

Stock Market Risk and Return: An Equilibrium Approach

Robert F. Whitelaw

Empirical evidence that expected stock returns are weakly related to volatility at the market level appears to contradict the intuition that risk and return are positively related. We investigate this issue in a general equilibrium exchange economy characterized by a regime-switching consumption process with time-varying transition probabilities between regimes. When estimated using consumption data, the model generates a complex, non-linear and time-varying relation between expected returns and volatility, duplicating the salient features of the risk/return trade-off in the data. The results emphasize the importance of time-varying investment opportunities and highlight the perils of relying on intuition from static models.

Understanding the risk/return trade-off is fundamental to equilibrium asset pricing. In this context, the stock market is one of the most natural starting points since it serves as a proxy for the wealth portfolio that is studied in finance theory. It is perhaps surprising, therefore, that there is still a good deal of controversy around the issue of how to measure risk at the market level. Recent empirical studies [e.g., Glosten, Jagannathan, and Runkle (1993), Whitelaw (1994), 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.² Second, they document significant time variation in this relation. Specifically, in a modified GARCH-M framework using post-World War II monthly data, Glosten, Jagannathan, and Runkle (1993) find that the estimated coefficient on volatility in the expected return regression is negative. In a similar dataset, when both conditional moments are estimated as functions of predetermined financial variables, Whitelaw (1994)

My thanks to an anonymous referee, the editor, Ravi Jagannathan, Kobi Boudoukh, Kent Daniel, Wayne Ferson, John Heaton, Anthony Lynch, Matt Richardson, and seminar participants at UCLA, Duke University, University of Minnesota, Northwestern University, University of Southern California, the 1997 AFA meetings, and the 1996 Utah Winter Finance Conference for helpful comments. Address correspondence to Robert F. Whitelaw, New York University, Stern School of Business, 44 West 4th St., Suite 9-190, New York, NY 10012, or email: rwhitela@stern.nyu.edu.

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

² Similar results are also reported in Nelson (1991), Pagan and Hong (1991), and Harrison and Zhang (1999). However, Harrison and Zhang (1999) show that at longer horizons (i.e., 1–2 years) there is a significantly positive relation between expected returns and conditional volatility. The shorter horizon phenomenon is also present in international data. For example, De Santis and Imrohorglu (1997) find a significant positive relation in only 2 countries out of a sample of 14 emerging and 3 developed markets, and Ang and Bekaert (1999) identify a low mean high volatility return regime using U.S., U.K., and German data.

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. 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?³ Second, what features are necessary to generate this counterintuitive behavior of expected returns and volatility?

This article 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) and Kandel and Stambaugh (1990), who attempt to duplicate various features of equity return data in an equilibrium setting. Consumption growth is modeled as an autoregressive process, with two regimes in which the parameters differ, an extension of Hamilton (1989). The probability of a regime shift is modeled as a function of the level of consumption growth, yielding time-varying transition probabilities as in 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 short-run and long-run relation between expected equity returns and the volatility of returns.

The two-regime specification is able to identify the expansionary and contractionary phases of the business cycle consistent with the NBER business cycle dating. More important, the model generates results that are broadly consistent with the empirical evidence. Expected returns and conditional volatility exhibit a complex, nonlinear relation. They are negatively related in the long run and this relation varies widely over time. 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. In marked contrast, a single-regime model calibrated to the same data generates a strong positive, and essentially linear, relation between expected returns and volatility.

In order to preserve tractability, the two-regime specification is kept simple. As a consequence, the reduced form model, while providing insight into the

³ 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 article is whether the more specific intertemporal patterns documented recently are consistent with economic data.

relation between risk and return, fails to match other features of the equity return data. For example, the magnitude of the equity premium is much too low. However, efforts to address this and other puzzles using tools such as habit persistence [see, e.g., Campbell and Cochrane (1999)] are beyond the scope of this article.

The major contribution of this article 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 also provides the relevant economic intuition. 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.

Put slightly differently, regime shifts introduce large movements in the investment opportunity set, and therefore induce a desire among investor's to hedge adverse changes [see Merton (1973)]. In some states of the world, the market claim provides such a hedge. Specifically, when a regime shift is likely, its value is high and its expected return is low as a consequence. These are also states of the world with high volatility, generating the required negative relation between volatility and expected returns. In other states of the world, regime shifts are less likely and the standard positive relation dominates.

The remainder of the article is organized as follows. Section 1 develops the asset pricing framework, provides the intuition behind the risk/return relation in this setting, and describes the two-regime specification. In Section 2, we estimate and analyze the two-regime model, contrasting the results to those from a single-regime model. The expected return and volatility patterns are analyzed and a sensitivity analysis is performed. Section 3 concludes.

1. Theory

1.1 The asset pricing framework

Consider a pure exchange economy with a single consumption good and a representative agent whose utility function exhibits constant relative risk aversion [Lucas (1978)]. The resulting pricing equation is

$$E_t \left[\beta \left(\frac{c_{t+1}}{c_t} \right)^{-\alpha} r_{t+1} \right] = 1, \quad (1)$$

where c_t is consumption, r_{t+1} is the asset return, α is the coefficient of risk aversion, β is the time preference parameter, and $E_t[\cdot]$ denotes the expectation conditional on information available at time t . The expected return on any asset in excess of the riskless rate (denoted r_{ft}) is proportional to the negative of the covariance of this return with the marginal rate of substitution (MRS), that is,

$$E_t[r_{t+1} - r_{ft}] = -r_{ft} \text{Cov}_t[m_{t+1}, r_{t+1}], \quad (2)$$

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, that is, to equate aggregate consumption and the aggregate stock market dividend [see, e.g., Mehra and Prescott (1985) and Cecchetti, Lam, and Mark (1990)]. Consequently, the market return is

$$r_{st+1} = \frac{s_{t+1} + c_{t+1}}{s_t} = \left(\frac{c_{t+1}}{c_t} \right) \frac{(s_{t+1}/c_{t+1}) + 1}{s_t/c_t}, \quad (3)$$

and the price:dividend ratio is

$$\frac{s_t}{c_t} = \sum_{s=1}^{\infty} E_t \left[\beta^s \left(\frac{c_{t+s}}{c_t} \right)^{1-\alpha} \right]. \quad (4)$$

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.⁴ 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.

It is not immediately clear how the expected excess equity return and the conditional volatility of this return will be related in this framework. Nevertheless, for many specifications, the variance of the market return and the covariance between the market return and the MRS will be closely linked. Specifically, for the stock market, Equation (2) can be rewritten as

$$E_t[r_{st+1} - r_{ft}] = -r_{ft} \text{Vol}_t[r_{st+1}] \text{Vol}_t[m_{t+1}] \text{Corr}_t[m_{t+1}, r_{st+1}], \quad (5)$$

where

$$\text{Corr}_t[m_{t+1}, r_{st+1}] = \text{Corr}_t \left[\beta \left(\frac{c_{t+1}}{c_t} \right)^{-\alpha}, \left(\frac{c_{t+1}}{c_t} \right) \frac{(s_{t+1}/c_{t+1}) + 1}{s_t/c_t} \right] \quad (6)$$

⁴ 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 nonseparability of preferences [Epstein and Zin (1989)], and consumption adjustment costs [Marshall (1993)].

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,⁵ this condition holds (for $\alpha > 0$) since

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

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, that is, that the correlation be high when volatility is low and vice versa. Again, this is impossible for fixed price:dividend ratios.

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.

For $\alpha > 1$, the price:dividend ratio is positively related to expectations of the inverse of future consumption growth [see Equation (4)], that is, 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. Therefore 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. Consequently, variation in the price:dividend ratio offsets variation in dividend growth in the equity return, and the correlation in Equation (5) is reduced.

There are two remaining issues. 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)

⁵ 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 trade-off.

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 ARMA 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.

Can alternative specifications of preferences achieve the desired result even when consumption growth follows an ARMA process? Two popular generalizations, habit persistence [Constantinides (1990) and Ferson and Constantinides (1991)] and recursive utility [Epstein and Zin (1989) and Hung (1994)], have been investigated for their ability to match other features of stock return data, particularly the magnitude and volatility of the equity premium. Both approaches permit a separation between the intertemporal elasticity of substitution and the inverse of the relative risk aversion coefficient; while under CRRA utility, these two quantities are equal. The additional flexibility may help in resolving the conflict between a relatively smooth consumption process and a large and variable equity premium, which is at the heart of the equity premium puzzle [Mehra and Prescott (1985)], although Kocherlakota (1990) argues that this flexibility does not substantially increase the explanatory power of the model.

The principle effects of these generalizations are on the volatility of the MRS, not on the correlation between the MRS and equity returns, which is the focus of this investigation. For example, under recursive utility, the MRS depends on both consumption growth and the market return; consequently, covariations with both these quantities determine the risk premium. Such a specification provides little or no additional help in generating time-varying conditional correlations between the MRS and the market return. Under habit persistence, the MRS is modified to depend not on consumption growth but on the growth of weighted differences in consumption, due to the dependence of utility on past levels of consumption. Again, however, time-varying correlations are not a natural feature of the model with standard consumption processes. For example, Campbell and Cochrane (1999) use a model with external habit persistence to match a wide variety of dynamic asset pricing phenomena. Nevertheless, they still generate a monotonic, albeit nonlinear, relation between expected returns and volatility. In many ways, the literature on more general preferences is complementary to the work in this article. Combining these preferences with the consumption process proposed in this article may simultaneously address a variety of puzzles regarding stock market returns.

1.2 Regime shifts

One simple and attractive way to overcome the problems outlined above 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 aggregate economic regimes has gathered increasing empirical support in the literature [see Hamilton (1994, chap. 22) for a survey]. In general, this research provides evidence of multiple regimes within the course of a single business cycle which then repeat in succeeding cycles. 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. This article focuses on consumption data due to the nature of the model, but GNP and industrial production, among other business cycle variables, have also been shown to conform to regime shift specifications. Recent evidence [Sichel (1994)] even suggests the existence of more than two regimes.

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” [see, e.g., Bekaert, Hodrick, and Marshall (1998) and Veronesi (1998)]. Moreover, we model the fundamental process, consumption growth, rather than modeling asset prices or returns directly. For example, Gray (1996), Bekaert, Hodrick, and Marshall (1998), and Ang and Bekaert (1998) estimate regime-shift models for interest rates and Ang and Bekaert (1999) estimate a model for stock returns. The approach in this article is very different in that the fundamental economic process is modeled in a regime-shift framework and asset returns are derived using rational expectations. A similar approach is applied to the bond market in Evans (1998) and Boudoukh et al. (1999). Given that agents rationally anticipate regime shifts in the underlying process, the behavior of equity returns can take on potentially complex, interesting, and realistic characteristics.

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

$$g_{t+1} = \begin{cases} a_1 + b_1 g_t + \epsilon_{1t+1} & \epsilon_{1t+1} \sim N(0, \sigma_1^2) \quad \text{for } I_{t+1} = 1 \\ a_2 + b_2 g_t + \epsilon_{2t+1} & \epsilon_{2t+1} \sim N(0, \sigma_2^2) \quad \text{for } I_{t+1} = 2, \end{cases} \quad (7)$$

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

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

that is, the probability of being in a given regime next period depends only on the current regime and the underlying state variable. This function is parameterized as

$$\begin{aligned} P_{t+1}(1, 1) &\equiv \Pr[I_{t+1} = 1 | I_t = 1, g_t] = \frac{\exp(p_0 + p_1 g_t)}{1 + \exp(p_0 + p_1 g_t)} \\ P_{t+1}(1, 2) &\equiv \Pr[I_{t+1} = 2 | I_t = 1, g_t] = 1 - P_{t+1}(1, 1) \\ P_{t+1}(2, 2) &\equiv \Pr[I_{t+1} = 2 | I_t = 2, g_t] = \frac{\exp(q_0 + q_1 g_t)}{1 + \exp(q_0 + q_1 g_t)} \\ P_{t+1}(2, 1) &\equiv \Pr[I_{t+1} = 1 | I_t = 2, g_t] = 1 - P_{t+1}(2, 2). \end{aligned} \quad (8)$$

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) and Hung (1994). It is similar to the specifications that Gray (1996) uses to estimate the process for short-term interest rates and that Filardo (1994) uses to model the business cycle dynamics of industrial production.

2. Empirical Results

2.1 Data

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 January 1959 to December 1996 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 article. 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 analysis later in the article.

Table 1
Descriptive statistics

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 (0.021)	-0.239 (0.046)	0.381 (0.015)

Descriptive statistics for monthly log consumption growth (in percent) for the period February 1959–December 1996. The AR(1) is estimated using GMM, with heteroscedasticity-consistent standard errors in parentheses. σ denotes the residual standard deviation.

Table 1 provides descriptive statistics for the monthly log consumption growth data (in percent) over the 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 broadly consistent with other results in the literature that study consumption data.

2.2 Estimating a two-regime model

The two-regime model [Equations (7) and (8)] is estimated using the maximum likelihood methodology in Gray (1996). 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.⁶ The parameter values are chosen to maximize this function in the standard manner.

Table 2 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 given that the parameters imply regimes with mean consumption growth of 0.323% and 0.146% . The parameters of both regimes are estimated with good precision, and they are significant at all conventional levels. Table 2 also presents tests for the equality of the parameters

⁶ See the appendix of Gray (1996) for the details concerning construction of the likelihood function. Thanks to Steve Gray for the estimation code that was modified for this application.

Table 2
Parameter estimates for the two-regime model

	a_1	b_1	σ_1	p_0	p_1
Expansion	0.422 (0.034)	-0.307 (0.066)	0.383 (0.017)	3.986 (0.749)	2.783 (3.403)
	a_2	b_2	σ_2	q_0	q_1
Contraction	0.191 (0.036)	-0.312 (0.081)	0.328 (0.020)	4.058 (1.068)	-0.741 (2.989)
Test statistic	20.533 [0.000]	0.002 [0.968]	3.893 [0.048]		

Parameter estimates for the two-regime AR(1) model of log consumption growth given in Equations (7) and (8). The model is estimated by maximum likelihood using monthly data on the consumption of nondurable goods and services from February 1959 to December 1996. Standard errors are in parentheses. Test statistics for the equality of the parameters across the two regimes are also reported, with p -values in brackets.

across the two regimes, with p -values in brackets. Both the mean and volatility of consumption growth are higher in expansions, but the level of mean reversion is almost identical across the regimes. Clearly the model is able to identify two distinct regimes within the time series of consumption data.

Of greatest interest are the parameters which control the regime shifts. While the constants are 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. To illustrate the magnitude of this implied time variation, Figure 1 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. There is a 58% unconditional probability of being in an expansion and a 42% probability of being in a contraction.

Before proceeding to the implications of the estimated parameters for stock market risk and return, we examine the estimation more closely since the results that follow are sensitive to the parameter values (see Section 2.5). First, note that the model is deliberately kept simple in order to illuminate the underlying economic intuition. Such a reduced form model is unlikely to capture all the complexities of the time series of consumption growth. From an economic perspective, the two key questions are whether the data truly indicate the existence of multiple regimes, and whether the transition probabilities are time varying.




Figure 1
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 (8), and the parameter estimates are in Table 2.

In partial answer to the first question, Table 2 shows that the parameter values are statistically different across the regimes. A direct test of a two-regime model versus a single regime model is difficult because, under the null hypothesis of a single regime, the regime shift parameters are not identified. As a result, the likelihood ratio test does not have the standard chi-square distribution. Nevertheless, this test statistic can be used informally to evaluate the specification as in Gray (1996).⁷ The statistic has a value of 19.93, with a corresponding p -value of 0.000 under the $\chi^2(3)$ distribution, which is supportive of the two-regime specification. Note that this informal rejection of the single-regime model is not due exclusively to the existence of different conditional volatilities across the regimes. The test statistic for a two-regime model with constant volatilities (i.e., $\sigma_1 = \sigma_2$) versus the single regime model is 16.82, with a corresponding p -value of 0.000.

A different way to evaluate the goodness-of-fit of the model is to examine the precision with which it identifies the regimes. Ideally the conditional probability of being in either regime should be close to zero or one most of the time, that is, the data should identify the state of the economy with a high degree of certainty. In addition, the identified regimes and transitions should

⁷ A formal test can be constructed using a grid search over the nuisance parameters [see Hansen (1992)], but it is prohibitively computationally intensive.




Figure 2
Regime identification and consumption growth
Panel A shows the estimated conditional probability of being in an expansion based on the two-regime model. Panel B shows 9-month moving averages of consumption growth. In both graphs, vertical solid and dashed lines mark NBER business cycle peaks and troughs, respectively.

correspond to economic intuition. Figure 2 presents evidence to this effect. Panel A graphs the time series of $\text{Pr}(I_t = 1|\Phi_{t-1})$, that is, the conditional probability of being in regime 1. For reference purposes, the NBER peaks and troughs of the business cycle are marked by solid and dashed vertical

lines, respectively. While the probability series is not exceptionally smooth, it identifies the NBER cycles accurately, with two exceptions.⁸

First, the estimation fails to pick up the recession of 1970. Second, the model has difficulty identifying the post-1991 period as an expansion. The explanation for both results is clear when looking at the underlying consumption growth data. Panel B shows 9-month moving averages (to smooth the data) against the same NBER turning points. Consumption growth around 1970 shows no contraction-like behavior, hence the estimation fails to isolate this period. Similarly the recent expansion is weak by historical standards (with mean monthly consumption growth of 0.18% relative to 0.32% in the five previous expansions), so it is again difficult to identify. Overall the regime-shift model performs excellently.

The issue of time-varying transition probabilities is somewhat less clear-cut. While the point estimates in Table 2 are consistent with this interpretation, the standard errors are large. One explanation is that the reduced form model may be too simple to fully capture the regime dynamics. For example, Filardo (1994) provides strong evidence of state-dependent transition probabilities in the cyclical process for industrial production using a more elaborate model. When the transition probabilities are allowed to depend on other exogenous variables such as the index of leading economic indicators, constant transition probabilities can be rejected statistically. Moreover, the resulting model provides a superior fit to the data.

2.3 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, multiregime specifications, closed-form solutions are no longer available; therefore 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 follow Tauchen and Hussey (1991).⁹ The key point, from the perspective of examining the risk/return trade-off, 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. Throughout the analysis we use nine consumption growth states within each regime.

To analyze stock market risk and return we also need to specify the degree of risk aversion and the time preference parameter. We use $\alpha = 2$ and $\beta = 0.997$, and the sensitivity of the results to these parameters is addressed later.

⁸ The probability can easily be smoothed by constructing $\Pr(I_t = 1 | \Phi_T)$, that is, the state probability given the full dataset. The resulting series is less jagged, but it generates similar inferences.

⁹ Thanks to George Tauchen for the discrete approximation code that has been modified for this application.

The two-regime specification has a total of 18 states of the world, 9 within each regime. The nine 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 3 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. The states are indexed from lowest to highest consumption growth. For ease of interpretation, 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 P:E ratios calculated using annual earnings. In addition, the risk premium is multiplied by 100 for presentation purposes.

The conditional moments of returns exhibit dramatic nonmonotonicities as shown in Figure 3, which 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 is the weak relation between volatility and the risk premium. In the contractionary regime, the risk premium and volatility are negatively related. In the expansion, a positive relation holds for states 5–9, but even in this limited set the relation is nonlinear.

Table 3
Risk and return in a two-regime model

State	g_t	s_t/c_t	$E_t[r_{st+1} - r_{ft}]$ (% $\times$ 100)	$\text{Vol}_t[r_{st+1}]$ (%)	$P_t(i, j)$	Probability
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

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 2. All values except for consumption growth are annualized.




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 2. Expansion and contraction states are marked by circles and squares, respectively.

These results are in marked contrast to those generated from a single-regime model. Figure 4 shows the the state-by-state volatility and risk premium for the model based on the AR(1) estimates in Table 1 (marked by triangles). 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 trade-off at the market level in a dynamic CAPM setting [see Merton (1980)]. The other major differences between the two models are the increase in the variability of the risk premium and the higher volatility associated with the two-regime model.

The unconditional relation is difficult to ascertain from the graph due to the differing probabilities associated with each state. For example, state 5 in the expansion has an unconditional probability of 22.5%, while states 1 and 9 in both regimes have probabilities of less than 0.01%. 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 -0.481 , which coincides with the empirical results in Glosten, Jagannathan, and Runkle (1993) and Whitelaw (1994). Both of these articles 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.




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 2 and 1, respectively. Expansion states are marked by circles, contraction states are marked by squares, and the single-regime states are marked by triangles.

It is tempting to attribute this negative relation between the mean and volatility of returns to the extreme values observed in certain states of the world. For example, states 1 and 2 in the expansion have high regime-shift probabilities, high volatilities, and low expected returns. Note, however, that the unconditional probability of those states is small so they contribute little to the unconditional moments. One way to verify this conjecture is to set the transition probabilities to zero in the four most extreme consumption growth states in each regime (states 1, 2, 8, and 9). The resulting unconditional correlation is -0.498 , little different from the previous result. In other words, the observed behavior is not being driven by the tails of the distribution.

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. 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 4 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 consumption growth in each state is identical to the values given in Table 3. 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. It is this simple feature that generates the key results. The market claim may act as a hedge against shifts in investment opportunities, that is, shifts from one regime to the other. Relative to dividends, prices are high when investment opportunities are poor, and vice versa.

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

Table 4
Risk and return in single-regime models

State	s_t/c_t	$E_t[r_{st+1} - r_{ft}]$ (% $\times 100$)	$\text{Vol}_t[r_{st+1}]$ (%)
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

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 2. All values are annualized.

zero probability of ever entering a contraction. In Table 3, 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 (4), 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 3. 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 4 is preserved, albeit in a weakened form.

The price:dividend ratio 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 (3)]. Note first that the variation in price:dividend ratios, especially across the regimes, tends to be larger than the variation in consumption growth. Table 3 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 3.

The second issue is the correlation between equity returns and the MRS (see Section 1.1). In other words, does the market claim provide a hedge against consumption risk? 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, that is, price:dividend ratios are higher in the contraction than in the expansion. Consequently, a shift from contraction to 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 more than 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.¹⁰ 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.

We are also interested in potential time variation in the relation between the risk premium and volatility. 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. This correlation captures both the sign and the magnitude of the relation between the conditional moments. 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 18 possible states) at time t . Next period (time $t + 1$), the economy can be in any of the 18 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.

The contrast between the single-regime model and the two-regime model is equally apparent when considering these conditional correlations between the risk premium and the volatility. For the single-regime model the conditional correlations are 1.000 in every state. The short-run relation exhibits no time variation. For the two-regime model, 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 range 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. 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, e.g., Whitelaw (1994) and Boudoukh, Richardson, and Whitelaw (1997)], but the absence of positive correlations is not. This question is addressed in more detail in Section 2.5.

2.4 Other implications of the model

While the focus of this article is on the relation between the mean and volatility of stock market returns, it is interesting to investigate other implications of

¹⁰ Boudoukh, Richardson, and Whitelaw (1997) make a similar point in the context of a simple, four-state, discrete economy.




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 2. The x -axis gives the number of the state using the scheme in Table 3. Expansion and contraction states are marked by circles and squares, respectively.

the model, both to consider further testable restrictions and to determine the reasonableness of the parameterization. This latter concern is well founded given that attempts to resolve the equity premium puzzle have sometimes led to models that produce startlingly unrealistic interest rate processes. Such an outcome is not totally surprising since a common approach is to attempt to increase the volatility of the MRS [see Equation (5)], which also determines the volatility of interest rates.

Intuitively, working on the correlation between current and future consumption growth, as in this article, will have a less dramatic effect on interest rates. This intuition is correct. The mean monthly real risk-free rate is essentially the same in the two-regime model as in the single-regime model. As expected, the volatility is somewhat higher (0.29% per month versus 0.19% per month), but not unreasonably so. The patterns in the volatility of excess returns are similar—0.67% in the two-regime model and 0.46% in the single-regime model. The ordering of the mean excess returns is reversed because much of the induced regime risk is unpriced.

This relative similarity in unconditional moments contrasts markedly with the dramatic differences in the conditional moments, as shown in Figures 3 and 4. Consequently it is in the conditional moments that we should look for the strongest implications of the model. One interesting implication is that the conditional moments of bond returns should exhibit patterns similar

to those seen in stock returns because the discount rate effect dominates the dividend effect (see Section 1.1). Counter to the intuition of the liquidity preference hypothesis, long-term bond returns should be volatile in periods of low term premia associated with business cycle phase transitions. Evidence in Boudoukh et al. (1999) is consistent with this prediction.

Finally, much of the evidence in the literature on the conditional moments of bond and stock returns deals with predictability using various financial variables. In the two-regime model, there are only two relevant state variables: the level of consumption growth and the state of the economy. As a result, such variables as price:dividend ratios (or equivalently dividend yields) and interest rates have predictive power for both the mean and volatility of stock and bond returns. Of equal importance are proxies for the current and future state of the economy, especially around regime transitions. While investigating possible variables is beyond the scope of this article, the intuition is consistent with the use in the literature of such forward-looking variables as the slope of the term structure or the credit yield spread.

2.5 Sensitivity analysis

The purpose of the sensitivity analysis is twofold: (1) to examine the robustness of the results of the two-regime model to changes in the parameters and (2) to find the parameters 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. These correlations conveniently summarize the direction and strength of the relation between the conditional moments, at least in a linear context.

Initially, consider the time preference parameter β . Increasing β toward 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.

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. Of course, raising the level of risk aversion also increases the magnitude of the risk premium, but it has less plausible effects on the properties of interest rates.

With respect to the within-regime time-series properties of consumption growth, there are six parameters and innumerable variations of 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, that is, 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 3 and 4. 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 upward 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 5 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

Table 5
Risk and return in a two-regime model

State	$E_t[r_{st+1} - r_{ft}]$ (% $\times$ 100)	$\text{Vol}_t[r_{st+1}]$ (%)	$P_t(i, j)$	Probability	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

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 2 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.

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.

3. Conclusion

This article 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 relation between these moments of returns and substantial time variation in this relation. This article 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.

These insights into the risk/return relation at the market level may be a precursor to a better understanding of asset pricing in a variety of markets. For example, the large changes in investment opportunities implied by the regime-shift model and the implied hedging demands are also likely to have strong implications for bond prices and returns. Moreover, given the somewhat surprising results at the market level, the model may have interesting implications for the cross-section of expected equity returns, a topic that has been studied extensively in the literature.

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 article is inconsistent with many of these empirical specifications. One potential correction is to model expected returns in a multifactor framework, with conditional volatility as one of the factors

and other factors proxying for changes in the investment opportunity set [e.g., Scruggs (1998)]. 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, 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.

Moreover, further research is needed on the interaction between asymmetric consumption processes and alternative preferences such as recursive utility. Hung (1994), for example, shows that a simple Markov model in combination with unexpected utility can match the unconditional moments of equity returns and interest rates when the dividend process of the stock market is permitted to differ from the aggregate consumption process.

Throughout the article 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. However, it would be worthwhile to estimate a multiregime 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.

References

- Abel, A. 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.
- Ang, A., and G. Bekaert, 1998, "Regime Switches in Interest Rates," Working Paper 6508, NBER.
- Ang, A., and G. Bekaert, 1999, "International Asset Allocation with Time-Varying Correlations," working paper, Stanford University.
- Backus, D., and A. Gregory, 1993, "Theoretical Relations Between Risk Premiums and Conditional Variances," *Journal of Business and Economic Statistics*, 11, 177–185.
- Bekaert, G., R. Hodrick, and D. Marshall, 1998, "Peso Problem Explanations for Term Structure Anomalies," working paper, Stanford University, Columbia University, and Federal Reserve Bank of Chicago.
- Boudoukh, J., M. Richardson, and R. F. Whitelaw, 1997, "Nonlinearities in the Relation Between the Equity Risk Premium and the Term Structure," *Management Science*, 43, 371–385.

- Boudoukh, J., M. Richardson, T. Smith, and R. F. Whitelaw, 1999, "Regime Shifts and Bond Returns," working paper, New York University.
- Campbell, J. Y., 1987, "Stock Returns and the Term Structure," *Journal of Financial Economics*, 18, 373–399.
- Campbell, J. Y., and J. H. Cochrane, 1999, "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107, 205–251.
- Cecchetti, S. G., P.-S. Lam, and N. C. Mark, 1990, "Mean Reversion in Equilibrium Asset Prices," *American Economic Review*, 80, 398–418.
- Constantinides, G., 1990, "Habit Formation: A Resolution of the Equity Premium Puzzle," *Journal of Political Economy*, 98, 519–543.
- De Santis, G., and S. Imrohroglu, 1997, "Stock Returns and Volatility in Emerging Financial Markets," *Journal of International Money and Finance*, 16, 561–579.
- Epstein, L., and S. Zin, 1989, "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57, 937–970.
- Evans, M. D. D., 1998, "Regime Shifts, Risk and the Term Structure," working paper, Georgetown University.
- Ferson, W., and G. Constantinides, 1991, "Habit Persistence and Durability in Aggregate Consumption: Empirical Tests," *Journal of Financial Economics*, 29, 199–240.
- Filardo, A. J., 1994, "Business Cycle Phases and Their Transitional Dynamics," *Journal of Business and Economic Statistics*, 12, 299–308.
- French, K. R., G. W. Schwert, and R. F. Stambaugh, 1987, "Expected Stock Returns and Volatility," *Journal of Financial Economics*, 19, 3–29.
- Glosten, L. R., R. Jagannathan, and D. 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.
- Gray, S. F., 1996, "Modeling the Conditional Distribution of Interest Rates as a Regime-Switching Process," *Journal of Financial Economics*, 42, 27–62.
- Hamilton, J. D., 1989, "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle," *Econometrica*, 57, 357–384.
- Hamilton, J. D., 1994, *Time Series Analysis*, Princeton University Press, Princeton, N.J.
- Hansen, B. E., 1992, "The Likelihood Ratio Test Under Non-Standard Conditions: Testing the Markov Switching Model of GNP," *Journal of Applied Econometrics*, 7, S61–S82.
- Hansen, L. P., 1982, "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 50, 1929–1954.
- Hansen, L. P., and K. Singleton, 1982, "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models," *Econometrica*, 50, 1269–1286.
- Harrison, P., and H. H. Zhang, 1999, "An Investigation of the Risk and Return Relation at Long Horizons," *Review of Economics and Statistics*, 81, 1–10.
- Hung, M.-W., 1994, "The Interaction Between Unexpected Utility and Asymmetric Market Fundamentals," *Journal of Finance*, 49, 325–343.
- Kandel, S., and R. F. Stambaugh, 1990, "Expectations and Volatility of consumption and Asset Returns," *Review of Financial Studies*, 3, 207–232.
- Kocherlakota, N. R., 1990, "Disentangling the Coefficient of Relative Risk Aversion from the Elasticity of Intertemporal Substitution: An Irrelevance Result," *Journal of Finance*, 45, 175–190.
- Lucas, R. E., Jr., 1978, "Asset Prices in an Exchange Economy," *Econometrica*, 46, 1429–1445.

- Marshall, D. A., 1993, "Asset Return Volatility with Extremely Small Costs of Consumption Adjustment," working paper, Federal Reserve Bank of Chicago.
- Mehra, R., and E. Prescott, 1985, "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 15, 145–161.
- Merton, R. C., 1973, "An Intertemporal Capital Asset Pricing Model," *Econometrica*, 41, 867–887.
- Merton, R. C., 1980, "On Estimating the Expected Return on the Market," *Journal of Financial Economics*, 8, 323–361.
- Nelson, D. B., 1991, "Conditional Heteroscedasticity in Asset Returns: A New Approach," *Econometrica*, 59, 347–370.
- Pagan, A. R., and Y. S. Hong, 1991, "Nonparametric Estimation and the Risk Premium," in W. Barnett, J. Powell, and G. Tauchen (eds.), *Nonparametric and Semiparametric Methods in Econometrics and Statistics*, Cambridge University Press, Cambridge, UK.
- Scruggs, J. T., 1998, "Resolving the Puzzling Relation Between the Market Risk Premium and Conditional Market Variance: A Two-Factor Approach," *Journal of Finance*, 53, 575–603.
- Sichel, D. E., 1994, "Inventories and the Three Phases of the Business Cycle," *Journal of Business and Economic Statistics*, 12, 269–277.
- Tauchen, G., and R. Hussey, 1991, "Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models," *Econometrica*, 59, 371–396.
- Veronesi, P., 1998, "Regime Shifts and Changing Volatility in Stock Returns: A Rational Expectations Equilibrium Model," working paper, University of Chicago.
- Whitelaw, R. F., 1994, "Time Variations and Covariations in the Expectation and Volatility of Stock Market Returns," *Journal of Finance*, 49, 515–541.