of return predictability increasing with the return horizon, as demonstrated in a single-equation context by Fama and French (1988a).

Looking at the Wald restrictions in more detail, if the constant expected real return model were true, a regression of δ_t' on δ_t and Δd_t would place a unit weight on δ_t and zero on Δd_t. Concentrating for the moment on the infinite-horizon case (final column, Table 5.2), the VAR estimates imply

$$\delta_t' = \underset{(0.118)}{0.158}\,\delta_t - \underset{(0.199)}{0.484}\,\Delta d_t$$

which is clearly a long way from the null hypothesis. Although the correlation between δ_t and δ_t' is quite high (0.722, s.e.=0.249), and only a little more than one standard error from unity, the ratio of their standard deviations is estimated to be 0.223 (s.e.=0.105). The theoretical dividend-price ratio is therefore less than a quarter as variable than the observed series. This is illustrated in Figure 5.1 where, despite some close correspondence in the *sign* of the movements in δ_t and δ_t', the former is much more volatile. This "excess volatility" is also apparent in Figure 5.2, which compares the actual and theoretical log real stock prices, p_t and p_t'. The findings here are very similar to those reported by Campbell and Shiller (1988, 1989) for US data.

5.9.5 Constant Expected Excess Returns

The same qualitative conclusions noted above apply to the constant *excess* returns model. Informally, this is obvious from the fact that Figures 5.3 and 5.4, which compare the actual and theoretical log dividend-price ratios and log real stock

Table 5.1: Comparison of Actual and Theoretical Log Dividend-Price Ratio
Constant Expected Real Returns

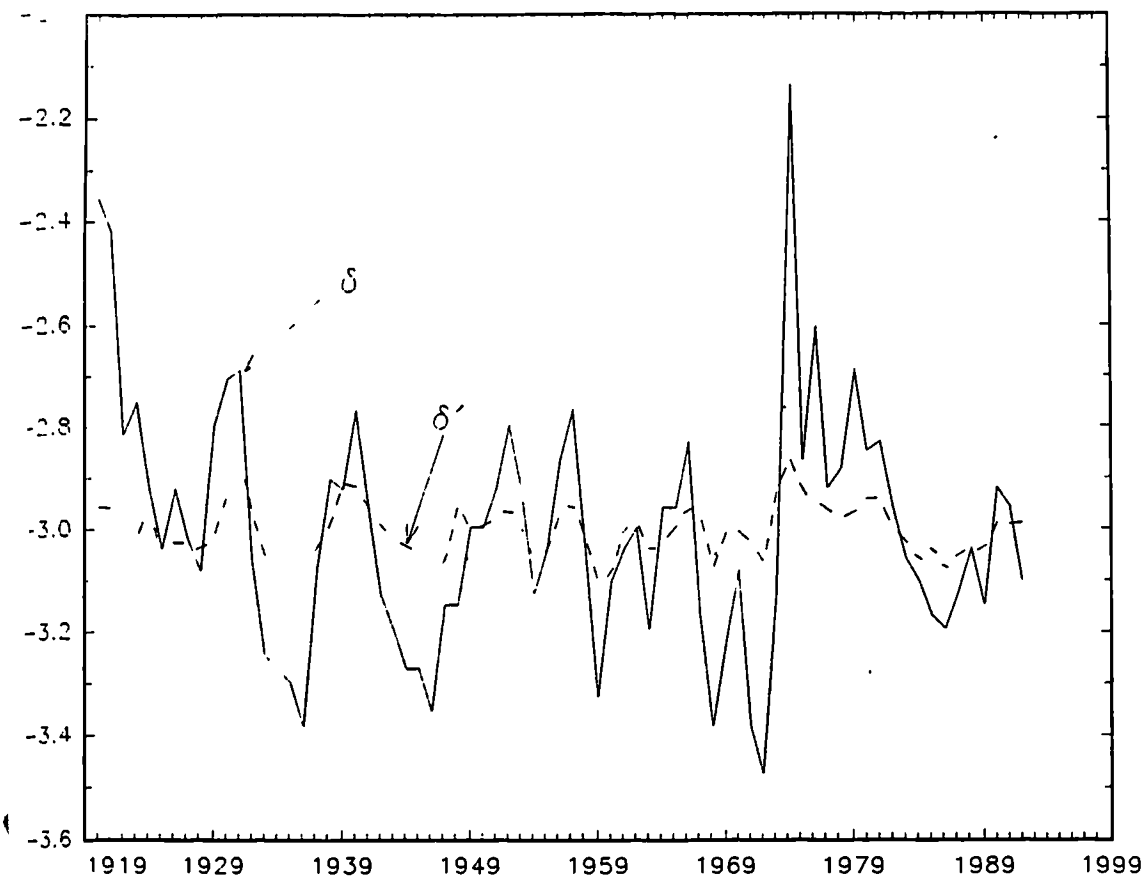

Table 5.2: Comparison of Actual and Theoretical Log Real Stock Price
Constant Expected Real Returns

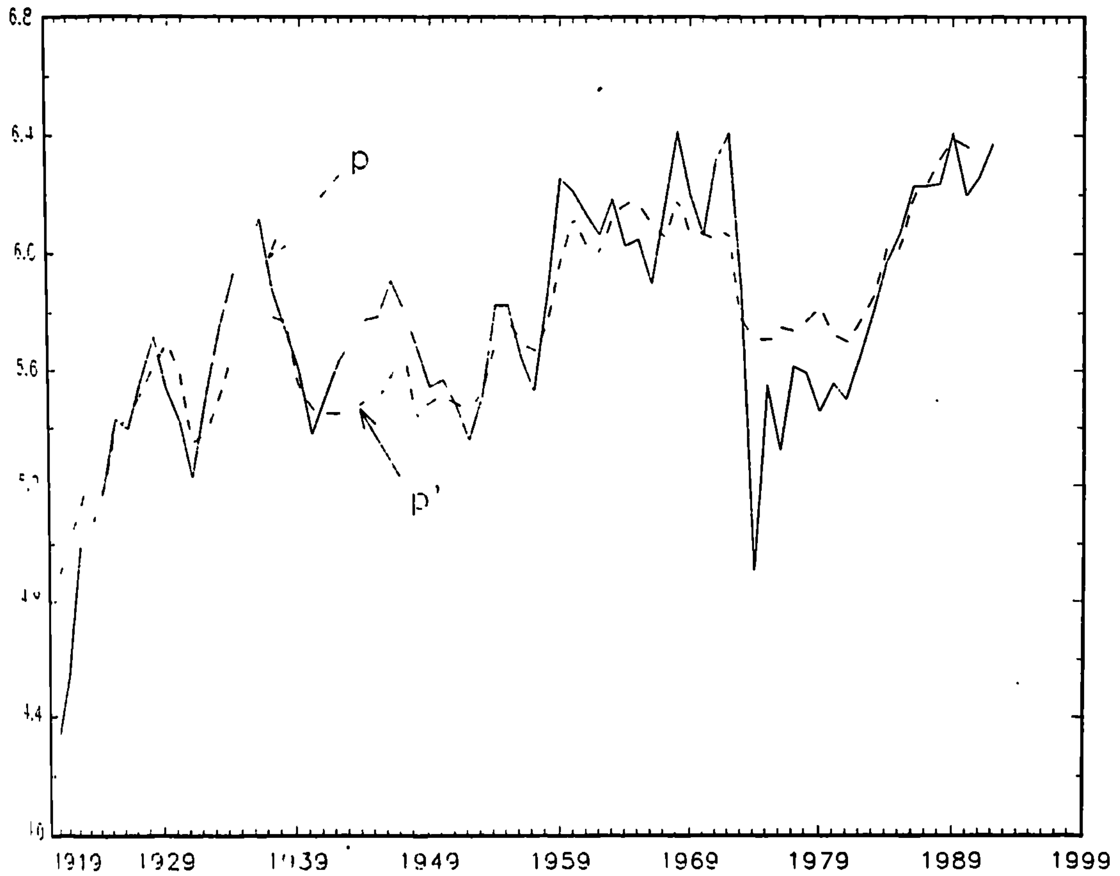

Table 5.3: Comparison of Actual and Theoretical Log Dividend-Price Ratio
Constant Expected Excess Returns

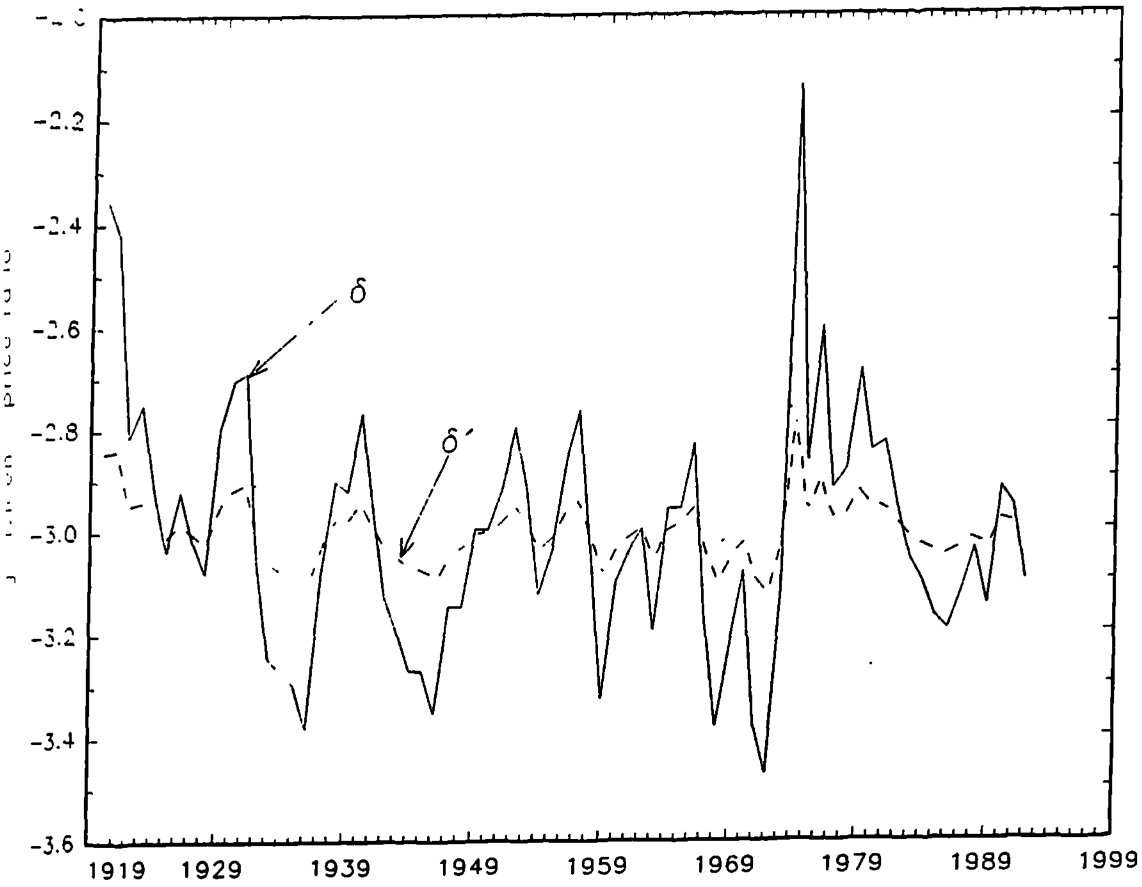

Table 5.4: Comparison of Actual and Theoretical Log Real Stock Price
Constant Expected Excess Returns

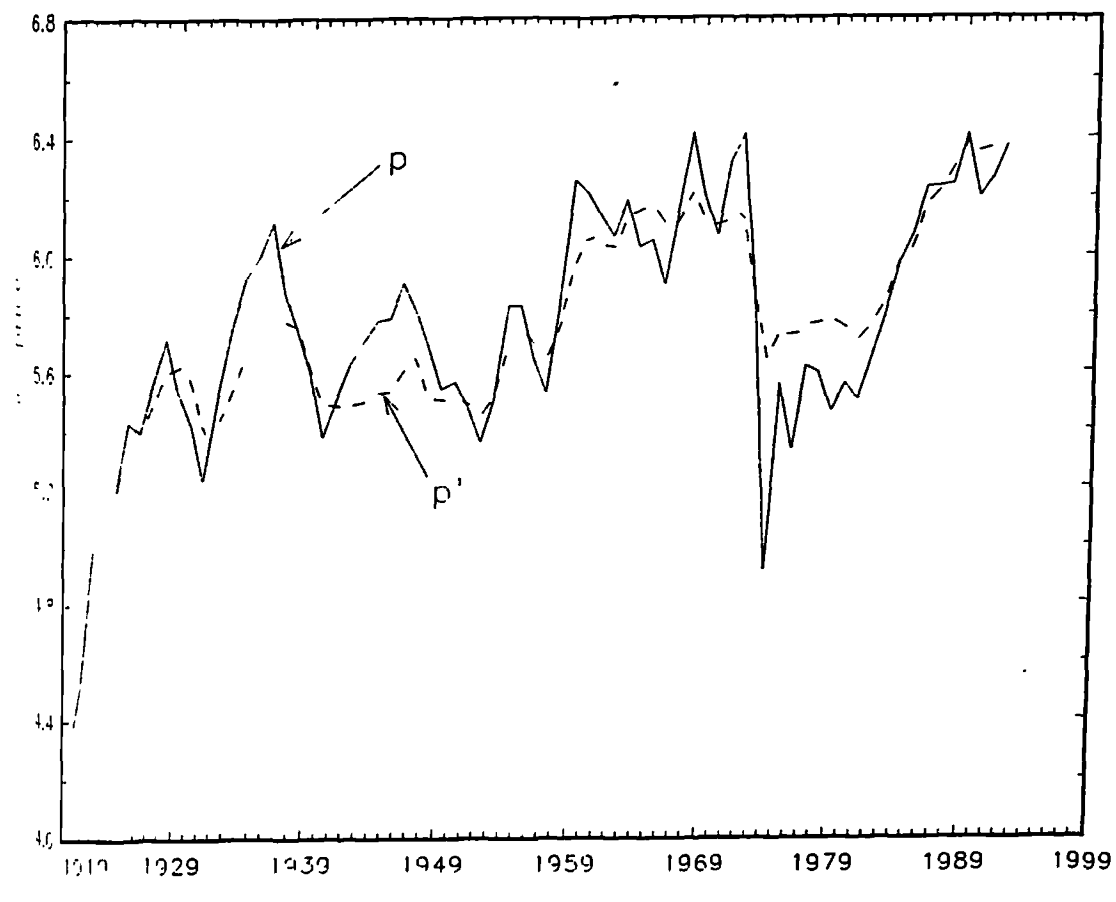

prices for the constant excess returns model, are almost identical to Figures 5.1 and 5.2. Using $\mathbf{z}_t=(\delta_t, r_t-\Delta d_t)'$ in the VAR (i.e. the impact of r_t is restricted to equal that of $-\Delta d_t$), the estimated equations are

$$\delta_t = \underset{(0.104)}{0.528}\,\delta_{t-1} + \underset{(0.405)}{0.175}\,(r_{t-1} - \Delta d_{t-1}) + e_{1t}$$

$$(r_t - \Delta d_t) = -\underset{(0.054)}{0.123}\,\delta_{t-1} + \underset{(0.124)}{0.088}\,(r_{t-1} - \Delta d_{t-1}) + e_{2t}$$

The log dividend-price ratio does have forecasting power over the combined discount rate and dividend growth series, so the weakest implication of the dividend-price ratio model is satisfied. However, direct formal comparison of the δ_t and δ_t' series, reported in Table 5.3, does not provide much support for this model. First of all, although the Wald statistics are uniformly lower than for the constant real returns model, the lowest statistics is 10.474 (for a 1-year horizon), which can be rejected at the 0.5% level. Once again, the Wald statistics increase uniformly with the return horizon. All of the estimated correlations between δ_t and δ_t' are high, ranging from 0.984 to 0.986, but once again there appears to be insufficient volatility in the δ_t' series. The standard deviation ratio $\sigma(\delta'_t)/\sigma(\delta_t)$ is estimated to be 0.624 (s.e.=0.134) for a 1-year return horizon, falling to 0.261 (s.e.=0.129) when the horizon is 10 years or more.

Although the two-variable VAR model benefits from the fact that the variable $(r_t-\Delta d_t)$ does not depend on a price deflator, more information regarding the lack of success of this model may be gleaned from entering the variables separately and performing the decomposition of δ_t into its dividend and discount rate components. When r_t and Δd_t are entered separately in the VAR (Table 5.4), the metrics are nearly identical to those in Table 5.3. However, for the "unrestricted" case we find that the standard deviation ratio $\sigma(\delta_t-\delta_{dt}')/\sigma(\delta_{\xi t}')$ is much greater than unity (4.056, s.e.=1.172), although $\mathrm{Corr}(\delta_t-\delta_{dt}', \delta_{\xi t})$ is close to one (0789, s.e.=0.356). The interpretation here is that the movement in the dividend-price ratio which is not attributable to real dividend growth $(\delta_t-\delta_{dt}')$, is not well approximated by a discount rate which moves only with the

Table 5.3: Constant Expected Excess Returns - Restricted

		Return Horizon (years)					
		1	2	3	5	10	infinity
(i)	Wald statistic (p-value)	10.474 (0.005)	15.667 (0.000)	21.194 (0.000)	29.233 (0.000)	33.180 (0.000)	33.318 (0.000)
(ii)	$\sigma(\delta_t')/\sigma(\delta_t)$ (s.e.)	0.624 (0.134)	0.427 (0.150)	0.336 (0.147)	0.276 (0.135)	0.261 (0.129)	0.261 (0.129)
(iii)	$Corr(\delta_t,\delta_t')$ (s.e.)	0.985 (0.010)	0.986 (0.005)	0.986 (0.004)	0.985 (0.013)	0.984 (0.017)	0.984 (0.017)
(iv)	$\sigma(\xi_{1t}')/\sigma(\xi_{1t})$ (s.e.)	0.693 (0.156)	0.529 (0.139)	0.456 (0.119)	0.412 (0.101)	0.402 (0.094)	0.401 (0.094)
(v)	$Corr(\xi_{1t},\xi_{1t}')$ (s.e.)	0.971 (0.026)	0.933 (0.056)	0.894 (0.087)	0.853 (0.114)	0.841 (0.120)	0.840 (0.120)

[1] The statistics in this table were calculated from a 2-variable VAR. The variables used were δ_t and $(r_t-\Delta d_t)$, so that the impact of r_t in the VAR is restricted to equal that of $-\Delta d_t$.

Table 5.4: Constant Expected Excess Returns - Unrestricted[1]

		Return Horizon (years)					
		1	2	3	5	10	infinity
(i)	Wald statistic (p-value)	16.185 (0.001)	24.808 (0.000)	33.341 (0.000)	43.403 (0.000)	48.589 (0.000)	48.829 (0.000)
(ii)	$\sigma(\delta_t')/\sigma(\delta_t)$ (s.e.)	0.644 (0.143)	0.446 (0.149)	0.353 (0.146)	0.291 (0.138)	0.275 (0.133)	0.274 (0.133)
(iii)	$Corr(\delta_t,\delta_t')$ (s.e.)	0.972 (0.020)	0.959 (0.049)	0.961 (0.052)	0.975 (0.033)	0.981 (0.021)	0.981 (0.021)
(iv)	$\sigma(\xi_{1t}')/\sigma(\xi_{1t})$ (s.e.)	0.695 (0.172)	0.517 (0.137)	0.446 (0.111)	0.413 (0.094)	0.409 (0.090)	0.409 (0.090)
(v)	$Corr(\xi_{1t},\xi_{1t}')$ (s.e.)	0.963 (0.027)	0.929 (0.048)	0.899 (0.076)	0.865 (0.106)	0.852 (0.113)	0.852 (0.113)

[1] The statistics in this table were derived from a 3-variable VAR, where δ_t, r_t and Δd_t were entered separately.

Decomposition of Dividend-price Ratio

When $i = \infty$:	$\sigma(\delta_t-\delta_{dt}')/\sigma(\delta_{rt}')$ (s.e.)	4.056 (1.172)
	$Corr(\delta_t-\delta_{dt}', \delta_{rt}')$ (s.e.)	0.789 (0.356)

safe rate of interest ($\delta_{\xi t}$). The correlation is high, but the safe rate does not move enough to explain movements in δ_t-δ_{dt}'.

5.9.6 *The CAPM*

The above gloomy assessments of market efficiency contrast with the results from our CAPM-type models. Table 5.5 presents the results for a constant safe rate but a time-varying risk premium. The estimated VAR equations for the market volatility model are

$$\delta_t = \underset{(0.108)}{0.502}\,\delta_{t-1} - \underset{(0.503)}{0.128}\,\Delta d_{t-1} + \underset{(0.115)}{0.259}\,V_{mt-1} + e_{1t}$$

$$\Delta d_t = -\underset{(0.042)}{0.066}\,\delta_{t-1} + \underset{(0.087)}{0.335}\,\Delta d_{t-1} + \underset{(0.062)}{0.151}\,V_{mt-1} + e_{2t}$$

$$V_{mt} = \underset{(0.140)}{0.282}\,\delta_{t-1} - \underset{(0.174)}{0.145}\,\Delta d_{t-1} + \underset{(0.201)}{0.153}\,V_{mt-1} + e_{3t}$$

The most important finding here is that the log dividend-price ratio Granger-causes the market return variance (the coefficient on δ_{t-1} in the third equation is significant at the 4.8% level). Also notice that the market return variance has significant predictive power over the dividend-price ratio. It is a possibility, therefore, that the ability of the dividend-price ratio to predict stock returns is due to a close relationship between this variable and market volatility.

Turning now to the formal statistical analysis, the first row of Table 5.5 indicates that the lowest p-value for the Wald test is 47.1% (one-period horizon), rising to over 60% as the return horizon increases. The clear implication is that, once time-variation in the expected market return variance has been accounted for, the variables in the VAR have no predictive power over returns. Thus, although the dividend-price ratio is not a significant predictor of real dividend growth, its role in predicting stock price movements may stem from the fact that it appears to proxy for investors' perceived risk.

Again looking directly at the relationship between δ_t and δ_t', in the infinite-horizon case, my estimates yield

Table 5.5: Constant Safe Rate, Time-varying Risk Premium (α=2)

		Return Horizon (years)					
		1	2	3	5	10	infinity
(i)	Wald statistic (p-value)	2.525 (0.471)	2.017 (0.569)	1.890 (0.596)	1.852 (0.604)	1.859 (0.602)	1.862 (0.602)
(ii)	$\sigma(\delta_t')/\sigma(\delta_t)$ (s.e.)	1.194 (0.257)	1.412 (0.493)	1.529 (0.637)	1.631 (0.775)	1.683 (0.847)	1.687 (0.853)
(iii)	$Corr(\delta_t,\delta_t')$ (s.e.)	0.946 (0.051)	0.945 (0.055)	0.945 (0.057)	0.945 (0.058)	0.945 (0.058)	0.945 (0.058)
(iv)	$\sigma(\xi_{1t}')/\sigma(\xi_{1t})$ (s.e.)	0.984 (0.204)	1.165 (0.367)	1.264 (0.479)	1.349 (0.591)	1.390 (0.650)	1.394 (0.655)
(v)	$Corr(\xi_{1t},\xi_{1t}')$ (s.e.)	0.907 (0.115)	0.895 (0.126)	0.890 (0.131)	0.887 (0.136)	0.885 (0.137)	0.885 (0.138)

Decomposition of Dividend-price Ratio

When $i = \infty$: $\sigma(\delta_t-\delta_{dt}')/\sigma(\delta_{rt}')$ 0.593
(s.e.) (0.316)

$Corr(\delta_t-\delta_{dt}', \delta_{rt}')$ 0.934
(s.e.) (0.078)

Table 5.6: Time-varying Safe Rate and Risk Premium (α=2)

		Return Horizon (years)					
		1	2	3	5	10	infinity
(i)	Wald statistic (p-value)	2.635 (0.451)	2.081 (0.556)	2.024 (0.568)	2.090 (0.554)	2.196 (0.533)	2.212 (0.530)
(ii)	$\sigma(\delta_t')/\sigma(\delta_t)$ (s.e.)	1.253 (0.225)	1.496 (0.418)	1.633 (0.537)	1.757 (0.647)	1.817 (0.670)	1.822 (0.704)
(iii)	$Corr(\delta_t,\delta_t')$ (s.e.)	0.964 (0.033)	0.965 (0.032)	0.965 (0.032)	0.965 (0.033)	0.965 (0.033)	0.965 (0.033)
(iv)	$\sigma(\xi_{1t}')/\sigma(\xi_{1t})$ (s.e.)	1.062 (0.201)	1.261 (0.324)	1.373 (0.415)	1.475 (0.506)	1.524 (0.553)	1.528 (0.557)
(v)	$Corr(\xi_{1t},\xi_{1t}')$ (s.e.)	0.944 (0.076)	0.937 (0.083)	0.933 (0.089)	0.929 (0.094)	0.927 (0.096)	0.927 (0.096)

Decomposition of Dividend-price Ratio

When $i = \infty$: $\sigma(\delta_t-\delta_{dt}')/\sigma(\delta_{rt}')$ 0.554
(s.e.) (0.216)

$Corr(\delta_t-\delta_{dt}', \delta_{rt}')$ 0.962
(s.e.) (0.050)

Table 5.5: Comparison of Actual and Theoretical Log Dividend-Price Ratio
Volatility Model, i=1 (α=2)

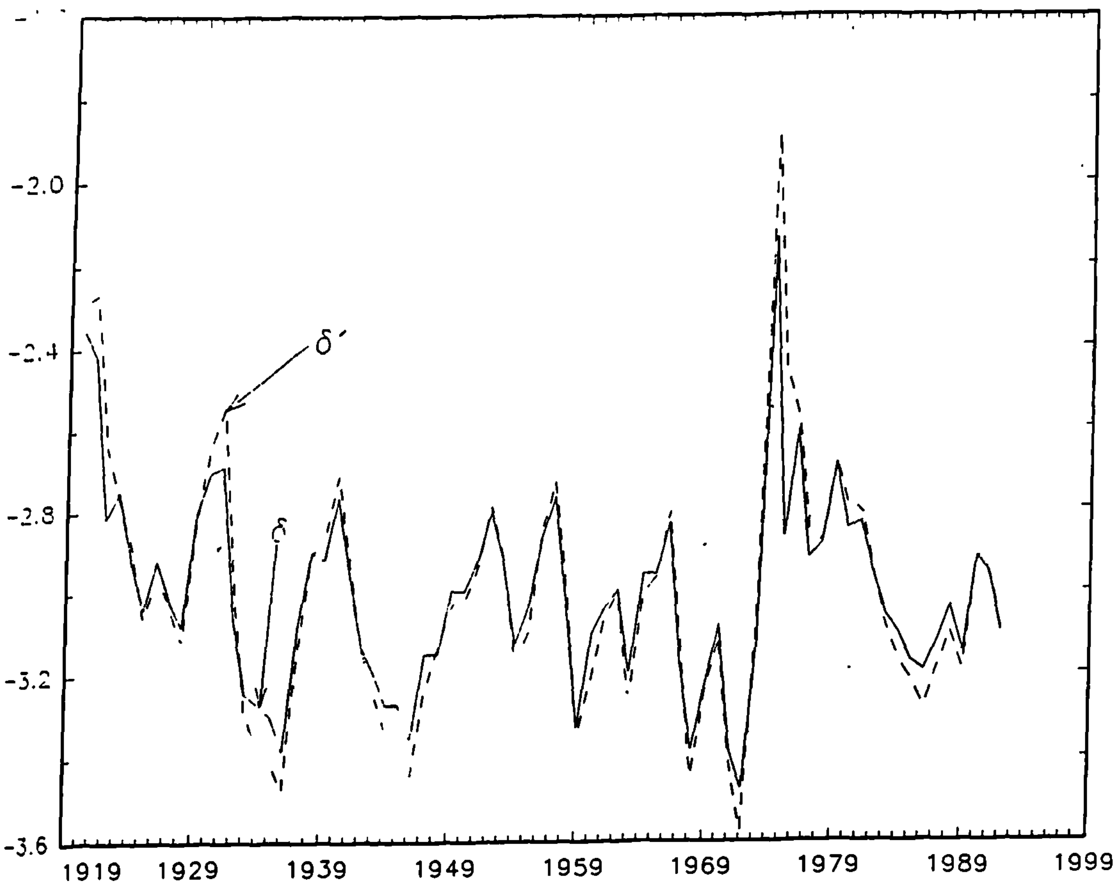

Table 5.6: Comparison of Actual and Theoretical Log Real Stock Price
Volatility Model, i=1 (α=2)

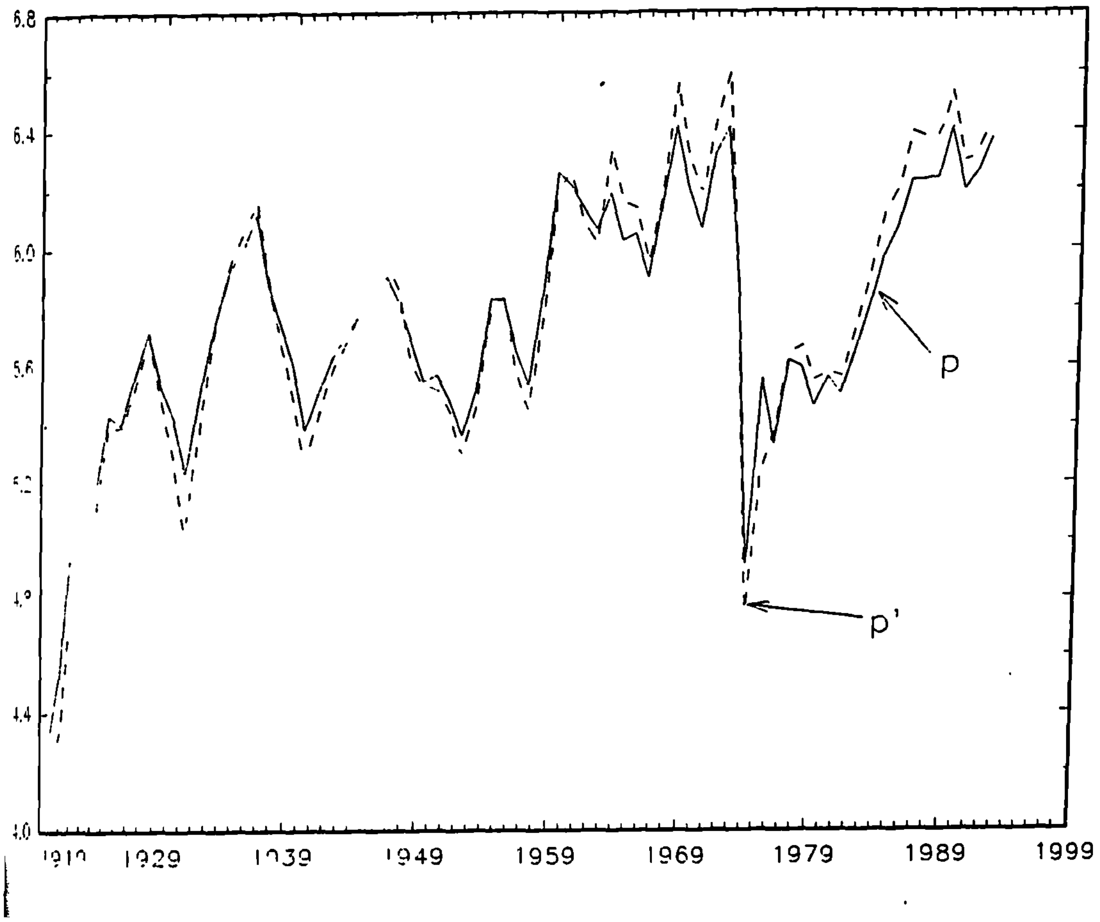

Table 5.7: Comparison of Actual and Theoretical Log Real Stock Price Volatility Model, $i=\infty$ ($\alpha=2$)

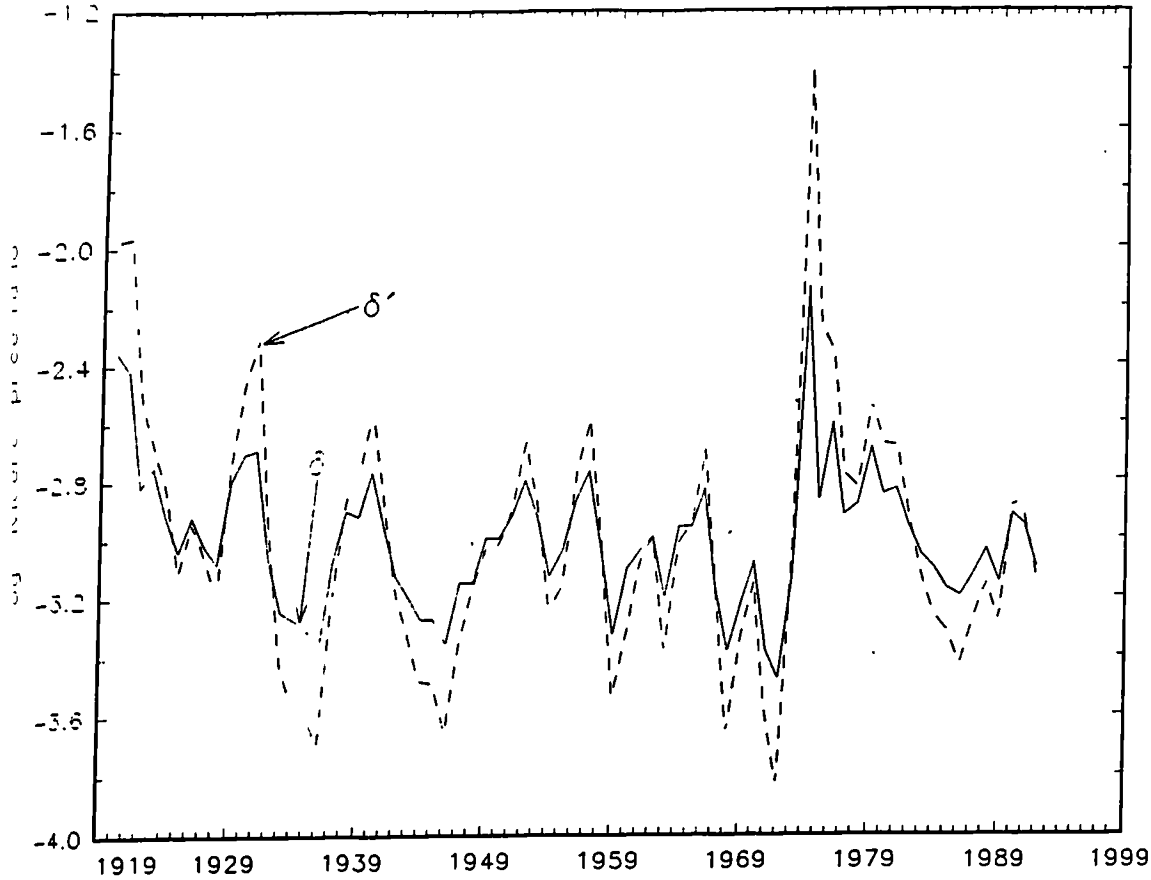

Table 5.8: Comparison of Actual and Theoretical Log Real Stock Price Volatility Model, $i=\infty$ ($\alpha=2$)

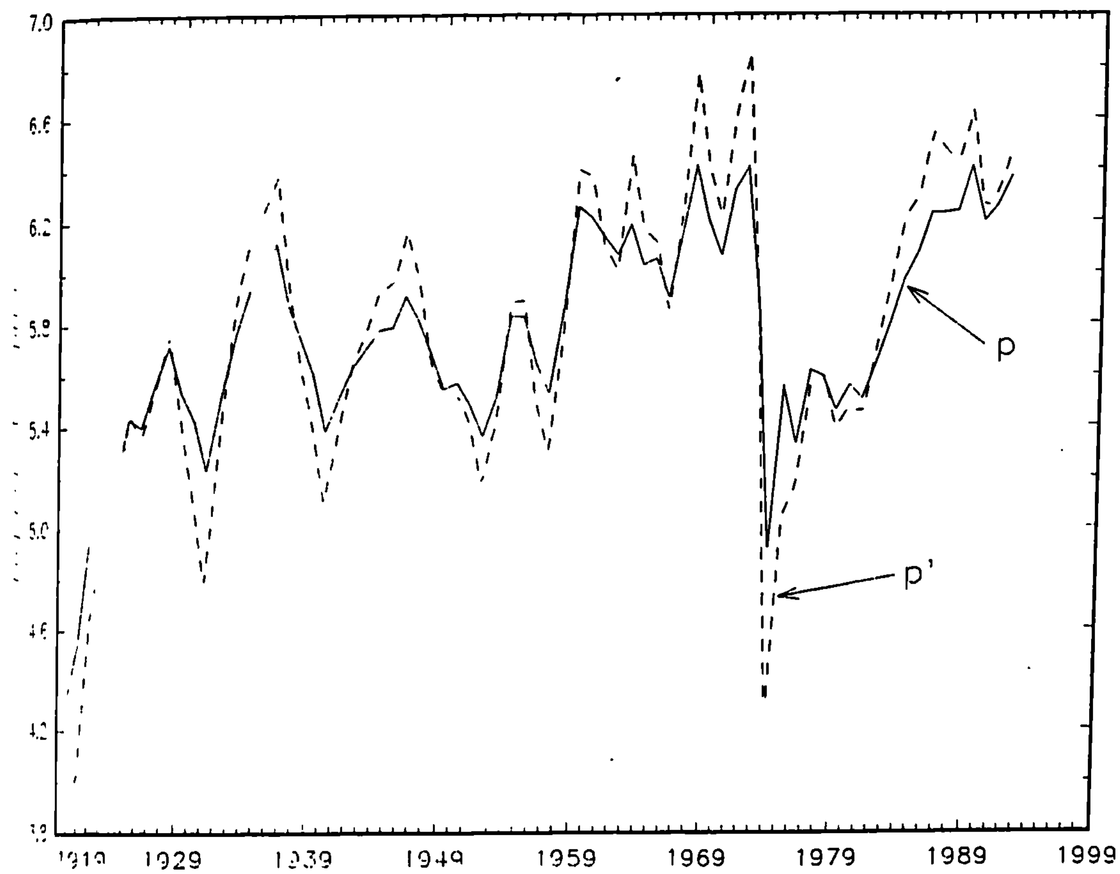

$$\delta_t' = \underset{(0.798)}{1.572}\,\delta_t - \underset{(1.418)}{1.281}\,\Delta d_t + \underset{(0.527)}{0.418}\,V_t$$

where each coefficient is within 1 standard error of its implied theoretical value. In addition, all of the standard deviation ratios $\sigma(\delta'_t)/\sigma(\delta_t)$ are within 1 standard error of unity, and $Corr(\delta'_t, \delta_t)$ are all very close to one (rows (ii) and (iii) in Table 5.5).

It is interesting to note that all of the standard deviation ratio point estimates in Table 5.5 are larger than unity: the *actual* dividend-price ratio is *less* volatile than that implied by the EMH. Since (with $i = \infty$) $\sigma(\delta_t - \delta_{dt}')/\sigma(\delta_{rt}')$ is also less than unity (see bottom of Table 5.5), the discount rate in the CAPM appears more variable than the residual movement in δ_t, once the effects of real dividend growth have been removed. Figures 5.5 to 5.8 plot the actual and theoretical log dividend-price ratios and log real stock prices when i=1 (Figures 5.5 and 5.6) and $i = \infty$ (Figures 5.7 and 5.8). It is very noticeable that when i=1 the actual and theoretical series are almost identical. When $i=\infty$, the theoretical series are more volatile than the actual series, although the degree of comovement between the two is still impressive. The volatility model appears to fit the data extremely well.

Allowing the safe rate to vary through time (Table 5.6) has little effect on the statistics, except to worsen slightly the correspondence between $\sigma(\delta_t)$ and $\sigma(\delta_t')$. I conclude, therefore, that the CAPM model with a constant safe rate provides an adequate representation of the data.

The CAPM (Merton 1973) suggests that the risk premium is determined by the variance of the market portfolio and the empirical literature using ARCH models indicates that this variance may be time varying and highly persistent (see for example Chou 1988 for the US and Hall et. al. 1990 for the UK). The VAR approach has an advantage over the Poterba-Summers (1986) model of the RVF in that I consider variation in both the risk premium and dividend growth, although my measure of the former is perhaps not as sophisticated as the GARCH model used by Chou. Nevertheless my results in Table 5.5 using the BZW annual data are consistent with those of Chou who finds that the RVF performs reasonably well when a time-varying risk premium is introduced into the model.

5.9.7 Comparison of Actual and Theoretical One-Period Returns

Notice first that for the constant real returns model, Corr(ξ_t, ξ_{1t}') is quite high, at least for the lower return horizons (Table 5.2, row (v)). However, once again the movement is not marked enough, resulting in low standard deviation ratios (Table 5.2, row (iv)).

The story is the same for the constant expected returns model (Table 5.3, rows (iv) and (v)). Tables 5.5 and 5.6, however, show that the theoretical returns derived from the CAPM models closely match the actual returns. All of the standard deviation ratios $\sigma(\xi_{1t}')/\sigma(\xi_{1t})$ are within 1 standard error of unity, and in Table 5.6 this is also true for the estimated correlations. When the safe rate is held constant, the correlations are somewhat lower, although still well within 2 standard errors of unity. Thus I again find reasonable of support for the time-varying risk premium specifications.

5.9.7 Some Variants

When the VAR lag length p is increased to 2, in general there is a rise in the Wald statistics (not reported). This does not, of course, affect my conclusions on the constant real and constant excess returns models. However, with the two CAPM specifications we find that the Wald test rejects the EMH at the 5% level for return horizons of 1 and 2 years. Thereafter (horizons 3, 5, 10 and infinity), the p-values are comfortably large (over 20%), and the previous non-rejections are sustained. In fact, even for the two lowest return horizons, the volatility models outperform the other specifications of equilibrium expected returns.

I also studied the effects of raising the coefficient of relative risk aversion in the volatility models. Table 5.7 and 5.8 present statistics for the coefficient of relative risk aversion varying from 1 to 5 when the return horizon is 1 year and infinite respectively. As α increases, we observed a slight increase in the Wald statistics (from 3.267 to 4.291 for i=1, and from 2.731 to 2.906 when $i\Rightarrow\infty$), together with a rise in the standard deviation ratio $\sigma(\delta_t')/\sigma(\delta_t)$ and a marginal fall in Corr(δ_t, δ_t'). However, standard errors on the latter also increase, so that the hypothesis that δ_t and δ_t' are not statistically different, and even with α=5 all of the p-values for the Wald statistics are well over 20%. Comparison of the actual and theoretical returns reveals a similar

Table 5.7: Sensitivity to α: Return Horizon i=1
Time-Varying Safe Rate and Risk Premium

		Coefficient of Relative Risk Aversion, α				
		1	2	3	4	5
(i)	Wald statistic (p-value)	3.276 (0.351)	2.635 (0.451)	3.255 (0.354)	3.860 (0.277)	4.291 (0.232)
(ii)	$\sigma(\delta_t')/\sigma(\delta_t)$ (s.e.)	0.941 (0.110)	1.253 (0.225)	1.566 (0.370)	1.878 (0.522)	2.190 (0.675)
(iii)	$Corr(\delta_t,\delta_t')$ (s.e.)	0.975 (0.026)	0.964 (0.033)	0.958 (0.042)	0.954 (0.049)	0.951 (0.055)
(iv)	$\sigma(\xi_{1t}')/\sigma(\xi_{1t})$ (s.e.)	0.857 (0.143)	1.062 (0.201)	1.287 (0.299)	1.519 (0.408)	1.754 (0.521)
(v)	$Corr(\xi_{1t},\xi_{1t}')$ (s.e.)	0.972 (0.042)	0.944 (0.076)	0.921 (0.106)	0.903 (0.130)	0.889 (0.149)

Table 5.8: Sensitivity to α: Return Horizon $i=\infty$
Time-Varying Safe Rate and Risk Premium

		Coefficient of Relative Risk Aversion, α				
		1	2	3	4	5
(i)	Wald statistic (p-value)	2.731 (0.435)	2.212 (0.530)	2.362 (0.501)	2.661 (0.447)	2.906 (0.406)
(ii)	$\sigma(\delta_t')/\sigma(\delta_t)$ (s.e.)	0.986 (0.295)	1.822 (0.704)	2.650 (1.143)	3.480 (1.589)	4.297 (2.309)
(iii)	$Corr(\delta_t,\delta_t')$ (s.e.)	0.975 (0.032)	0.965 (0.033)	0.960 (0.036)	0.957 (0.037)	0.955 (0.039)
(iv)	$\sigma(\xi_{1t}')/\sigma(\xi_{1t})$ (s.e.)	0.901 (0.252)	1.528 (0.557)	2.171 (0.888)	2.823 (1.228)	3.486 (1.563)
(v)	$Corr(\xi_{1t},\xi_{1t}')$ (s.e.)	0.972 (0.047)	0.927 (0.096)	0.898 (0.120)	0.879 (0.132)	0.867 (0.140)

Table 5.9: Results from Post-War Data (α=2)
Time-Varying Safe Rate and Risk Premium

		Return Horizon (years)					
		1	2	3	5	10	infinity
(i)	Wald statistic (p-value)	2.174 (0.537)	1.729 (0.631)	1.584 (0.663)	1.596 (0.660)	1.655 (0.647)	1.659 (0.646)
(ii)	$\sigma(\delta_t')/\sigma(\delta_t)$ (s.e.)	1.272 (0.359)	1.536 (0.689)	1.661 (0.885)	1.756 (1.057)	1.785 (1.114)	1.785 (1.115)
(iii)	$Corr(\delta_t,\delta_t')$ (s.e.)	0.919 (0.088)	0.922 (0.080)	0.918 (0.087)	0.915 (0.091)	0.915 (0.092)	0.915 (0.092)
(iv)	$\sigma(\xi_{1t}')/\sigma(\xi_{1t})$ (s.e.)	0.996 (0.294)	1.213 (0.515)	1.308 (0.657)	1.383 (0.794)	1.406 (0.843)	1.407 (0.844)
(v)	$Corr(\xi_{1t},\xi_{1t}')$ (s.e.)	0.874 (0.185)	0.868 (0.174)	0.853 (0.193)	0.844 (0.203)	0.841 (0.204)	0.841 (0.204)

pattern, with the standard deviation ratios increasing and the correlation between the two series falling (the lowest correlation being 0.867 for α=5 and i=∞). I conclude that the results are robust to changes in the coefficient of relative risk aversion within a plausible range.

In Table 5.9 I report the results of the VAR tests over the Post War data period 1945-1994 for the CAPM model with a time-varying safe interest rate. The results are remarkably similar to those in Table 5.6 for the whole data period. For every return horizon, the non-linear Wald test that $\delta_t=\delta_t'$ indicates strong non-rejection, with p-values ranging from 0.537 (i=1) to 0.663 (i=3). The point estimates of the standard deviation ratio $\sigma(\delta_t')/\sigma(\delta_t)$ are all above unity, although not significantly so, and the correlation between δ_t and δ_t' ranges from 0.915 (i=5) to 0.922 (i=2). Looking at the comparison of actual and theoretical one-period returns, it is notable that the estimated correlations between ξ_{1t} and ξ_{1t}' are somewhat lower over the Post War period than in Table 5.6. For example, Corr(ξ_{1t}, ξ_{1t}') is 0.874 for a one-year return horizon over the more recent data period, as compared with 0.944 over the full sample. However, the standard deviation ratios in Table 5.9 are not much different from those in Table 5.6. The overall impression is that there is a remarkable consistency of results over the two sample periods.

5.10 Conclusions

In this Chapter I applied the Campbell-Shiller VAR methodology to an annual UK stock index to test for market efficiency under four different assumptions regarding equilibrium expected returns. Over the period 1918-1993, I found no evidence supportive of the view that expected returns are constant, or that they depend only on a safe rate of interest. However, I was unable to reject specifications for equilibrium returns which include a time-varying risk premium based on the variance of past returns. These findings appear to be robust to changes in the VAR lag length and the coefficient of relative risk aversion.

In the final chapter of their elegant macro-theory text, Blanchard and Fischer (1989) present two models of asset pricing that they consider "useful" for macroeconomists: the consumption CAPM, and the standard CAPM as studied here. Lund and Engsted (1993) use the VAR methodology to test the consumption CAPM on the same data as we use, and reject its implications. Combining this evidence with that presented above confirms Blanchard and Fisher's conclusion that, whilst the consumption CAPM is theoretically more general, and so, to some, more appealing, empirically the standard CAPM outperforms it.

There are of course many caveats to add to the supportive evidence reported above. My model is linear in the variables, and the Campbell-Shiller (1989) proxy for a time-varying risk premium may be subject to measurement error. Given the balance of the evidence in this Chapter, it remains my view that there may yet be life in the EMH applied to the stock market providing a time-varying risk premium based on the variance of returns is incorporated in the determination of equilibrium returns. At present my results do not warrant a stronger inference, but further work using this methodology on alternative data sets may allow more definitive conclusions to be made.

Chapter 6: MARKET SEGMENTATION AND STOCK PRICE BEHAVIOUR

6.1 *Introduction*

With few exceptions, studies of aggregate stock market indices, whether based on stock prices (e.g. variance bounds tests) or one-period and multiperiod returns, have found evidence against the Efficient Markets Hypothesis (EMH)[1]. The consensus seems to be that existing models of expected returns do not adequately capture the behaviour of discount rates in the Rational Valuation Formula (RVF). Clearly, rejection of the EMH based on aggregate data may be due to mispricing in only a sub-sector of the market. Also, it is well documented that market analysts concentrate on specific sub-sectors of the market (see Francis 1986, ch.17, and Bing 1971), and therefore investors may be more skilled at eliminating mispricing in sub-sectors of the market rather than in the market as a whole. If market segmentation on an industry basis is prevalent then it is possible that we may find more support for the EMH at these lower levels of aggregation.

In this chapter I use the VAR methodology pioneered by Campbell and Shiller (1987, 1988, 1989) to test the efficiency of sub-sectors of the UK stock market. Previous studies using this methodology have used aggregated data under various models of equilibrium returns, namely either that real returns are constant, or they depend on real interest rates or the growth in consumption. The VAR approach has not been applied to the (standard) CAPM at a sectoral level. I take Merton's intertemporal CAPM as my theoretical anchor for expected returns, and using both monthly and quarterly data from 1965 to 1993, I find little evidence that the conditional market variance is the sole factor in the determination of equilibrium market returns. I then turn my attention to six portfolios based on industry sub-sectors, and use the VAR/CAPM framework to analyse two issues. First, under the CAPM-EMH, expected returns in market sub-sectors depend on the conditional covariance of sector returns with the market. The question then is whether one can trace the rejection of this model at the aggregate level to a failure of the CAPM to hold in specific sub-sectors of the market. I therefore test the conditional covariance specification on each of the market sub-sectors. The second issue is that of market segmentation. If investors concentrate their attention only on market sub-sectors, one should not expect any measure of covariance to be relevant in the determination

[1] See Sections 3.2 and 4.2-4.4.

of expected sub-sector returns. In fact, in this case the sub-sector *is* the market, and so the expected return should depend only on the conditional return variance within that sub-sector. Thus, in the final part of my analysis, I compare the results obtained from the covariance specification with those when expected returns are assumed to depend only on the expected sub-sector return variance.

There have been several notable attempts to study the implications of the CAPM on time-series data. French, Schwert and Stambaugh (1987) modelled the market return variance as an ARMA process, and whilst they found only a weak relationship between the resultant variance forecasts and ex post returns, they concluded nevertheless that the conditional variance is an important factor in the determination of equilibrium returns. French et al (1987), Chou (1988), Bollerslev, Engle and Wooldridge (1988) and Hall, Miles and Taylor (1990) used GARCH-M models to capture conditional variances and covariances, the general conclusion being that such effects are present and important[2]. Although I am unable to incorporate such complex models of conditional variance into the VAR analysis, my study does have the important advantage that I require not only that the conditional variance/covariance measure determines expected returns, but that it does so in a way which is consonant with the RVF. In particular, other variables which have commonly been found to have predictive power over returns should have no incremental effect, except in so far as they provide information about future return variability.

The remainder of this chapter is as follows. Section 6.2 outlines the models of expected returns under consideration in this chapter. In Section 6.3 I briefly review previous attempts to test the implications of the CAPM using time-series data, and 6.4 re-introduces the Campbell-Shiller dividend-price ratio model, which is the basis of my testing procedure. Section 6.5 discusses various data issues and the construction of the VAR for monthly data, whilst 6.6 presents the empirical results for the monthly data set. Section 6.7 presents the results using quarterly data. Section 6.8 discusses some variants to the basic model, and 6.9 concludes.

---

2 Clare, O'Brien, Thomas and Wickens (1993) is an exception.

6.2 *Expected Returns and Discount Rates*

Tests of the EMH are conditional on a model of equilibrium returns, and my analysis of market efficiency relies on the implications of Merton's intertemporal CAPM. This specifies the expected excess return (over a short interest rate, r_t) on asset i as being proportional to the conditional covariance of the return on asset i with the return on the market portfolio (C_{imt}):

$$(6.1) \quad E_t h_{it+1} = E_t(r_{t+1} + \alpha C_{imt+1})$$

where h_{it+1} is the one-period real holding period return on asset i. Under certain conditions (Merton 1980), α is the harmonic mean of investors' coefficient of relative risk aversion. The expected excess return on the *market* portfolio is linearly related to the conditional market return variance (V_{mt}):

$$(6.2) \quad E_t h_{mt+1} = E_t(r_{t+1} + \alpha V_{mt+1})$$

Finally, when considering market segmentation, the expected excess return on asset i depends only on the conditional return variance of asset i[3]:

$$(6.3) \quad E_t h_{it+1} = E_t(r_{t+1} + \alpha V_{it+1})$$

6.3 *Time-Series Studies of the CAPM*

Most studies of the CAPM using time-series data have focused on the main prediction of the model, that returns on any asset i are linearly related to the conditional covariance of asset-i returns with market returns. As discussed in Chapter 3, French, Schwert and Stambaugh (1987) (FSS) and Chou (1988) employed univariate GARCH-M models in order to relate the expected return on an aggregate stock market index to the conditional return error variance. The coefficient on the conditional variance term was taken as an estimate of the aggregate coefficient of relative risk aversion, or the "market price of risk", α. If the CAPM is an

---

[3] It is important to notice at this point that the tests of the within-sector variance model have no power against the case where investors do consider the covariances with other sectors, but these are either negligibly small or happen to sum to zero, since in these cases (6.1) and (6.3) will be approximately equal.

accurate model of required returns, market risk should be significantly priced. Chou estimated α to be in the range 4.5-6.15 whilst FSS estimated a much wider range (2.41-7.22). Both studies generally found that α was significantly different from zero.

Bollerslev, Engle and Wooldridge (1988) argued that studies such as these may be omitting important covariances by taking the aggregate stock market return to be the "market" return in the CAPM. Bollerslev et al took the market to consist of stocks, bills and bonds, and set up a multivariate GARCH model in order to address two questions: first, are the conditional covariances between these asset returns time-varying; and second, does the inclusion of the extra covariances in the model improve the specification over the univariate models used in previous studies.

If e_t is a vector of excess returns on N assets, e_{it}, together with the market portfolio excess return, e_{mt}, the multivariate GARCH-M(p,q) model may be written as

$$e_t = \alpha H_t \iota_m + \varepsilon_t$$

$$\mathrm{vech}(H_t) = A_0 + \sum_{i=1}^{p} A_i \mathrm{vech}(\varepsilon_{t-i}\varepsilon_{t-i}') + \sum_{j=1}^{q} B_j \mathrm{vech}(H_{t-j})$$

where $H_t = \begin{bmatrix} V_{1t} & C_{12t} & . & C_{1mt} \\ C_{12t} & V_{2t} & . & C_{2mt} \\ . & . & . & . \\ C_{1mt} & C_{2mt} & . & V_{mt} \end{bmatrix}$, ι_m picks out e_{mt} from e_t and vech(.) denotes the column-stacking operator of the lower triangular portion of a symmetric matrix. Assuming that each covariance depends only on its own past values and surprises, setting p=q=1, and using quarterly data from 1959(1) to 1984(2), Bollerslev et al estimated α to be 0.499 with a standard error of 0.160. Thus, although expected returns were found to be significantly related to conditional covariances, the estimated market price of risk was much lower than previous studies had obtained. With regards to the time-variation in conditional covariances, none of the variance or covariances in H_t was individually statistically significant, but the joint hypothesis that all of the elements of H_t were zero was rejected. They concluded that models which do not allow for time-variation in conditional covariances are likely to be misspecified. However, support for the CAPM was qualified by the fact that lagged returns and consumption growth were found to have marginal explanatory power for returns, even after the inclusion of the conditional covariance.

Hall, Miles and Taylor (1990) took the FT500 index to be the "market", and studied the covariances of returns from four industrial sectors - mechanical engineering, financial, electrical and chemical - with the market return. Also employing a multivariate GARCH-M(1,1) specification, estimation of the unconstrained model produced an estimate of α of 3.24 (s.e.=1.05), which was statistically significant, and of a magnitude which compared closely with previous estimates. However, Hall et al noted that the off-diagonal terms in the constant matrix $\mathbf{A}_0$, which are the unconditional asset covariances, were mostly insignificant. They re-estimated the model, imposing zero unconditional covariances, and found little change in the estimates of the remaining parameters. They therefore posited that, whilst conditional variances and covariances were time-varying, and are significant predictors of returns, from an unconditional perspective only variances were important.

This last result is, perhaps, suggestive of market segmentation. However, it is not possible to determine from either of the multivariate GARCH-M studies what is the contribution of the conditional covariances to the estimates ofα. If covariances are not important, one would expect that their omission from*all* of the terms in H_t (i.e. both conditional and unconditional) would affect neither the magnitude nor the significance of the estimated market price of risk. This is one of the motivations of the following analysis, albeit in a different modelling framework.

6.4 *The Dividend-Price Ratio Model*

The methodology used here is the same as that in Chapter 5, although the discount rate models are slightly different. Referring back to equation (5.44), one can replace $E_t\xi_{t+j}$ with one of the models of equilibrium returns (6.1)-(6.3). For example, in the case of expected excess market returns depending on market volatility (equation (6.2)), equation (4.44) becomes

$$(6.4)\quad \delta_t = \sum_{j=0}^{\infty} \rho^j E_t(r_{t+1+j} + \alpha V_{mt+1+j} - \Delta d_{t+1+j}) - \frac{k}{1-\rho}$$

Replacing the expectations in (6.4) with rational forecasts produces the *theoretical* dividend-price ratio δ_t'. δ_t' is the dividend-price ratio that would obtain if the CAPM were true. Once again, forecasts of the RHS variables in (6.4) are obtained from a VAR, and the discussion of the VAR methodology in Section 5.6 is relevant here, the only difference being that I now concentrate only on one-period returns and infinite-horizon returns.

6.5 *Data Issues and VAR Construction*

Returns data are calculated from the Datastream Equities Indices. The Total Market data are disaggregated into six sub-sector portfolios: Industrials (both including and excluding oil), Financial Services, Capital Goods, Consumer Goods, and Other Sectors. Observations are taken on the last day of each month from January 1965 to January 1993. Following Campbell and Shiller (1988, 1989), the instantaneous variances V_{it} are calculated as the squared ex post monthly real return in sector i, whilst the covariances C_{imt} are the product of the real ex post return in sector i and the real ex post market return. As always, the following VAR analysis assumes that the series are stationary. Table 6.1 presents Phillips-Perron (1988) unit root tests for the log divided-price ratios, real dividend growth, and the variance and covariance terms. The PP tests are not completely supportive of the stationarity assumption. In particular, the t-statistics do not reject the null of non-stationarity for the log dividend-price ratios for two of the industrial portfolios, Consumer Goods and Other Sectors, although for the former, the $T(1-\phi)$ statistics indicate stationarity at the 10% level. This might illustrate a difference in the power of the two tests. I continue under the assumption that the series are stationary[4].

As far as the quality of the approximation (3.14) is concerned, Table 6.2 presents summary statistics for the series h_{it} and ξ_{it}. The correlation between the actual and approximate log real returns is reasonably high for all series (ranging from 0.908 to 0.975), and their standard deviations are similar[5]. The linearisation constant ρ was estimated as $1/(1+e^{\delta})$, where δ is the mean log dividend-price ratio for the Total

[4] In fact, PP tests indicate that the actual dividend-price ratio (as opposed to the log) is stationary for all sectors. That a logarithmic transformation reverses the inference perhaps tells us more about the tests for stationarity than the nature of the series.

[5] I do find that the means of the two series often differ, but this does not matter since the means of the series are unconstrained in the current framework.

Table 6.1: Phillips-Perron Unit Root Tests (Monthly Data)

Variable	PP statistic[1]	
	t-statistic	$T(1-\phi)$
Total Market		
δ_t	-2.913	-16.707
Δd_t	-18.595	-383.575
V_{mt}	-15.615	-298.474
Capital Goods		
δ_t	-2.734	-14.711
Δd_t	-18.675	-363.412
C_{imt}	-15.297	-290.417
V_{it}	-15.008	-282.470
Consumer Goods		
δ_t	-2.532	-12.912
Δd_t	-16.531	-348.266
C_{imt}	-15.831	-304.330
V_{it}	-16.093	-302.909
Industrials (including Oil)		
δ_t	-2.884	-16.378
Δd_t	-18.683	-394.007
C_{imt}	-15.598	-298.417
V_{it}	-15.585	-298.593
Industrials (excluding Oil)		
δ_t	-2.676	-14.243
Δd_t	-18.510	-400.342
C_{imt}	-15.683	-299.703
V_{it}	-15.776	-301.650
Financial Services		
δ_t	-2.704	-14.386
Δd_t	-18.298	-333.274
C_{imt}	-15.666	-298.457
V_{it}	-15.741	-298.655
Other Sectors		
δ_t	-2.425	-11.102
Δd_t	-19.488	-405.117
C_{imt}	-15.804	-299.918
V_{it}	-16.122	-305.280

[1] All reported statistics are calculated using a truncation lag of 4. 10% critical values are -2.57 for the t-statistics and -11.1 for the $T(1-\phi)$ statistics (Fuller 1976). The PP statistics for the real interest rate are -7.224 (t-statistic), -285.514 ($T(1-\phi)$).

Table 6.2: Comparison of Actual and Thoeretical Log Real Returns (Monthly Data)

Variable	Standard Deviation	Correlation with h_{it}
Total Market (ρ=0.996)[1]		
h_{mt}	0.061	1.000
ξ_{mt}	0.067	0.944
Capital Goods (ρ=0.996)		
h_{it}	0.070	1.000
ξ_{it}	0.077	0.946
Consumer Goods (ρ=0.997)		
h_{it}	0.065	1.000
ξ_{it}	0.068	0.975
Industrials (including Oil) (ρ=0.996)		
h_{it}	0.061	1.000
ξ_{it}	0.068	0.940
Industrials (excluding Oil) (ρ=0.996)		
h_{it}	0.063	1.000
ξ_{it}	0.067	0.960
Financial Services (ρ=0.996)		
h_{it}	0.066	1.000
ξ_{it}	0.074	0.910
Other Sectors (ρ=0.995)		
h_{t}	0.062	1.000
ξ_{it}	0.069	0.908

[1] The figure in parentheses after each sector label is the estimated value of ρ for that sector. However, since there is little variation in these values, and I found the results to be invariant to changes in this parameter, all statistics have been calculated using the value of ρ for the Total Market, i.e. 0.996.

Market (see footnote, Table 6.2). The resultant figure of 0.996 corresponds to an annualised mean dividend-price ratio of 4.742%.

Seasonality in real dividend growth caused some concern over the appropriate choice of VAR lag length. In particular, I found that, in order to remove serial correlation from the residuals of the real dividend growth equation, it was generally necessary to include lagged dependent variables at 3, 6, 9, and 12 months. This could, of course, be achieved by specifying a 12-lag VAR. However, 3 lags were found to be sufficient for all other equations, and even ignoring the substantial loss of degrees of freedom from estimating a VAR(12) model, the large number of parameters is likely to give a markedly better in-sample fit than forecasting performance. This is not uncommon in large VAR systems since "the parameters fit not only the systematic relationships...but also the random variation" (Litterman 1986). The model that I actually employ is a 3-lag VAR but with extra lagged dependent variables in the real dividend growth equation. In practice, therefore, to maintain a square **A** matrix, I have a 12-lag VAR, but with zero constraints on all parameters relating to variables beyond lag 3, except for the aforementioned lags of real dividend growth. The VAR parameters are estimated by GMM using White's (1984) heteroscedasticity-consistent covariance matrix for the entire set of parameters (see Section 2.3). The vector of statistics of interest, γ, is a non-linear function of the **A** matrix. If we denote this as γ=g(**A**) then the standard errors of the estimated statistics are calculated as $[g_\gamma(\mathbf{A})'\,\Psi\,g_\gamma(\mathbf{A})]^{1/2}$, where Ψ is the parameter covariance matrix.

For all sector models, given the estimates of Hall et al, results are initially reported for α=3. The effects of variations in this parameter are discussed in Section 6.9.

6.6 *Empirical Results with Monthly Data*

6.6.1 *Results for the Total Market Data*

Table 6.3 contains the linear and non-linear Wald statistics, and the statistics comparing the time-series behaviour of δ_t and δ_t', for the CAPM model in which the market excess return depends on the conditional market return variance. Results are presented for the coefficient of relative risk aversion, α, varying from 1 to 5. The VAR

Table 6.3: Total Market, Variance Model
Monthly Data

	Coefficient of Relative Risk Aversion				
	α=1	α=2	α=3	α=4	α=5
Non-linear Wald (p-value)	66.989 (0.000)	25.975 (0.100)	13.710 (0.748)	10.409 (0.918)	9.749 (0.940)
Linear Wald (p-value)	68.541 (0.000)	90.737 (0.000)	98.576 (0.000)	93.725 (0.000)	89.208 (0.000)
$\sigma(\delta')/\sigma(\delta)$ (s.e)	1.453 (1.098)	1.992 (1.481)	2.498 (1.886)	3.001 (2.296)	3.503 (2.661)
Corr(δ,δ') (s.e.)	0.597 (0.315)	0.657 (0.305)	0.696 (0.299)	0.722 (0.482)	0.741 (0.293)

Figure 6.1a: Actual and Theoretical Log Dividend-Price Ratio
Total Market, Monthly Data (α=3)

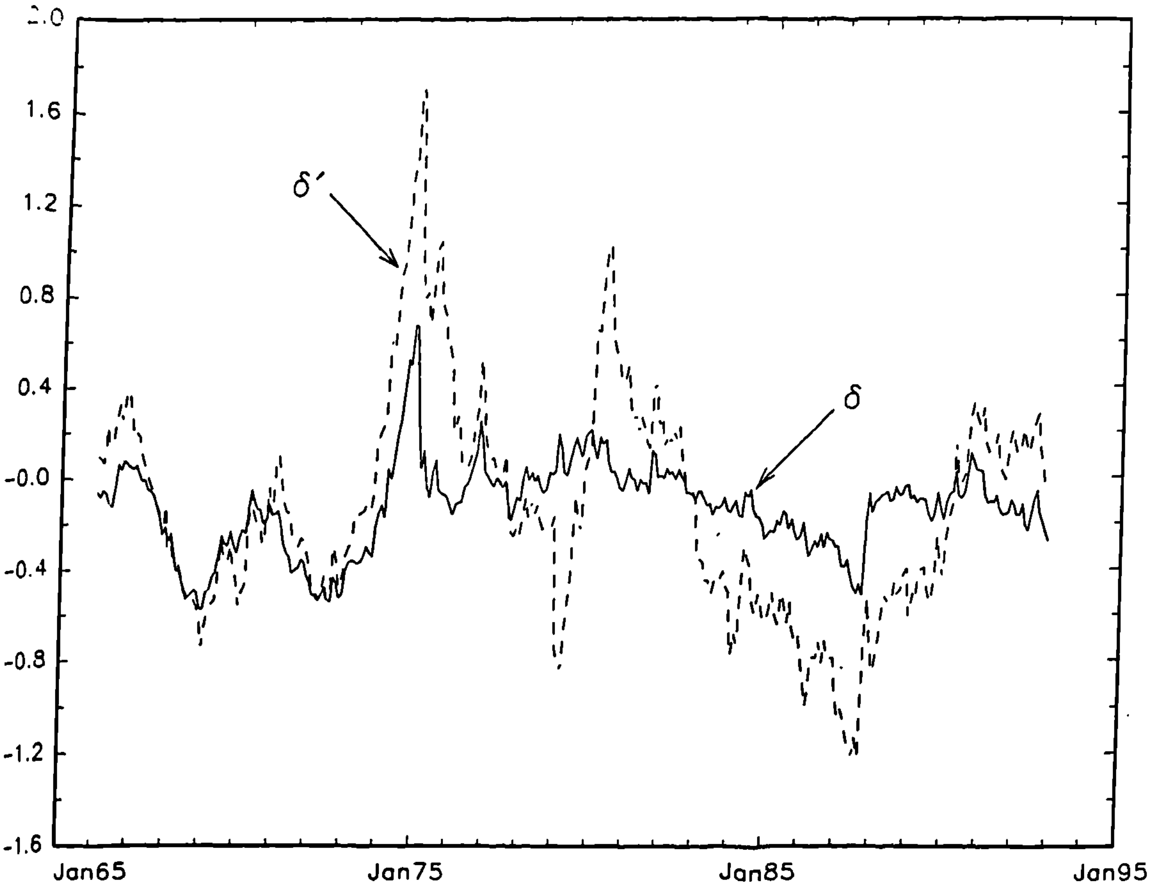

Figure 6.1b: Actual and Theoretical Log Real Stock Price
Total Market, Monthly Data (α=3)

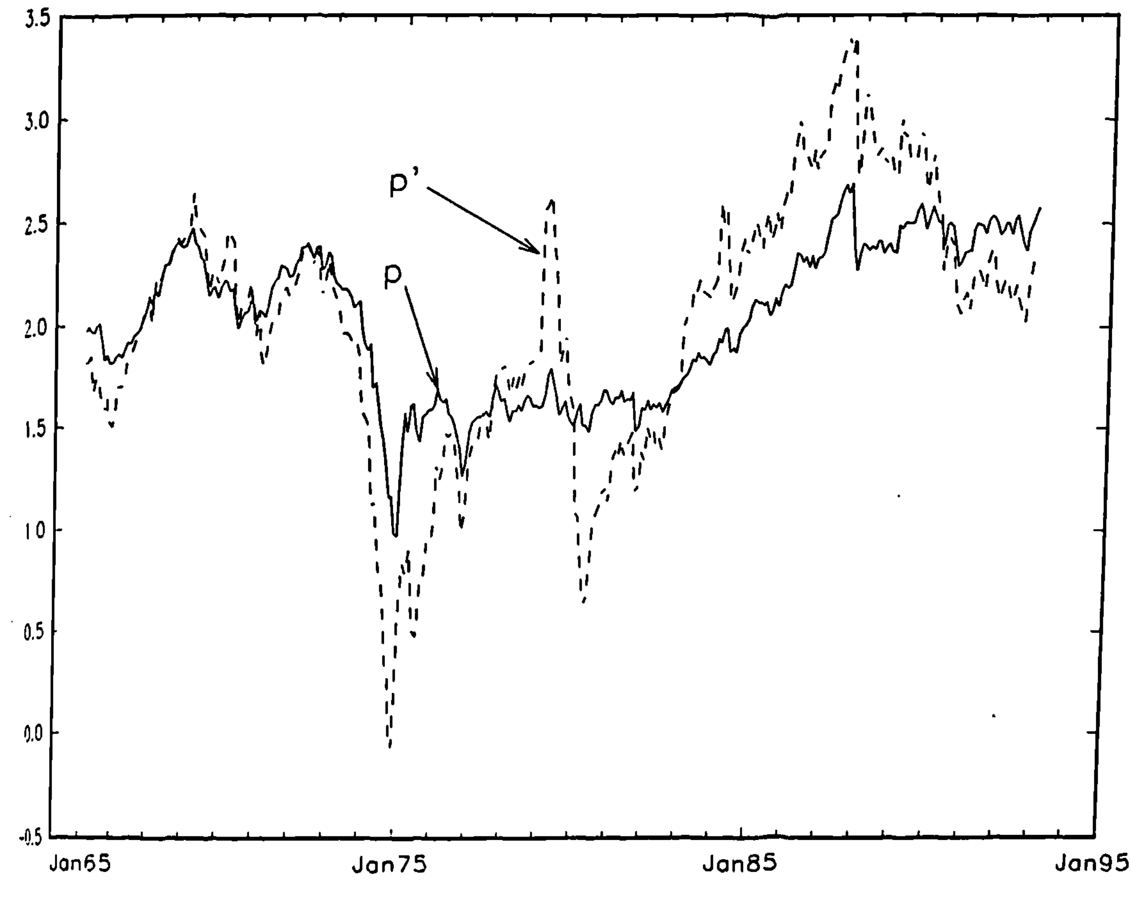

estimates (with α=3) imply the following relationship between the actual and theoretical log dividend-price ratio[6]:

$$\delta_t' = \underset{(1.10)}{1.67}\,\delta_t - \underset{(5.67)}{6.33}\,\Delta d_t - \underset{(0.71)}{0.04}\,(r-\alpha V)_t + \ldots$$

The large standard errors on these estimates mean that one cannot reject the null hypothesis that $\delta_t=\delta_t'$, but clearly, on the basis of the point estimates the two series will behave quite differently. This is illustrated in Figure 6.1a, where one can see that, although the series move broadly together, δ_t' is clearly more variable than δ_t. Much the same can be seen in Figure 6.1b, which compares the actual and theoretical log real stock prices. Looking now at the direct comparison of the δ_t and δ_t' series in Table 6.3, the correlation between the two ranges from 0.597 to 0.741, and is always within two standard errors of unity. Thus there is a high degree of comovement between the actual and theoretical dividend-price ratios. The standard deviation ratio $\sigma(\delta')/\sigma(\delta)$ is always greater than unity (1.453 to 3.503), but the large standard errors on the point estimates mean that none of them is significantly different from unity.

These points are reflected in the non-linear Wald statistics, which indicate non-rejection of the null hypothesis that $\delta_t=\delta_t'$ for $\alpha>1$. However, notice that, in contrast, the linear Wald test (which tests the hypothesis that one-period returns are unpredictable once the conditional market variance has been accounted for) always rejects at a high level of significance.

There are two points to make which help one to come to a conclusion about these contradictory results. First, there is a growing body of opinion that inference based on non-linear Wald statistics is not sound. Phillips and Park (1988) demonstrated that for a given data set and a given set of restriction, any (non-negative) value of the Wald statistic may be obtained by reparameterising the restrictions. Davidson and MacKinnon (1993) added that the common hope that inference based on a "natural" specification of non-linear restrictions (as, say, in equation (5.51b)) will prove to be more valid is, in fact, without foundation. One might

---

6 There are, in fact, 36 terms in this expression, but the first three are sufficient to illustrate the point being made.

therefore wish to place more emphasis on inference from the linear Wald test. Second, recall from Chapter 5 that perhaps the weakest implication of the dividend-price ratio model is that the dividend-price ratio should Granger-cause the discount rate measure, in this case $(r+\alpha V_m)_t$. The VAR equation for this series is

$$\begin{aligned}(r+\alpha V)_t = &-\underset{(0.044)}{0.016}\,\delta_{t-1}-\underset{(0.059)}{0.039}\,\Delta d_{t-1}+\underset{(0.056)}{0.078}\,(r+\alpha V)_{t-1}\\ &+\underset{(0.172)}{0.189}\,\delta_{t-2}-\underset{(0.165)}{0.165}\,\Delta d_{t-2}+\underset{(0.063)}{0.083}\,(r+\alpha V)_{t-2}\\ &-\underset{(0.164)}{0.133}\,\delta_{t-3}-\underset{(0.073)}{0.197}\,\Delta d_{t-3}-\underset{0.088}{0.008}\,(r+\alpha V)_{t-3}\end{aligned}$$

Clearly lagged δ_t's do not impact significantly on the measured discount rate[7]. I am therefore led to the conclusion that the balance of evidence is against the CAPM: the conditional market return variance is not the sole determinant of expected market returns. The question now is whether this rejection can be traced to specific sub-sectors of the market.

6.6.2 *Results for Sector Portfolios: The Covariance Model*

Table 6.4 presents the statistics of interest when the sector portfolio returns are assumed to depend only on the conditional covariance of the sector return with the market return, whilst Figures 6.2-6.7 show graphical comparisons of δ_t and δ_t', and p_t and p_t' for each sector. The figures highlight some cross-sectional variation in the closeness of δ_t and δ_t'. For example, for Financial Services (Figure 6.6a) the two series move quite closely together, whereas Consumer Goods (Figure 6.3a) and the Industrials portfolios (Figures 6.4 and 6.5) have the theoretical log dividend-price ratio much more variable than the actual series. These differences are reflected in the point estimates of the various test statistics. For Financial Services, the correlation between δ_t and δ_t' is 0.850 and the standard deviation ratio $\sigma(\delta_t')/\sigma(\delta_t)$ is 0.801, whereas for Consumer Goods the figures are 0.600 and 3.178 respectively. It seems rather odd, then, that the non-linear Wald statistics indicate rejection of the hypothesis that $\delta_t=\delta_t'$ for Financial Services, but not for any of the other sector portfolios. However, as the linear Wald statistics reject the unpredictability null for all

7 An F-test of the hypothesis that the coefficients on the lagged log dividend-price ratios were jointly zero returned a marginal p-value of 0.41.

Table 6.4: Covariance Model, Monthly Data (α=3)

	Capital Goods	Consumer Goods	Industrials (incl. Oil)	Industrials (excl. Oil)	Financial Services	Other Sectors
Non-linear Wald (p-value)	7.229 (0.988)	2.857 (0.999)	8.926 (0.961)	4.623 (0.999)	47.122 (0.000)	7.703 (0.983)
Linear Wald (p-value)	77.903 (0.000)	57.877 (0.000)	88.729 (0.000)	69.133 (0.000)	82.840 (0.000)	51.777 (0.000)
$\sigma(\delta')/\sigma(\delta)$ (s.e.)	2.360 (3.458)	3.178 (8.457)	2.588 (4.429)	2.913 (5.601)	0.801 (0.677)	1.831 (2.678)
$Corr(\delta,\delta')$ (s.e.)	0.887 (0.119)	0.600 (0.348)	0.840 (0.170)	0.809 (0.186)	0.850 (0.228)	0.802 (0.231)
$\sigma(\delta-\delta_d')/\sigma(\delta_r')$ (s.e.)	0.499 (0.630)	0.894 (1.083)	0.548 (0.688)	0.644 (0.732)	1.642 (1.561)	0.864 (1.320)
$Corr(\delta-\delta_d',\delta_r')$ (s.e.)	-0.147 (2.471)	-0.524 (2.354)	-0.152 (2.893)	-0.452 (2.381)	0.831 (0.363)	0.106 (1.950)
$\sigma(\xi')/\sigma(\xi)$ (s.e.)	1.969 (2.986)	2.760 (7.432)	2.241 (3.912)	2.396 (4.704)	0.962 (0.746)	1.587 (2.358)
$Corr(\xi,\xi')$ (s.e.)	0.807 (0.141)	0.452 (0.438)	0.744 (0.200)	0.689 (0.260)	0.854 (0.125)	0.702 (0.298)

Figure 6.2a: Actual and Theoretical Log Dividend-Price Ratio
Capital Goods, Monthly Data (α=3)

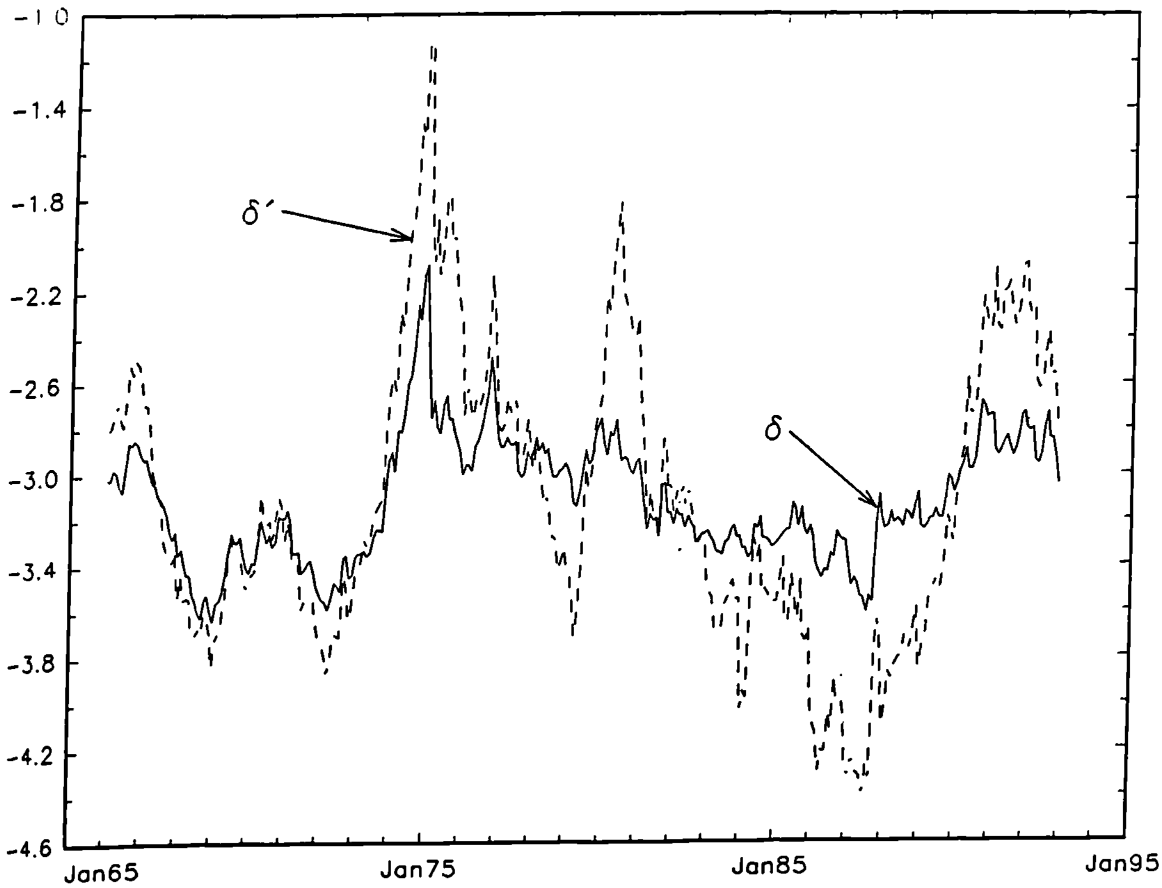

Figure 6.2b: Actual and Theoretical Log Real Stock Price
Capital Goods, Monthly Data (α=3)

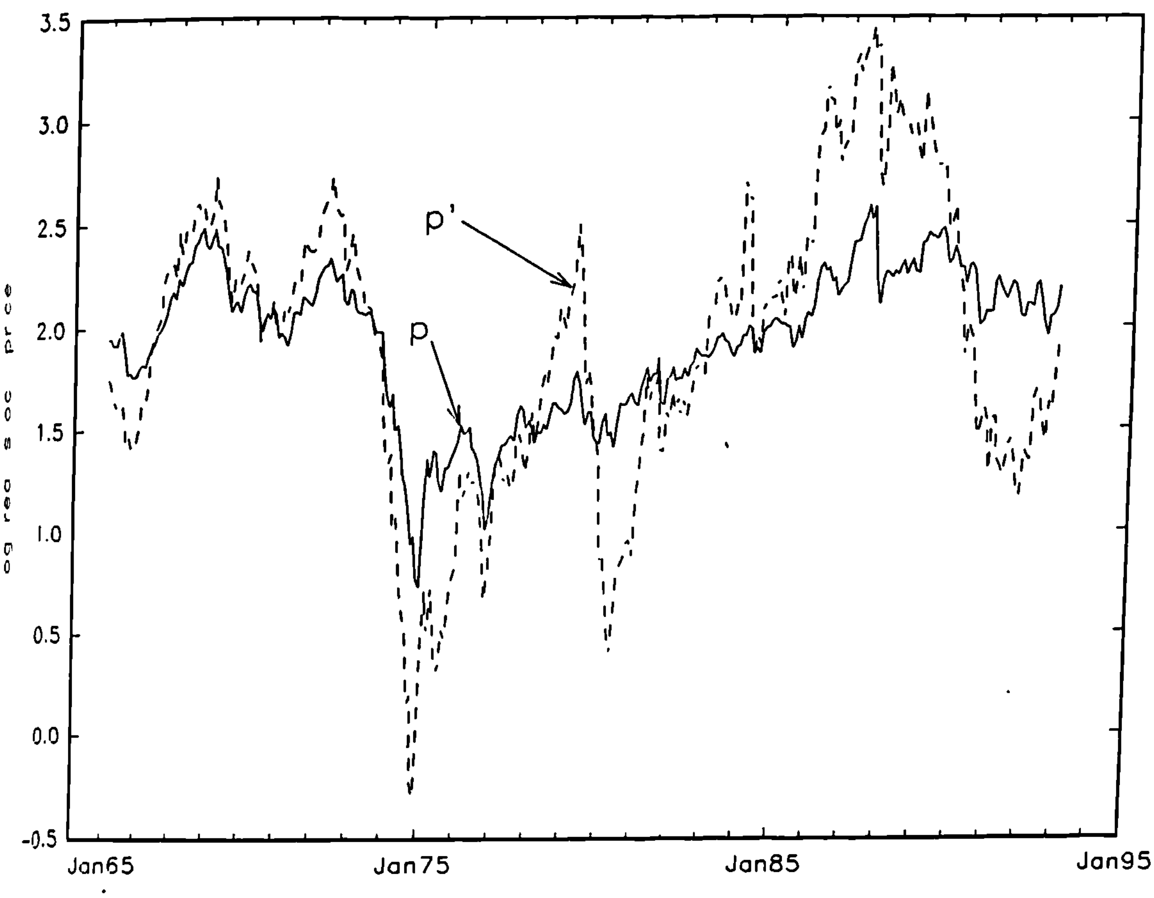

Figure 6.3a: Actual and Theoretical Log Dividend-Price Ratio
Consumer Goods, Monthly Data (α=3)

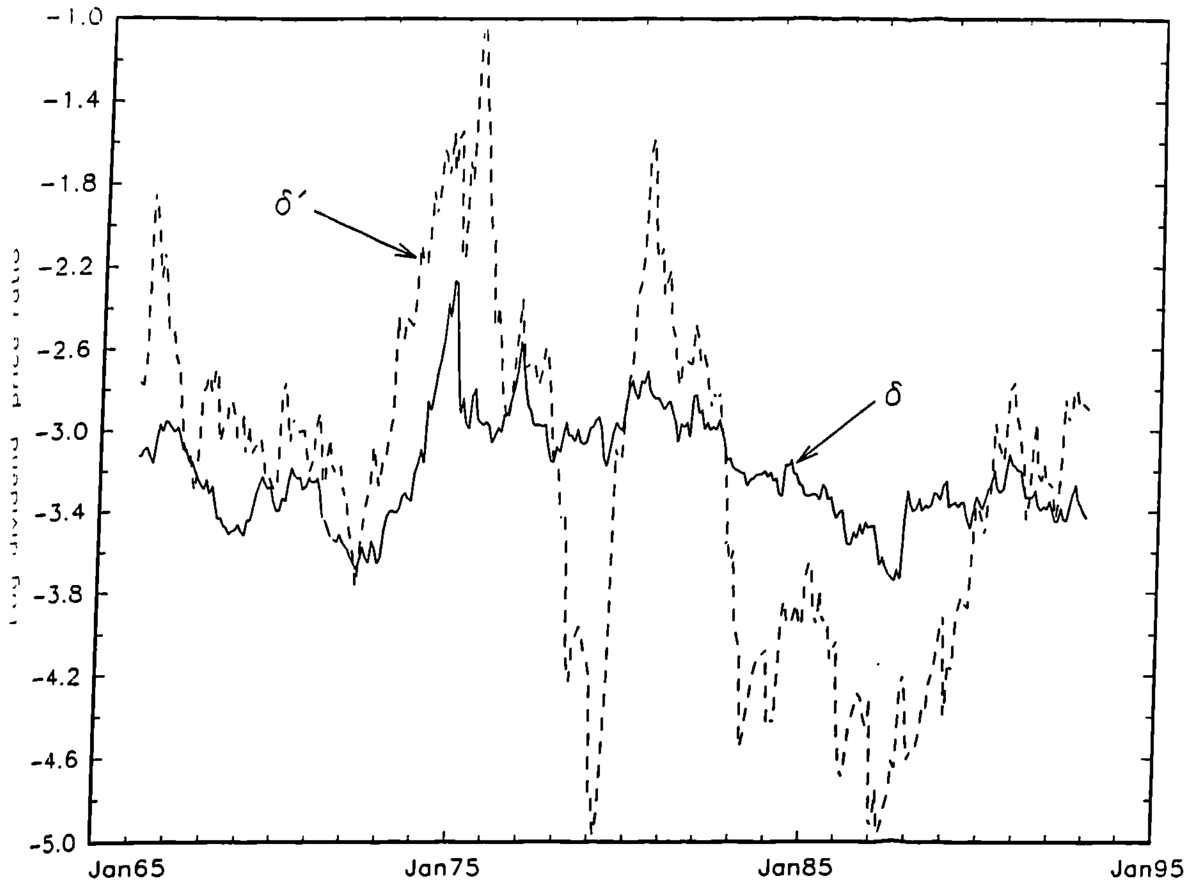

Figure 6.3b: Actual and Theoretical Log Real Stock Price
Consumer Goods, Monthly Data (α=3)

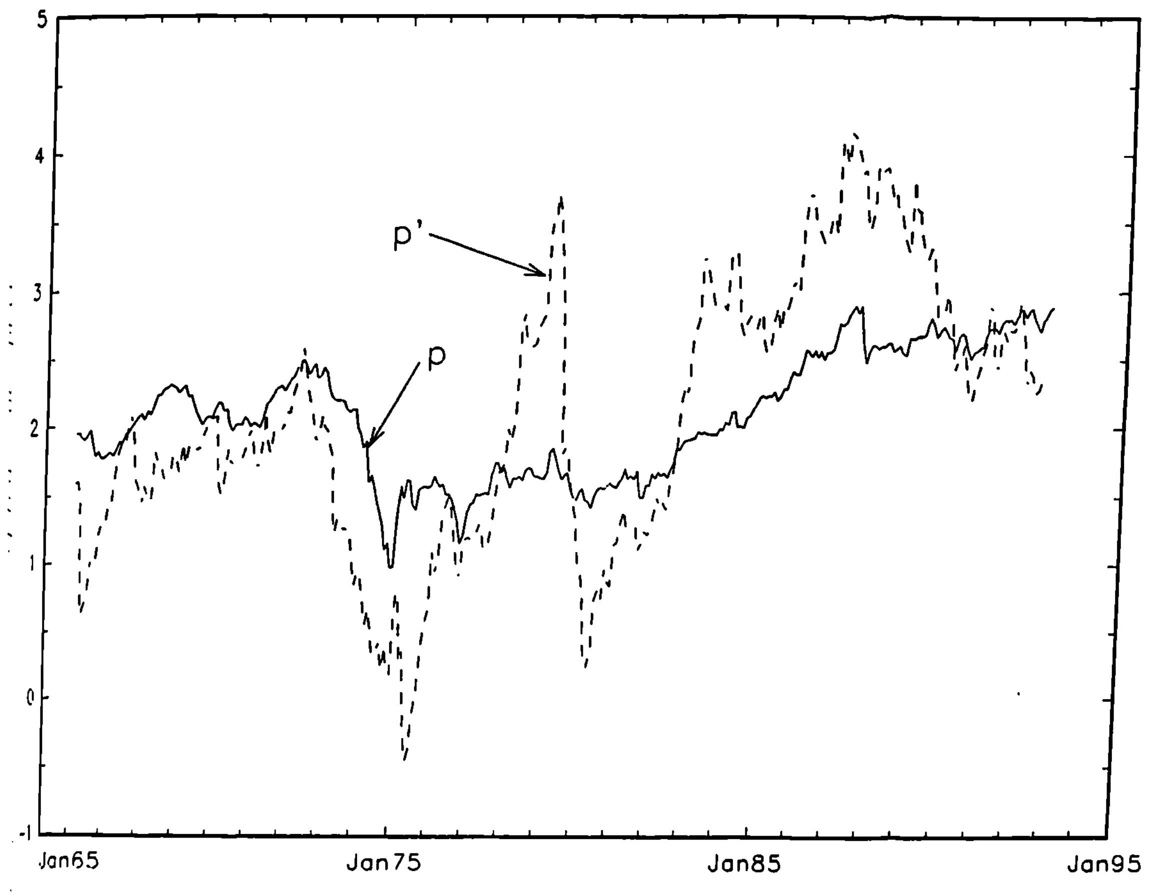

Figure 6.4a: Actual and Theoretical Log Dividend-Price Ratio
Industrials Including Oil, Monthly Data (α=3)

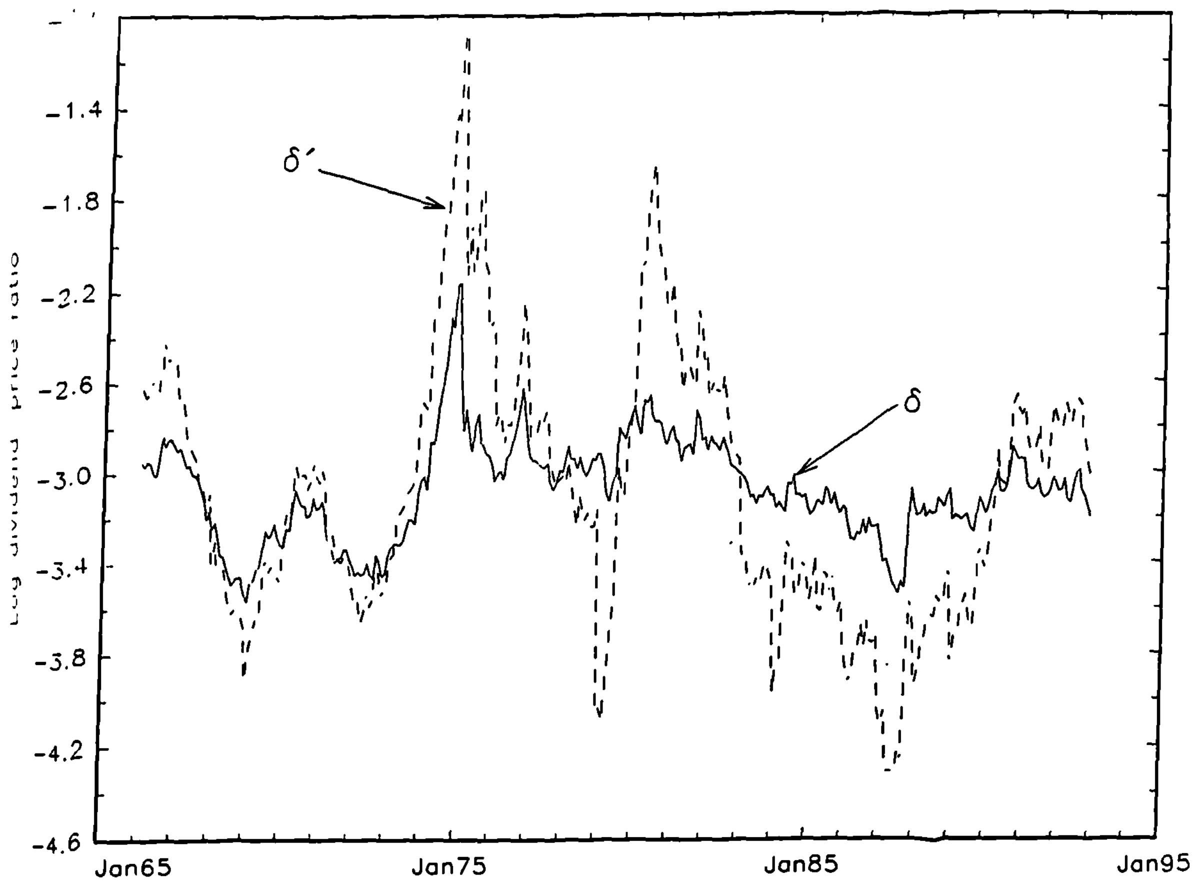

Figure 6.4b: Actual and Theoretical Log Real Stock Price
Industrials Including Oil, Monthly Data (α=3)

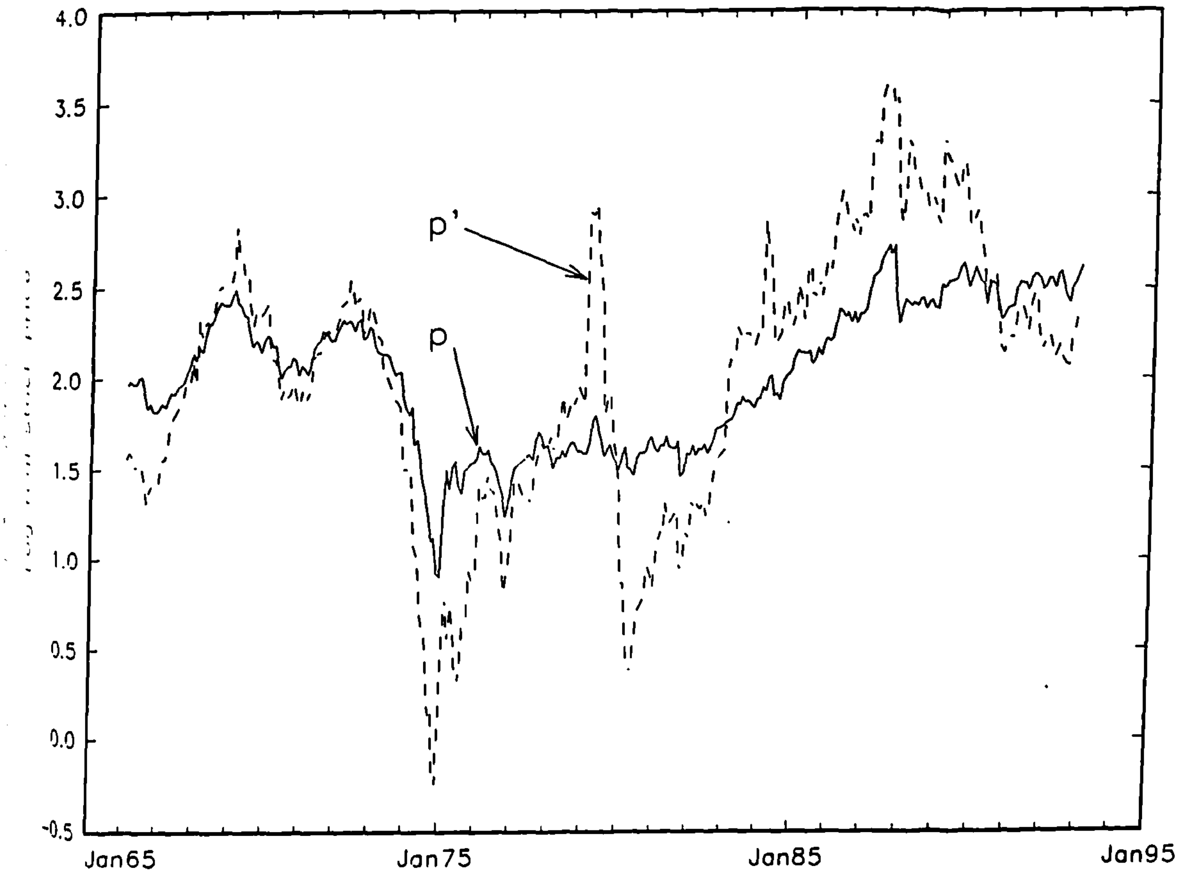

Figure 6.5a: Actual and Theoretical Log Dividend-Price Ratio
Industrials Excluding Oil, Monthly Data (α=3)

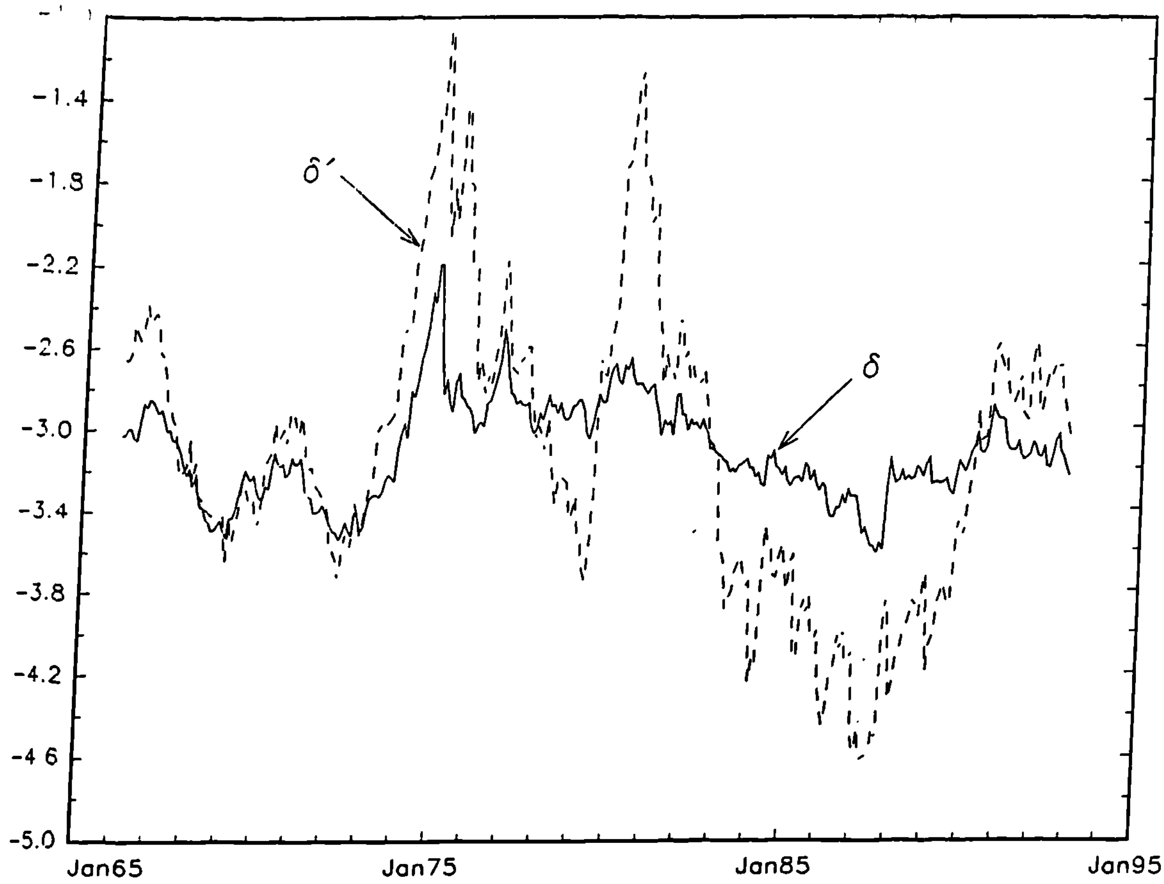

Figure 6.5b: Actual and Theoretical Log Real Stock Price
Industrials Excluding Oil, Monthly Data (α=3)

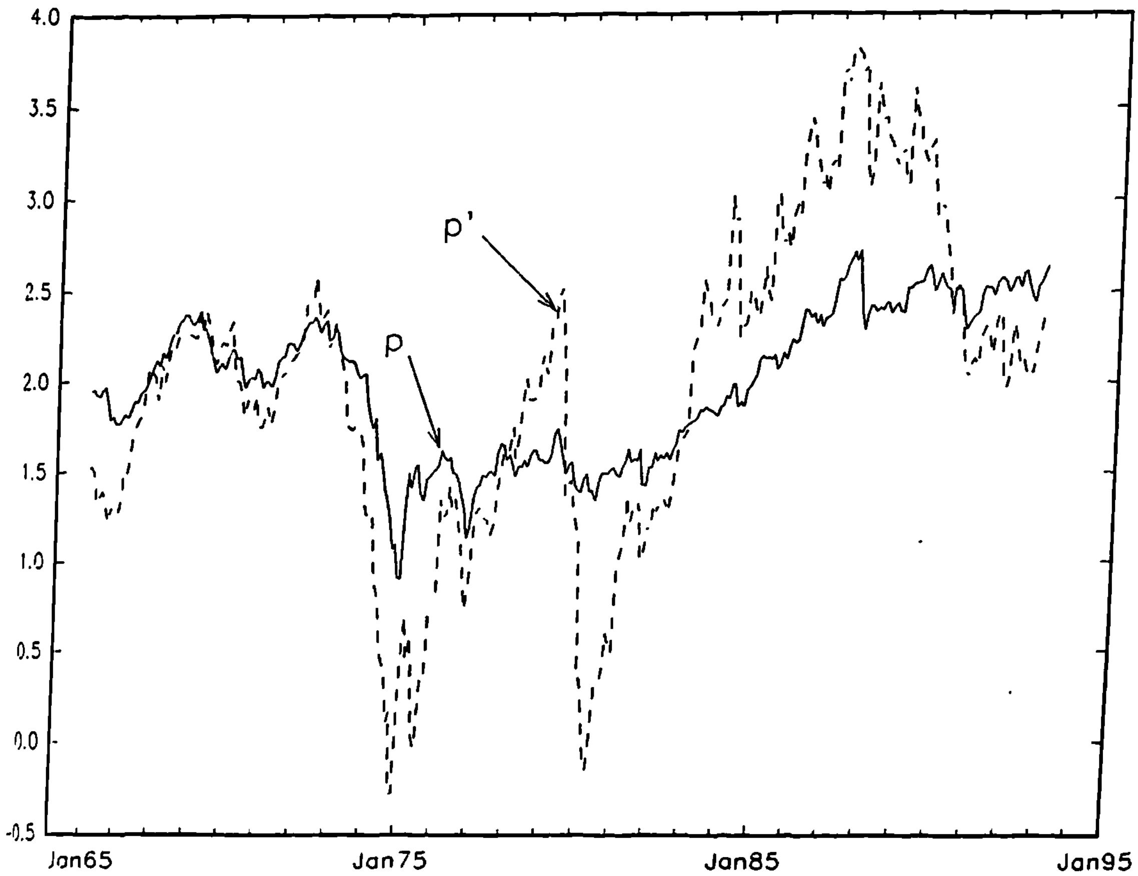

Figure 6.6a: Actual and Theoretical Log Dividend-Price Ratio
Financial Services, Monthly Data (α=3)

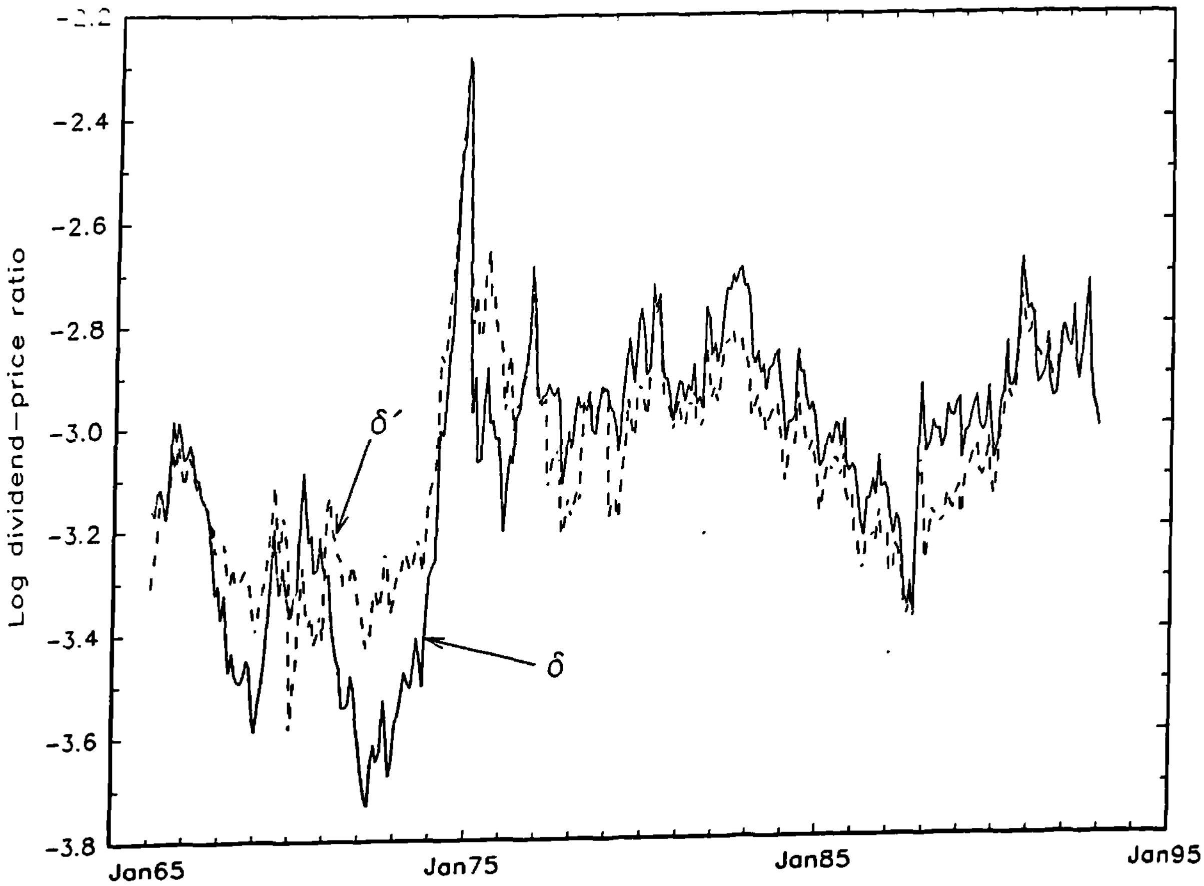

Figure 6.6b: Actual and Theoretical Log Real Stock Price
Financial Services, Monthly Data (α=3)

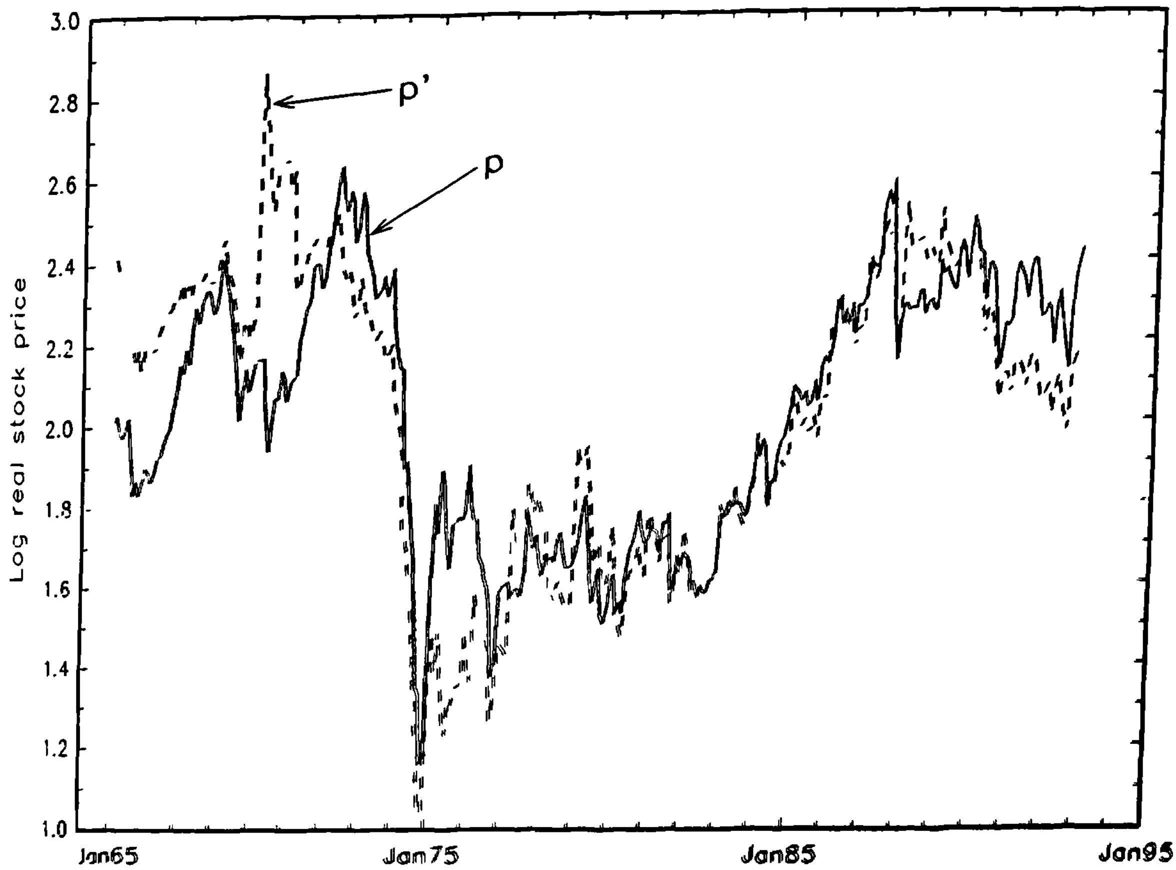

Figure 6.7a: Actual and Theoretical Log Dividend-Price Ratio
Other Sectors, Monthly Data (α=3)

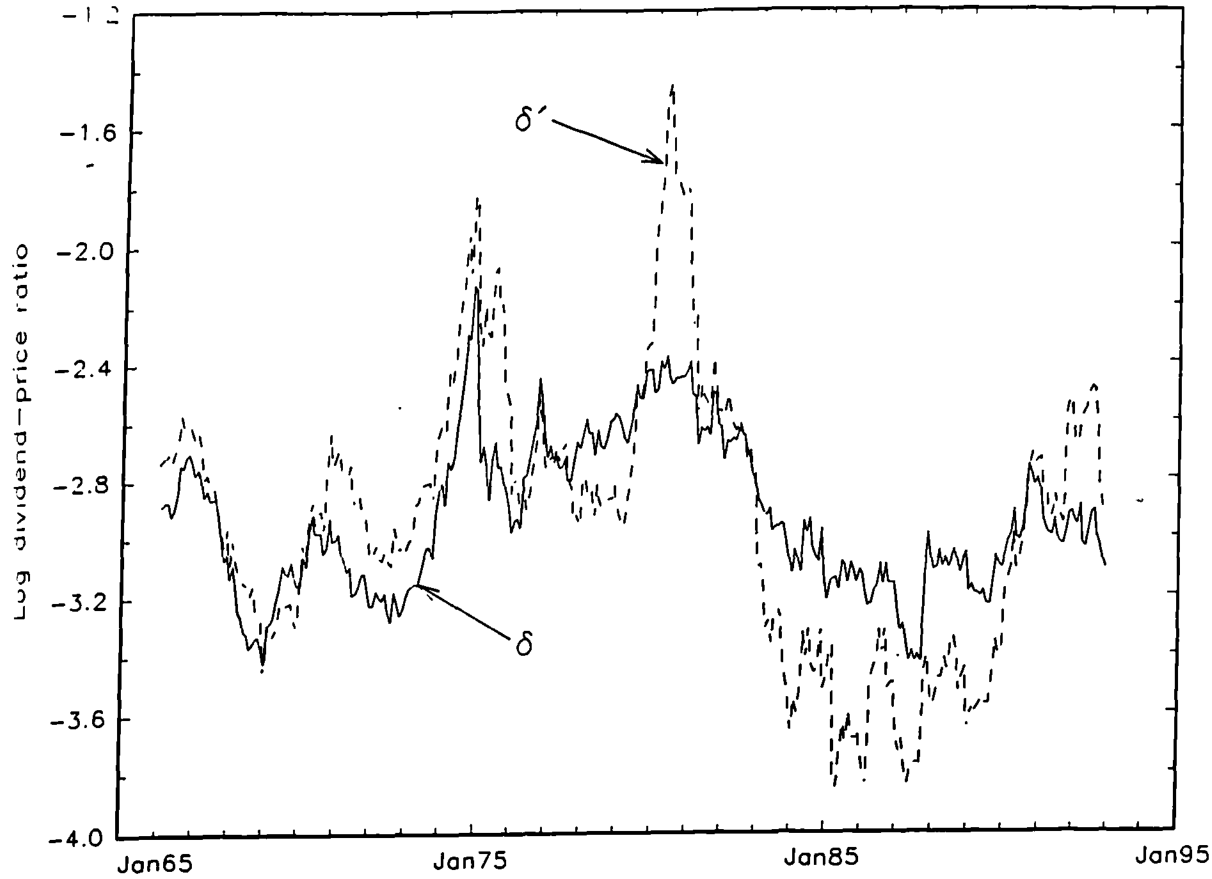

Figure 6.7b: Actual and Theoretical Log Real Stock Price
Other Sectors, Monthly Data (α=3)

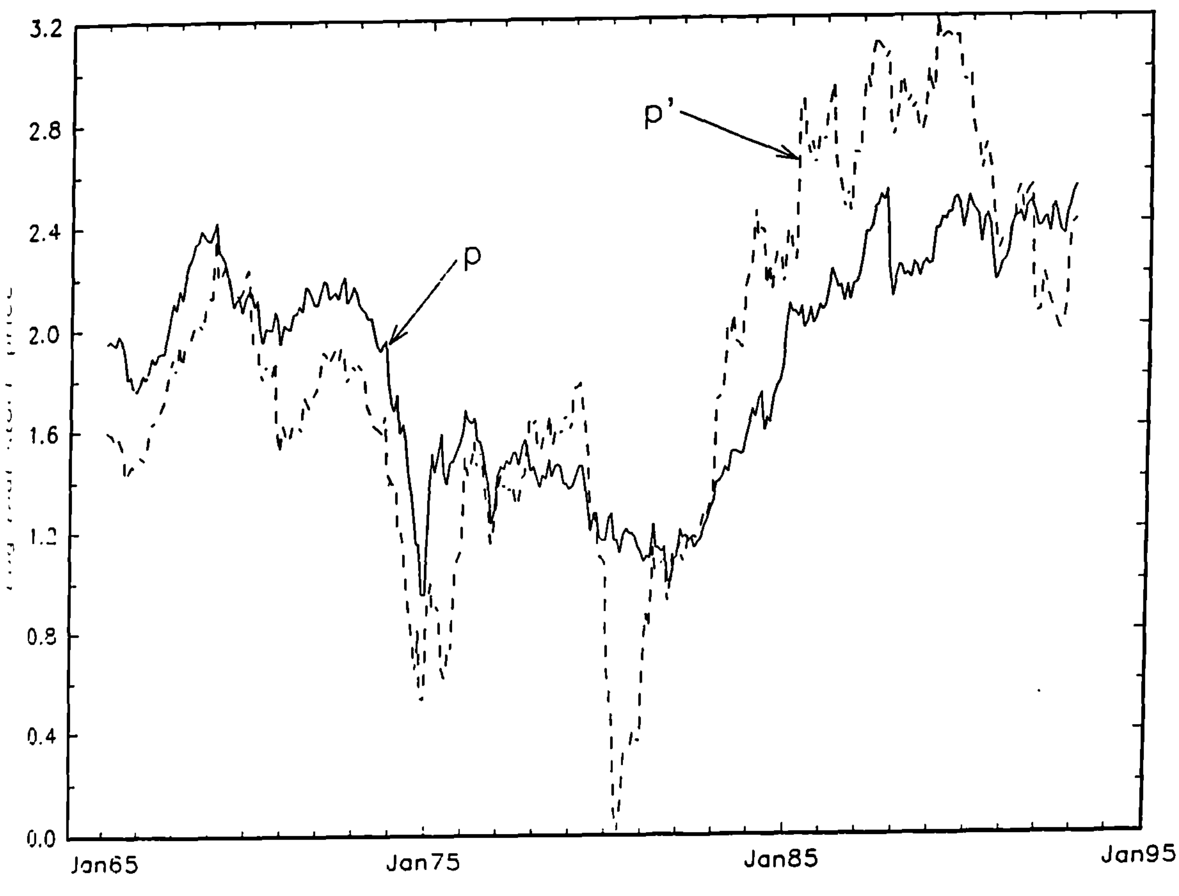

Table 6.5: Granger-causality, Covariance Model (α=3)

	Capital Goods	Consumer Goods	Industrials (incl. Oil)	Industrials (excl. Oil)	Financial Services	Other Sectors
δ_{t-1}	-0.008 (0.032)	-0.024 (0.054)	-0.010 (0.045)	-0.019 (0.045)	-0.015 (0.055)	0.000 (0.038)
Δd_{t-1}	-0.005 (0.103)	0.046 (0.234)	-0.019 (0.106)	0.047 (0.207)	-0.055 (0.040)	-0.037 (0.069)
$(r+\alpha C_{im})_{t-1}$	0.100 (0.347)	0.071 (0.441)	0.081 (0.349)	0.078 (0.441)	0.087 (0.4085)	0.101 (0.352)
δ_{t-2}	0.167 (0.153)	0.198 (0.178)	0.188 (0.177)	0.191 (0.182)	0.158 (0.153)	0.148 (0.149)
Δd_{t-2}	-0.118 (0.120)	-0.159 (0.172)	-0.109 (0.102)	-0.204 (0.216)	-0.183 (0.079)	0.018 (0.109)
$(r+\alpha C_{m})_{t-2}$	0.077 (0.323)	0.089 (0.268)	0.089 (0.242)	0.089 (0.272)	0.074 (0.215)	0.091 (0.229)
δ_{t-3}	-0.129 (0.178)	-0.143 (0.176)	-0.139 (0.167)	-0.140 (0.181)	-0.107 (0.100)	-0.121 (0.132)
Δd_{t-3}	-0.183 (0.093)	-0.151 (0.109)	-0.140 (0.133)	-0.135 (0.174)	-0.219 (0.140)	-0.167 (0.110)
$(r+\alpha C_{im})_{t-3}$	-0.012 (0.314)	-0.004 (0.265)	-0.014 (0.303)	-0.016 (0.277)	0.005 (0.253)	-0.016 (0.428)

sectors, I am inclined to attribute the non-linear Wald test results to a combination of the very large standard errors on many of the point estimates in Table 6.4, and the problems discussed above with the construction of non-linear Wald statistics.

Turning now to the comparison of (δ-δ_d') and δ_r', the problem with large standard errors is even more acute. This is because although the estimated correlations between the two are actually *negative* for four of the six sectors, the lack of precision in estimation means that they are all within one standard error of unity.

Finally, comparison of actual and theoretical returns, ξ_{1t} and ξ_{1t}', shows that there is significant positive correlation between the two for all but one sector (Consumer Goods), and none of these is outside two standard errors of unity. However, once again most of the point estimates indicate that the theoretical real return is much more volatile than the actual, although not significantly so.

As before, it is worth looking at the VAR estimates for the discount rate equation, which indicate whether δ_t Granger-causes the discount rate measure. These are presented in Table 6.5, and it is clear that lagged dividend-price ratios are not significant predictors of conditional covariances[8].

6.6.3 *Results for Sector Portfolio: The Own Variance Model*

When the covariance terms are omitted from the models for sector returns, so that each return is assumed to depend only on its own conditional variance, there are two questions to address. First, does the own-variance model provide a complete description of expected sector returns? Second, does the omission of the covariance terms have a marked effect on the performance of the model? The answers are, respectively, "probably not" and "no". Table 6.6 contains the test statistics for the own-variance model. There is, in fact, very little difference - both in quantitative and qualitative terms - between these results and those in Table 6.4. The linear and non-linear Wald statistics give conflicting signals, and the large standard errors on the remaining statistics mean that the null hypotheses which are implied by a successful model cannot be rejected, despite the fact that the point estimates are often far from their hypothesised values.

---

8 F-tests once again supported this conclusion, with marginal p-values ranging from 0.45 to 0.72.

Table 6.6: Variance Model, Monthly Data (α=3)

	Capital Goods	Consumer Goods	Industrials (incl. Oil)	Industrials (excl. Oil)	Financial Services	Other Sectors
Non-linear Wald (p-value)	8.128 (0.977)	4.996 (0.999)	12.074 (0.843)	8.382 (0.972)	65.185 (0.000)	8.694 (0.966)
Linear Wald (p-value)	98.117 (0.000)	71.961 (0.000)	125.291 (0.000)	97.238 (0.000)	102.172 (0.000)	58.979 (0.000)
$\sigma(\delta')/\sigma(\delta)$ (s.e.)	2.421 (1.520)	3.331 (3.598)	2.550 (1.794)	2.920 (2.573)	0.824 (0.374)	1.892 (1.411)
$Corr(\delta,\delta')$ (s.e.)	0.887 (0.099)	0.609 (0.342)	0.841 (0.177)	0.812 (0.170)	0.849 (0.241)	0.806 (0.202)
$\sigma(\delta-\delta_d')/\sigma(\delta_r')$ (s.e.)	0.474 (0.265)	0.825 (0.572)	0.562 (0.434)	0.642 (0.413)	1.572 (0.851)	0.809 (0.645)
$Corr(\delta-\delta_d',\delta_r')$ (s.e.)	-0.142 (1.214)	-0.535 (0.937)	-0.147 (1.126)	-0.451 (1.105)	0.823 (0.340)	0.101 (1.088)
$\sigma(\xi')/\sigma(\xi)$ (s.e.)	2.044 (1.187)	2.819 (2.943)	2.210 (1.429)	2.389 (1.960)	0.993 (0.406)	1.631 (1.140)
$Corr(\xi,\xi')$ (s.e.)	0.806 (0.135)	0.429 (0.456)	0.739 (0.206)	0.687 (0.242)	0.841 (0.131)	0.701 (0.279)

6.7 Results Using Quarterly Data

The failure of the dividend-price ratio to forecast volatility with monthly data contrasts sharply with the notable success found in Chapter 5 using annual data. It may well be, therefore, that the dividend-price ratio is more useful for tracking lower-frequency components of expected returns. If this is the case, one would expect the VAR models used above, which rely in large part on the dividend-price ratio to model expected returns, to perform poorly, and this may account in part for the disappointing performance of the dividend-price ratio model. I therefore now apply the same tests to quarterly data. I constructed three non-overlapping quarterly data series for each of the original monthly series, with quarters beginning at the end of January, March and April. As far as the quality of the approximation (3.14) is concerned, Table 6.7 presents summary statistics for the series h_{it} and ξ_{it}. The correlation between the actual and approximate log real returns is high for all series, ranging from 0.927 to 0.982, and their standard deviations are similar. The linearisation constant ρ was estimated as $1/(1+e^{\delta})$, where δ is the mean log dividend-price ratio for the Total Market (see footnote, Table 6.7). The resultant figure of 0.988 corresponds, as before, to an annualised mean dividend-price ratio of 4.742%. The PP tests presented in Table 6.8 give similar inference to those for monthly data reported above. In order to take account of the stochastic seasonal component of real dividend growth, four lags are included in the VAR. Since there was practically no difference in the results using the three different quarterly series, I concentrate only on the findings from the data from January.

6.7.1 Results for the Total Market data

Table 6.9 contains the results for the Total Market data, with $\mathbf{z}_t=[\delta_t, \Delta d_t, r_t+\alpha V_{mt}]$ and the coefficient of relative risk aversion, α, varying from 1 to 5. There is a steady improvement in the non-linear Wald statistic as α increases, so that with α=5 the null hypothesis that $\delta_t=\delta_t'$ cannot be rejected at the 5% level. There is also a gradual improvement in the linear Wald statistic, although the largest marginal probability value for this is only 0.035. The initial impression is then that even if we are willing to accept values of α of 5 then there is only limited support for the CAPM-EMH[9]. This

[9] As reported in Section 4.5.3, this qualitative conclusion was also found by Grossman and

Table 6.7: Comparison of Actual and Thoeretical Log Real Returns (Quarterly Data)

Variable	Standard Deviation	Correlation with h_{it}
Total Market (ρ=0.988)[1]		
h_{mt}	0.110	1.000
ξ_{mt}	0.125	0.965
Capital Goods (ρ=0.989)		
h_{it}	0.130	1.000
ξ_{it}	0.143	0.961
Consumer Goods (ρ=0.990)		
h_{it}	0.117	1.000
ξ_{it}	0.125	0.982
Industrials (including Oil) (ρ=0.988)		
h_{it}	0.109	1.000
ξ_{it}	0.124	0.961
Industrials (excluding Oil) (ρ=0.989)		
h_{it}	0.115	1.000
ξ_{it}	0.124	0.976
Financial Services (ρ=0.988)		
h_{it}	0.121	1.000
ξ_{it}	0.138	0.927
Other Sectors (ρ=0.987)		
h_{it}	0.112	1.000
ξ_{it}	0.122	0.943

[1] The figure in parentheses after each sector label is the estimated value of ρ for that sector. However, since there is little variation in these values, and I found the results to be invariant to changes in this parameter, all statistics have been calculated using the value of ρ for the Total Market, i.e. 0.988.

Table 6.8: Phillips-Perron Unit Root Tests (Quarterly Data)

Variable	PP statistic[1]	
	t-statistic	$T(1-\phi)$
Total Market		
δ_t	-2.990	-16.750
Δd_t	-7.491	-77.482
V_{mt}	-7.139	-70.242
Capital Goods		
δ_t	-2.802	-14.728
Δd_t	-8.581	-96.557
C_{imt}	-8.039	-84.616
V_{it}	-8.977	-97.215
Consumer Goods		
δ_t	-2.534	-12.488
Δd_t	-6.498	-64.940
C_{imt}	-7.265	-72.265
V_{it}	-7.658	-78.494
Industrials (including Oil)		
δ_t	-2.943	-16.281
Δd_t	-7.212	-72.607
C_{imt}	-7.509	-75.348
V_{it}	-7.866	-80.157
Industrials (excluding Oil)		
δ_t	-2.713	-13.579
Δd_t	-6.338	-60.392
C_{imt}	-7.494	-75.265
V_{it}	-7.923	-81.359
Financial Services		
δ_t	-2.772	-14.404
Δd_t	-10.823	-118.691
C_{imt}	-6.372	-59.269
V_{it}	-5.983	-53.953
Other Sectors		
δ_t	-2.431	-11.530
Δd_t	-7.419	-75.681
C_{imt}	-7.881	-80.348
V_{it}	-9.187	-95.428

[1] All reported statistics are calculated using a truncation lag of 4. 10% critical values are -2.58 for the t-statistics and -11 for the $T(1-\phi)$ statistics (Fuller 1976). The PP statistics for the real interest rate are -6.950 (t-statistic), -72.785 ($T(1-\phi)$).

Table 6.9: Total Market Data, Variance Model
Quarterly Data

	Coefficient of Relative Risk Aversion				
	α=1	α=2	α=3	α=4	α=5
Non-linear Wald (p-value)	167.854 (0.000)	74.939 (0.000)	37.736 (0.000)	24.562 (0.017)	18.776 (0.094)
Linear Wald (p-value)	37.279 (0.000)	30.214 (0.003)	25.583 (0.012)	22.990 (0.028)	22.187 (0.035)
$\sigma(\delta')/\sigma(\delta)$ (s.e)	0.419 (0.277)	0.503 (0.360)	0.583 (0.438)	0.666 (0.519)	0.753 (0.602)
Corr(δ,δ') (s.e.)	-0.230 (0.400)	-0.117 (0.472)	-0.056 (0.542)	-0.015 (0.600)	0.015 (0.646)

Figure 6.8a: Actual and Theoretical Log Dividend-Price Ratio Total Market, Quarterly Data (α=3)

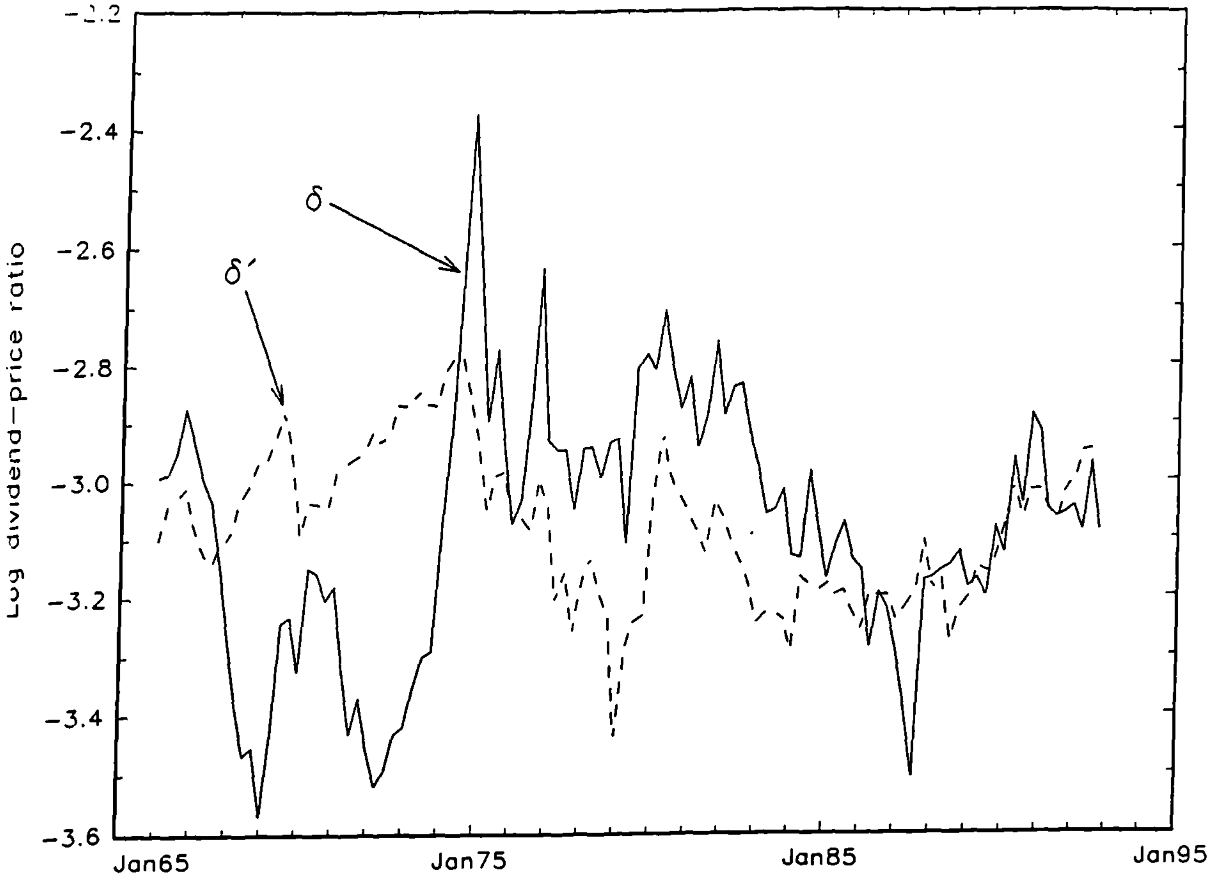

Figure 6.8b: Actual and Theoretical Log Real Stock Price Total Market, Quarterly Data (α=3)

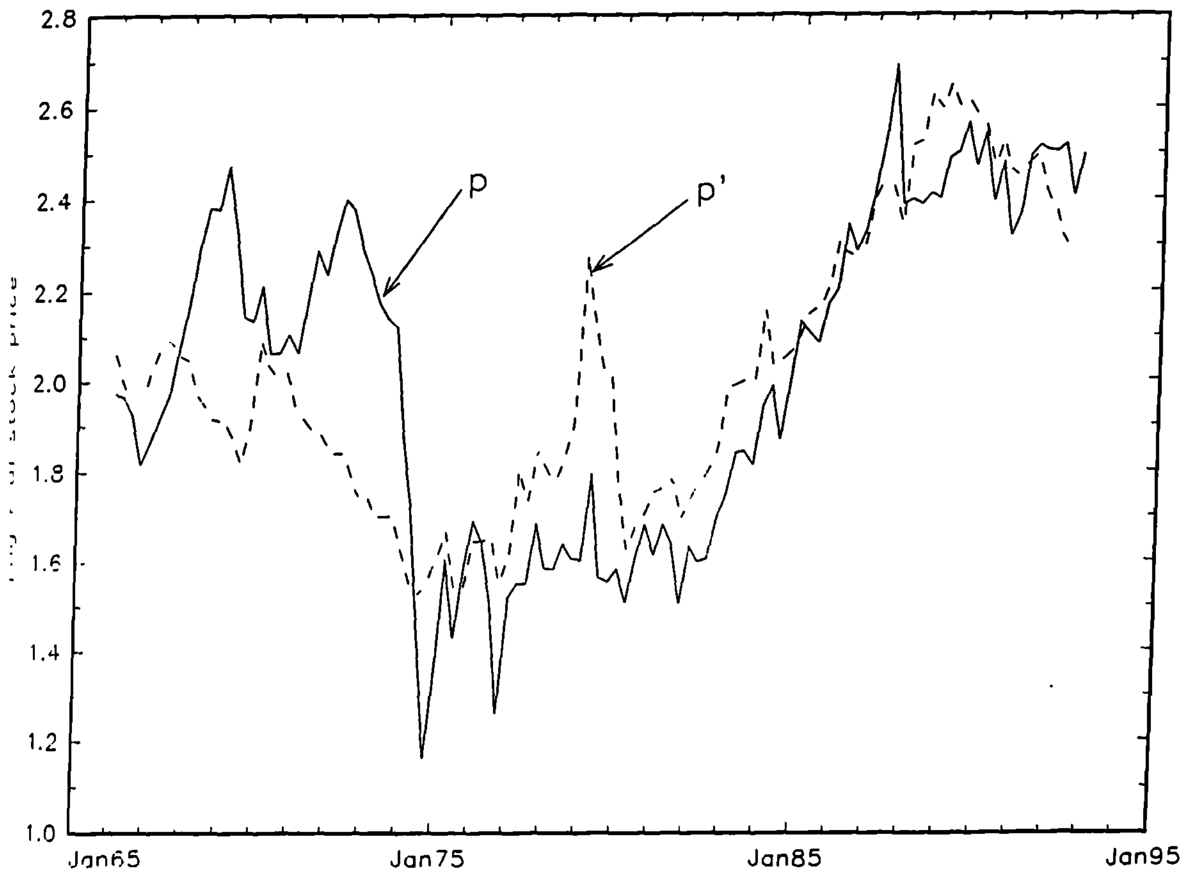

conclusion applies a fortiori when we examine the time-series behaviour of δ_t and δ_t'. Although the estimated standard deviation ratio is within one standard error of unity, the estimated correlations between δ_t and δ_t' are mostly *negative* and always insignificantly different from zero: there is no statistically discernible relationship between the actual and theoretical dividend-price ratios. This is illustrated in Figures 6.8a and 6.8b which compare δ_t and δ_t' and p_t and p_t' respectively when α=3.

Looking at the estimates of $f_1(\mathbf{A})$, $f_2(\mathbf{A})$ and $f_3(\mathbf{A})$[10], we have,

$$\delta_t' = \underset{(0.335)}{0.160}\,\delta_t - \underset{(1.751)}{1.973}\,\Delta d_t + \underset{(0.546)}{0.282}\,(r_t + \alpha V_t) + \ldots$$

with GMM standard errors in parentheses. The direct comparison of δ_t and δ_t' in Figure 6.8a depends only on these point estimates, which are quite a way from their hypothesised values of 1, 0 and 0 respectively, so that the two series do not move closely together. However, the non-linear Wald test takes sampling error into account, and the relatively large standard errors on the point estimates mean that the null cannot statistically be rejected. The same argument applies to the correlation between δ_t and δ_t' for which even the negative point estimate with α=3 is within two standard errors of +1. Again, my overall interpretation is that there is little support for the CAPM-EMH applied directly to the market return.

6.7.2 Results for Sector Portfolios: The Covariance Model

Table 6.10 presents the results of the application of the CAPM to the market sub-sectors, assuming no market segmentation exists. Again, for brevity, the results are presented only for α=3. Figures 6.9-6.14 compare graphically the actual and theoretical dividend-price ratios and real stock prices.

Looking first at the Figures, for most sectors, the story is the same as before: the actual and theoretical series move closely together, but the actual is always more variable than the theoretical. The exceptions appear to be Financial Services (Figure

---

Shiller (1981) using variance bounds tests on US data. The larger is α, the closer is the variance bound to its theoretically-determined value.

[10] There are actually 12 restrictions (corresponding to the 4 lags of the VAR), rather than the three reported here, but inclusion of the extra nine coefficients would not add anything to the point being made.

Table 6.10: Covariance Model, Quarterly Data (α=3)

	Capital Goods	Consumer Goods	Industrials (excl. Oil)	Industrials (incl. Oil)	Financial Services	Other Sectors
Non-linear Wald (p-value)	49.464 (0.000)	44.058 (0.000)	58.690 (0.000)	58.136 (0.000)	105.501 (0.000)	69.499 (0.000)
Linear Wald (p-value)	20.760 (0.054)	27.243 (0.007)	24.461 (0.018)	46.246 (0.000)	55.268 (0.000)	14.549 (0.267)
$\sigma(\delta')/\sigma(\delta)$ (s.e.)	0.553 (0.303)	0.456 (0.279)	0.468 (0.288)	0.354 (0.263)	0.533 (0.320)	0.270 (0.200)
$Corr(\delta,\delta')$ (s.e.)	0.678 (0.425)	0.499 (0.685)	0.599 (0.553)	0.294 (0.873)	-0.322 (0.776)	0.438 (0.823)
$\sigma(\delta-\delta_d')/\sigma(\delta_r')$ (s.e.)	3.503 (2.818)	3.734 (3.057)	3.538 (2.742)	4.309 (3.433)	3.978 (2.078)	3.673 (2.405)
$Corr(\delta-\delta_d',\delta_r')$ (s.e.)	0.912 (0.156)	0.869 (0.250)	0.877 (0.201)	0.819 (0.269)	-0.400 (0.726)	0.563 (0.610)
$\sigma(\xi')/\sigma(\xi)$ (s.e.)	0.603 (0.223)	0.586 (0.259)	0.567 (0.226)	0.472 (0.152)	0.633 (0.208)	0.418 (0.167)
$Corr(\xi,\xi')$ (s.e.)	0.721 (0.231)	0.670 (0.348)	0.681 (0.292)	0.598 (0.468)	0.213 (0.429)	0.751 (0.262)

Table 6.11: Variance Model, Quarterly Data (α=3)

	Capital Goods	Consumer Goods	Industrials (excl. Oil)	Industrials (incl. Oil)	Financial Services	Other Sectors
Non-linear Wald (p-value)	41.006 (0.000)	36.927 (0.000)	51.518 (0.000)	61.999 (0.000)	99.024 (0.000)	94.245 (0.000)
Linear Wald (p-value)	19.103 (0.086)	36.613 (0.000)	23.608 (0.023)	45.622 (0.000)	58.656 (0.000)	12.763 (0.387)
$\sigma(\delta')/\sigma(\delta)$ (s.e.)	0.615 (0.311)	0.549 (0.311)	0.510 (0.292)	0.328 (0.241)	0.677 (0.430)	0.212 (0.177)
$Corr(\delta,\delta')$ (s.e.)	0.741 (0.384)	0.636 (0.574)	0.660 (0.503)	0.349 (0.903)	-0.423 (0.641)	0.606 (0.778)
$\sigma(\delta-\delta_d')/\sigma(\delta_r')$ (s.e.)	2.767 (2.105)	2.643 (1.941)	3.078 (2.319)	3.993 (2.991)	2.560 (1.261)	2.583 (1.697)
$Corr(\delta-\delta_d',\delta_r')$ (s.e.)	0.913 (0.169)	0.894 (0.235)	0.917 (0.112)	0.829 (0.242)	-0.556 (0.490)	0.672 (0.405)
$\sigma(\xi')/\sigma(\xi)$ (s.e.)	0.662 (0.242)	0.698 (0.294)	0.608 (0.233)	0.454 (0.146)	0.739 (0.255)	0.398 (0.133)
$Corr(\xi,\xi')$ (s.e.)	0.762 (0.211)	0.731 (0.284)	0.723 (0.266)	0.637 (0.437)	0.105 (0.457)	0.808 (0.194)

Figure 6.9a: Actual and Theoretical Log Dividend-Price Ratio
Capital Goods, Quarterly Data (α=3)

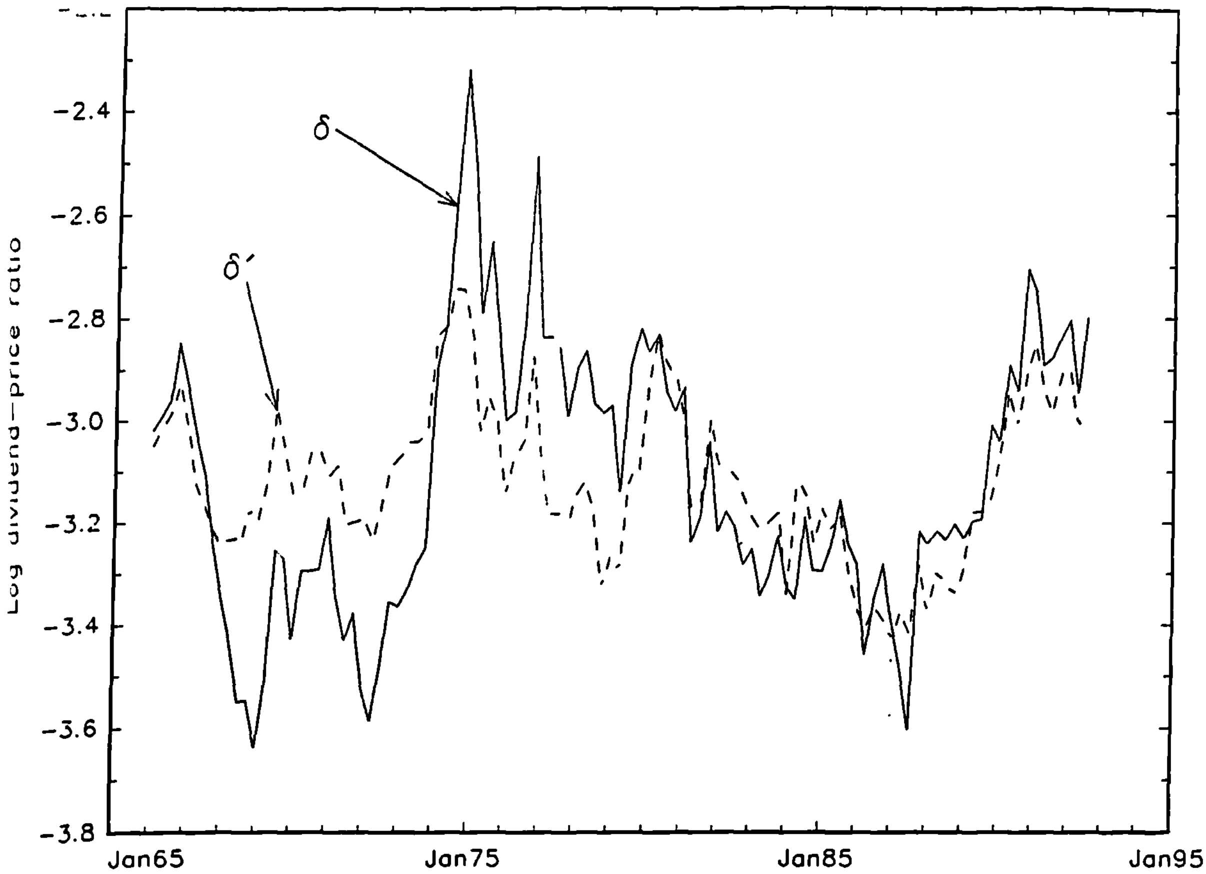

Figure 6.9b: Actual and Theoretical Log Real Stock Price
Capital Goods, Quarterly Data (α=3)

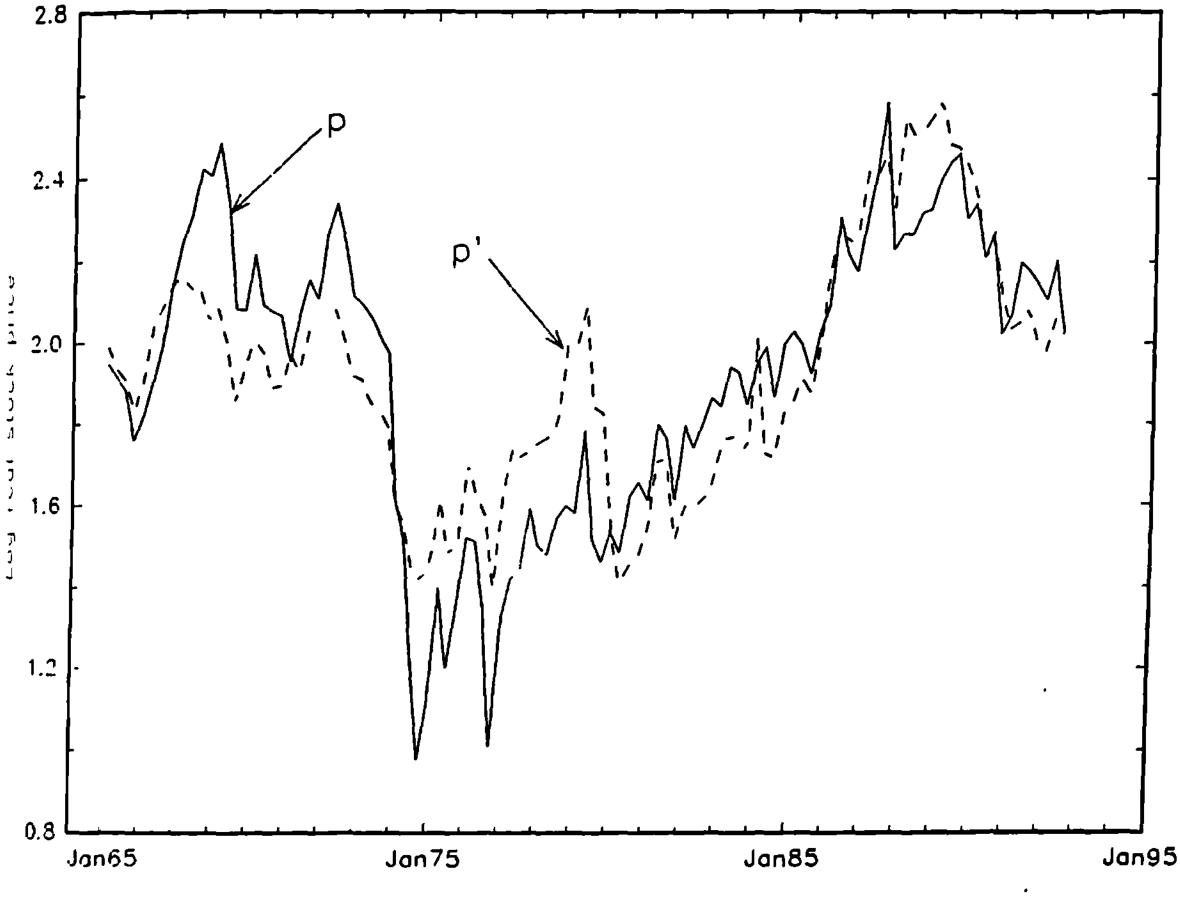

Figure 6.10a: Actual and Theoretical Log Dividend-Price Ratio
Consumer Goods, Quarterly Data (α=3)

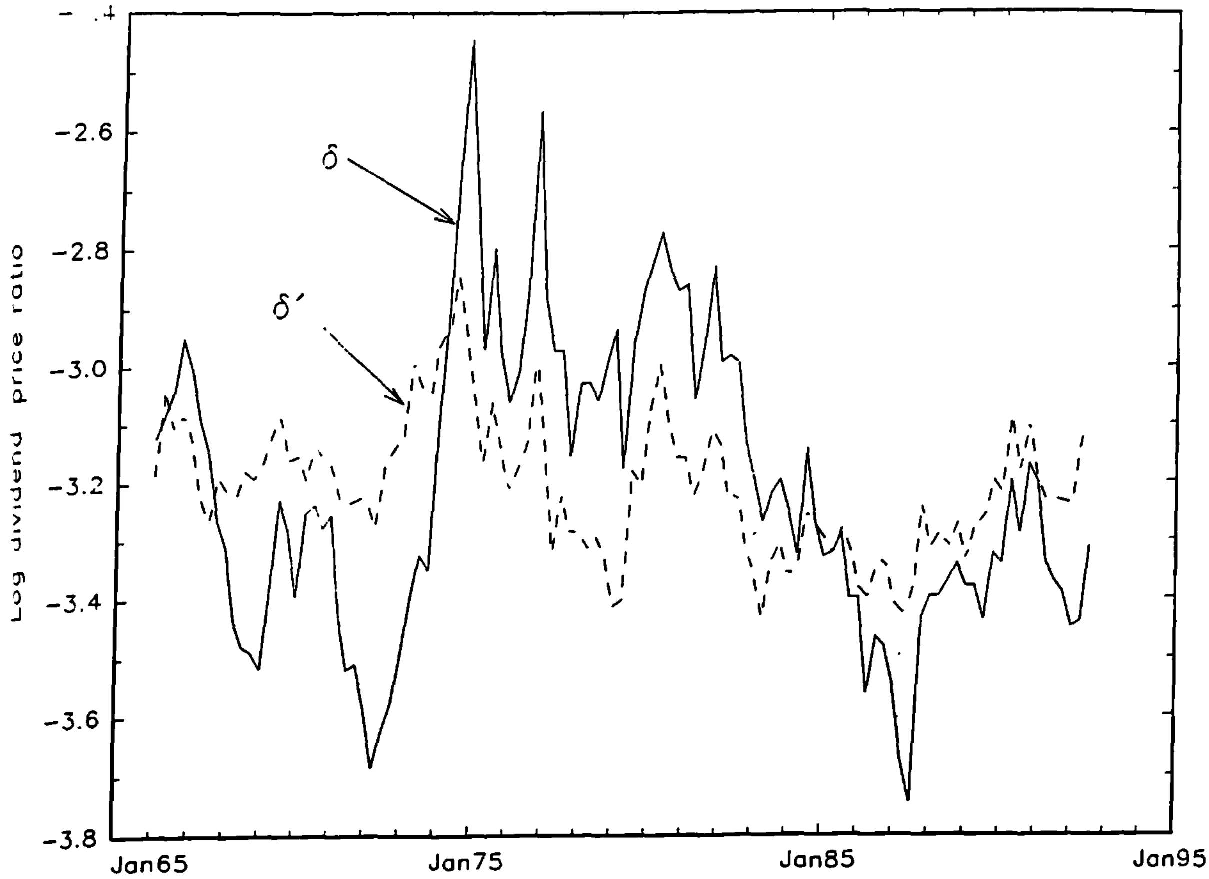

Figure 6.10b: Actual and Theoretical Log Real Stock Price
Consumer Goods, Quarterly Data (α=3)

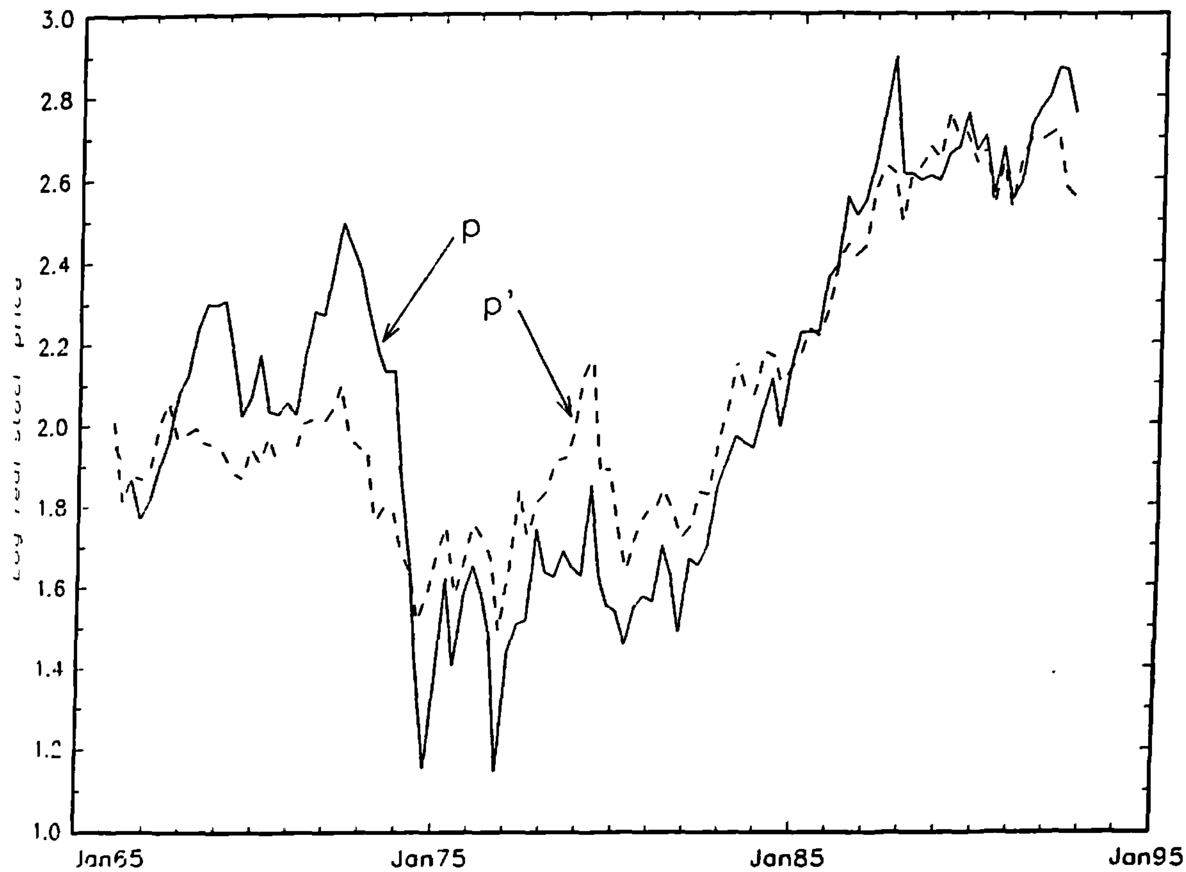

Figure 6.11a: Actual and Theoretical Log Dividend-Price Ratio
Industrials Including Oil, Quarterly Data (α=3)

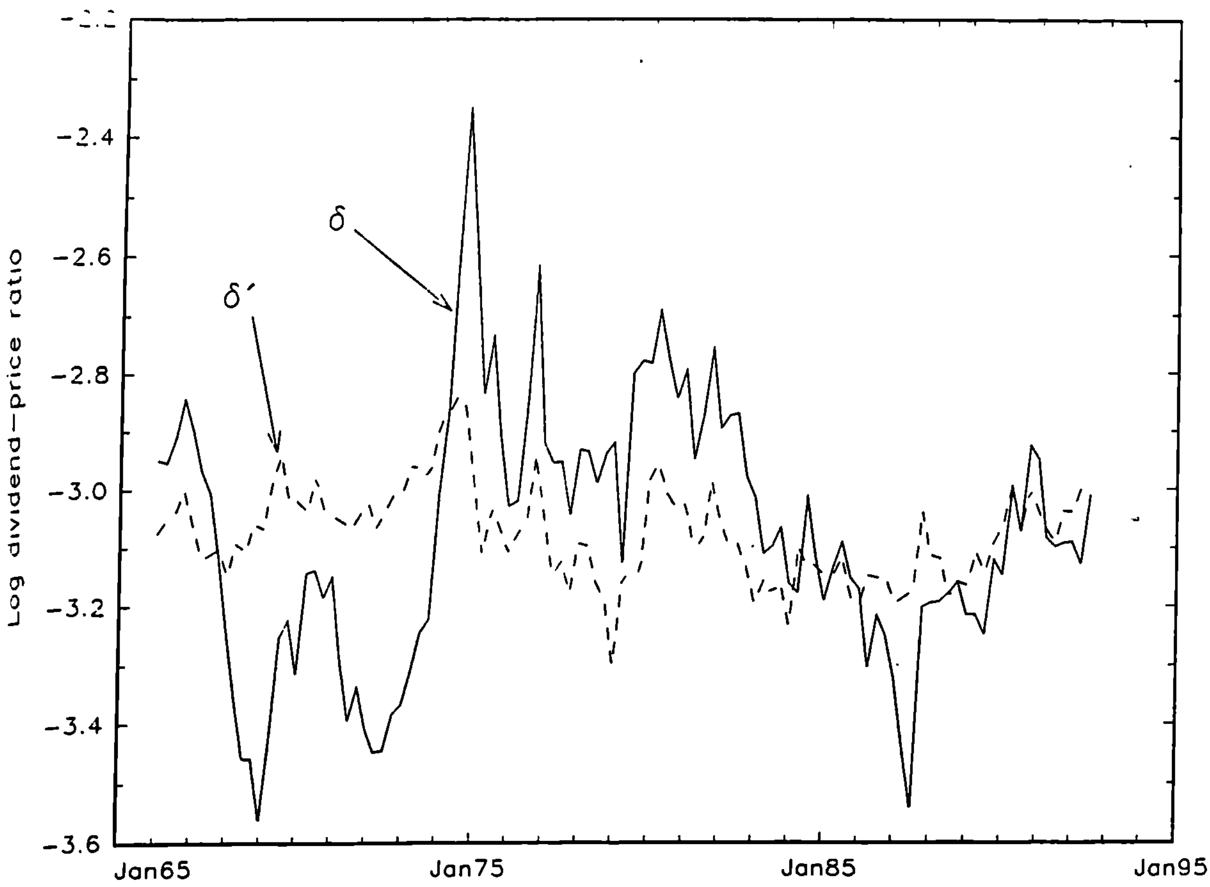

Figure 6.11b: Actual and Theoretical Log Real Stock Price
Industrials Including Oil, Quarterly Data (α=3)

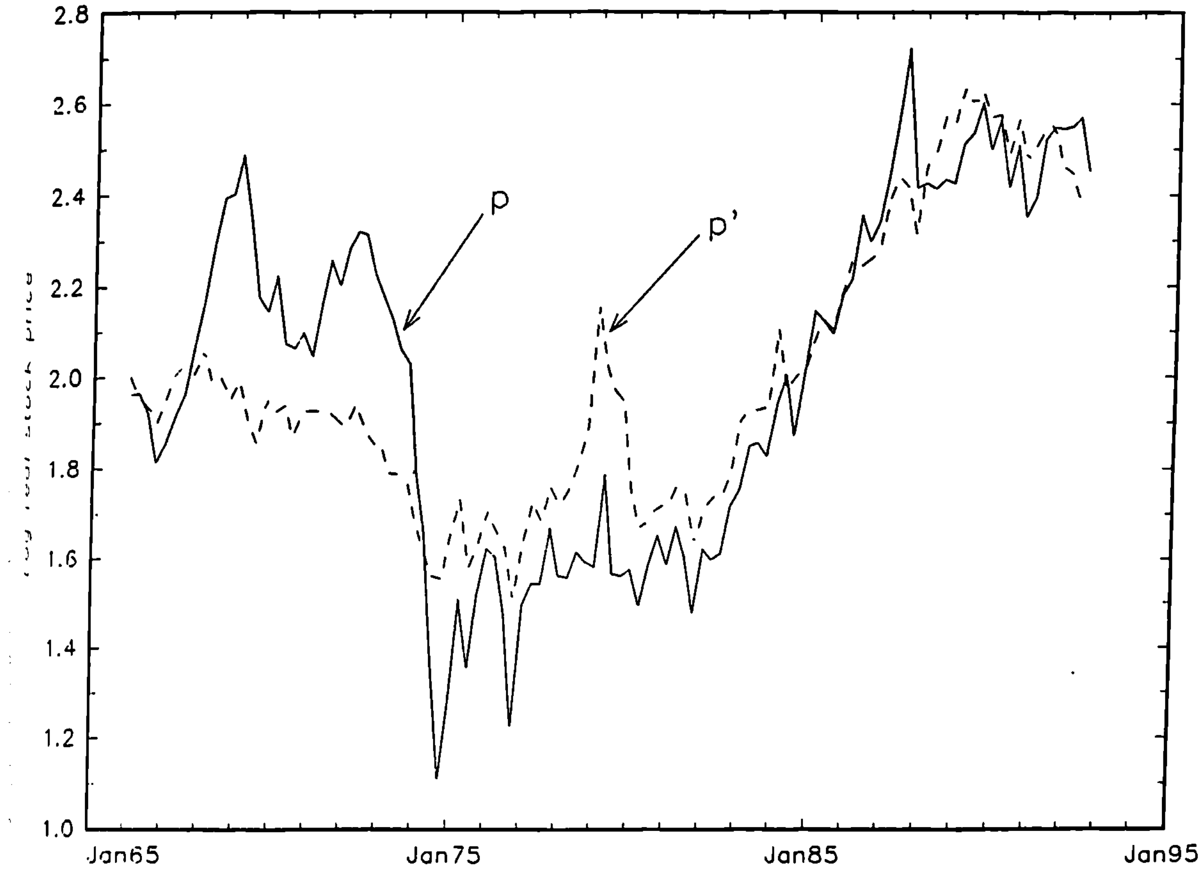

Figure 6.12a: Actual and Theoretical Log Dividend-Price Ratio Industrials Excluding Oil, Quarterly Data (α=3)

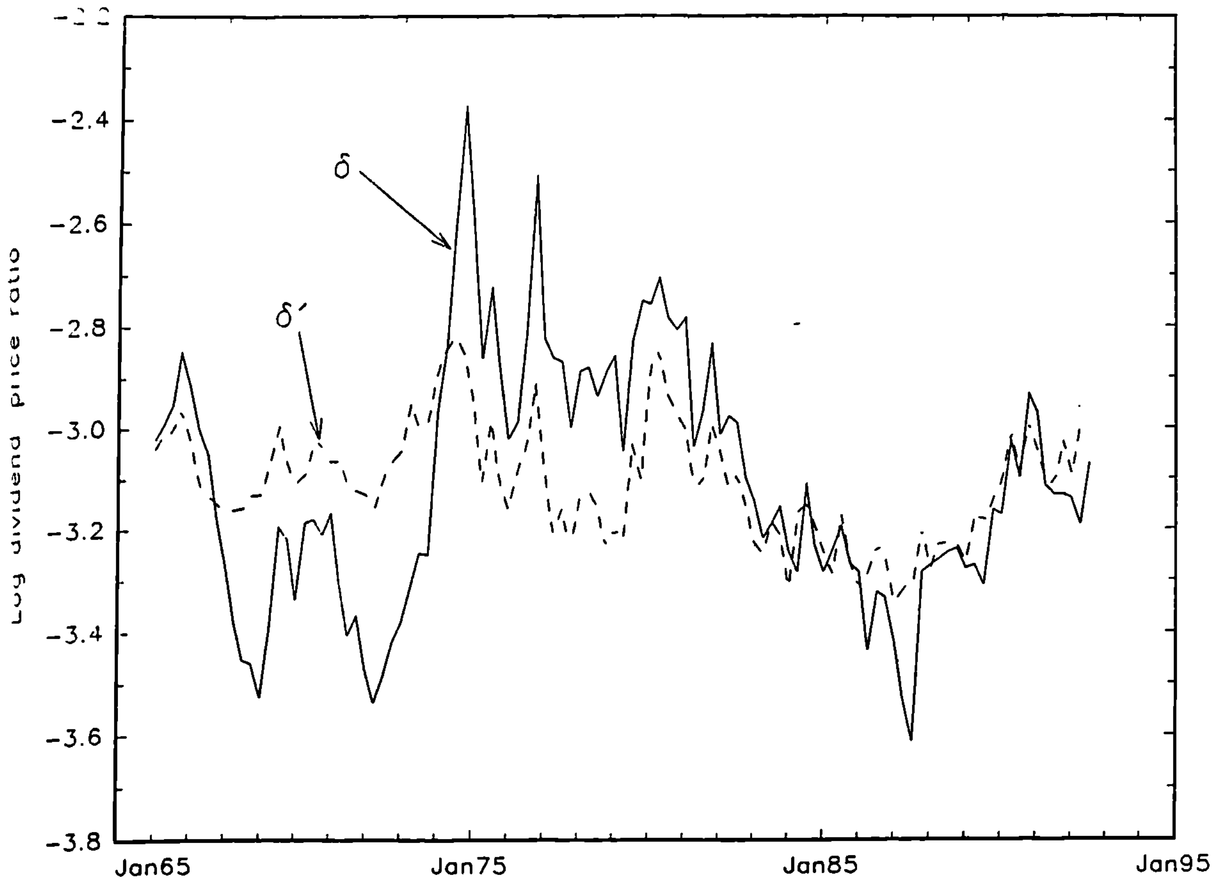

Figure 6.12b: Actual and Theoretical Log Real Stock Price Industrials Excluding Oil, Quarterly Data (α=3)

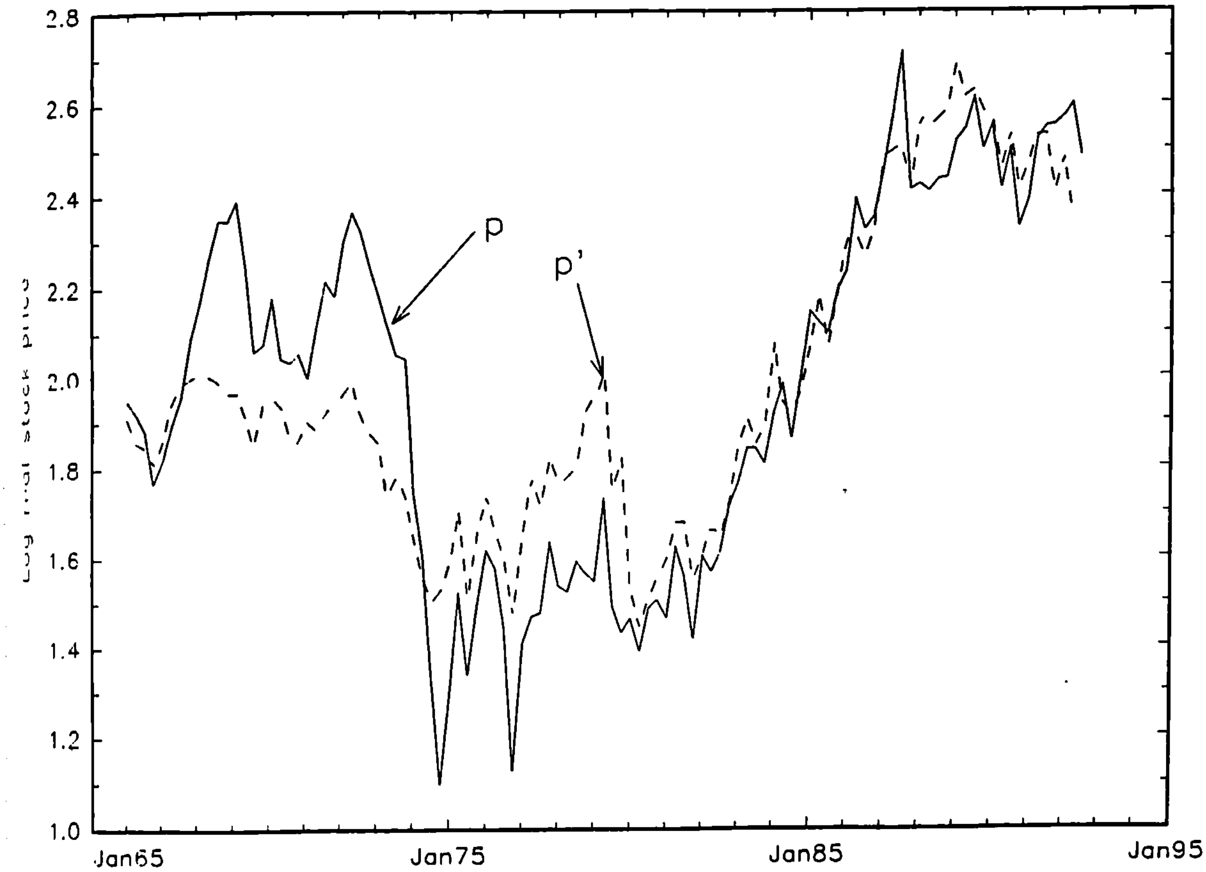

Figure 6.13a: Actual and Theoretical Log Dividend-Price Ratio
Financial Services, Quarterly Data (α=3)

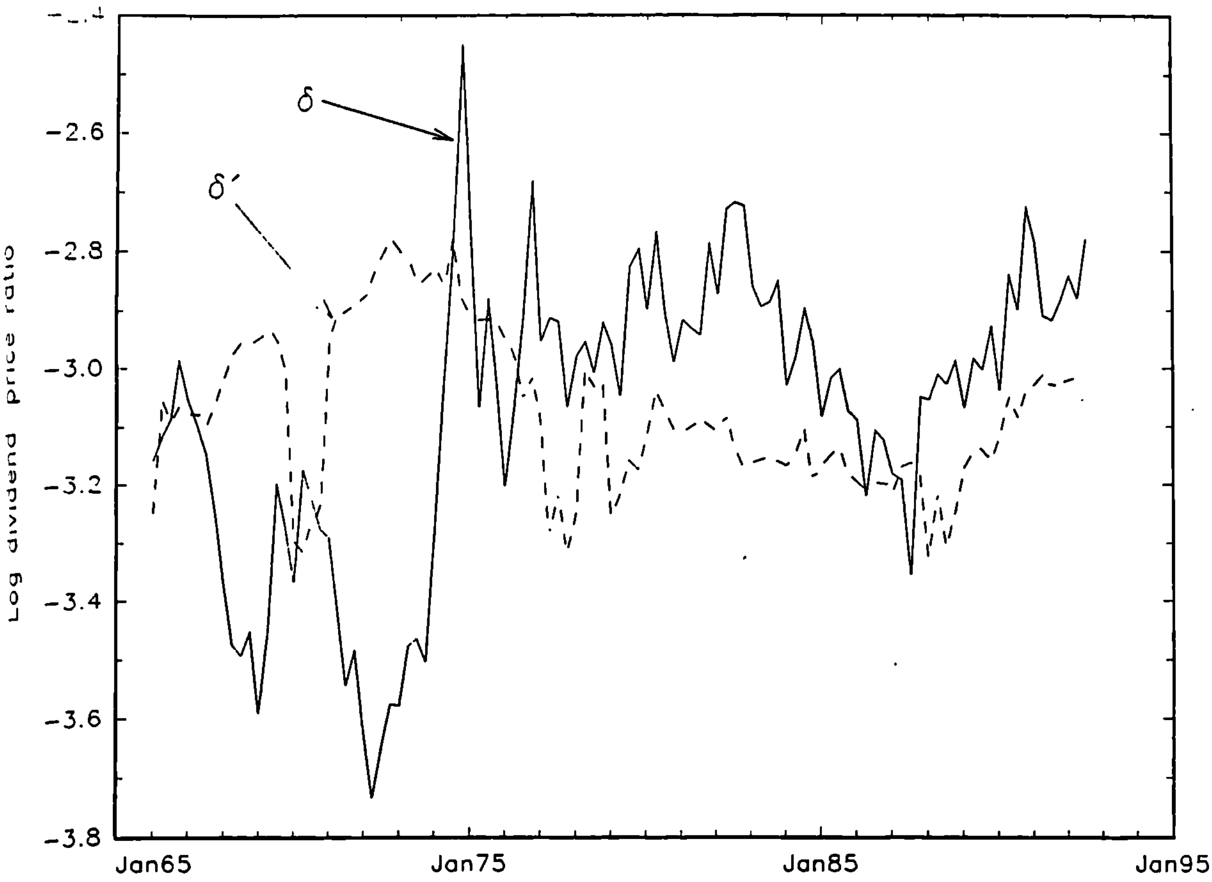

Figure 6.13b: Actual and Theoretical Log Real Stock Price
Financial Services, Quarterly Data (α=3)

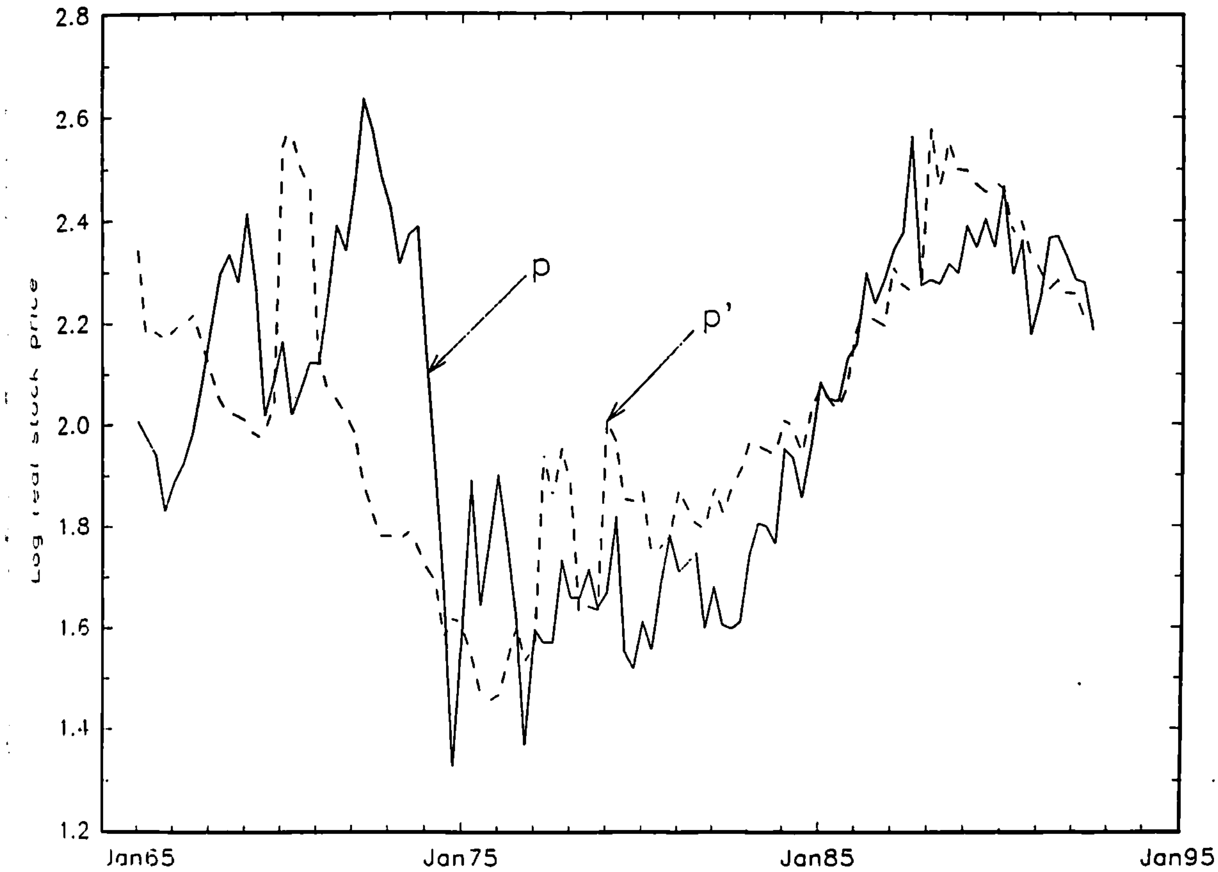

Figure 6.14a: Actual and Theoretical Log Dividend-Price Ratio
Other Sectors, Quarterly Data (α=3)

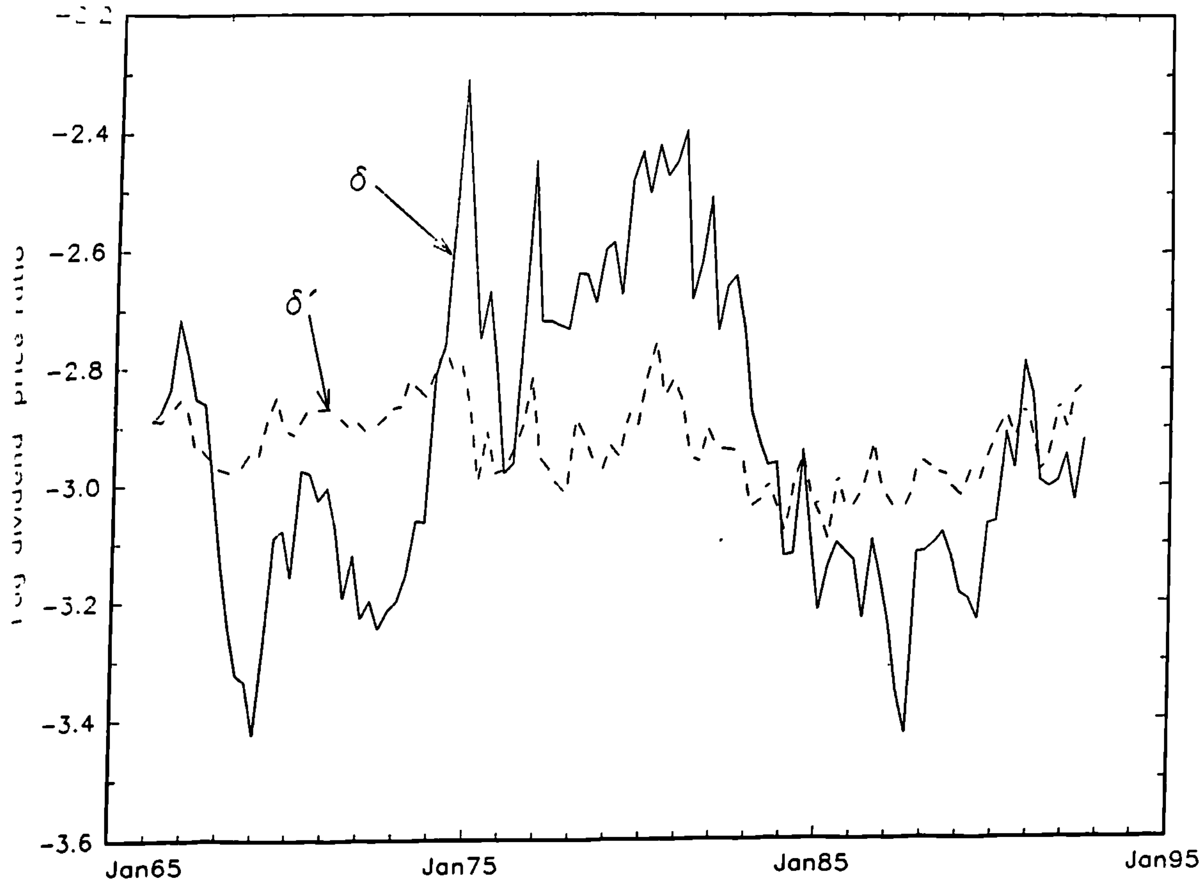

Figure 6.14b: Actual and Theoretical Log Real Stock Price
Other Sectors, Quarterly Data (α=3)

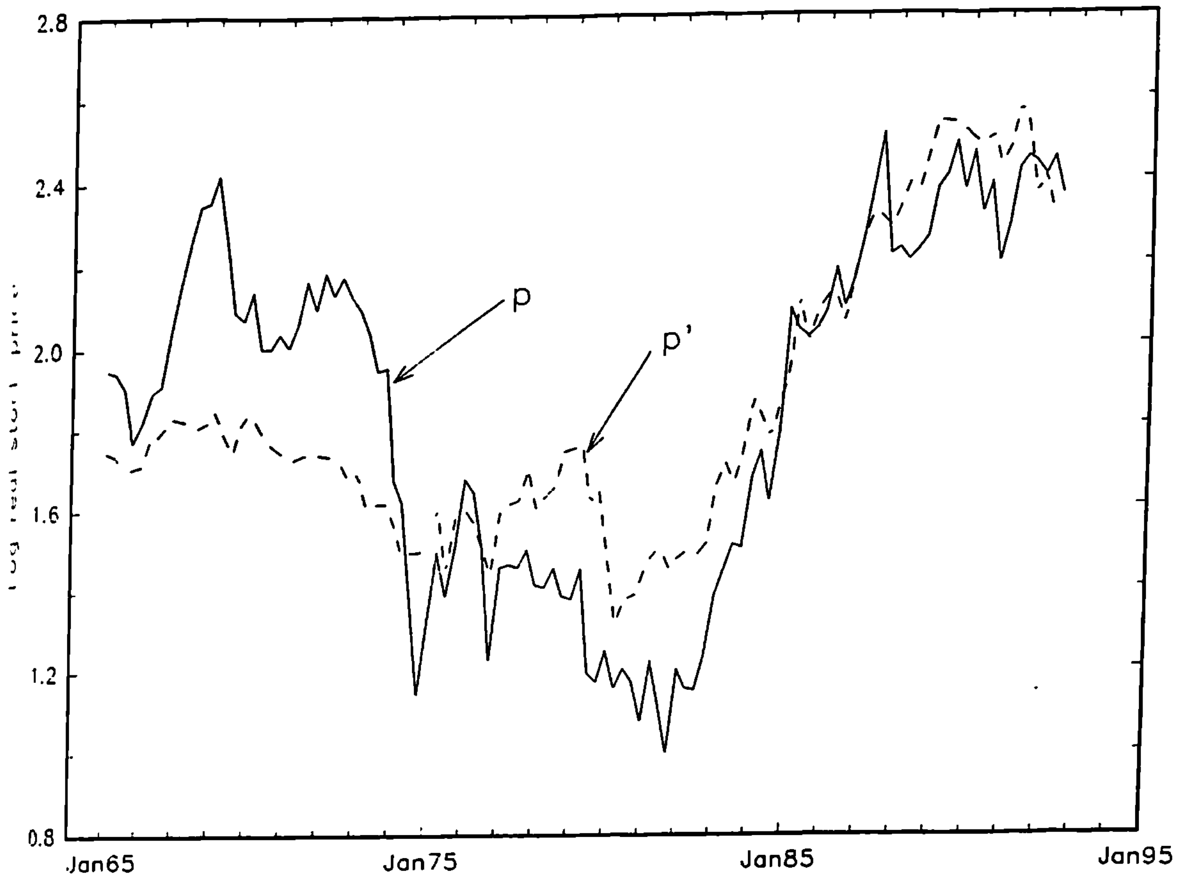

6.13) and Other Sectors (Figure 6.14). For both of these portfolios there are periods when the actual and theoretical series move markedly in the opposite direction. I therefore do not expect to find much support for the CAPM from these sectors.

The first row of results in Table 6.10 shows that all of the non-linear Wald statistics point towards rejection of the covariance model. However, as I have just discussed, I do not wish to put too much weight on inference from these statistics. Turning to the linear Wald test, I find that one-period returns are unpredictable for both the Capital Goods portfolio and for Other Sectors. The implication is that the forecasting variables in the VAR, in particular the log dividend-price ratio, are relevant only in so far as they forecast the covariance between portfolio and market returns. However, for none of the other sectors does the linear Wald test support this conclusion.

One notable result is that the log dividend-price ratio is now found significantly to Granger cause the discount rate measure. F-tests of the joint significance of the lagged δ_t's in the covariance equation produce the following marginal probability values:

Capital Goods	0.09
Consumer Goods	0.01
Industrials Including Oil	0.03
Industrials Excluding Oil	0.03
Financial Services	0.03
Other Sectors	0.21

This highlights a major difference between the monthly and quarterly data. When monthly data were used, the log dividend-price ratio did not Granger-cause the discount rate measure. However, with quarterly data, in five out of the six sector portfolios, the dividend-price ratio does significantly Granger-cause discount rates at the 10% level. One might therefore expect to find more supportive evidence for market efficiency from the analysis of the lower-frequency data.

Direct comparison of δ_t and δ_t' reveals more positive results for most sectors in comparison with those reported for the aggregate market. The standard deviation ratios (row 3) all indicate that the theoretical dividend-price ratio is less variable than is actually observed, but for all sectors except Industrials (including oil) and Other Sectors, the point estimates are within two standard errors of unity. However, the negative estimated correlation between δ_t and δ_t' reported for the Financial Services portfolio is strongly suggestive of a rejection of the model for this sector (although, again notice the large standard error). The correlation is largest for the Capital Goods sector (0.678, s.e.=0.425), but none of the other estimates is outside of one standard error of unity.

Comparison of the actual and theoretical discount rate components (rows 5 and 6) indicates a very high correlation between the two for most of the portfolios (for Capital Goods the correlation is highest at 0.912, s.e.=0.156), and although all of the estimated standard deviation ratios are very large, so are the standard errors, so that the null of unity cannot be rejected.

Finally, in rows 8 and 9 the behaviour of actual and theoretical returns,ξ_t and ξ_t' is compared. For every sector the theoretical return is less variable than the actual return, but only for Industrials Including Oil and Other Sectors is the difference statistically significant. None of the estimated correlations is beyond two standard errors of unity, but the point estimate for Financial Services is particularly low (0.213).

Thus although inference is once again clouded by sizeable standard errors, there is mixed but more positive support for the CAPM-EMH at this lower level of aggregation, with evidence for the Capital Goods sector being fairly strong, and that for Financial Services quite weak.

6.7.3 Results for Sector Portfolios: The Own Variance model

One's initial impression from Table 6.11 is that the point estimates for most statistics for the own-variance model are a marginal improvement on those for the covariance model, although the sizeable standard errors mean that the improvement is not statistically significant. Once again, all of the non-linear Wald statistics are so large as to reject the efficient markets null. As with the covariance model, both Capital Goods and Other Sectors do not reject return unpredictability (linear Wald test) at the 5% level, whilst the marginal probability value for Industrials Excluding Oil is 2.3%. None of the other sectors provides supportive statistics. This suggests that the question of whether returns over and above equilibrium returns are predictable is unaffected by the inclusion of the covariances.

The vast majority of point estimates of the remaining statistics move closer to, and are within one standard error of, their hypothesised value of unity. The notable exception is again the Financial Services sector, for which I estimate negative correlations between δ_t and δ_t' and the discount rate components δ_t-δ_{dt}' and δ_{rt}'. Capital Goods and Industrials Excluding Oil fare particularly well.

Although the large standard errors on some of the estimates make precise conclusions difficult, my general finding is that the own-variance model performs marginally better than the covariance model. However, neither model provides a complete explanation of movements in the dividend-price ratio for most of the industrial portfolios studied.

6.8 Some Variants

In this section I discuss the effects of varying the coefficient of relative risk aversion, α, on the VAR test results. Looking first at the results using monthly data, Tables 6.12 to 6.19 present the results for the covariance and own-variance model for α=1,2,4 and 5. The most striking result is that the linear Wald test rejects the null of unpredictability of one-period returns for all values of α. Thus, neither the covariance of the sector return with the market return, nor the sector return variance, is found to be a sufficient statistic to summarise systematic movements in sectoral returns. The point estimates of the standard deviation ratios, $\sigma(\delta')/\sigma(\delta)$, get larger as α increases, but so do the standard errors, so that their difference from unity is never significant. The correlations between δ_t and δ_t' tend to increase with α, although the rise is never marked. Much the same can be said when comparing the actual and theoretical log real returns, ξ_t and ξ_t', although the increase in the correlation between the two series as α increases is perhaps more prominent.

The results using quarterly data appear in Tables 6.20 to 6.27. As far as the linear Wald test is concerned, the most supportive results are found for Capital Goods and Other Sectors, which do not reject the null of unpredictability for α >1. For all sectors, the non-linear Wald statistic falls as α rises, but the null hypothesis

Table 6.12: Covariance Model, Monthly Data (α=1)

	Capital Goods	Consumer Goods	Industrials (incl. Oil)	Industrials (excl. Oil)	Financial Services	Other Sectors
Non-linear Wald (p-value)	24.577 (0.137)	6.622 (0.993)	23.456 (0.174)	14.398 (0.703)	252.618 (0.000)	35.529 (0.008)
Linear Wald (p-value)	38.646 (0.000)	23.871 (0.021)	46.437 (0.000)	29.117 (0.004)	60.450 (0.000)	36.943 (0.000)
$\sigma(\delta')/\sigma(\delta)$ (s.e.)	1.315 (0.969)	1.844 (2.758)	1.397 (1.260)	1.561 (1.535)	0.463 (0.276)	1.115 (0.956)
Corr(δ,δ') (s.e.)	0.835 (0.115)	0.515 (0.337)	0.780 (0.201)	0.765 (0.177)	0.722 (0.274)	0.707 (0.237)
$\sigma(\delta-\delta_d')/\sigma(\delta_r')$ (s.e.)	2.233 (1.612)	4.701 (5.397)	2.224 (2.127)	2.612 (2.814)	5.124 (2.776)	3.689 (3.751)
Corr($\delta-\delta_d'$,δ_r') (s.e.)	-0.210 (0.239)	-0.483 (1.322)	-0.085 (1.378)	-0.358 (1.297)	0.838 (0.279)	0.022 (1.171)
$\sigma(\xi')/\sigma(\xi)$ (s.e.)	1.053 (0.749)	1.574 (2.368)	1.176 (1.031)	1.248 (1.189)	0.598 (0.238)	0.962 (0.751)
Corr(ξ,ξ') (s.e.)	0.757 (0.158)	0.328 (0.404)	0.681 (0.249)	0.627 (0.266)	0.768 (0.105)	0.566 (0.326)

Table 6.13: Variance Model, Monthly Data (α=1)

	Capital Goods	Consumer Goods	Industrials (incl. Oil)	Industrials (excl. Oil)	Financial Services	Other Sectors
Non-linear Wald (p-value)	41.249 (0.001)	10.142 (0.927)	47.760 (0.000)	41.847 (0.001)	331.286 (0.000)	39.360 (0.003)
Linear Wald (p-value)	48.998 (0.000)	79.573 (0.000)	63.429 (0.000)	51.010 (0.000)	91.138 (0.000)	44.595 (0.000)
$\sigma(\delta')/\sigma(\delta)$ (s.e.)	1.331 (0.857)	1.869 (2.134)	1.384 (1.062)	1.556 (1.394)	0.464 (0.283)	1.132 (0.868)
Corr(δ,δ') (s.e.)	0.836 (0.118)	0.530 (0.342)	0.779 (0.218)	0.768 (0.183)	0.724 (0.294)	0.713 (0.234)
$\sigma(\delta-\delta_d')/\sigma(\delta_r')$ (s.e.)	2.065 (1.345)	4.214 (4.037)	2.301 (2.059)	2.959 (2.613)	4.963 (2.570)	3.373 (3.078)
Corr($\delta-\delta_d'$,δ_r') (s.e.)	-0.197 (1.170)	-0.473 (1.052)	-0.076 (1.187)	-0.342 (1.225)	0.837 (0.278)	0.027 (1.089)
$\sigma(\xi')/\sigma(\xi)$ (s.e.)	1.081 (0.639)	1.566 (1.758)	1.170 (0.818)	1.245 (1.033)	0.615 (0.174)	0.974 (0.679)
Corr(ξ,ξ') (s.e.)	0.757 (0.158)	0.317 (0.427)	0.678 (0.263)	0.627 (0.260)	0.766 (0.117)	0.572 (0.319)

Table 6.14: Covariance Model, Monthly Data (α=2)

	Capital Goods	Consumer Goods	Industrials (incl. Oil)	Industrials (excl. Oil)	Financial Services	Other Sectors
Non-linear Wald (p-value)	11.043 (0.893)	3.427 (0.999)	12.259 (0.834)	6.446 (0.994)	88.889 (0.000)	13.824 (0.740)
Linear Wald (p-value)	46.353 (0.000)	26.922 (0.008)	64.655 (0.000)	37.815 (0.000)	51.833 (0.000)	32.174 (0.001)
$\sigma(\delta')/\sigma(\delta)$ (s.e.)	1.841 (1.967)	2.552 (5.322)	2.006 (2.586)	2.268 (3.277)	0.630 (0.442)	1.477 (1.695)
$Corr(\delta,\delta')$ (s.e.)	0.869 (0.116)	0.562 (0.344)	0.818 (0.180)	0.791 (0.183)	0.806 (0.234)	0.765 (0.233)
$\sigma(\delta-\delta_d')/\sigma(\delta_r')$ (s.e.)	0.816 (0.767)	1.513 (1.591)	0.881 (0.920)	1.059 (1.025)	2.488 (1.794)	1.397 (1.680)
$Corr(\delta-\delta_d',\delta_r')$ (s.e.)	-0.152 (1.824)	-0.528 (1.800)	-0.139 (2.089)	-0.433 (1.801)	0.832 (0.327)	0.092 (1.560)
$\sigma(\xi')/\sigma(\xi)$ (s.e.)	1.506 (1.643)	2.197 (4.629)	1.713 (2.230)	1.839 (2.686)	0.770 (0.456)	1.269 (1.433)
$Corr(\xi,\xi')$ (s.e.)	0.792 (0.148)	0.403 (0.430)	0.723 (0.217)	0.666 (0.264)	0.833 (0.113)	0.653 (0.314)

Table 6.15: Variance Model, Monthly Data (α=2)

	Capital Goods	Consumer Goods	Industrials (incl. Oil)	Industrials (excl. Oil)	Financial Services	Other Sectors
Non-linear Wald (p-value)	14.519 (0.695)	6.570 (0.993)	19.187 (0.380)	18.994 (0.392)	135.386 (0.000)	18.028 (0.454)
Linear Wald (p-value)	57.575 (0.000)	52.793 (0.000)	99.295 (0.000)	61.354 (0.000)	75.662 (0.000)	36.294 (0.000)
$\sigma(\delta')/\sigma(\delta)$ (s.e.)	1.879 (1.193)	2.643 (2.937)	1.980 (1.433)	2.269 (2.031)	0.641 (0.306)	1.516 (1.147)
$Corr(\delta,\delta')$ (s.e.)	0.869 (0.104)	0.577 (0.343)	0.818 (0.192)	0.794 (0.174)	0.807 (0.255)	0.771 (0.214)
$\sigma(\delta-\delta_d')/\sigma(\delta_r')$ (s.e.)	0.772 (0.440)	1.390 (1.013)	0.906 (0.716)	1.057 (0.719)	2.392 (1.275)	1.304 (1.057)
$Corr(\delta-\delta_d',\delta_r')$ (s.e.)	-0.146 (1.212)	-0.534 (0.956)	-0.133 (1.150)	-0.428 (1.141)	0.827 (0.320)	0.088 (1.099)
$\sigma(\xi')/\sigma(\xi)$ (s.e.)	1.559 (0.918)	2.223 (2.412)	1.695 (1.127)	1.834 (1.524)	0.793 (0.282)	1.297 (0.918)
$Corr(\xi,\xi')$ (s.e.)	0.791 (0.141)	0.384 (0.449)	0.719 (0.227)	0.664 (0.249)	0.825 (0.124)	0.656 (0.296)

Table 6.16: Covariance Model, Monthly Data (α=4)

	Capital Goods	Consumer Goods	Industrials (incl. Oil)	Industrials (excl. Oil)	Financial Services	Other Sectors
Non-linear Wald (p-value)	6.045 (0.996)	2.799 (0.999)	7.812 (0.981)	4.261 (0.999)	31.543 (0.025)	5.534 (0.998)
Linear Wald (p-value)	104.885 (0.000)	81.556 (0.000)	108.314 (0.000)	98.112 (0.000)	91.575 (0.000)	70.824 (0.000)
$\sigma(\delta')/\sigma(\delta)$ (s.e.)	2.878 (5.395)	3.801 (12.243)	3.165 (6.771)	3.544 (8.464)	0.976 (0.977)	2.187 (3.884)
$Corr(\delta,\delta')$ (s.e.)	0.898 (0.121)	0.620 (0.349)	0.854 (0.165)	0.822 (0.188)	0.876 (0.225)	0.826 (0.227)
$\sigma(\delta-\delta_d')/\sigma(\delta_r')$ (s.e.)	0.359 (0.576)	0.634 (0.899)	0.398 (0.600)	0.463 (0.617)	1.226 (1.447)	0.625 (1.171)
$Corr(\delta-\delta_d',\delta_r')$ (s.e.)	-0.146 (3.116)	-0.518 (2.910)	-0.157 (3.729)	-0.460 (2.955)	0.830 (0.392)	0.113 (2.337)
$\sigma(\xi')/\sigma(\xi)$ (s.e.)	2.436 (4.734)	3.323 (10.835)	2.771 (6.055)	2.946 (7.197)	1.165 (1.105)	1.915 (3.502)
$Corr(\xi,\xi')$ (s.e.)	0.815 (0.136)	0.484 (0.438)	0.756 (0.189)	0.704 (0.254)	0.859 (0.134)	0.731 (0.280)

Table 6.17: Variance Model, Monthly Data (α=4)

	Capital Goods	Consumer Goods	Industrials (incl. Oil)	Industrials (excl. Oil)	Financial Services	Other Sectors
Non-linear Wald (p-value)	6.805 (0.992)	4.290 (0.999)	10.303 (0.922)	5.951 (0.996)	41.607 (0.001)	6.061 (0.996)
Linear Wald (p-value)	115.959 (0.000)	92.960 (0.000)	130.812 (0.000)	109.021 (0.000)	130.533 (0.000)	72.746 (0.000)
$\sigma(\delta')/\sigma(\delta)$ (s.e.)	2.962 (1.850)	4.012 (4.258)	3.117 (2.163)	3.557 (3.101)	1.011 (0.444)	2.269 (1.680)
$Corr(\delta,\delta')$ (s.e.)	0.899 (0.097)	0.632 (0.342)	0.855 (0.168)	0.825 (0.168)	0.873 (0.233)	0.829 (0.194)
$\sigma(\delta-\delta_d')/\sigma(\delta_r')$ (s.e.)	0.342 (0.190)	0.586 (0.399)	0.407 (0.312)	0.461 (0.290)	1.171 (0.639)	0.587 (0.465)
$Corr(\delta-\delta_d',\delta_r')$ (s.e.)	-0.142 (1.214)	-0.532 (0.932)	-0.152 (1.144)	-0.459 (1.088)	0.821 (0.351)	0.107 (1.081)
$\sigma(\xi')/\sigma(\xi)$ (s.e.)	2.534 (1.456)	3.413 (3.472)	2.728 (1.737)	2.938 (2.343)	1.206 (0.537)	1.974 (1.366)
$Corr(\xi,\xi')$ (s.e.)	0.813 (0.130)	0.459 (0.458)	0.751 (0.194)	0.701 (0.238)	0.843 (0.135)	0.728 (0.267)

Table 6.18: Covariance Model, Monthly Data (α=5)

	Capital Goods	Consumer Goods	Industrials (incl. Oil)	Industrials (excl. Oil)	Financial Services	Other Sectors
Non-linear Wald (p-value)	5.783 (0.997)	2.900 (0.999)	7.471 (0.986)	4.333 (0.999)	24.697 (0.134)	4.668 (0.999)
Linear Wald (p-value)	117.420 (0.000)	92.941 (0.000)	117.240 (0.000)	113.555 (0.000)	104.535 (0.000)	75.396 (0.000)
$\sigma(\delta')/\sigma(\delta)$ (s.e.)	3.395 (7.733)	4.427 (16.597)	3.742 (9.568)	4.171 (11.795)	1.153 (1.331)	2.544 (5.281)
Corr(δ,δ') (s.e.)	0.906 (0.122)	0.637 (0.348)	0.864 (0.161)	0.830 (0.188)	0.892 (0.222)	0.842 (0.223)
$\sigma(\delta-\delta_d')/\sigma(\delta_r')$ (s.e.)	0.281 (0.545)	0.491 (0.804)	0.312 (0.555)	0.361 (0.556)	0.978 (1.364)	0.490 (1.081)
Corr($\delta-\delta_d',\delta_r'$) (s.e.)	-0.146 (3.730)	-0.513 (3.440)	-0.159 (4.563)	-0.463 (3.494)	0.830 (0.417)	0.118 (2.709)
$\sigma(\xi')/\sigma(\xi)$ (s.e.)	2.906 (6.846)	3.890 (14.758)	3.302 (8.619)	3.495 (10.104)	1.376 (1.527)	2.250 (4.833)
Corr(ξ,ξ') (s.e.)	0.819 (0.131)	0.506 (0.435)	0.763 (0.182)	0.714 (0.248)	0.858 (0.140)	0.750 (0.265)

Table 6.19: Variance Model, Monthly Data (α=5)

	Capital Goods	Consumer Goods	Industrials (incl. Oil)	Industrials (excl. Oil)	Financial Services	Other Sectors
Non-linear Wald (p-value)	6.955 (0.990)	4.109 (0.999)	10.086 (0.929)	5.787 (0.997)	35.325 (0.009)	6.064 (0.996)
Linear Wald (p-value)	131.490 (0.000)	123.595 (0.000)	181.345 (0.000)	129.446 (0.000)	160.874 (0.000)	79.028 (0.000)
$\sigma(\delta')/\sigma(\delta)$ (s.e.)	3.502 (2.182)	4.695 (4.925)	3.682 (2.538)	4.190 (3.626)	1.199 (0.517)	2.649 (1.953)
Corr(δ,δ') (s.e.)	0.906 (0.095)	0.647 (0.342)	0.865 (0.161)	0.833 (0.167)	0.888 (0.228)	0.845 (0.188)
$\sigma(\delta-\delta_d')/\sigma(\delta_r')$ (s.e.)	0.268 (0.149)	0.454 (0.306)	0.319 (0.243)	0.359 (0.223)	0.932 (0.511)	0.461 (0.364)
Corr($\delta-\delta_d',\delta_r'$) (s.e.)	-0.142 (1.214)	-0.528 (0.930)	-0.154 (1.108)	-0.464 (1.079)	0.820 (0.357)	0.111 (1.076)
$\sigma(\xi')/\sigma(\xi)$ (s.e.)	3.025 (1.726)	4.009 (4.006)	3.247 (2.050)	3.487 (2.745)	1.426 (0.669)	2.324 (1.596)
Corr(ξ,ξ') (s.e.)	0.817 (0.128)	0.480 (0.459)	0.758 (0.185)	0.711 (0.234)	0.840 (0.136)	0.745 (0.256)

Table 6.20: Covariance Model, Quarterly Data (α=1)

	Capital Goods	Consumer Goods	Industrials (excl. Oil)	Industrials (incl. Oil)	Financial Services	Other Sectors
Non-linear Wald (p-value)	158.711 (0.000)	152.750 (0.000)	151.463 (0.000)	188.970 (0.000)	311.183 (0.000)	143.737 (0.000)
Linear Wald (p-value)	36.765 (0.000)	39.292 (0.000)	36.957 (0.000)	62.109 (0.000)	52.734 (0.000)	35.198 (0.000)
$\sigma(\delta')/\sigma(\delta)$ (s.e.)	0.379 (0.219)	0.285 (0.152)	0.336 (0.189)	0.282 (0.211)	0.302 (0.154)	0.288 (0.170)
$Corr(\delta,\delta')$ (s.e.)	0.363 (0.454)	0.021 (0.645)	0.220 (0.556)	-0.152 (0.558)	-0.261 (0.845)	0.113 (0.703)
$\sigma(\delta-\delta_d')/\sigma(\delta_r')$ (s.e.)	9.039 (5.557)	7.434 (4.048)	5.781 (2.842)	6.754 (3.390)	18.003 (13.912)	4.699 (2.412)
$Corr(\delta-\delta_d',\delta_r')$ (s.e.)	0.594 (0.457)	0.544 (0.471)	0.460 (0.432)	0.550 (0.401)	0.684 (0.825)	0.370 (0.479)
$\sigma(\xi')/\sigma(\xi)$ (s.e.)	0.413 (0.137)	0.310 (0.112)	0.356 (0.135)	0.321 (0.103)	0.483 (0.101)	0.363 (0.120)
$Corr(\xi,\xi')$ (s.e.)	0.496 (0.292)	0.302 (0.472)	0.390 (0.391)	0.331 (0.479)	0.167 (0.324)	0.500 (0.321)

Table 6.21: Variance Model, Quarterly Data (α=1)

	Capital Goods	Consumer Goods	Industrials (excl. Oil)	Industrials (incl. Oil)	Financial Services	Other Sectors
Non-linear Wald (p-value)	138.292 (0.000)	129.678 (0.000)	142.099 (0.000)	199.757 (0.000)	359.144 (0.000)	138.140 (0.000)
Linear Wald (p-value)	37.608 (0.000)	37.718 (0.154)	37.048 (0.000)	60.780 (0.000)	57.066 (0.000)	33.794 (0.001)
$\sigma(\delta')/\sigma(\delta)$ (s.e.)	0.403 (0.224)	0.311 (0.154)	0.353 (0.196)	0.273 (0.208)	0.325 (0.167)	0.274 (0.155)
$Corr(\delta,\delta')$ (s.e.)	0.432 (0.429)	0.119 (0.608)	0.267 (0.520)	-0.150 (0.552)	-0.299 (0.839)	0.088 (0.685)
$\sigma(\delta-\delta_d')/\sigma(\delta_r')$ (s.e.)	8.345 (5.328)	7.173 (3.951)	5.663 (2.780)	6.411 (3.155)	17.086 (7.734)	3.806 (2.112)
$Corr(\delta-\delta_d',\delta_r')$ (s.e.)	0.756 (0.326)	0.787 (0.281)	0.576 (0.369)	0.561 (0.381)	0.107 (1.329)	0.526 (0.385)
$\sigma(\xi')/\sigma(\xi)$ (s.e.)	0.429 (0.140)	0.332 (0.112)	0.366 (0.138)	0.316 (0.098)	0.493 (0.107)	0.356 (0.100)
$Corr(\xi,\xi')$ (s.e.)	0.531 (0.289)	0.373 (0.455)	0.417 (0.382)	0.343 (0.477)	0.128 (0.349)	0.481 (0.313)

Table 6.22: Covariance Model, Quarterly Data (α=2)

	Capital Goods	Consumer Goods	Industrials (excl. Oil)	Industrials (incl. Oil)	Financial Services	Other Sectors
Non-linear Wald (p-value)	87.825 (0.000)	74.057 (0.000)	95.943 (0.000)	98.573 (0.000)	195.976 (0.000)	108.097 (0.000)
Linear Wald (p-value)	29.828 (0.003)	27.830 (0.006)	31.159 (0.002)	49.950 (0.000)	43.871 (0.000)	24.335 (0.018)
$\sigma(\delta')/\sigma(\delta)$ (s.e.)	0.466 (0.257)	0.363 (0.207)	0.404 (0.235)	0.314 (0.244)	0.396 (0.226)	0.282 (0.178)
Corr(δ,δ') (s.e.)	0.562 (0.439)	0.326 (0.696)	0.451 (0.554)	0.112 (0.748)	-0.272 (0.873)	0.291 (0.762)
$\sigma(\delta-\delta_d')/\sigma(\delta_r')$ (s.e.)	5.568 (4.245)	5.789 (4.465)	4.892 (3.234)	5.613 (3.734)	7.837 (3.981)	4.439 (2.558)
Corr($\delta-\delta_d',\delta_r'$) (s.e.)	0.912 (0.125)	0.885 (0.191)	0.800 (0.305)	0.763 (0.296)	-0.186 (1.013)	0.486 (0.554)
$\sigma(\xi')/\sigma(\xi)$ (s.e.)	0.498 (0.174)	0.428 (0.172)	0.449 (0.170)	0.381 (0.113)	0.537 (0.148)	0.385 (0.141)
Corr(ξ,ξ') (s.e.)	0.643 (0.266)	0.569 (0.426)	0.583 (0.352)	0.509 (0.515)	0.216 (0.393)	0.642 (0.298)

Table 6.23: Variance Model, Quarterly Data (α=2)

	Capital Goods	Consumer Goods	Industrials (excl. Oil)	Industrials (incl. Oil)	Financial Services	Other Sectors
Non-linear Wald (p-value)	69.941 (0.000)	60.827 (0.000)	85.272 (0.000)	103.007 (0.000)	197.642 (0.000)	145.322 (0.000)
Linear Wald (p-value)	29.627 (0.003)	30.741 (0.002)	31.074 (0.002)	50.039 (0.000)	49.607 (0.000)	21.931 (0.038)
$\sigma(\delta')/\sigma(\delta)$ (s.e.)	0.505 (0.261)	0.414 (0.217)	0.431 (0.239)	0.295 (0.231)	0.477 (0.290)	0.231 (0.147)
Corr(δ,δ') (s.e.)	0.628 (0.413)	0.461 (0.634)	0.506 (0.520)	0.138 (0.771)	-0.368 (0.758)	0.328 (0.775)
$\sigma(\delta-\delta_d')/\sigma(\delta_r')$ (s.e.)	4.422 (3.278)	4.239 (2.966)	4.401 (2.920)	5.208 (3.250)	4.826 (2.440)	3.235 (1.987)
Corr($\delta-\delta_d',\delta_r'$) (s.e.)	0.929 (0.090)	0.938 (0.101)	0.871 (0.197)	0.764 (0.276)	-0.459 (0.632)	0.617 (0.400)
$\sigma(\xi')/\sigma(\xi)$ (s.e.)	0.533 (0.180)	0.491 (0.187)	0.473 (0.173)	0.370 (0.276)	0.590 (0.169)	0.365 (0.106)
Corr(ξ,ξ') (s.e.)	0.685 (0.253)	0.649 (0.363)	0.624 (0.331)	0.537 (0.499)	0.129 (0.426)	0.668 (0.252)

Table 6.24: Covariance Model, Quarterly Data (α=4)

	Capital Goods	Consumer Goods	Industrials (excl. Oil)	Industrials (incl. Oil)	Financial Services	Other Sectors
Non-linear Wald (p-value)	32.918 (0.001)	29.927 (0.003)	37.626 (0.000)	39.395 (0.000)	64.663 (0.000)	46.806 (0.000)
Linear Wald (p-value)	14.954 (0.244)	28.095 (0.005)	21.990 (0.038)	35.446 (0.000)	49.742 (0.000)	13.113 (0.361)
$\sigma(\delta')/\sigma(\delta)$ (s.e.)	0.647 (0.359)	0.563 (0.360)	0.543 (0.350)	0.408 (0.284)	0.684 (0.418)	0.273 (0.230)
Corr(δ,δ') (s.e.)	0.751 (0.400)	0.598 (0.647)	0.695 (0.523)	0.420 (0.906)	-0.358 (0.701)	0.554 (0.938)
$\sigma(\delta-\delta_d')/\sigma(\delta_r')$ (s.e.)	2.510 (2.044)	2.656 (2.1863)	2.666 (2.189)	3.395 (2.909)	2.583 (1.334)	3.049 (2.187)
Corr($\delta-\delta_d',\delta_r'$) (s.e.)	0.893 (0.224)	0.830 (0.353)	0.884 (0.178)	0.822 (0.285)	-0.477 (0.611)	0.604 (0.648)
$\sigma(\xi')/\sigma(\xi)$ (s.e.)	0.722 (0.285)	0.764 (0.354)	0.705 (0.297)	0.583 (0.211)	0.754 (0.277)	0.473 (0.205)
Corr(ξ,ξ') (s.e.)	0.762 (0.200)	0.708 (0.299)	0.725 (0.244)	0.635 (0.417)	0.199 (0.447)	0.804 (0.222)

Table 6.25: Variance Model, Quarterly Data (α=4)

	Capital Goods	Consumer Goods	Industrials (excl. Oil)	Industrials (incl. Oil)	Financial Services	Other Sectors
Non-linear Wald (p-value)	28.485 (0.005)	26.790 (0.008)	33.681 (0.001)	42.666 (0.000)	57.946 (0.000)	55.858 (0.000)
Linear Wald (p-value)	17.443 (0.134)	25.617 (0.012)	22.263 (0.035)	33.257 (0.001)	56.230 (0.000)	11.918 (0.452)
$\sigma(\delta')/\sigma(\delta)$ (s.e.)	0.736 (0.377)	0.704 (0.419)	0.604 (0.357)	0.377 (0.257)	0.894 (0.569)	0.238 (0.230)
Corr(δ,δ') (s.e.)	0.808 (0.345)	0.724 (0.507)	0.754 (0.457)	0.495 (0.906)	-0.455 (0.569)	0.757 (0.541)
$\sigma(\delta-\delta_d')/\sigma(\delta_r')$ (s.e.)	1.991 (1.522)	1.885 (1.405)	2.297 (1.808)	3.166 (2.598)	1.711 (0.829)	2.116 (1.455)
Corr($\delta-\delta_d',\delta_r'$) (s.e.)	0.895 (0.224)	0.861 (0.302)	0.916 (0.129)	0.847 (0.229)	-0.594 (0.434)	0.703 (0.409)
$\sigma(\xi')/\sigma(\xi)$ (s.e.)	0.805 (0.318)	0.926 (0.408)	0.762 (0.309)	0.559 (0.209)	0.914 (0.350)	0.459 (0.178)
Corr(ξ,ξ') (s.e.)	0.800 (0.176)	0.758 (0.242)	0.766 (0.217)	0.679 (0.377)	0.083 (0.466)	0.870 (0.131)

Table 6.26: Covariance Model, Quarterly Data (α=5)

	Capital Goods	Consumer Goods	Industrials (excl. Oil)	Industrials (incl. Oil)	Financial Services	Other Sectors
Non-linear Wald (p-value)	24.551 (0.017)	22.597 (0.031)	26.777 (0.008)	29.584 (0.003)	44.509 (0.000)	33.439 (0.001)
Linear Wald (p-value)	14.383 (0.277)	25.203 (0.014)	22.297 (0.034)	30.067 (0.003)	42.901 (0.000)	13.765 (0.316)
$\sigma(\delta')/\sigma(\delta)$ (s.e.)	0.747 (0.424)	0.679 (0.445)	0.629 (0.417)	0.473 (0.312)	0.842 (0.516)	0.291 (0.266)
Corr(δ,δ') (s.e.)	0.799 (0.371)	0.658 (0.605)	0.755 (0.478)	0.505 (0.884)	-0.381 (0.649)	0.627 (0.933)
$\sigma(\delta-\delta_d')/\sigma(\delta_r')$ (s.e.)	1.947 (1.589)	2.040 (1.674)	2.113 (1.775)	2.763 (2.428)	1.894 (0.966)	2.578 (1.963)
Corr($\delta-\delta_d'$,δ_r') (s.e.)	0.878 (0.268)	0.800 (0.413)	0.878 (0.201)	0.810 (0.326)	-0.516 (0.552)	0.627 (0.674)
$\sigma(\xi')/\sigma(\xi)$ (s.e.)	0.850 (0.353)	0.952 (0.451)	0.854 (0.374)	0.706 (0.277)	0.890 (0.350)	0.544 (0.523)
Corr(ξ,ξ') (s.e.)	0.782 (0.177)	0.723 (0.270)	0.744 (0.213)	0.650 (0.381)	0.185 (0.454)	0.819 (0.198)

Table 6.27: Variance Model, Quarterly Data (α=5)

	Capital Goods	Consumer Goods	Industrials (excl. Oil)	Industrials (incl. Oil)	Financial Services	Other Sectors
Non-linear Wald (p-value)	21.853 (0.039)	21.750 (0.040)	24.478 (0.018)	32.299 (0.001)	38.698 (0.000)	35.441 (0.000)
Linear Wald (p-value)	17.999 (0.116)	20.707 (0.533)	22.351 (0.034)	27.656 (0.006)	49.670 (0.000)	12.517 (0.405)
$\sigma(\delta')/\sigma(\delta)$ (s.e.)	0.864 (0.455)	0.869 (0.533)	0.708 (0.431)	0.439 (0.287)	1.118 (0.708)	0.300 (0.306)
Corr(δ,δ') (s.e.)	0.849 (0.308)	0.772 (0.453)	0.811 (0.402)	0.589 (0.842)	-0.474 (0.525)	0.775 (0.507)
$\sigma(\delta-\delta_d')/\sigma(\delta_r')$ (s.e.)	1.550 (1.187)	1.457 (1.092)	1.816 (1.455)	2.596 (2.225)	1.278 (0.611)	1.783 (1.265)
Corr($\delta-\delta_d'$,δ_r') (s.e.)	0.882 (0.258)	0.839 (0.340)	0.907 (0.169)	0.847 (0.237)	-0.613 (0.405)	0.723 (0.411)
$\sigma(\xi')/\sigma(\xi)$ (s.e.)	0.958 (0.401)	1.163 (0.525)	0.927 (0.393)	0.677 (0.278)	1.104 (0.446)	0.540 (0.234)
Corr(ξ,ξ') (s.e.)	0.819 (0.150)	0.767 (0.219)	0.785 (0.184)	0.694 (0.337)	0.066 (0.467)	0.883 (0.117)

that $\delta_t=\delta_t'$ is always firmly rejected. Comparing the behaviour of δ_t and δ_t' directly, when $\alpha=5$ the findings are reasonably supportive, except for Financial Services (for which the correlation is negative) and Other Sectors (for which the theoretical log dividend-price ratio is much less variable than the actual).

Finally, it is worth noting that, for both monthly and quarterly data, the own-variance model usually performs marginally better, and certainly no worse, than the covariance model.

6.9 Conclusions

If one accepts the results from the linear and non-linear Wald statistics then most of the models tend to reject the EMH for $\alpha \leq 3$. However, Campbell and Shiller (1987) noted that this statistical rejection may not be *economically* important, if, say, δ_t and δ_t' move closely together. This is because, although small deviations of the VAR parameters from their theoretical values imply some predictability in one-period, and hence multi-period, returns, such predictability might not lead to large deviations of prices from their theoretical values. In terms of the behaviour of δ_t and δ_t', and of one-period returns ξ_t and ξ_t', the sectoral variance model performs marginally better than the covariance model. The latter would imply a poor performance for the market model, since if $E_t h_{it+1}=\alpha E_t V_{it+1}$[11], then the market return, which is a weighted sum of the returns on the s industrial portfolios, is given by

$$E_t h_{mt+1} = \sum_{i=1}^{s} w_i E_t h_{it+1} = \alpha \sum_{i=1}^{s} w_i E_t V_{it+1}$$

whereas the CAPM implies that the expected market return is given by

$$E_t h_{mt+1} = \alpha E_t V_{mt+1} = \alpha\left(\sum_{i=1}^{s} w_i^2 E_t V_{it+1} + \sum_{i=1}^{s}\sum_{j=1}^{s} w_i w_j C_{ij}\right)$$

If covariances do not play a part in determining $E_t h_{it+1}$ then the market model is misspecified, since it depends on V_{mt+1} which contains covariance terms. The evidence is consistent with this view, since for the market model the correlation

11 I assume a zero risk-free rate here for ease of exposition.

between δ_t and δ_t' is either negative or zero, a feature found for only one of the sectors (Financial Services). The methodology adopted here provides a framework within which the issues of market efficiency and market segmentation may be addressed. However, given the large standard errors on many of the statistics of interest, my conclusions must be considered extremely tentative.

Chapter 7: Discussion and Conclusions

7.1 Summary of Findings

This thesis is concerned with detailing the salient features of UK stock price behaviour. The Rational Valuation Formula, which states that the current price of a stock equals the expected discounted present value of future dividends, provides the theoretical basis for two distinct areas of research. In Chapters 3 and 4, the emphasis is on analysing some of the most important influences on stock price movements. In Chapter 3, I study the relative contributions of revisions to expected dividends and revisions to expected returns to unexpected changes in stock prices. In Chapter 4, the analysis turns to the question of which financial and macroeconomic news events appear to induce expectations of future dividends, interest rates and risk premia to be revised. Chapters 5 and 6 are concerned with testing the implications of particular models for investors' equilibrium expected returns. I attempt to determine empirically the extent to which observed movements in stock prices coincide with the predictions of various theories of asset pricing.

The Efficient Markets Hypothesis is basically the application of the theory of competitive markets to financial markets. Whatever the finer details of the definition of stock market efficiency, the condition that investors must not be able to obtain returns in excess of expected returns is paramount. Thus, stock returns in excess of equilibrium expected returns must not be predictable. The stickier issue is the question of what determines investors' required returns, and how may they be modelled empirically. In much of the early stock market literature, the assumption was made that expected returns were constant. In this case, any predictability in returns could be taken as evidence against market efficiency. However, there also exist models for required returns which imply that they vary through time in a systematic fashion. In Chapters 3 and 4, it is presumed that the stock market is, in fact, efficient, so that any predictability in returns is taken as systematic movement in required returns. In most cases, significant return predictability is found. The question then is whether this predictability is consistent with any existing theories on the time-variation in equilibrium expected returns. The link between such predictability and the formal implications of specific economic models for expected returns is addressed in Chapter 5 and 6.

All four chapters share a common methodological feature: the Campbell-Shiller dividend-price ratio model is combined with VAR forecasting methods to produce the relevant statistics. Previous studies have used return autocorrelations and single-equation return regressions to study the predictability of stock returns, and variance bounds tests to address the issue of excess volatility in stock prices. Also, many researchers have found it necessary to assume either that expected returns are constant or that dividends grow at a constant rate. The Campbell-Shiller dividend-price ratio model provides an alternative methodology. In return for a degree of approximation error, this model allows one to study various aspects of stock price movements without assuming the constancy of either dividend growth or expected returns. In addition, because the dividend-price ratio model involves only stationary variables, statistical analysis of various asset pricing models can be conducted using traditional tools of statistical inference.

One version of the Campbell-Shiller model states that the current unexpected real return on a stock can be separated into two components: revisions to expectations of future real dividend growth (which have a positive effect on stock prices), and revisions to expectations of future real returns. In Chapter 3, stock returns are allowed to depend upon a number of financial variables such as the dividend yield and the gilt-equity yield ratio. A VAR is used to provide forecasts of these variables, and so predict stock returns into the distant future. These forecasts provide that portion of unexpected returns which is down to changes in expected real returns, whilst revisions to expected future dividends are initially taken as the residual. The variance of unexpected returns is then apportioned into the variance of revisions to expectations of future dividends, the variance of revisions to expectations of future expected returns, and the covariance between the two. Using annual observations of an aggregate UK stock market index over the period 1918 to 1994, I find that the portion of the unexpected real return variance accounted for by changes in expected future returns outweighs that of changes in expected dividends by a ratio of 3:1. The covariance between revisions in expectations of these two fundamentals is found to be negligible. I then look at the effects of including a measure of return volatility in the VAR. I find that the dividend-price ratio Granger-causes volatility, and there is some evidence that volatility is "persistent", in that a current shock to volatility has a protracted effect on expectations of future volatility. However, the inclusion of volatility does not have any major effect on the results of the variance decomposition. Next, I apply the variance decomposition to excess stock returns, which introduces changes in expectations of future real interest rates in to the analysis. The covariance between changes in expectations of future dividends and future real interest rates is found to be positive. Since the two elements have opposite effects on returns, with dividends having a positive effect whilst real interest rates have a negative effect, they tend to offset each other. Consequently, it is changes in expectations of future excess returns that is found to have the most important impact on current unexpected returns.

There are four main conclusions that can be drawn from the analysis in Chapter 3. First, since returns are found to be predictable, the models of constant expected real and excess returns are not supported. Second, assuming that this predictability is rational, so that the RVF continues to hold but with a time-varying discount factor, changes in the latter have the major influence on changes in stock prices. Models which allow discount rates to vary but assume dividend growth to be constant (such as those used by Poterba and Summers 1986, and Chou 1988) may be more successful in capturing stock price movements than constant returns models. However, it should not be forgotten that although the contribution of changes in expected dividends was small, it was found to be statistically significant. Third, the finding that volatility is persistent raises the possibility that it is changing expectations of volatility which cause the large movements in expected future returns. Finally, given that the dividend-price ratio has been shown to be a significant predictor of future returns, the evidence here that the dividend-price ratio has predictive power over volatility raises that possibility that a more formal specification of expected returns, which involves expected volatility, might have some success in explaining movements in annual UK stock prices.

Chapter 4 is a first attempt at disentangling the effects that news about macroeconomic factors have on investors' required returns in the stock market. According to the Linear Factor Model (of which the APT and the CAPM are special cases), the unexpected return on any asset is a linear function of innovations in risk factors. These factors are sources of non-diversifiable risk which impact upon investors' required returns. The coefficients which relate each individual factor innovation to the unexpected stock return are known as the factor "betas", and are scaled covariances between factor innovations and unexpected returns. According to the RVF, each factor can affect returns only be affecting investors' expectations of future dividends, future interest rates, or future risk premia. In Chapter 4, I take a number of macroeconomic and financial factors, and attempt to calculate how each of them affects expectations of these fundamental components. The methodology is straightforward. A VAR is used to decompose the factors into expected and unexpected components. The covariances between the latter and revisions to expectations of future dividends, future real interest rates and future excess returns are then calculated directly, and scaled to make them into factor betas.

The decomposition is applied to monthly observations of 27 industry-based stock portfolios over the period 1970 to 1993. When the return on the market portfolio is taken as a risk factor (as implied by the CAPM), I find that for most sectors the positive market beta is due primarily to the fact that revisions to market returns coincide with revisions to future required sector portfolio returns. Many of the betas between portfolio dividend expectations and the market return are not statistically significant, and are, in any case, much smaller than the return betas. I also use the beta decomposition to test formally the implications of the CAPM for the behaviour of the portfolio returns. Whilst some statistical support is found for the restrictions placed on the LFM by the CAPM, informal analysis of the implications of the CAPM for the cross-sectional behaviour of portfolio returns is not supportive of the model.

The dividend-price ratio betas are all found to be significantly negative, so that a positive shock to the dividend-price ratio coincides with a fall in the current excess return. This is due, in the main, to the large positive future excess return component. The latter arises because higher future expected excess returns, through the RVF, cause current stock prices to fall, and so the dividend-price ratio to rise. No consistent pattern is found in the signs of the dividend betas.

The beta decomposition allows me to study the relationship between stock returns and inflation in more detail than previous researchers. Much of the previous work in this area has looked at the relationship between the real return on stocks and the expected rate of inflation. The puzzle is that, whilst the Fisher hypothesis implies that there should be no relationship between the two, a significant negative correlation has been consistently observed. I demonstrate that the coefficient linking unexpected real stock returns to unexpected inflation has four separate components: a coefficient relating shocks to inflation and shocks to the current real interest rate; a coefficient relating shocks to inflation with revisions to future expected excess stock returns; a coefficient relating shocks to inflation with shocks to expected future dividends; and a coefficient relating shocks to inflation with expectations of future real interest rates. Moreover these effects are i) potentially offsetting and ii) likely to vary across industrial sectors, so that there is a great deal of scope for the net relationship between inflation and real portfolio returns to vary widely. This decomposition allows the source of any such variation to be identified more accurately than in previous analyses.

For the market portfolio, I find all four components to be negative. The negative relation between shocks to inflation and revisions to expectations of future dividends has been attributed to the fact that inflation, through the demand and/or supply of money, is negatively related to expected future real activity, which in turn affects dividend expectations (the "proxy hypothesis"). However, I am unable to explain the negative link between inflation shocks and revisions to future expected excess returns. If anything, given the proxy hypothesis, the relationship should be positive. My estimates result in a positive relation between shocks to inflation and unexpected current excess returns, but a negative relation between shocks to inflation and unexpected real returns. The difference between the two is accounted for by the strong negative relation between innovations in inflation and innovations in the current real rate of interest. With regards to the cross-sectional variation of inflation betas, I do not find support for the hypothesis of Boudoukh, Richardson and Whitelaw (1994) that such differences are accounted for by differences in the degree of cyclicality of an industry's output.

My other findings are as follows. Innovations in the real interest rate have a negative impact on current unexpected returns, as changes in real interest rates lead expectations of all future real interest rates to be revised, resulting in a negative effect on current stock prices which is reinforced by higher future excess returns. Similarly, a higher real exchange rate has its major impact though lower expected future real interest rates. Finally, I find that since shocks to real output growth do not appear to affect expectations of fundamentals, they have little or no effect on most asset returns.

The main benefit derived from this beta decomposition is a greater understanding of the net influences of systematic risk factors observed in single-equation return regressions. One might well have prior views on how a particular factor might impact upon expected dividends *or* real interest rates *or* future excess returns. For example, one might expect shocks to the exchange rate to affect dividend expectations for exporting firms, or shocks to the dividend-price ratio (which may proxy for expected volatility) to affect expected returns on all assets. However, only by combining the three can one begin to understand the net effect of factors on asset prices. The case of the inflation beta serves as a useful example. The Fisher hypothesis means that one does not expect investors to revise their required excess returns in direct response to shocks to inflation. The proxy hypothesis points out that the two may be correlated via a mutual relation with expected future output. However, this line of argument only explains why inflation may be found to be related to future dividends. The fact that inflation and stock returns are found to be related might also be taken to imply a link (direct or indirect) between inflation and future required returns, and my analysis suggests that such a link exists. As the focus hitherto has been solely on dividends, there exists no explanation for the correlation between expected future returns and inflation. Moreover, it appears that the latter relationship has as much impact on the cross-sectional pattern of the correlation between stock returns and inflation as does the degree of cyclicality in dividends (which is the only existing explanation of this pattern).

Chapter 5 presents formal tests of three popular asset pricing models: constant real returns, constant excess returns and the CAPM. The Campbell-Shiller dividend-price ratio model is found to be particularly useful as it clarified the role of the dividend-price ratio model in asset pricing. According to this model, the dividend-price ratio is an optimal forecast of future real dividend growth and future expected real returns. Consequently, if we combine optimal forecasts of dividends and real returns, the resultant series should behave just like the observed dividend-price ratio. The question is whether economic theory points towards an empirical proxy for expected returns which will satisfy this consistency condition. If we take first the constant expected real returns model, all of the movement in the dividend-price ratio should be accounted for by changing forecasts of future real dividend growth. The tests are applied to the same data employed in Chapter 3. Using a VAR model to obtain forecasts of real dividend growth, I find that a large portion of movements in the dividend-price ratio remain unexplained. This is a similar finding to those which led to Shiller's (1981) claim that stock prices are too volatile to be accounted for by changes in expected dividends. The difference is that the VAR methodology is robust to non-stationary real stock prices and dividends. I next tested the hypothesis that expected excess returns are constant. Although, in conformity with the Campbell-Shiller model, the dividend-price ratio did have some power to predict real interest rates, more formal tests of the constant excess returns model find this to be seriously deficient as a model of expected returns. In contrast, as suggested by the analysis in Chapter 3, not only does the dividend-price ratio forecast volatility, but it does so in a way which is consonant with the predictions of the CAPM. That is, when time-variation in expected volatility was combined with forecasts of real dividend growth, the resultant series was statistically indistinguishable from the actual dividend-price ratio. It has often been observed that the dividend-price ratio forecasts stock returns. The results in Chapter 5 imply that this ability to forecast returns is due to the fact that the dividend-price ratio tracks expected volatility, and it is the latter that causes investors to adjust their required returns on risky assets.

Finally, Chapter 6 studied the cross-sectional implications of the CAPM for monthly and quarterly UK data. Over the period 1965-1993, and using the aggregate stock market return variance as "market risk", the implications of the CAPM were easily rejected. The analysis of the market sub-sectors highlighted some cross-sectional differences in efficiency, but support for the covariance model was by no means overwhelming. However, in an apparent contradiction to the implications of the CAPM, the omission of the covariance terms, so that the expected return in each sub-sector depended only on the expected variance of the return in that sector, had, if anything, a positive effect on the results. This might be indicative of a degree of market segmentation, where investors are more skilled at eliminating idiosyncratic risk within industrial sectors than across the market as a whole.

The contradictory evidence presented in Chapters 5 and 6 with regards to the efficacy of the CAPM for explaining stock price behaviour could be down to two effects. First, because of data availability, the tests were performed over very different time periods. Moreover, there were too few annual observations over the shorter period for a rigorous sensitivity analysis of sample horizon to be made. However, I believe that the more important issue is to do with data frequency. In broad terms, the CAPM received more support the lower the frequency of the data. Now it is well documented that the dividend-price ratio forecasts a greater portion of the return variance the longer is the return horizon. It is quite possible, therefore, that the dividend-price ratio is able to pick up long-term swings in expected volatility but is not such a good proxy for short-term volatility. Indeed, the dividend-price ratio is an exceptionally smooth series, with much more important low-frequency components than high-frequency components. In contrast, there are certain patterns in volatility (such as volatility clustering) which tend to show up more in high-frequency data than low frequency data (which is why ARCH-type models are so popular for modelling high-frequency returns data). It may well be, therefore, that it is the empirical proxy for volatility which is at fault for the tests of higher-frequency data, rather than the hypothesis that expected returns depend on expected volatility.

Finally, one important lesson which can be learnt from this thesis is the importance of studying a wide range of metrics when testing a hypothesis. For example the linear Wald test suggested that the quarterly returns data for Other Sectors in Chapter 6 were unpredictable. If I had concentrated solely on return predictability to make inference about market efficiency, I might have concluded that the CAPM was a suitable model for this portfolio. However, a comparison of the actual and theoretical log dividend-price ratios clearly indicated that such a conclusion would have been erroneous. Clearly, conclusions drawn from a large amount of information are much more persuasive than studies that base inference on only one or two test statistics.

7.2 Further Research

Most of the research into stock price behaviour has been at the micro level. However, in recent years, the importance of studying risk pricing for macroeconomic purposes has become more apparent. In particular, the failure of investment in the UK to pick up after the recent recession, despite historically low borrowing rates, has been cited as one of the most worrying aspects of the "recovery". The stock market provides a forum in which some of the most basic aspects of risk pricing can be studied, since the data are usually of better quality than other macroeconomic indicators, and there is a relatively high degree of homogeneity across alternative investments as compared with investments in physical assets.

There are quite a number of areas for further research suggested by this thesis. First of all, it would be useful if ARCH-type models for return volatility could be combined with the RVF to provide a more formal test of the implications of the CAPM than do the standard tests of the significance of the conditional variance in the conditional mean equation. As noted above, ARCH models are more successful at modelling conditional volatility at high frequencies than VAR equations.

Although not addressed directly in this thesis, the relationship between stock returns and consumption is clearly not understood. All of the studies cited here find that "consumption has the wrong sign", and it is evident that the relationship between returns on financial assets and particular measures of aggregate consumption needs to be refined. I would speculate that this is likely to depend on the proportion of national income generated by quoted firms.

I think that the most obvious manifestation of our ignorance about investor behaviour is what Mehra and Prescott (1985) termed the "equity premium puzzle". The basic problem is that ex post excess stock returns seem implausibly high, given the perceived level of uncertainty. This may be taken either as evidence that stock prices are not closely related to fundamentals, which may occur if investors do not assess rationally the risks and opportunity costs of investments, or that our current understanding of rational risk assessment is deficient. Dixit and Pindyck (1994) have noted that returns to marginal physical investments appear to be excessively high, based on traditional net present value calculations. They developed a theory in which the irreversibility of an investment gives extra value to the option of postponing the investment, when postponement may lead to the revelation of more information regarding the likely payoff. Of course, the traditional view of the stock market is that the liquidity of investments is high, so that such an explanation could not be applied. However, transactions charges mean that reversing a sale or purchase is not costless, and this might make it worthwhile for investors to postpone trading on information until the quality of that information is more firmly established. Also, liquidity is not so high for some thinly-traded stocks, and institutional factors may impose costs on fund managers who constantly adjust their positions.

Although I have concentrated on models of rational investor behaviour, I believe that recent developments in the identification and modelling of irrational behaviour may prove fruitful. For example, the well-known model of De Long, Schleifer, Summers and Waldmann (1990), which studied the interaction of so-called "noise traders" with fundamentalist traders, demonstrated that even a small amount of irrational trading can cause prices to diverge from fundamentals as the rational investors have to allow for the greater uncertainty surrounding future prices. Kirman (1993) argued that herding behaviour can explain the waves of volatility and tranquillity observed in financial markets, whilst Shiller (1984) and Cutler, Poterba and Summers (1990) developed simple alternative models of stock price determination when not all traders are assumed to be fully rational. Of course, it can be argued that any apparent irrationality on the part of economic agents is more to do with our definitions of rationality than genuinely inexplicable behaviour. For example, it is often noted that agents base decisions on rules of thumb which may consistently lead to sub-optimal outcomes. However, if the costs of a) assessing the payoffs to such decisions, and b) developing more accurate decision-making processes, are prohibitive, the use of such informal decision rules fits in well with broader concepts of rational economic behaviour. The problem is that, just as it is difficult to believe that individuals, as products of natural social processes, can be ascribed with the degree of rationality required by standard economic models, it seems implausible that people behave in a way which does not maximise their personal well-being. However, I do believe that the highly rigorous scientific analyses which have been a notable characteristic of the debate on stock market efficiency will substantially improve our understanding of these fundamental issues.

References

Akgiray, V. (1989). "Conditional heteroscedasticity in time series of stock returns: Evidence and forecasts." *Journal of Business* 62, 55-80.

Attanasio, O. and Wadhwani, S. (1990). "Does the CAPM explain why the dividend yield helps predict returns?" LSE Financial Markets Group Discussion Paper, no. 049.

Bachelier, L. (1900). "Theory of speculation." In Cootner (1964), 17-78.

Barr, D.G. and Pesaran, B. (1994). "An assessment of the relative importance of real interest rates, inflation and term premia in determining the returns on real and nominal UK bonds." Mimeo, Brunel University, Uxbridge, Middlesex, UK.

Beenstock, M. and Chan, K. (1988). "Economic forces in the London stock market." *Oxford Bulletin of Economics and Statistics* 50, 27-39.

Bing, R. (1971). "Survey of practitioners' stock evaluation methods." *Financial Anlysts' Journal* 27, 55-69.

Black, F. (1976). "Studies of stock price volatility changes." *Proceedings of the American Statistical Association*, Business and Economic Statistics Section, 177-186.

Blanchard, O.J. and Fischer, S. (1989). *Lectures on Macroeconomics.* MIT Press, Cambridge, Massachusetts.

Bodie, Z. (1976). "Common stock as a hedge against inflation." *Journal of Finance* 31, 459-470.

Bollerslev, T., Engle, R.F. and Wooldridge, J.M. (1988). "A Capital Asset Pricing Model with time-varying covariances." *Journal of Political Economy* 96, 116-131.

Boudoukh, J., Richardson, R. and Whitelaw, R.F. (1994). "Industry returns and the Fisher effect." *Journal of Finance* 49, 1595-1615.

Box, G.E.P. and Jenkins, G.M. (1970). *Time-Series Analysis: Forecasting and Control.* Holden-Day, San Fransisco.

Bulkley, G. and Tonks, I. (1989). "Are UK stock prices excessively volatile? Trading rules and variance bounds tests." *Economic Journal* 99, 1083-1098.

Burmeister, E. and McElroy, M.B. (1988). "Joint estimation of factor sensitivities and risk premia for the Arbitrage Pricing Theory." *Journal of Finance* 43, 721-735.

Campbell, J.Y. (1991). "A variance decomposition for stock returns." *Economic Journal* 101, 157-179.

Campbell, J.Y. (1987). "Stock returns and the term structure." *Journal of Finance*, vol. 18, pp. 373-399.

Campbell, J.Y. and Ammer, J. (1993). "What moves the stock and bond markets? A variance decomposition for long-term asset returns." *Journal of Finance* 48, 3-37.

Campbell, J.Y. and Cochrane, J.H. (1994). "By force of habit: A consumption-based explanation of aggregate stock market behaviour." Mimeo, Department of Economics, Harvard University, USA.

Campbell, J.Y. and Mei, J. (1993). "Where do betas come from? Asset price dynamics and the sources of systematic risk." *Review of Financial Studies* 6, 567-592.

Campbell, J.Y and Shiller, R.J. (1989). "The dividend-price ratio and expectations of future dividends and discount factors." *Review of Financial Studies* 1, 195-228.

Campbell, J.Y. and Shiller, R.J. (1988). "Stock prices, earnings and expected dividends." *Journal of Finance* 43, 661-676.

Campbell, J.Y. and Shiller, R.J. (1987). "Cointegration and tests of present value models." *Journal of Political Economy* 95, 1062-1088.

Chen, N. (1991). "Financial investment opportunities and the macroeconomy." *Journal of Finance* 46, 529-554.

Chen, N., Roll, R. and Ross S.A. (1986). "Economic forces and the stock market." *Journal of Business* 59, 383-403.

Chou, R.Y. (1988). "Volatility persistence and stock valuations: some empirical evidence using GARCH." *Journal of Applied Econometrics* 3, 279-294.

Clare, A. and Thomas, S.H. (1994). "Macroeconomic factors, the APT and the UK stock market." *Journal of Business, Finance and Accounting* 21, 309-330.

Clare, A., O'Brien, R., Thomas, S. and Wickens, M. (1993). "Macroeconomic shocks and the domestic CAPM: Evidence from the UK stock market." Discussion Paper 93-02, Department of Economics, Brunel University, UK.

Clare, A.D., Thomas, S.H. and Wickens, M.R. (1994). "Is the gilt-equity yield ratio useful for predicting UK stock returns?" *Economic Journal* 104, 303-316.

Cochrane, J.H. (1991). "Production-based asset pricing and the link between stock returns and evonomic fluctuations." *Review of Financial Studies*, 195-228.

Conrad, J. and Kaul, G. (1988). "Time-variation in expected returns." *Journal of Business* 61, 409-425

Cooley, T.F. and LeRoy, S.F. (1985). "Atheoretical macroeconomics: a critique." *Journal of Monetary Economics* 16, 283-308.

Cootner, P.H., ed. (1964). *The Random Character of Stock Market Prices.* MIT Press, Cambridge, Massachusetts..

Cootner, P.H. (1962). "Stock prices: Random versus systematic changes." *Industrial Management Review* 3, 24-45.

Cowles, A. (1933). "Can stock market forecasters forecast?" *Econometrica* 53, 385-407.

Cutler, D.M., Poterba, J.M. and Summers, L.H. (1991). "Speculative dynamics." *Review of Economic Studies* 58, 529-546.

Cutler, D.M., Poterba, J.M. and Summers, L.H. (1990). "Speculative dynamics and the role of feedback traders." *American Economic Review*, Papers and Proceedings, 63-68.

Cutler, D.M., Poterba, J.M. and Summers, L.H. (1989). "What moves stock prices?" *Journal of Portfolio Management* 15, 4-12.

Davidson, R. and MacKinnon, J.G. (1993). *Estimation and Inference in Econometrics.* Oxford University Press, New York.

De Long, J.B., Schleifer, A., Summers, L.H. and Waldmann, R.J. (1990). "Noise trader risk in financial markets." Journal of Political Economy 98, 703-738.

Dixit, A.K. and Pindyck, R.S. (1994). *Investment Under Uncertainty.* Princeton University Press, New Jersey.

Durlauf, S.N. and Hall, R.E. (1989). "Measuring noise in stock prices." Working paper, Stanford University, USA.

Engle R.F., Lilien, D.M. and Robins, R.P. (1987). "Estimating time varying risk premia in the term structure: The ARCH-M model." *Econometrica* 55, 391-407.

Fama, E.F. (1990). "Stock returns, expected returns and real activity." *Journal of Finance* 45, 1089-1108.

Fama, E.F. (1981). "Stock returns, real activity, inflation and money." *American Economic Review* 71, 545-565.

Fama, E.F. (1970). "Efficient capital markets: A review of theory and empirical work." *Journal of Finance* 25, 383-417.

Fama, E.F. and French, K. (1989). "Business conditions and expected returns on stocks and bonds." *Journal of Financial Economics* 25, 23-49.

Fama, E.F. and French, K.R. (1988a). "Permanent and temporary components of stock prices." *Journal of Political Economy* 96, 246-273.

Fama, E.F. and French, K. (1988b). "Dividend yields and expected stock returns." *Journal of Financial Economics* 22, 3-25.

Fama, E.F. and Schwert, G.W. (1977). "Asset returns and inflation." *Journal of FInancial Economics* 5, 115-146.

Feldstein, M. (1980). "Inflation and the stock market." American Economic Review 70, 636-651.

Firth, M. (1979). "The relationship between stock market returns and rates of inflation." *Journal of Finance* 34, 743-749.

Fisher, R. (1966). "Some new stock market indexes." *Journal of Business* 39, 191-225.

Flavin, M.A. (1983). "Excess volatility in the financial markets: A reassessment of the empirical evidence." *Journal of Political Economy* 91, 246-273.

Francis, J.C. (1986). *Investments: Analysis and Management.* 4th edition, McGraw-Hill, New York.

French, K.R., Schwert, G.W. and Stambaugh, R.F. (1987). "Expected stock returns and volatility." *Journal of Financial Economics* 19, 3-29.

Friedman, B.M. (1979). "Optimal expectations of the extreme information assumptions of 'Rational Expectations' macromodels." *Journal of Monetary Economics* 5, 23-41.

Fuller, W. (1976). *Introduction to Statistcal Time Series.* Wiley, New York.

Geske, R. and Roll, R. (1983). "The fiscal and monetary linkage between stock returns and inflation." *Journal of Finance* 38, 1-33.

Gilles, C. and LeRoy, S.F. (1991). "Econometric aspects of the variance bounds tests: A survey." *Review of Financial Studies* 4, 753-791.

Gordon, M.J. (1962). *The Investment, Financing and Valuation of the Corporation.* Irwin, Homewood Ill.

Granger, C.W.J. and Morgenstern, O. (1963). "Spectral analysis of New York Stock Market prices." In Cootner (1964), 162-188.

Gregory, A.W. and Veall, M.R. (1985). "Formulating Wald tests of non-linear restrictions." *Econometrica* 53, 1465-1468.

Grossman, S.J. and Shiller, R.J. (1981). "The determinants of the variability of stock market prices." *American Economic Review* 71, 222-227.

Gultekin, N.B. (1983). "Stock market returns and inflation: Evidence from other countries." *Journal of Finance* 38, 49-65.

Hall, S.G., Miles, D.K. and Taylor, M.P (1990). "A CAPM with time-varying betas: some results." In *Economic Modelling at the Bank of England* (eds S.G.B. Henry and K.D. Patterson), Chapman and Hall, London.

Hall, S.G., Miles, D. and Taylor, M.P. (1989). "Modelling asset prices with time-varying risk." *The Manchester School* 57, 340-356.

Hansen, L. (1982). "Large sample properties of Generalised Method of Moments estimators." *Econometrica* 50, 1029-1054.

Holden, K. (1994). "Vector autoregression modelling and forecasting." Discussion Paper 5\94, Centre for International Banking, Economics and Finance, Liverpool Business School, Liverpool John Moores University.

Jaffe, J.F. and Mandelker, G. (1976). "The Fisher effect for risky assets: an empirical investigation." *Journal of Finance* 31, 447-458.

Kaul, G. (1987). "Stock returns and inflation. The role of the monetary sector." *Journal of Financial Economics* 18, 253-276.

Keim, D.B. and Stambaugh, R.F. (1986). "Predicting returns in the stock and bond markets." *Journal of Financial Economics* 17, 357-390.

Kendall, M.G. (1953). "The analysis of economic time series, Part 1: Prices." In Cootner (1964), 85-99.

Kirman, A. (1993). "Ants, rationality and recruitment." *Quarterly Journal of Economics*, February, 137-156.

Kleidon, A.W. (1986). "Variance bounds tests and stock price evaluation models." *Journal of Political Economy* 94, 953-1001.

Lamoureux, C.G. and Lastrapes, W.D. (1990a). "Heteroscedasticity in stock return data: Volume versus GARCH effects." *Journal of Finance* 45, 221-229.

Lamoureux, C.G. and Lastrapes, W.D. (1990b). "Persistence in variance, structural change and the GARCH model." *Journal of Business and Economic Statistics* 8, 225-234.

LeRoy, S.F. and Porter, R.D. (1981). "The present-value relation: Tests based on implied variance bounds." *Econometrica* 49, 555-574.

Litterman, R.B. (1986). "A statistical approach to economic forecasting." *Journal of Business and Economic Statistics* 4, 1-4.

Lo, A.W. and MacKinlay, A.C. (1988). "Stock market prices do not follow random walks: Evidence from a simple specification test." *Review of Financial Studies* 1, 41-66.

Lucas, R.E. (1978). "Asset prices in an exchange economy." *Econometrica* 46, 1426-1446.

Lund, J. and Engsted, T. (1993). "GMM and present value tests of the C-CAPM: Some evidence from the Danish, German, Swedish and UK stock markets." Mimeo, Department of Information Science, Aarhus School of Business, Denmark.

MacDonald, R. and Power, D. (1991). "Persistence in UK stock market returns." In *Money and Financial Markets* (ed M.P. Taylor). Basil Blackwell, Oxford.

Malkiel, B.G. (1979). "The capital formation problem in the United States." *Journal of Finance* 34, 291-306.

Mandelbrot, B. (1963). "The variation of certain speculative prices." *Journal of Buisness* 36, 394-419.

Mankiw, N.G., Romer, D. and Shapiro, M.D. (1991). "Stock market forecastability and volatility: A statistical appraisal." *Review of Economic Studies* 58, 455-477.

Mankiw, N.G., Romer, D. and Shapiro, M.D. (1985). "An unbiased reexamination of stock market volatility." *Journal of Finance* 40, 677-687.

Marsh, T.A. and Merton, R.C. (1986). "Dividend variability and variance bounds tests for the rationality of stock market prices." *American Economic Review* 76, 483-498.

Mehra, R. and Prescott, E.C. (1985). "The equity premium: A puzzle." *Journal of Monetary Economics* 15, 145-161.

Merton, R.C.(1987). "On the current state of the stock market rationality hypothesis." In S. Fischer (ed.), *Macroeconomics and Finance: Essays in Honor of Franco Modigliani.* MIT Press, Cambridge, Massachusets.

Merton, R.C. (1980). "On estimating the expected return on the market." *Journal of Financial Economics* 8, 323-361.

Merton, R.C. (1973). "An intertemporal capital asset pricing model." *Econometrica* 41, 867-887.

Miles, D. (1993). "Testing for short-termism in the UK stock market." *Economic Journal* 103, 1379-1396.

Modigliani, F. and Cohn, R.A. (1979). "Inflation, rational valuation and the market." *Financial Analysts Journal*, March/April, 24-44.

Muth, J.F. (1961). "Rational expectations and the theory of price movements." *Econometrica* 29, 315-335.

Nelson, C.R. (1976). "Inflation and rates of return on common stocks." *Journal of Finance* 31, 471-483.

Nelson, C.R. and Plosser, C.I. (1982). "Trends and random walks in macroeconomic time series: Some evidence and implications." *Journal of Monetary Economics* 10, 139-162.

Newey, W. and West, K. (1987). "A simple positive semi-definite, heteroscedasticity and autocorrelation consistent covariance matrix." *Econometrica* 55, 703-708.

Perron, P. (1988). "Trends and random walks in macroeconomic time series: Further evidence from a new approach." *Journal of Economic Dynamics and Control* 12, 297-332.

Pesaran, M.H. and Timmermann, A. (1992). "Forecasting stock returns: An examination of stock market trading in the presence of transaction costs." Mimeo, University of Cambridge.

Phillips, P.C.B. (1987). "Time series regressions with a unit root." *Econometrica* 55, 277-301.

Phillips, P.C.B. and Park, J.Y. (1988). "On the formulation of Wald tests of nonlinear restrictions." *Econometrica* 56, 1065-1083.

Phillips, P.C.B. and Perron, P. (1988). "Testing for a unit root in time series regressions." *Biometrica* 75, 335-346.

Pindyck, R.S. (1986). "Risk aversion and determinants of stock market behaviour." NBER Working Paper, no.21.

Pindyck, R.S. (1984). "Risk, inflation and the stock market." *American Economic Review* 74, 335-351.

Poon, S. and Taylor, S.J. (1991). "Macroeconomic factors and the UK stock market." *Journal of Business, Finance and Accounting* 18, 619-636.

Poterba, J.M. and Summers, L.H. (1988). "Mean reversion in stock prices: Evidence and implications." *Journal of Financial Economics* 22, 26-59.

Poterba, J.M. and Summers, L.H. (1986). "The persistence of volatility and stock market fluctuations." *American Economic Review* 76, 1142-1151.

Roll, R. (1988). "R^2" *Journal of Finance* 43, 541 566.

Ross, S.A.(1976). "The arbitrage theory of capital asset pricing." *Journal of Economic Theory* 13, 341-360.

Said, S.E. and Dickey, D.A. (1984). "Testing for unit roots in autoregressive-moving average models of unknown order." *Biometrica* 71, 599-607.

Samuelson, P.A. (1965). "Proof that properly anticipated prices fluctuate randomly." *Industrial Management Review* 6, 41-49.

Samuelson, P.A. (1973). "Proof that properly discounted present values of assets vibrate randomly." *Bell Journal of Economics* 4, 369-374.

Scott, L.O. (1991). "Financial market volatility: A survey." *IMF Staff Papers* 38, 582-625.

Scott, L.O. (1985). "The present value model of stock market prices: Regression tests and Monte Carlo results." *Review of Economics and Statistics* 57, 599-605.

Shah, M. and Wadhwani, S. (1993). "The effect of the term spread, dividend yield and real activity on stock returns: Evidence from 15 countries." LSE Financial Markets Group, Discussion Paper 98.

Shea, G.S. (1989). "Ex post rational price approximations and the empirical reliability of the present value relation." *Journal of Applied Econometrics* 4, 139-159.

Shiller, R.J. (1989). *Market Volatility.* MIT Press, Cambridge, Massachusetts.

Shiller, R.J. (1984). "Stock prices and social dynamics." *Brookings Papers on Economic Activity* 2, 457-498.

Shiller, R.J. (1981). "Do stock prices move too much to be justified by subsequent changes in dividends?" *American Economic Review* 71, 421-436.

Shiller, R.J. and Beltratti, A.E. (1992). "Stock prices and bond yields." *Journal of Monetary Economics* 30, 25-46.

Sims, C.A. (1980). "Macroeconomics and Reality." *Econometrica* 48, 1-48.

Summers, L.H. (1986). "Does the stock market rationally reflect fundamental values?" *Journal of Finance* 41, 591-601.

Summers, L.H. (1981). "Inflation, the stock market and owner-occupied housing." *American Economic Review Papers and Proceedings* 71, 429-434.

Thaler, R.H. (1994). *Quasi Rational Economics.* Russell Sage Foundation, New York.

Tiao, G.C. and Box, G.E.P. (1981). "Modelling multiple time series with applications." *Journal of the American Statistical Association* 76, 802-816.

West, K.D. (1988). "Dividend innovations and stock price volatility." *Econometrica* 56, 37-61.

White, H. (1984). *Asymptotic Theory for Econometricians.* Academic Press, Orlando, Florida, USA.

Wold, H. (1938). *A Study in the Analysis of Stationary Time Series.* Alquist and Wiksell, Upsala.

Working, H. (1934). "A random difference series for use in the analysis of time series." *Journal of the American Statistical Association* 29, 11-24.

Zellner, A. and Palm, F. (1974). "Time series analysis and simultaneous equation econometric models." *Journal of Econometrics* 2, 17-54.