Department of Finance
Working Paper Series 1998

FIN-98-073

Stock Market Risk and Return: An Equilibrium Approach

Robert F. Whitelaw

October 1997

This working paper series has been generously supported by a grant from

CDC Investment Management Corporation

Stock Market Risk and Return: An Equilibrium Approach

Robert F. Whitelaw *

NYU, Stern School of Business

October 31, 1997

*New York University, Stern School of Business. (44 West 4th Street, Suite 9-190, New York, NY 10012, (212) 998-0338, FAX: (212) 995-4233, email: rwhitela@stern.nyu.edu). Thanks to Kobi Boudoukh, Kent Daniel, Wayne Ferson, John Heaton, Anthony Lynch, Matt Richardson, and seminar participants at the University of California, Los Angeles, Duke University, the University of Minnesota, Northwestern University, the University of Southern California, the 1997 AFA meetings, and the 1996 Utah Winter Finance Conference for helpful comments.

STOCK MARKET RISK AND RETURN: AN EQUILIBRIUM APPROACH

Abstract

Recent empirical evidence suggests that expected stock returns are weakly, or even negatively, related to the volatility of stock returns at the market level, and that this relation varies substantially over time. This evidence contradicts the apparently reliable intuition that risk and return are positively related and that stock market volatility is a good proxy for risk. This paper investigates the relation between volatility and expected returns in a general equilibrium, exchange economy. A relatively simple model, estimated using aggregate consumption data, is able to duplicate the salient features of the observed expected return/volatility relation. The key features of the model are the existence of two regimes with different consumption growth processes and time-varying transition probabilities between regimes. This structure generates time-varying correlations between stock returns and the marginal rate of substitution; thus inducing variability in the short-run relation between expected returns and volatility and a weakening of the long-run relation. These results highlight the perils of relying on intuition from static models. They also have important implications for the empirical modeling of returns.

1 Introduction

Recent empirical studies (e.g., Glosten, Jagannathan, and Runkle (1993), Whitelaw (1994,1997), and Boudoukh, Richardson, and Whitelaw (1997)) document two puzzling results with regard to the intertemporal relation between equity risk and return at the market level.[1] First, they provide evidence of a weak, or even negative, relation between conditional expected returns and the conditional volatility of returns. Second, they document significant time-variation in this relation. Specifically, in a modified GARCH-M framework, Glosten, Jagannathan, and Runkle (1993) find that the estimated coefficient on volatility in the expected return regression is negative using post-WWII monthly data. In a similar dataset, when both conditional moments are estimated as functions of predetermined financial variables, Whitelaw (1994) finds that the long-run correlation between the fitted moments is negative. Moreover, the short-run correlation varies substantially from approximately -0.8 to 0.8 when measured over 17-month horizons. This evidence is supported by the significant in- and out-of-sample predictability of market Sharpe ratios in Whitelaw (1997). Finally, in a nonparametric estimation using almost two centuries of annual data, Boudoukh, Richardson, and Whitelaw (1997) find that time-variation in expected returns and the variance of returns, as functions of slope of the term structure, do not coincide.

These empirical results are especially interesting because they run counter to the strong intuition of a positive relation between volatility and expected returns at the market level that comes from such models as the dynamic CAPM (Merton (1980)). Two questions arise naturally. First, are these results consistent both with general equilibrium models and with the time series properties of variables such as consumption growth which drive equity returns in these models?[2] Second, what features are necessary to generate this counterintuitive behavior of expected returns and volatility?

This paper addresses these two questions in the context of a representative agent, exchange economy (Lucas (1978)). As such, the exercise is similar in spirit to that of Cecchetti, Lam, and Mark (1990), who attempt to duplicate the serial correlation patterns in equity returns in an equilibrium setting. Consumption growth is modeled as an autoregressive process, with two

[1] These papers extend earlier work on the subject by Campbell (1987) and French, Schwert, and Stambaugh (1987), among many others.

[2] It is known that equilibrium models can generate a wide variety of relations between the mean and volatility of returns (e.g., Abel (1988) and Backus and Gregory (1993)). The question addressed in this paper is whether the more specific intertemporal patterns documented recently are consistent with economic data.

regimes in which the parameters differ (similar to Hamilton (1989)). The probability of a regime shift is modeled as a function of the level of consumption growth (similar to Filardo (1994)). The parameters are estimated by maximum likelihood using monthly consumption data over the period 1959-1996. The stock market is modeled as a claim on aggregate consumption, and the quantities of interest are the unconditional correlation between expected equity returns and the volatility of returns and the conditional correlations between these moments of returns in various states of the world. The economy is approximated by a discrete state space economy in which the state variables can take on only a fixed number of values (see Tauchen and Hussey (1991)). The discrete state space methodology permits closed-form calculations of expected returns and conditional volatility, while generating excellent approximations to the continuous state space.

A single-regime model is not able to generate the correlation patterns between expected returns and volatility that are apparent in the data. In such a model, these variables are strongly positively correlated both unconditionally and in every state of the world. In contrast, a two-regime specification generates both a negative unconditional correlation and conditional correlations that vary widely. The key features of the specification are regime parameters that imply different means of consumption growth across the regimes and state dependent regime switching probabilities.

The major contribution of this paper is in establishing the fact that the recent empirical evidence is consistent with reasonable parameterizations of a relatively simple equilibrium model. This finding adds credibility to these empirical results, and the model provides the economic intuition behind these results. Specifically, the possibility of shifts between regimes that exhibit different consumption growth processes increases volatility while simultaneously reducing the equity risk premium in certain states of the world. The equity risk premium is a function of the correlation between equity returns and the marginal rate of substitution. However, the marginal rate of substitution depends only on next period's consumption growth, while the equity return depends on the infinite future via its dependence on the stock price next period. In states in which a regime shift is likely, this divergence of horizons weakens the link between market returns and the marginal rate of substitution. As a result, the risk premium is low but the volatility of returns is high.

The remainder of the paper is organized as follows. Section 2 develops the asset pricing framework, provides the intuition behind the risk/return relation in this setting, and describes the two-regime AR specification. A single-regime model is estimated in Section 3. The discrete approximation technique is illustrated and the results provide a benchmark for the later analysis. In Section 4, we estimate and analyze the two-regime model. The expected return and volatility patterns are presented and a sensitivity analysis is performed. Section 5 concludes.

2 The Theory

2.1 The Asset Pricing Framework

Consider a pure exchange economy with a single consumption good (Lucas (1978)). Assume further, the existence of a representative agent whose utility function exhibits constant relative risk aversion. All assets will be priced according to the first order conditions of this agent, giving the standard pricing equation

$$p_t = \beta \, \mathrm{E}_t \left[\left(\frac{c_{t+1}}{c_t} \right)^{-\alpha} (p_{t+1} + y_{t+1}) \right], \tag{1}$$

where c_t is consumption, p_t is the ex-payoff price of the asset, y_t is the asset's payoff, α is the coefficient of risk aversion, β is the time preference parameter, and $\mathrm{E}_t[\cdot]$ denotes the expectation conditional on information available at time t.

Denote the one-period, gross return on an asset by

$$r_{t+1} \equiv \frac{p_{t+1} + y_{t+1}}{p_t},$$

and the one-period, gross, riskless rate by

$$r_{ft} \equiv \left(\beta \, \mathrm{E}_t \left[\left(\frac{c_{t+1}}{c_t} \right)^{-\alpha} \right] \right)^{-1}.$$

The expected return on any asset in excess of the riskless rate is proportional to the negative of the covariance of this return with the marginal rate of substitution (MRS), i.e.,

$$\mathrm{E}_t[r_{t+1} - r_{ft}] = -r_{ft} \, \mathrm{Cov}_t[m_{t+1}, r_{t+1}], \tag{2}$$

where $m_{t+1} \equiv \beta (c_{t+1}/c_t)^{-\alpha}$ is the MRS.

In this setting it is standard to identify the stock market as the claim on the aggregate consumption stream, i.e., to equate aggregate consumption and the aggregate stock market dividend (see, for example, Mehra and Prescott (1985) and Cecchetti, Lam, and Mark (1990)). Applying equation (1), the value of equity is

$$s_t = \beta \, \mathrm{E}_t \left[\left(\frac{c_{t+1}}{c_t} \right)^{-\alpha} (s_{t+1} + c_{t+1}) \right], \tag{3}$$

where c_{t+1} is the dividend. Substituting for future equity prices, the current value can be written as the sum of discounted future dividends (consumption), i.e.,

$$s_t = \sum_{s=1}^{\infty} \beta^s \mathrm{E}_t \left[\left(\frac{c_{t+s}}{c_t} \right)^{-\alpha} c_{t+s} \right].$$

The gross, one-period, equity return is

$$\begin{aligned} r_{st+1} &= \frac{s_{t+1} + c_{t+1}}{s_t} \\ &= \left(\frac{c_{t+1}}{c_t} \right) \frac{s_{t+1}/c_{t+1} + 1}{s_t/c_t}, \end{aligned} \tag{4}$$

where s_t/c_t is the price-dividend ratio.

This model is not intended to capture all the complexities inherent in equity returns. In fact, similar models have been rejected on the grounds that they cannot match the observed equity premium or other features of the joint time series of equity returns and consumption data.[3] Nevertheless, as will become apparent, this model is both sufficiently complex to produce insight into the time-variation of the mean and volatility of equity returns and sufficiently simple to preserve tractability.

2.2 Volatility and Expected Returns

It is not immediately clear how the expected excess equity return and the conditional volatility of this return will be related in the framework above, even for the aggregate stock market. The covariance in equation (2) is not necessarily proportional to the variance of returns. Nevertheless, for many specifications, the variance of the market return and the covariance between the market return and the MRS will be closely linked.

[3]For early examples, see Hansen and Singleton (1982) and Mehra and Prescott (1985). Numerous attempts have been made to modify the model to better fit the data. These include introducing habit persistence and durability (Constantinides (1990) and Ferson and Constantinides (1991)), time non-separability of preferences (Epstein and Zin (1989)), and consumption adjustment costs (Marshall (1993)).

For the stock market, equation (2) can be rewritten as

$$\mathrm{E}_t[r_{st+1} - r_{ft}] \quad = \quad -r_{ft} \ \mathrm{Vol}_t[r_{st+1}] \ \mathrm{Vol}_t[m_{t+1}] \ \mathrm{Corr}_t[m_{t+1}, r_{st+1}] \ , \tag{5}$$

where

$$\begin{aligned} \mathrm{Corr}_t[m_{t+1}, r_{st+1}] &= \mathrm{Corr}_t \left[\beta \left(\frac{c_{t+1}}{c_t} \right)^{-\alpha}, \left(\frac{c_{t+1}}{c_t} \right) \frac{s_{t+1}/c_{t+1} + 1}{s_t/c_t} \right] \\ &= \mathrm{Corr}_t \left[\left(\frac{c_{t+1}}{c_t} \right)^{-\alpha}, \left(\frac{c_{t+1}}{c_t} \right) (s_{t+1}/c_{t+1} + 1) \right] \ , \end{aligned} \tag{6}$$

and Vol_t and Corr_t are the conditional volatility and conditional correlation, respectively. The conditional moments of returns will be positively related (period by period) as long as the correlation between the MRS and the equity return is negative. Holding the price-dividend ratio constant,[4] this condition holds (for $\alpha > 0$) since

$$\mathrm{Corr}_t \left[\left(\frac{c_{t+1}}{c_t} \right)^{-\alpha}, \left(\frac{c_{t+1}}{c_t} \right) \right] < 0.$$

The long-run relation between expected returns and volatility is less obvious because of potential time-variation in the conditional correlation, the conditional volatility of the marginal rate of substitution, and the riskless rate. However, a negative long-run relation between the moments of equity returns would generally require that time-variation in the correlation offset movements in the conditional volatility, i.e., that the correlation be high when volatility is low and vice versa. Again, this is impossible for fixed price-dividend ratios.

The above discussion makes it clear that the only way to duplicate the salient features of the data (i.e., weak or negative short-run and long-run relations between expected returns and volatility) is to formulate a model in which variation in the price-dividend ratio partially offsets the variation in the dividend growth component of the equity return in some states of the world. In other words, the price-dividend ratio must either covary positively with the MRS or covary weakly, but be volatile enough to reduce the overall correlation between the MRS and the return on equity. In these states of the world, the magnitude of the correlation will be reduced, and high volatility will no longer correspond to high expected returns.

[4]Of course, it may be difficult to imagine a world in which the price-dividend ratio is literally constant, yet there is time-variation in the moments of equity returns, since both depend on future consumption/dividend growth. This thought experiment is intended simply to illustrate the intuition behind the standard risk/return tradeoff.

To understand the behavior of the price-dividend ratio, rewrite equation (3) as

$$\frac{s_{t+1}}{c_{t+1}} = \sum_{s=1}^{\infty} \beta^s \mathrm{E}_{t+1} \left[\left(\frac{c_{t+1+s}}{c_{t+1}} \right)^{1-\alpha} \right]. \tag{7}$$

For $\alpha > 1$, this ratio is positively related to expectations of the inverse of future consumption growth. In other words, for suitably high risk aversion, the dominant effect is through the discount rate, not the growth in future dividends. High expected consumption growth implies low price-dividend ratios and vice versa. This predominance of the discount rate provides a link between this paper and research on long-term Treasury bond returns, where there is no dividend effect at all. More specifically, it supports the conjecture in Boudoukh, Richardson, Smith, and Whitelaw (1996) that it is the long-term nature of both stocks and bonds that is responsible for their similar volatility and expected return patterns.

Equation (7) shows that the price-dividend effect depends on the relation between consumption growth and expected future consumption growth. For example, if high consumption growth today implies high expected consumption growth in the future, then high consumption growth states will be associated with low price-dividend ratios. Variation in the price-dividend ratio therefore offsets variation in dividend growth in the equity return, and the correlation in equation (5) is reduced. Alternatively, if the price-dividend ratio is variable enough and unrelated to current dividend growth, the correlation will also be reduced.

There are two remaining problems. First, the magnitude of the variation in the correlation must be sufficiently large to offset variation in volatility. Second, time-variation in the short-run relation between expected returns and volatility (i.e., the existence of both positive and negative short-run correlations) requires that the correlation be strongly time-varying in some periods and much less so in others. Both of these problems are difficult, if not impossible, to overcome if consumption growth follows a simple autoregressive process. The correlation will vary little over time because the price-dividend ratio, which is an expectation of future consumption growth, will be less variable than consumption growth itself. Moreover, correlations will be relatively stable because both the immediate and distant future depend on a limited number of past values of consumption growth.

2.3 The Role of Regime Shifts

One simple and attractive way to overcome these problems is to consider a model with regime shifts and transition probabilities between regimes that are state dependent. For regimes that are sufficiently far apart in terms of the time series behavior of consumption growth, the regime switching probability will control the conditional volatility of returns. That is, states with a high probability of switching to a new regime will have high volatility. At the same time, however, increasing the probability of a regime switch may decrease the correlation between equity returns and the marginal rate of substitution, thus reducing the risk premium. This second effect will occur because the price-dividend ratios, which depend on expected future consumption growth, will be related to the regime not to short-run consumption growth.

The idea of shifts in regimes has gathered increasing empirical support in the literature. For example, Hamilton (1989) develops and estimates a two-regime model of the business cycle with constant switching probabilities. Filardo (1994) extends this model to time-varying transition probabilities, and he shows that allowing the probabilities to depend on economic state variables improves the goodness of fit.[5] There are numerous possible specifications, but for simplicity we consider a two-regime model. In particular, we assume that, at any point in time, the natural logarithm of consumption growth follows an autoregressive process of order 1 (AR(1)) with normally distributed errors and a constant variance. However, we also allow for the possibility of two different AR regimes. The state process follows a specified AR until a regime switch is triggered. This process then follows an AR with different parameters until another switch occurs. In particular, using the notation $g_{t+1} \equiv \ln(c_{t+1}/c_t)$, the two-regime economy is parameterized as

$$g_{t+1} = \begin{cases} a_1 + b_1 g_t + \epsilon_{1t+1} & \epsilon_{1t+1} \sim N(0, \sigma_1^2) \quad \text{for } I_{t+1} = 1 \\ a_2 + b_2 g_t + \epsilon_{2t+1} & \epsilon_{2t+1} \sim N(0, \sigma_2^2) \quad \text{for } I_{t+1} = 2 \end{cases} \tag{8}$$

where I_{t+1} indexes the regime. The evolution of this sequence of random variables is governed by the regime transition probabilities

$$\Pr[I_{t+1} = 1,\ 2|\Phi_t] = f(I_t, g_t),$$

i.e., the probability of being in a given regime next period depends only on the current regime and

[5] Gray (1995b) estimates a similar model for short-term interest rates, and the results indicate that the transition probabilities depend on the level of interest rates.

the underlying state variable. This function is parameterized as

$$\begin{aligned} P_{t+1}(1,1) &\equiv \Pr[I_{t+1}=1|I_t=1,g_t] = \tfrac{\exp(p_0+p_1 g_t)}{1+\exp(p_0+p_1 g_t)} \\ P_{t+1}(1,2) &\equiv \Pr[I_{t+1}=2|I_t=1,g_t] = 1-P_{t+1}(1,1) \\ P_{t+1}(2,2) &\equiv \Pr[I_{t+1}=2|I_t=2,g_t] = \tfrac{\exp(q_0+q_1 g_t)}{1+\exp(q_0+q_1 g_t)} \\ P_{t+1}(2,1) &\equiv \Pr[I_{t+1}=1|I_t=2,g_t] = 1-P_{t+1}(2,2) \end{aligned} \tag{9}$$

The parameterization of the regime switching model is a generalization of the switching model in Hamilton (1989), which is also studied in the context of stock returns in Cecchetti, Lam, and Mark (1990). It is similar to the specifications that Gray (1995b) uses to estimate the process for short-term interest rates and that Filardo (1994) uses to model the business cycle dynamics of industrial production.

3 A Single-Regime Model

Before proceeding to the estimation of the two-regime model described in Section 2.3, it is worthwhile to examine a single-regime model, i.e., a model in which consumption growth follows a single AR(1). This simple model provides a base case for comparison to the two-regime model, a setting for developing intuition about the risk/return tradeoff, and a context in which to discuss the methodology for computing the mean and volatility of equity returns.

3.1 Estimation

The model is estimated using monthly data on real, aggregate, chain-weighted consumption of nondurable goods and services from the Basic Economics database (series GMCNQ and GMCSQ). The monthly series starts in January 1959, but the late start date relative to the quarterly series is more than compensated for by the higher frequency of the data. Using data from 1/59-12/96 yields 455 observations for consumption growth. There are numerous issues with respect to the quality of the data, problems of time aggregation, etc., which are beyond the scope of this paper. Fortunately, the implied intertemporal relation between expected returns and volatility is relatively insensitive to the precise time series properties of the data. This issue is addressed in more detail in the sensitivity analyses later in the paper.

Table 1 provides descriptive statistics for the monthly, log consumption growth data (in percent) over the 2/59-12/96 sample period. Consumption growth varies from a low of -1.138% to a high of 1.696%, with a mean of 0.260%. The table also provides results from a generalized method of moments (GMM) estimation (Hansen (1982)) of an AR(1) on the same data. Heteroscedasticity-consistent standard errors are in parentheses, and the residual standard deviation (denoted σ) is also given. The coefficient indicates that consumption growth is negatively autocorrelated, but that lagged consumption growth does not explain a great deal of the variation in consumption growth. Note that the residual standard deviation of 0.381% is only slightly lower than the sample standard deviation of 0.392%. These results are consistent with other results in the literature that study consumption data.

3.2 Risk and Return

Using the pricing equations and the law of motion for consumption growth, it is sometimes possible to calculate the conditional moments of equity returns in closed form. For more complex, multi-regime specifications, closed-form solutions are no longer available; therefore, throughout the paper, we employ a discrete state space methodology that provides accurate numerical solutions. The continuous state variable (consumption growth) is approximated by a variable that takes on only a finite number of values. The dynamics are described by a transition matrix that gives the probabilities of moving between the various discrete states. The details of the discretization methodology, which follows Tauchen and Hussey (1991), and the technique for pricing in the discrete state space economy are given in Appendix A. The key point, from the perspective of examining the risk/return tradeoff, is that the discrete approximation converges quickly to the true model, and that the results are essentially identical to those from the continuous state space model.

As an illustration of the discrete approximation, Table 2 presents data for a 9-state approximation to the AR(1) process estimated in Table 1. The top panel shows the level of log consumption growth in each state and the unconditional probability of that state occurring. Note that the distribution is symmetric, centered around the mean of 0.260%, with extreme values of -1.459% and 1.979%. The probabilities are high near the center of the distribution and decrease rapidly at the extremes. As expected, the unconditional distribution looks much like a discrete approximation to the normal distribution. The bottom panel of the table shows the transition matrix, which captures the dynamics of the AR(1) specification. Specifically, each row of this matrix gives the probability of moving from a given state at time t to each of the nine states at time $t+1$. If consumption growth were i.i.d. over time, then the rows would be identical; however, the AR(1) estimation in Table 1 indicates the existence of negative autocorrelation. Consequently, from low consumption growth states (states 1-4), there is higher probability of going to a high consumption growth state (states 6-9) than there is of remaining in a low consumption growth state. The reverse applies to high consumption growth states (states 6-9). The discrete approximation captures all the salient features of the continuous state space AR(1) specification.

To analyze stock market risk and return we need to specify the degree of risk aversion and the time preference parameter. Initially we use $\alpha = 2$ and $\beta = 0.997$, and the sensitivity of the results to these parameters is addressed in the next section. Table 3 presents the state-by-state price-dividend ratio, risk premium, and volatility of the stock return. Note that the states are indexed from lowest to highest consumption growth using the same scheme as Table 2. For ease of understanding, all the values are annualized. The monthly risk premium and variance are multiplied by 12 and the price-dividend ratio is divided by 12. This latter adjustment makes the magnitudes of the ratios comparable to PE ratios calculated using annual earnings. In addition, the risk premium is multiplied by 100 for presentation purposes.

There are three major points worth noting about the results in Table 3. First, volatility is relatively low and not very variable. The lack of variation in the volatility is not surprising given the homoscedastic specification for log consumption growth; however, the log specification does induce some heteroscedasticity in consumption growth.

Second, the equity risk premium is small and also not very variable – this specification exhibits the standard equity risk premium puzzle. Within the context of the theoretical discussion in Section 2, one interpretation is that the volatility of the MRS in equation (5) is too small. This issue is basically orthogonal to the question of the relation between the risk premium and volatility. For example, it is possible to boost both the magnitude and the variability of the risk premium by increasing the level of risk aversion as discussed in Section 3.3. However, this change does not affect the basic risk/return relation.

Third, the price-dividend ratio is also not very variable, but it does increase as the level of consumption increases. This relation results from the negative autocorrelation in consumption growth and a level of risk aversion greater than one. As a result, high consumption growth in the current month implies an expectation of low consumption growth during the next month and thereafter, and a high price-dividend ratio (see equation (7)). From the discussion in Section 2.2, it is clear that a positive correlation between consumption growth and the price-dividend ratio is unlikely to generate the desired risk/return results.

From Table 3, it is also possible to see the risk premium/volatility relation; however, this relation is best illustrated graphically. Figure 1 plots the conditional risk premium against the conditional volatility of returns. Each point on the graph represents one of the nine states of the world. The graph shows a strong, positive, and essentially linear relation between the risk premium and the volatility of returns. This result coincides with the intuition of the risk/return tradeoff at the market level in a dynamic CAPM setting (see Merton (1980)). For this particular example, the implications of the graph are obvious, but it should be noted for future reference that state-by-state plots can sometimes be deceptive because each point does not have an equal probability. For example, state 5 has a probability of 39.5%, while the two extreme states (states 1 and 9) both have unconditional probabilities of 0.004% (see Table 2). Perhaps a more accurate idea of the unconditional relation between the expected risk premium and the volatility is given by the correlation between these conditional moments of returns. For this model, the unconditional correlation is 1.000, suggesting a very strong positive and linear relation.

We are also interested in potential time-variation in this relation. In the context of the discrete economy, time-variation is equivalent to variation across different states of the world. The most natural state-by-state measure is the conditional correlation between the conditional expected risk premium and the conditional volatility. Of course, at time t, the conditional moments based on time t information are known. Therefore, we consider the correlation at time t between the conditional expected risk premium and conditional volatility at time $t+1$. For example, suppose the economy is in a particular state (out of the 9 possible states) at time t. Next period (time $t+1$), the economy can be in any of the 9 states (with different probabilities), with corresponding conditional risk premiums and volatilities. The question we want to answer is whether high risk premiums are associated with high volatilities in these subsequent states. Conditional correlations will vary across states because transition probabilities vary across states. These conditional correlations are analogous to the short-run correlations between the estimated conditional moments that are reported in the empirical literature. As expected, given Figure 1, the conditional correlations are all 1.000. The short-run relation exhibits no time-variation.

3.3 Sensitivity Analysis

The risk/return relation for the single-regime model is not very interesting, but it is still worthwhile to consider how the various parameters affect this relation and the other values of interest.

The degree of risk aversion controls, among other things, the magnitude and variation of both volatility and the risk premium. When $\alpha = 40$, for example, the annualized risk premium varies from 5.2% to 7.2%, and annualized volatility varies from 9.8% to 13.7%; however, the positive, linear relation between the moments is preserved. For risk aversion lower than the base case value of 2, both moments exhibit even lower levels and less variation, but again they are essentially perfectly correlated. The only notable effect of lowering risk aversion is to change the relation between consumption growth and the price-dividend ratio. For $\alpha = 1$, the price-dividend ratio is constant across states, and for values less than one, the ratio decreases as consumption growth increases (see equation (7)).

The time preference parameter β has no effect on the risk/return relation, but it does control the magnitude of the price-dividend ratio. As β increases, the ratio increases, and vice versa.

The specification of the consumption growth process has important implications. Consider first the case where consumption growth is i.i.d., or equivalently, the AR(1) coefficient is zero. In this model, expectations do not depend on the current state of the world and consequently the price-dividend ratio, risk premium, and volatility are constant across states. For a positive AR(1) coefficient, time-variation is restored but reversed. The risk premium and volatility increase with consumption growth, and the price-dividend ratio decreases. The latter effect is a result of the positive autocorrelation, while the former effect arises because consumption growth volatility now increases with the level of consumption growth rather than decreasing, as in the base case. Unsurprisingly, the almost perfect correlation between the risk premium and volatility is preserved. Finally, the conditional volatility of consumption growth has an effect similar to that of the level of risk aversion. Higher consumption volatility generates a larger risk premium and higher return volatility, but the risk/return relation is again unaffected.

The clear conclusion from the above analysis is that, while the specific values are sensitive to the parameterization of the model, the major results concerning the relation between volatility and the risk premium are unaffected. In short, a single-regime model will not generate either a weak long-run relation or time-variation in the short-run relation.

4 A Two-Regime Model

Given the failure of the single-regime model, we now turn to a two-regime model to see if this generalization can produce a more complex risk/return relation. In recent years, there has been increasing empirical support for the hypothesis that the evolution of the economy is described by two or more distinct regimes with different parameters.[6] In general, this research provides evidence of multiple regimes within the course of a single business cycle which then repeat in succeeding cycles. This type of model should not be confused with models of one-time structural shifts, such as those used to model interest rates during the Fed experiment in 1979-1982. It also does not rely on extreme events that occur with small probability, as in the "peso problem".[7] Moreover, we model the fundamental process, consumption growth, rather than modeling equity returns directly. This approach contrasts with that of Gray (1995b), for example, who models interest rates as a regime switching process. Given that agents rationally anticipate regime shifts in the fundamental underlying process, the behavior of equity returns can take on potentially more complex, interesting, and realistic characteristics.

4.1 Estimation

Recall the specification of the two-regime model in equations (8) and (9), wherein log consumption growth can follow one of two AR(1) processes, depending upon the current regime, and the probability of a regime switch depends on the current regime and the level of consumption growth. There are a number of potential approaches to estimating this model, but we follow the maximum likelihood methodology in Gray (1995a). Using this approach, the model is reparameterized in terms of the probability of being in a given state at time t rather than in terms of the regime transition probabilities. This reparameterization allows for the construction of a recursive likelihood function much like the one used for GARCH estimation. The parameter values are chosen to maximize this

[6]See, for example, Hamilton (1989) and Filardo (1994).

[7]See, for example, Bekaert, Hodrick, and Marshall (1995) and Veronesi (1996).

function in the standard manner. See Appendix B for the details.

Table 4 presents the parameter estimates from this estimation, with standard errors in parentheses. Note that the regimes have been denoted as "expansion" and "contraction", which coincides with the regime shift business cycle literature. For example, Filardo (1994) estimates a model using industrial production data and attempts to match the identified regimes with the NBER business cycle dates. The reduced form model in this paper is not intended to match the observed cycles precisely; nevertheless, given that the parameters imply regimes with means of consumption growth of 0.323% and 0.146%, it makes intuitive sense to identify them with the phases of the cycle. The parameters of both regimes are estimated with good precision, and they are significant at all conventional levels. Both the mean and volatility of consumption growth are higher in expansions, but the level of mean reversion is almost identical across regimes.

Of greatest interest are the parameters which control the regime shifts. While the constants are both positive and significant, the coefficient on consumption growth is positive in expansions and negative in contractions. The standard errors on both estimates are large, but the point estimates suggest that regime persistence is positively related to consumption growth in expansions and negatively related to consumption growth in contractions. The existence of time-varying transition probabilities between regimes is consistent with the evidence in Filardo (1994). To illustrate the magnitude of this implied time-variation, Figure 2 plots the regime shift probabilities against log consumption growth. The graph shows $P(1,2)$ (solid line) and $P(2,1)$ (dashed line) – the probability of going from regime 1 to regime 2 and vice versa – as log consumption growth varies from -1.2% to 2.0%. The regime switch probabilities are relatively small for most reasonable levels of consumption growth. For example, at the within-regime means of the two regimes, $P(1,2)$ and $P(2,1)$ are 0.75% and 1.9%, respectively. If the probabilities were constant at these levels, then the regime half-lives would be approximately 92 months and 36 months, respectively. A slightly different perspective on the regime shift probabilities comes from considering the unconditional probability of being in each regime. At these estimated values, there is a 58% probability of being in an expansion and a 42% probability of being in a contraction.

4.2 Risk and Return

The two-regime specification has a total of 18 states of the world, 9 within each regime. The 9 states in each regime are identical in terms of their levels of consumption growth, but they differ with respect to their transition probabilities, both because of the differing AR parameters and the differing regime switching probabilities. Consequently, the conditional expected risk premium and the conditional volatility of returns can take on 18 different values. Table 5 reports log consumption growth, the price-dividend ratio, the risk premium, the volatility of returns, the probability of a regime shift, and the unconditional probability for each state. All the variables are annualized except for consumption growth, consistent with the presentation of the single-regime results in Tables 2 and 3. The discrete levels of consumption growth are slightly different than for the single-regime because they are based on the AR parameters of the expansion. Note, also, that the distributions within each regime are no longer unconditionally symmetric. For the expansion, this effect results from the asymmetry in the regime shift probabilities. For the contraction, there is the additional factor that the within-regime mean no longer equals consumption growth in state 5 because the means differ across regimes. Nevertheless, the transition matrix is calculated to produce a good approximation to the continuous state space model.

The conditional moments of returns do not exhibit the same monotonic patterns seen in the single-regime model. Figure 3 graphs the risk premium and volatility for each of the states. Expansion states are marked by circles and contraction states are marked by squares. The most notable feature of Figure 3, relative to the single-regime results in Figure 1, is the weakening of the relation between volatility and the risk premium. In the contractionary regime, the original positive relation has been reversed, and the risk premium and volatility are negatively related. In the expansion, the positive relation still holds for states 5-9, but, even in this limited set, the relation is no longer linear.

The overall relation is difficult to ascertain from the graph due to the differing probabilities associated with each state, but, in general, the points suggest a negative relation between the risk premium and volatility. In fact, the unconditional correlation between the first two conditional moments of returns is -0.481. This contrasts markedly with the 1.000 correlation in the single-regime model. Of greater importance, it coincides with the empirical results in Glosten, Jagannathan and Runkle (1993) and Whitelaw (1994). Both of these papers report a negative relation between conditional expected returns and conditional volatility. The analysis here shows that this negative relation is consistent with both general equilibrium and the fundamental time series properties of consumption growth.

How can the relatively straightforward two-regime specification generate such striking results? One perspective on the role of regime shifts can be gained by looking at the two regimes individually, as if they were each single-regime economies. In other words, consider the expansion or contraction with zero probability of a regime shift. Figure 4 graphs the state by state levels of the risk premium and the volatility for these two economies. Again, expansion states are marked by circles and contraction states are marked by squares. For comparison purposes, the single-regime premium and volatility are also plotted (marked by triangles). As expected, each regime individually bears a strong resemblance to the single-regime economy. The apparent clustering of points relative to Figure 1 is attributable to the differences in the scales of the graphs. If there are no regime switches, then there is a strong positive relation between the risk premium and the volatility in both regimes. The differences in the levels of risk premiums and volatilities across the three economies is due to the differences in the conditional volatility and autocorrelation of consumption growth. As the parameters change, so do the volatility and the risk premium. However, it is clearly not the parameters of the individual regimes but the existence of time-varying probabilities of regime shifts that creates the complex dynamics in the two-regime economy as plotted in Figure 3.

To understand these dynamics better, we start with the results underlying Figure 4. Table 6 presents the state-by-state values of the price-dividend ratio, the equity risk premium, and the volatility of stock returns for these two economies. Note that the consumption growth in each state is identical to the values given in Table 5. The only difference between the tables is that the regime shift probabilities have been set to zero in the latter table. As a result, price-dividend ratios are low and positively related to consumption growth in the expansion, and high and positively related to consumption growth in the contraction.

What happens when the possibility of a regime shift is introduced? Consider state 5 in the expansion. The price-dividend ratio is 13.337 if there is zero probability of ever entering a contraction. In Table 5, there is a 0.75% probability of an immediate switch of regimes, and a positive probability that a switch will occur in any subsequent period conditional on still remaining in the expansion. Consequently, the new price-dividend ratio accounts for the expectation that a switch to the contraction will occur, resulting in lower consumption growth in the future. From equation (7), lower consumption growth implies a higher price-dividend ratio; therefore, permitting regime shifts raises the price-dividend ratio from 13.337 to 14.879. A similar effect occurs in each state in the expansion, but the magnitude depends on the relative probability of a regime shift. In combination with the original consumption growth effect, state dependent probabilities lead to the U-shaped pattern for the expansion states in Table 5. For states in the contraction, the possibility of a shift to a high consumption growth regime lowers the price-dividend ratios, but the pattern from Table 6 is preserved, albeit in a weakened form.

The price-dividend ratios and consumption growth in each state, in turn, determine the behavior of equity returns. The return is a combination of two components: dividend (consumption) growth, and the change in the price-dividend ratio (see equation (4)). Note first that the variation in price-dividend ratios, especially across the regimes, tends to be larger than the variation in consumption growth. Table 5 is slightly deceptive in this respect because log consumption growth is given in percent. The implications are that the conditional volatility of returns is increasing in the probability of a regime shift and that volatility is larger than in the single-regime models. These patterns are clearly evident in the fifth column of Table 5.

The second issue is the correlation between equity returns and the MRS (see Section 2.2). If this correlation is strong and negative, as in the single-regime model, then expected returns will be positively related to volatility. However, the magnitude of this correlation also depends on the regime shift probability. Recall that consumption growth and price-dividend ratios are negatively correlated across regimes, i.e., price-dividend ratios are higher in the contraction than in the expansion. Consequently, a shift from the contraction to the expansion results, on average, not only in higher consumption growth and a lower MRS, but also in a lower price-dividend ratio and a lower equity return. Equity returns and the MRS tend to be positively correlated over regime transitions. This effect is sufficient to partially offset the standard negative correlation between the MRS and dividend growth. As a result, the correlation and the equity risk premium are low in states with high regime shift probabilities. For a sufficiently high regime shift probability, the correlation between the MRS and the return on equity may be positive, yielding a negative risk premium. This extreme case occurs in state 1 of the expansion with a regime shift probability of over 48%. While the unconditional probability of being in this state is low, the model does serve to illustrate the possibility of negative risk premiums at the stock market level.[8] Given the positive relation between regime shift probabilities and volatility noted above, the net result is a negative relation between the equity risk premium and the volatility of stock returns.

The contrast between the single-regime model and the two-regime model is equally apparent when considering the conditional correlation between the risk premium and the volatility. Given the graphical results in Figure 3, it is perhaps not surprising that the conditional correlation is negative in every state of the world. These correlations are plotted in Figure 5. Correlations in the expansion range from -0.99 in state 9 to -0.36 in state 3, while those in the contraction vary from -0.96 in state 1 to -0.37 in state 7. These patterns in the two regimes result from a combination of the within-regime transition probabilities, which look similar in both regimes, and the regime switch probabilities, which vary inversely. These results contrast with the essentially perfect positive conditional correlation in all states for the single-regime specification. The same effects that generate the long-run results discussed above are responsible for this short-run behavior. The existence of time-variation is consistent with results in the empirical literature (see, for example, Boudoukh, Richardson, and Whitelaw (1996), and Whitelaw (1994)), but the absence of positive correlations is not. This question is addressed in more detail in Section 4.3.

4.3 Sensitivity Analysis

For the two-regime model the purpose of the sensitivity analysis is twofold: to examine the robustness of the results to changes in the parameters and to find the conditions that yield the desired long-run and short-run behavior. For expositional clarity and brevity the discussion will focus primarily on the unconditional and conditional correlations between the risk premium and the volatility.

Initially, consider the time preference parameter β. Increasing β towards 1 increases the magnitude of the negative correlations because the price-dividend ratio is more sensitive to future consumption growth and hence, more sensitive to the regime. Decreasing β has the opposite effect. At a value of approximately 0.96, the unconditional correlation becomes positive and the conditional correlations are both negative and positive. Of course, a monthly value this low is difficult to justify.

[8]Boudoukh, Richardson, and Whitelaw (1997) make a similar point in the context of a simple, 4-state, discrete economy.

In some ways, the effect of risk aversion is similar. As α increases, the magnitudes of the negative correlations also increase. The reverse is true as α decreases. At $\alpha = 1$, price-dividend ratios are constant across both regimes, and the correlation becomes 1.000, unconditionally and conditionally. There are intermediate values for which the unconditional correlation is negative and the conditional correlations are both positive and negative.

With respect to the within-regime time series properties of consumption growth, there are six parameters and innumerable variations in these parameters that could be considered. We focus on three effects that illustrate how the short-run and long-run behavior result from a delicate balance between variation in consumption growth and variation in price-dividend ratios. First, consider altering the degree of autocorrelation in both regimes by changing b_1 and b_2 simultaneously. As these coefficients move toward zero, the short-run and long-run correlations between the mean and volatility become more negative because the current state of consumption growth has less influence on future expected consumption growth. When consumption growth is i.i.d. in both regimes, the expected dividend growth component of expected equity returns exhibits no within-regime variation. On the other hand, the variation in price-dividend ratios across regimes is still large. This cross-regime variation dominates expected equity returns. Second, consider varying the relative levels of mean consumption growth in the two states by changing a_1 and a_2. If the means are pushed further apart, the cross-regime variation in price-dividend ratios increases, and the correlations become more negative. Again, the issue is the relative variation in consumption growth and price-dividend ratios, especially across regimes. Third, think of changing the conditional volatility of consumption growth. Increasing volatility pushes the balance toward variation in consumption growth rather than price-dividend ratios, and the correlations increase.

As a final exercise, consider the critical role of the regime switch probabilities. In many ways, this is the most important analysis because these parameters are identified less accurately in the estimation. Consequently, from a statistical perspective, there is a wider range of plausible values, especially for the coefficients on consumption growth p_1 and q_1. The direct effects of changes in the parameters are relatively straightforward. For example, decreasing the constants p_0 and q_0, reduces regime persistence and increases the probability of a regime shift in every state. Decreasing only one of the constants reduces the unconditional probability of being in that regime. Decreasing the coefficients on consumption growth has similar effects, i.e., reducing regime persistence and reducing the unconditional probability of the regime. In addition, these coefficients also control the sensitivity of regime shift probabilities to the level of consumption growth. Large magnitudes, either negative or positive, generate larger variations across states.

The indirect effects on the moments of equity returns are less obvious and depend on the levels of these and the other parameters. For example, decreasing any of the parameters increases the probability of a regime shift, but it does not necessarily increase volatility. There is an offsetting effect on the distance between the price-dividend ratios across the regimes, as illustrated in Tables 5 and 6. The more persistent the regime, the more important is the mean consumption growth level in that regime for determining the price-dividend ratio. When regimes shifts are sufficiently likely, the current regime has little effect on the ratio. Of course, there is also an effect on the risk premium and on the conditional and unconditional correlations. A final point worth noting is that almost all the significant within-regime variation in risk premiums and volatilities comes from the variability in the regime shift probabilities. If either of the coefficients on consumption growth is set to zero, then all the points in that regime cluster in risk premium/volatility space. The direction of this effect can be seen in Figure 3, wherein the states within the contraction are more tightly clustered than those within the expansion.

Starting from the estimated parameter values, if either of the expansion parameters p_0 and p_1 decrease, both the unconditional and conditional correlations move upwards toward zero. In contrast, decreasing the contraction constant q_0 makes the correlation more negative. Finally, decreasing the coefficient q_1 has a small but positive effect on the correlations. The first three effects are all driven by the relative probability of the two regimes. Moving the weight more toward the contraction increases the correlation between the risk premium and the volatility. In the final case, this effect is offset by the negative effect associated with increasing variation across the states within the contraction.

Perhaps the easiest way to illustrate the complexity of the interactions and the magnitude of the effects is to look at a single interesting example. Table 7 gives the state-by-state values for the regime shift parameters $p_0 = 3.5$, $p_1 = 0.5$, $q_0 = 3.0$, and $q_1 = -1.4$. All the parameters have been reduced, so both regimes are less persistent. However, the unconditional probability of being in an expansion is increased to 71%. The effects on the moments of equity returns are quite dramatic. The risk premium/volatility patterns have changed in both regimes, and the overall variation in these moments has been reduced, especially for the expansion. The unconditional correlation is now weak and positive, taking on a value of 0.05. Finally, there has been a dramatic shift in the conditional correlations. These correlations exhibit extreme time-variation, achieving both high positive and high negative values.

From the above analysis, four conditions emerge as necessary to generate a weak or negative unconditional correlation and large time-variation in the conditional correlation between the risk premium and the volatility of stock returns. First, the AR parameters of the regimes must be sufficiently far apart to generate significant cross-regime variation in expected consumption growth. Second, risk aversion must be high enough to generate corresponding variation in price-dividend ratios. Third, regime shift probabilities must be relatively small to preserve the distinction between the regimes. Finally, these probabilities must also be state dependent to generate meaningful time-variation in the conditional correlations. However, given these conditions, there are numerous parameterizations that will generate results that are broadly consistent with the empirical evidence, but in direct contradiction to the standard risk/return intuition.

5 Conclusion

This paper shows that a two-regime exchange economy, estimated using consumption growth data, is able to duplicate two interesting features of the empirical relation between expected returns and volatility at the market level. Specifically, the model generates a negative unconditional correlation between these moments of returns and substantial time-variation in this relation. This paper demonstrates not only that a negative and time-varying relation between expected returns and volatility is consistent with rational expectations, but also that such a relation is consistent with aggregate consumption data in a representative agent framework.

An important implication of the results is that empirical models that impose a strong, often linear, relation between expected returns and volatility, such as GARCH-M, need to be employed with caution. The time series behavior implied by the model in this paper is inconsistent with many of these empirical specifications. One potential correction is to model expected returns in a multi-factor framework, with conditional volatility as one of the factors (e.g., Scruggs (1996)). Another promising approach is to model both expected returns and volatility nonparametrically, as functions of predetermined financial variables, thus allowing the data to tell the story (e.g., Boudoukh, Richardson, and Whitelaw (1997)).

Given the importance of regime shifts to the results, this paper indicates that further research in this area is clearly warranted. The sensitivity of the results to changes in the probability structure of regime shifts is both good and bad news in this respect. On the negative side, this sensitivity means that it is difficult to extract strong implications from models with parameters that are not estimated precisely. On the positive side, the time series properties of equity returns may provide a powerful information set with which to estimate these parameters.

Throughout the paper we have dealt with real equity returns. In contrast, the empirical literature works with nominal returns. Obviously, if inflation is constant, then all the results will carry through. More generally, adding stochastic inflation does not qualitatively affect the results as long as it has no real effects. The intuition behind this result is that inflation, while influencing the nominal marginal rate of substitution and equity returns, does not affect price-dividend ratios. Consequently, the dominant cross-regime dynamics are preserved. It would, however, be worthwhile to estimate a multi-regime model that permits a link between inflation and consumption growth and in which inflation is allowed to influence regime shift probabilities. Such a model presents a number of challenges, not the least of which is the fact that the time series of inflation appears to exhibit not only business cycle dynamics but also structural shifts.

A The Discretization Methodology

A.1 Formulating the Discrete State Space Model

Consider a variable x_t that follows a stationary AR process with a single lag:

$$x_{t+1} = a + bx_t + \epsilon_{t+1} \qquad \epsilon_t \sim \mathrm{N}(0, \sigma^2).$$

The assumption of one lag is made for the convenience of exposition only; longer lags can be handled in the same fashion by simply augmenting the vector of state variables. x_t can be approximated by a variable $\widehat{x}_t$ that takes on m discrete values. The evolution of $\widehat{x}_t$ through time can be described by a $m \times m$ transition matrix Π whose (i, j) entry is the probability of moving from state i at time t to state j at time $t+1$. The problem, of course, is choosing the discrete values of $\widehat{x}_t$ and the transition probabilities such that $\widehat{x}_t$ best approximates x_t. Tauchen and Hussey (1991) develop such a scheme based on numerical quadrature methods.[9] They choose the discrete points and transition probabilities such that the discretization matches the moments of $\widehat{x}_t$ with those of x_t. They also present an extensive discussion of the convergence of $\widehat{x}_t$ and functions of $\widehat{x}_t$ to their continuous state space counterparts as the number of quadrature points goes to infinity.

Denote the conditional mean of x_{t+1} as μ_t and the unconditional mean as μ:

$$\begin{aligned} \mathrm{E}_t[x_{t+1}] &\equiv \mu_t = a + bx_t \\ \mathrm{E}[x_{t+1}] &\equiv \mu = \frac{a}{1-b}\,. \end{aligned}$$

The first step is to decide on an appropriate discrete approximation to a standard normal random variable. This approximation amounts to choosing a set of discrete values and a set of corresponding weights (probabilities). A natural choice are values and weights which match as many moments of the standard normal as possible. For example, for two discrete values, the states $(1, -1)$ and the weights $(0.5, 0.5)$ match the mean, variance, and skewness of a standard normal. Consequently, define

$$z \equiv \begin{bmatrix} 1 \\ -1 \end{bmatrix},$$

[9] I would like to thank George Tauchen for the discrete approximation code that has been modified for this application.

and

$$w \equiv \begin{bmatrix} 0.5 \\ 0.5 \end{bmatrix} .$$

The discrete approximation to the original vector of state variables is computed as

$$\hat{x} = \mu + \sigma z = \begin{bmatrix} \mu + \sigma \\ \mu - \sigma \end{bmatrix} .$$

The final step is to calculate a 2×2 transition matrix Π, which will capture the dynamics of the original AR specification. The (i, j) element of this matrix, which is the probability of going from state i at time t to state j at time $t + 1$, is computed as

$$\begin{aligned} p_{ij} &= w_j * \phi_t(j)/\phi(j) \\ \pi_{ij} &= p_{ij}/\sum_j p_{ij} , \end{aligned}$$

where w_j is the weight on state j, $\phi_t(j)$ is the normal pdf with mean $\mu_t(i)$ and variance σ^2 evaluated at state j, and $\phi(j)$ is the normal pdf with mean μ and variance σ^2 evaluated at state j.

A.2 Pricing Assets

Given $\hat{x}_t$ and the transition matrix Π, the solutions to certain expectation equations become relatively easy to calculate. Assume, for example, that log consumption (dividend) growth follows an AR(1) and will be approximated by an m point discretization. Let $\hat{l}$ and $\hat{d}$ denote the $m \times 1$ vectors which contain the values that the MRS and consumption growth take on in each of the states:[10]

$$\begin{aligned} \hat{l} &= \left\{ \beta \left(\frac{c_{t+1}}{c_t} \right)^{-\alpha} \right\} \\ \hat{d} &= \left\{ \frac{c_{t+1}}{c_t} \right\} . \end{aligned}$$

Equation (5) from Section 5 can be rewritten in terms of the price-dividend ratio

$$\frac{s_t}{c_t} = \beta \mathrm{E}_t \left[\left(\frac{c_{t+1}}{c_t} \right)^{1-\alpha} \left(\frac{s_{t+1}}{c_{t+1}} + 1 \right) \right] .$$

[10]The time subscript is dropped because the state vectors are the same in every period.

In the discrete state space world, the expectation is simply a summation, and this equation can be written as

$$\widehat{v} = \Pi(\widehat{l}. * \widehat{d}. * \widehat{v}) + \Pi(\widehat{l}. * \widehat{d}) ,$$

which has the solution

$$\widehat{v} = \left[I_m - (1_m(\widehat{l}. * \widehat{d})^T). * \Pi\right]^{-1} \left[\Pi(\widehat{l}. * \widehat{d})\right] , \tag{10}$$

where $\widehat{v}$ denotes the vector of price–dividend ratios (one entry for each state of the world), I_i is an $i \times i$ identity matrix, 1_i is an i-vector of ones, superscript T denotes transpose, and $.*$ is element-by-element matrix multiplication.

The solution method for pricing finitely-lived securities such as bonds is somewhat different. For example, the price of a 1-period, riskless bond is the expectation of the MRS

$$q_{ft} \;=\; \beta \mathrm{E}_t \left[\left(\frac{c_{t+1}}{c_t} \right)^{-\alpha} \right],$$

which has the solution (in the discrete state space economy)

$$\widehat{q} = \Pi * \widehat{l}, \tag{11}$$

where $\widehat{q}$ is an m-vector of bond prices, one for each state of the world.

In addition to permitting solutions for asset prices in complex economies, the discrete state space pricing technique allows for relative simple computations of the conditional mean and variance of asset returns. The conditional moments are conditional on the specific state of the world as described by the state variables. The stationary distribution of the states can be found as the solution to $\Pi^T P = P$, where P is the vector of stationary probabilities.

The discussion above focuses on an economy described by a single AR process on the state variable. The extension to multiple regimes of individual AR processes is relatively straightforward. Within-regime transition probabilities, conditional on remaining within the regime, are derived as in Section A.1. The probabilities of moving between regimes, conditional on the state, must also be defined. The result is an augmented transition probability matrix and an extended matrix of state variables that can be used for asset pricing.

As an illustration, consider again the problem of estimating the price-dividend ratio, this time in an economy with two regimes. $\widehat{l}$ and $\widehat{d}$ are $m \times 1$ vectors which contain the values of the MRS and consumption growth in each of the m states. Let Π_1 denote the transition probabilities between these states, using the AR parameters of regime 1 and assuming there is only a single regime. In other words, $(\widehat{l}, \widehat{d}, \Pi_1)$ defines a single-regime model, and the rows of Π_1 sum to one. Similarly, let $(\widehat{l}, \widehat{d}, \Pi_2)$ define a single-regime model under the AR parameters of regime 2. The state vectors are identical across regimes, but the transition probabilities depend on the parameters of the two ARs. Assume that the probability of moving to regime 2 next period, conditional on being in regime 1, is state independent and equal to p. Similarly, assume the probability of moving from regime 2 to regime 1 is state independent and equal to q. Construct the augmented matrices $\widehat{l}^*$, $\widehat{d}^*$, and Π^* as follows:

$$\Pi^* = \begin{bmatrix} \Pi_1(1-p) & \Pi_2 p \\ \Pi_1 q & \Pi_2(1-q) \end{bmatrix} \quad \widehat{l}^* = \begin{bmatrix} \widehat{l} \\ \widehat{l} \end{bmatrix} \quad \widehat{d}^* = \begin{bmatrix} \widehat{d} \\ \widehat{d} \end{bmatrix}.$$

There are now $2m$ states of the world to reflect the fact that for each value of the MRS and consumption growth, the necessary information also includes the current regime. Π^* is a valid transition matrix since its rows sum to one. The augmented matrices are used for pricing in equations (10) and (11).

B Estimating a Two-Regime Model

Consider the model from equations (8) and (9):

$$\begin{aligned}
g_{t+1} &= \begin{cases} a_1 + b_1 g_t + \epsilon_{1t+1} & \epsilon_{1t+1} \sim N(0, \sigma_1^2) \quad \text{for } I_{t+1} = 1 \\ a_2 + b_2 g_t + \epsilon_{2t+1} & \epsilon_{2t+1} \sim N(0, \sigma_2^2) \quad \text{for } I_{t+1} = 2 \end{cases} \\
P_{t+1}(1,1) &= \frac{\exp(p_0 + p_1 g_t)}{1 + \exp(p_0 + p_1 g_t)} \\
P_{t+1}(1,2) &= 1 - P_{t+1}(1,1) \\
P_{t+1}(2,2) &= \frac{\exp(q_0 + q_1 g_t)}{1 + \exp(q_0 + q_1 g_t)} \\
P_{t+1}(2,1) &= 1 - P_{t+1}(2,2)
\end{aligned}$$

Following Gray (1995a), note that the conditional density of g_{t+1} can be written as

$$f(g_{t+1}|\Phi_t) = \sum_{i=1}^{2} f(g_{t+1}, I_{t+1} = i | I_t, g_t)$$

$$
\begin{aligned}
&= \sum_{i=1}^{2} f(g_{t+1}|I_{t+1}=i,g_t)\Pr[I_{t+1}=i|I_t,g_t] \\
&= \sum_{i=1}^{2} f(g_{t+1}|I_{t+1}=i,g_t)p_{it+1}
\end{aligned}
$$

where $p_{it+1} \equiv \Pr[I_{t+1}=i|I_t,g_t]$. Given conditional normality:

$$
f_{it+1} \equiv f(g_{t+1}|I_{t+1}=i,g_t) = \frac{1}{\sqrt{2\pi}\sigma_i}\exp\left(-\frac{g_{t+1}-a_i-b_i g_t}{\sigma_i^2}\right) \tag{12}
$$

The next, and critical step, is to reformulate the model in terms of p_{it+1} instead of $P_{t+1}(i,j)$. Noting that $p_{2t+1} = 1 - p_{1t+1}$, we focus on p_{1t+1}. The only conditioning information available is lagged values of consumption growth (denoted g_t^*), so p_{1t+1} can be rewritten as

$$
\begin{aligned}
\Pr[I_{t+1}=1|g_t^*] &= \sum_{i=1}^{2}\Pr[I_{t+1}=1|I_t=i,g_t^*]\Pr[I_t=i|g_t^*] \\
&= P_{t+1}(1,1)\Pr[I_t=1|g_t^*] + P_{t+1}(2,1)(1-\Pr[I_t=1|g_t^*])
\end{aligned} \tag{13}
$$

By Bayes rule,

$$
\Pr(I_t=1|g_t^*) = \frac{f(g_t|I_t=1,g_{t-1})\Pr[I_t=1,g_{t-1}]}{f(g_t|I_t=1,g_{t-1})\Pr[I_t=1,g_{t-1}] + f(g_t|I_t=2,g_{t-1})(1-\Pr[I_t=1,g_{t-1}])} \tag{14}
$$

Therefore, substituting (14) into (13),

$$
p_{1t+1} = P_{t+1}(1,1)\frac{f_{1t}p_{1t}}{f_{1t}p_{1t}+f_{2t}(1-p_{1t})} + P_{t+1}(2,1)\frac{f_{2t}(1-p_{1t})}{f_{1t}p_{1t}+f_{2t}(1-p_{1t})}
$$

where f_{it} is defined in equation (12). The log likelihood function is

$$
L = \sum_{t=1}^{T}\log[p_{1t}f_{1t} + (1-p_{1t})f_{2t}]
$$

which can be constructed recursively in the same way as in a GARCH model.[11]

[11] I would like to thank Steve Gray for the estimation code that has been modified for this application.

Mean	Std. Dev.	Min.	Max.
0.260	0.392	-1.138	1.696

AR(1) Estimation

	Constant	g_t	σ
g_{t+1}	0.322	-0.239	0.381
	(0.021)	(0.046)	(0.015)

Table 1: Descriptive Statistics

Descriptive statistics for monthly log consumption growth (in percent) for the period 2/59-12/96. The AR(1) is estimated using GMM, with heteroscedasticity-consistent standard errors in parentheses. σ denotes the residual standard deviation.

State	$g_t(\%)$	Prob.(%)
1	-1.459	0.004
2	-0.961	0.363
3	-0.531	5.482
4	-0.130	24.422
5	0.260	39.457
6	0.650	24.422
7	1.051	5.482
8	1.481	0.363
9	1.979	0.004

Transition Matrix(%)

	$t+1$								
t	1	2	3	4	5	6	7	8	9
1	0.00	0.00	0.30	4.53	22.71	41.14	26.21	4.95	0.16
2	0.00	0.02	0.76	8.31	30.30	39.86	18.27	2.42	0.05
3	0.00	0.05	1.57	12.99	35.93	35.86	12.37	1.21	0.02
4	0.00	0.12	2.92	18.45	39.44	30.43	8.05	0.59	0.01
5	0.00	0.28	4.99	24.41	40.63	24.41	4.99	0.28	0.00
6	0.01	0.59	8.05	30.43	39.44	18.45	2.92	0.12	0.00
7	0.02	1.21	12.37	35.86	35.93	12.99	1.57	0.05	0.00
8	0.05	2.42	18.27	39.86	30.30	8.31	0.76	0.02	0.00
9	0.16	4.95	26.21	41.14	22.71	4.53	0.30	0.00	0.00

Table 2: The 9-State Discrete Approximation

The top panel shows log consumption growth in each of the 9 states and the unconditional probability of the state. The bottom panel shows the transition matrix which determines the probability of moving from a given state at time t to any other state at time $t+1$.

State	s_t/c_t	$\mathrm{E}_t[r_{st+1} - r_{ft}]$ (% × 100)	$\mathrm{Vol}_t[r_{st+1}]$ (%)
1	14.794	4.222	1.599
2	14.808	4.212	1.596
3	14.820	4.203	1.592
4	14.832	4.195	1.589
5	14.843	4.187	1.586
6	14.854	4.179	1.583
7	14.866	4.171	1.580
8	14.878	4.163	1.577
9	14.892	4.153	1.573

Table 3: Risk and Return in a Single-Regime Model

State-by-state values for the price-dividend ratio, the risk premium, and the volatility of stock returns in a single-regime model based on the AR(1) specification in Table 1. All values are annualized.

	a_1	b_1	σ_1	p_0	p_1
Expansion	0.422	-0.307	0.383	3.986	2.783
	(0.034)	(0.066)	(0.017)	(0.749)	(3.403)
	a_2	b_2	σ_2	q_0	q_1
Contraction	0.191	-0.312	0.328	4.058	-0.741
	(0.036)	(0.081)	(0.020)	(1.068)	(2.989)

Table 4: Parameter Estimates for the Two-Regime Model

Parameter estimates for the two-regime, AR(1) model of log consumption growth given in equations (8) and (9). The model is estimated by maximum likelihood (see Appendix B), using monthly consumption of nondurable goods and services data from 2/59-12/96. Standard errors are in parentheses.

State	g_t	s_t/c_t	$\mathrm{E}_t[r_{st+1} - r_{ft}]$ (% × 100)	$\mathrm{Vol}_t[r_{st+1}]$ (%)	$P_t(i,j)$	Prob.
			Expansion			
1	-1.406	15.148	-1.859	7.529	48.139	0.003
2	-0.905	14.959	0.586	6.046	18.726	0.251
3	-0.473	14.889	2.766	4.023	6.471	3.400
4	-0.069	14.873	3.594	2.696	2.201	14.253
5	0.323	14.879	3.762	1.977	0.751	22.549
6	0.715	14.892	3.648	1.605	0.253	14.183
7	1.118	14.909	3.380	1.386	0.083	3.363
8	1.551	14.930	2.968	1.212	0.025	0.247
9	2.051	14.959	2.333	1.024	0.006	0.003
			Contraction			
1	-1.406	15.507	2.966	1.755	0.606	0.001
2	-0.905	15.526	2.880	1.904	0.876	0.212
3	-0.473	15.541	2.817	2.072	1.203	3.999
4	-0.069	15.554	2.736	2.261	1.616	15.674
5	0.323	15.566	2.626	2.477	2.148	16.530
6	0.715	15.576	2.478	2.727	2.852	4.937
7	1.118	15.585	2.273	3.018	3.808	0.389
8	1.551	15.592	1.981	3.365	5.171	0.004
9	2.051	15.597	1.531	3.799	7.324	0.000

Table 5: Risk and Return in a Two-Regime Model

State-by-state values for log consumption growth, the price-dividend ratio, the risk premium, the volatility of stock returns, the probability of a regime shift, and the unconditional state probability in a two-regime model based on the parameter values in Table 4. All values except for consumption growth are annualized.

State	s_t/c_t	$\mathrm{E}_t[r_{st+1} - r_{ft}]$ (% × 100)	$\mathrm{Vol}_t[r_{st+1}]$ (%)
		Expansion	
1	13.283	4.432	1.670
2	13.298	4.418	1.665
3	13.312	4.406	1.661
4	13.324	4.395	1.656
5	13.337	4.385	1.652
6	13.349	4.374	1.649
7	13.362	4.363	1.644
8	13.375	4.352	1.640
9	13.391	4.338	1.635
		Contraction	
1	18.586	3.296	1.439
2	18.608	3.237	1.424
3	18.627	3.225	1.419
4	18.645	3.217	1.416
5	18.663	3.209	1.412
6	18.680	3.201	1.409
7	18.698	3.194	1.405
8	18.717	3.187	1.402
9	18.740	3.190	1.400

Table 6: Risk and Return in Two Single-Regime Models

State-by-state values for the price-dividend ratio, the risk premium, and the volatility of stock returns for the individual regimes within the two-regime model, assuming zero probability of a regime shift. The parameter values are given in Table 4. All values are annualized.

State	$E_t[r_{st+1} - r_{ft}]$ (% × 100)	$Vol_t[r_{st+1}]$ (%)	$P_t(i,j)$	Prob.	ρ_t
		Expansion			
1	3.880	1.981	5.747	0.004	-0.816
2	3.959	1.925	4.532	0.306	-0.668
3	4.015	1.879	3.684	4.137	-0.441
4	4.057	1.839	3.031	17.350	-0.156
5	4.089	1.803	2.505	27.481	0.109
6	4.113	1.771	2.069	17.322	0.295
7	4.132	1.740	1.697	4.123	0.400
8	4.145	1.712	1.372	0.304	0.445
9	4.154	1.683	1.071	0.004	0.440
		Contraction			
1	2.611	1.229	0.691	0.001	-0.841
2	2.646	1.334	1.383	0.149	-0.703
3	2.656	1.464	2.505	2.785	-0.539
4	2.612	1.630	4.325	10.887	-0.353
5	2.496	1.842	7.256	11.461	-0.134
6	2.293	2.105	11.928	3.415	0.139
7	2.004	2.407	19.242	0.268	0.455
8	1.695	2.704	30.381	0.003	0.727
9	1.580	2.893	46.798	0.000	0.884

Table 7: Risk and Return in a Two-Regime Model

State-by-state values for the risk premium, the volatility of stock returns, the probability of a regime shift, the unconditional state probability, and the conditional correlation in a two-regime model based on the AR parameter values in Table 4 and the regime shift parameters $p_0 = 3.5$, $p_1 = 0.5$, $q_0 = 3.0$, and $q_1 = -1.4$. All values are annualized.

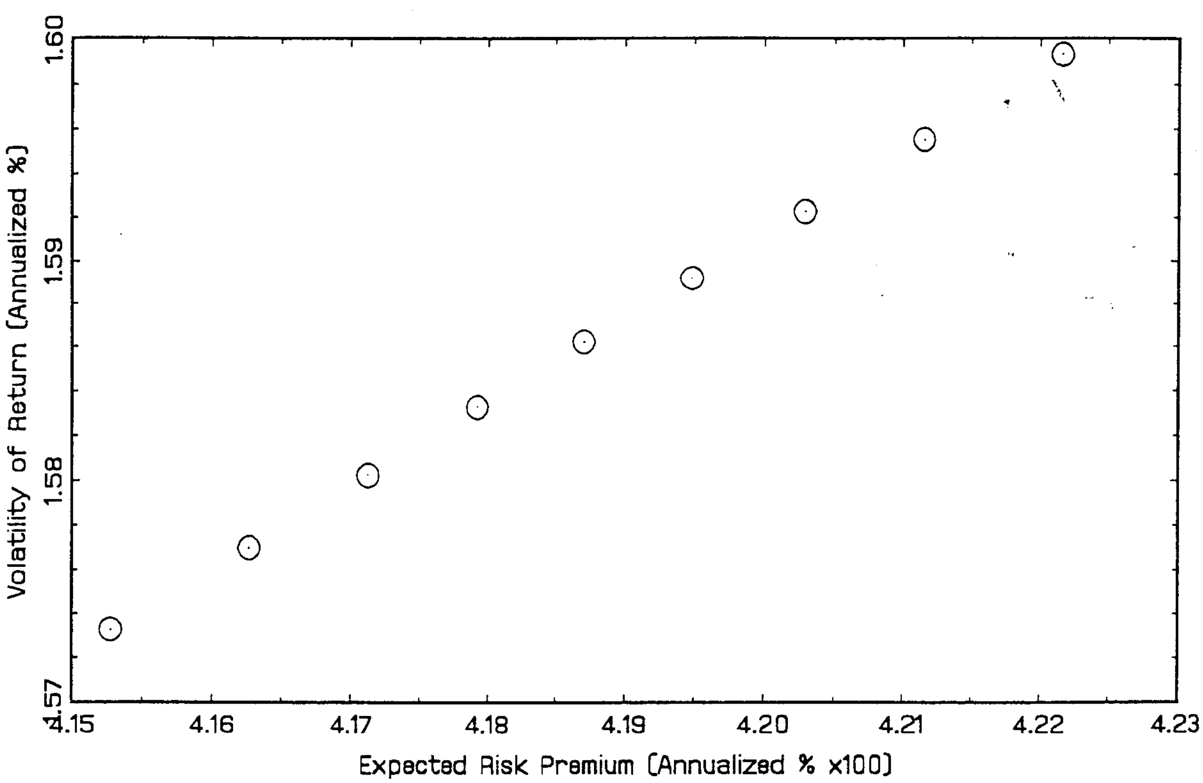

Figure 1: Risk and Return in a Single-Regime Model

State-by-state values of the conditional equity risk premium (times 100) and the conditional volatility of equity returns (both in percent, annualized) for a single-regime model. The parameter values are given in Table 2.

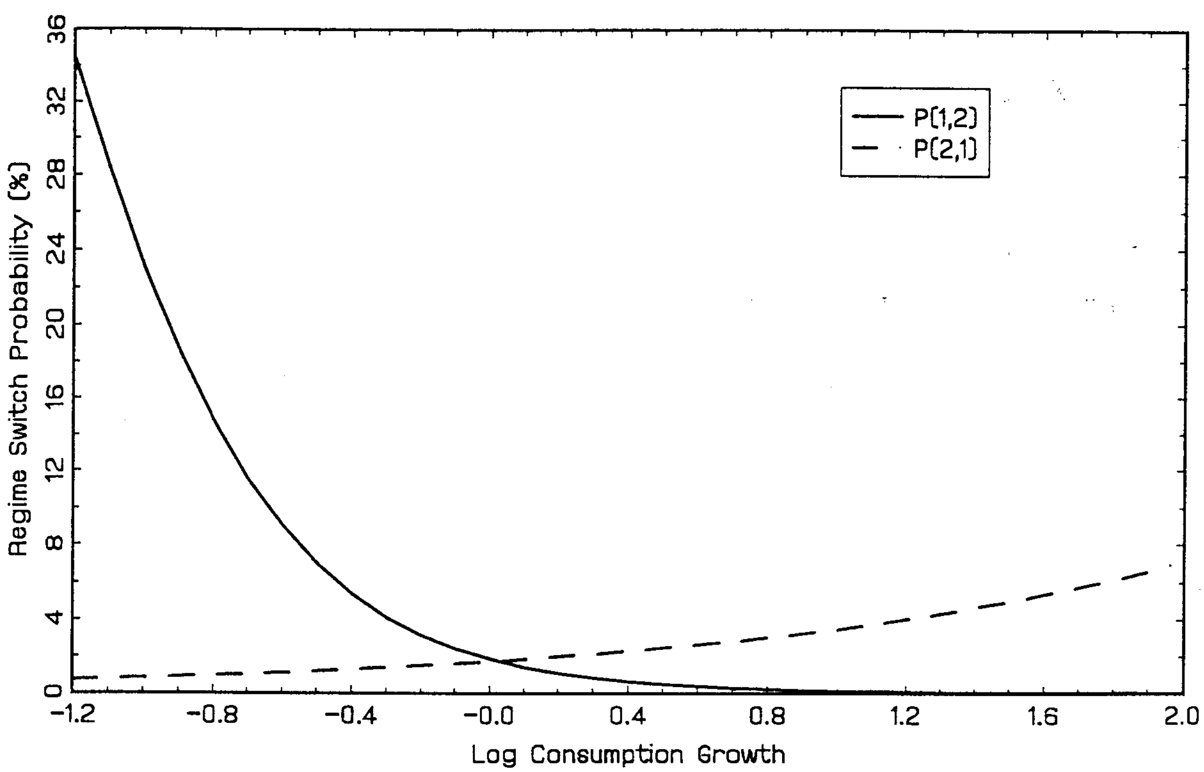

Figure 2: Regime Shift Probabilities

The estimated probabilities of shifting from an expansion to a contraction ($P(1,2)$, solid line) and from a contraction to an expansion ($P(2,1)$, dashed line) as a function of log consumption growth. The functions are given in equation (9), and the parameter estimates are in Table 4.

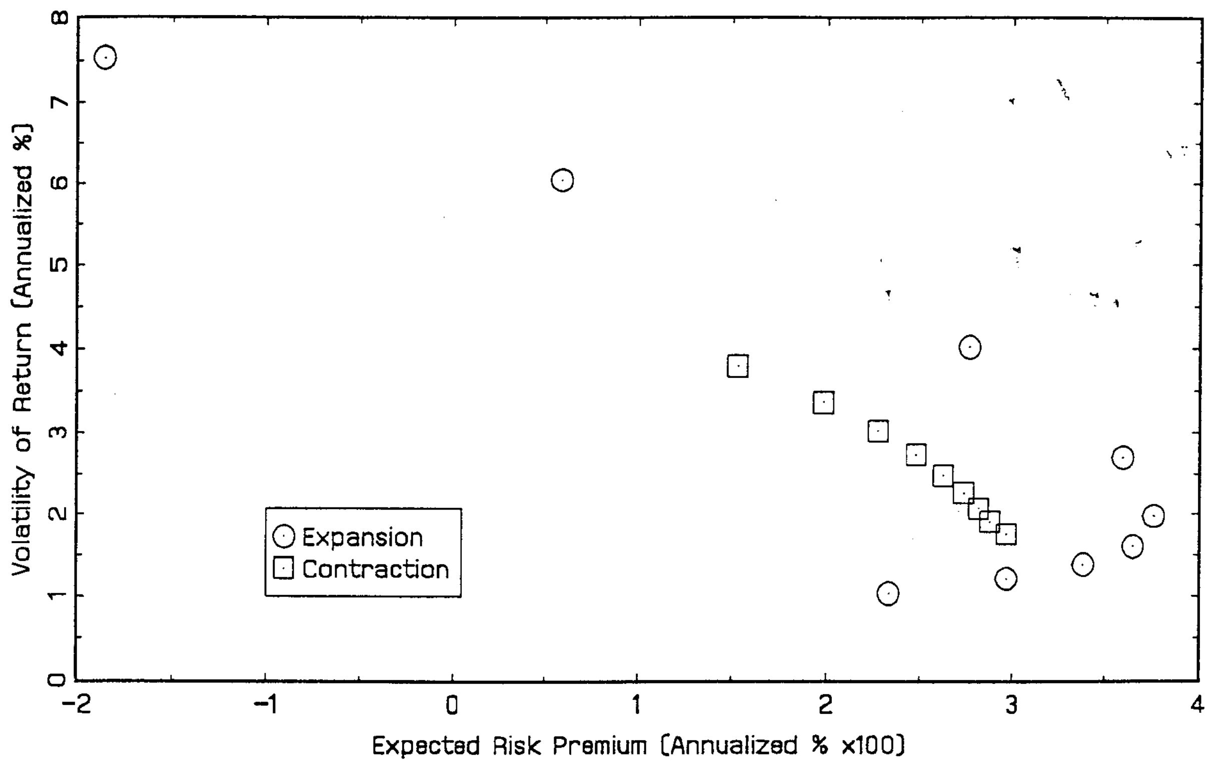

Figure 3: Risk and Return in a Two-Regime Model

State-by-state values of the conditional equity risk premium (times 100) and the conditional volatility of equity returns (both in percent, annualized) for a two-regime model. The model parameters are given in Table 5. Expansion and contraction states are marked by circles and squares, respectively.

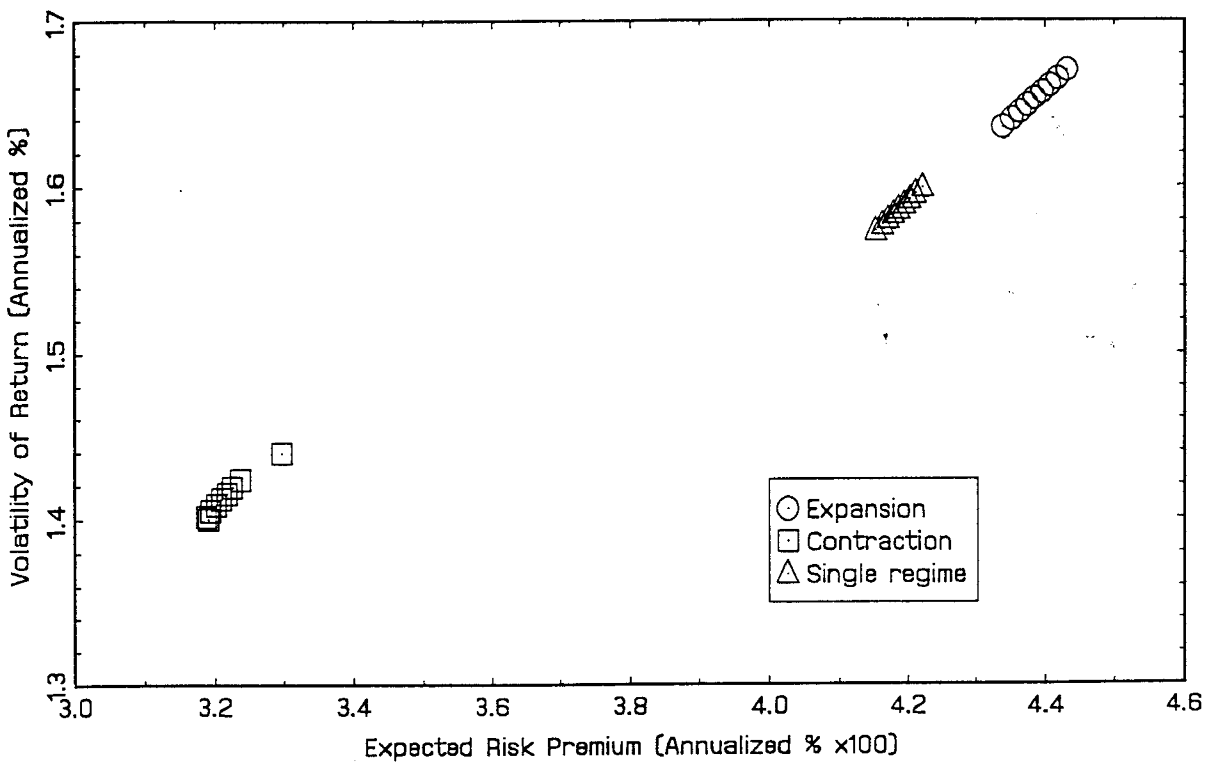

Figure 4: Risk and Return in Two Single-Regime Models

State-by-state values of the conditional equity risk premium (times 100) and the conditional volatility of equity returns (both in percent, annualized) for the individual regimes within the two-regime model, assuming zero probability of a regime shift, and for the single-regime model. The parameters for the two-regime and single-regime models are given in Tables 4 and 1, respectively. Expansion states are marked by circles, contraction states are marked by squares, and the single-regime states are marked by triangles.

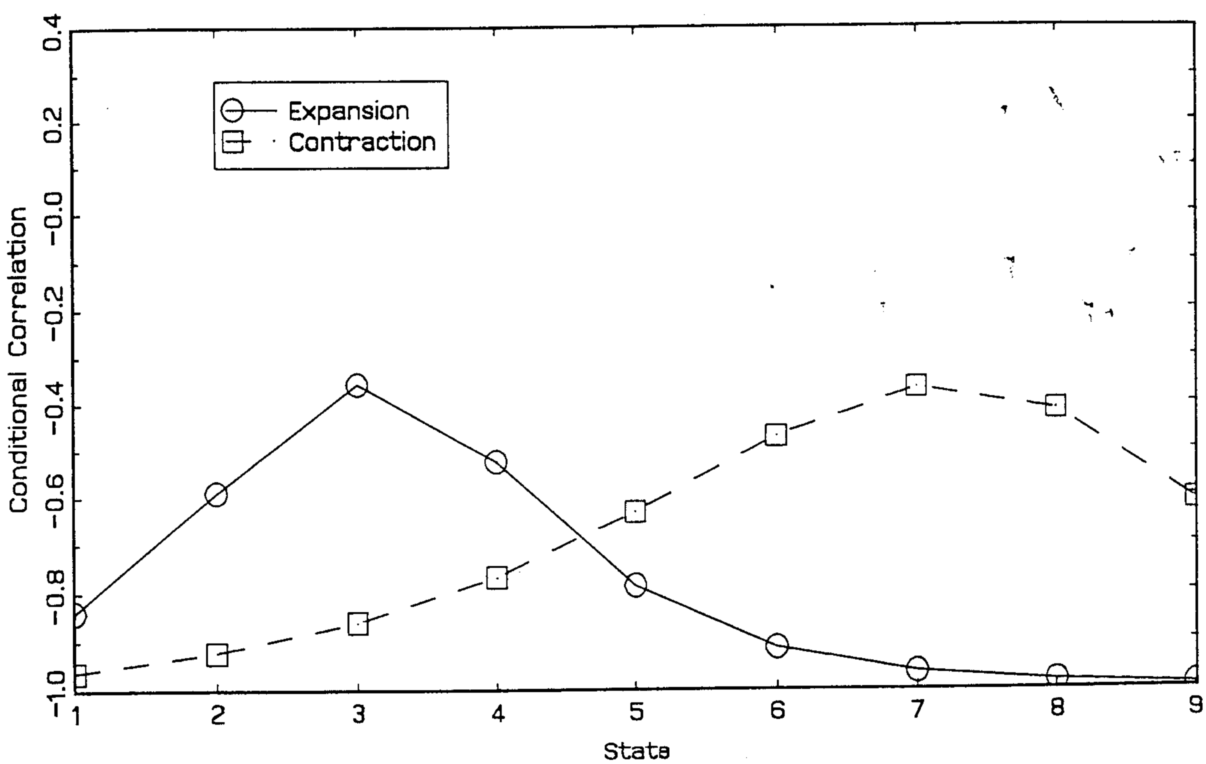

Figure 5: Conditional Correlations in a Two-Regime Model

State-by-state values of the conditional correlation between the equity risk premium and the conditional volatility of equity returns for a two-regime model. The model parameters are given in Table 5. The x-axis gives the number of the state using the scheme in Table 5. Expansion and contraction states are marked by circles and squares, respectively.

References

[1] Abel, Andrew B., 1988, "Stock Prices Under Time-Varying Dividend Risk: An Exact Solution in an Infinite Horizon General Equilibrium Model," *Journal of Monetary Economics*, 22, 375-395.

[2] Backus, David, and Allan Gregory, 1993, "Theoretical Relations Between Risk Premiums and Conditional Variances," *Journal of Business and Economic Statistics*, 11, 177-185.

[3] Bekaert, Geert, Robert Hodrick, and David Marshall, 1995, "Peso Problem Explanations for Term Structure Anomalies," unpublished working paper.

[4] Boudoukh, Jacob, Matthew Richardson, Tom Smith, and Robert F. Whitelaw, 1996, "Ex Ante Bond Returns and the Yield Curve," working paper, New York University.

[5] Boudoukh, Jacob, Matthew Richardson, and Robert F. Whitelaw, 1997, "Nonlinearities in the Relation Between the Equity Risk Premium and the Term Structure," *Management Science*, 43, 371-385.

[6] Campbell, John Y., 1987, "Stock Returns and the Term Structure," *Journal of Financial Economics*, 18, 373-399.

[7] Cecchetti, Stephen G., Pok-sang Lam, and Nelson C. Mark, 1990, "Mean Reversion in Equilibrium Asset Prices," *American Economic Review*, 80, 398-418.

[8] Cochrane, John H., 1991, "Production-Based Asset Pricing and the Link Between Stock Returns and Economic Fluctuations," *Journal of Finance*, 46, 209-236.

[9] Constantinides, George, 1990, "Habit Formation: A Resolution of the Equity Premium Puzzle," *Journal of Political Economy*, 98, 519-543.

[10] Epstein, Larry, and Stanley Zin, 1989, "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57, 937-970.

[11] Ferson, Wayne, and George Constantinides, 1991, "Habit Persistence and Durability in Aggregate Consumption: Empirical Tests," *Journal of Financial Economics*, 29, 199-240.

[12] Filardo, A.J., 1994, "Business Cycle Phases and Their Transitional Dynamics," *Journal of Business and Economic Statistics*, 12, 299-308.

[13] French, Kenneth R., G. William Schwert, and Robert F. Stambaugh, 1987,"Expected Stock Returns and Volatility," *Journal of Financial Economics*, 19, 3-29.

[14] Glosten, Lawrence R., Ravi Jagannathan, and David Runkle, 1993, "On the Relation Between the Expected Value and the Volatility of the Nominal Excess Return on Stocks," *Journal of Finance*, 48, 1779-1801.

[15] Gray, Stephen, 1995a, "Regime-Switching Models: A New Approach," working paper, Duke University.

[16] Gray, Stephen, 1995b, "Modeling the Conditional Distribution of Interest Rates as a Regime-Switching Process," working paper, Duke University.

[17] Hamilton, James D., 1989, "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle," *Econometrica*, 57, 357-384.

[18] Hansen, Lars P., 1982, "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 50, 1929-1954.

[19] Hansen, Lars, and Kenneth Singleton, 1982, "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models," *Econometrica*, 50, 1269-1286.

[20] Lucas, Robert E., Jr., 1978, "Asset Prices in an Exchange Economy," *Econometrica*, 46, 1429-1445.

[21] Marshall, David A., 1993, "Asset Return Volatility with Extremely Small Costs of Consumption Adjustment," unpublished working paper.

[22] Mehra, Rajnish, and Edward Prescott, 1985, "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 15, 145-161.

[23] Merton, Robert C., 1980, "On Estimating the Expected Return on the Market," *Journal of Financial Economics*, 8, 323-361.

[24] Scruggs, John T., 1996, "Resolving the Puzzling Relation Between the Market Risk Premium and Conditional Market Variance: A Two-Factor Approach," unpublished working paper.

[25] Tauchen, George, and Robert Hussey, 1991, "Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models," *Econometrica*, 59, 371-396.

[26] Veronesi, Pietro, 1996, "Regime Shifts and Changing Volatility in Stock Returns," working paper, Harvard University.

[27] Whitelaw, Robert F., 1994, "Time Variations and Covariations in the Expectation and Volatility of Stock Market Returns," *Journal of Finance*, 49, 515-541.

[28] Whitelaw, Robert F., 1997, "Time-Varying Sharpe Ratios and Market Timing," working paper, New York University.