Regime Shifts and Bond Returns

Jacob Boudoukh[a], Matthew Richardson[a], Tom Smith[b],

and Robert F. Whitelaw[c*]

November 1, 1999

*[a]Stern School of Business, New York University and the NBER; [b]Australian Graduate School of Management, University of New South Wales; [c]Stern School of Business, New York University. Thanks to Frank Diebold and seminar participants at the WFA meetings and Cornell University for helpful comments. Contact: Robert F. Whitelaw, 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. The current version of this paper can be found at http://www.stern.nyu.edu/~rwhitela/.

Regime Shifts and Bond Returns

Abstract

This paper investigates the implications of a 2-regime model of the business cycle for term premiums and volatilities in the bond market. The model, which is estimated via maximum likelihood using GDP, consumption and production data, has two key features – mean growth rates that vary across regimes and time-varying transition probabilities between regimes. The implied dynamics of term premiums and volatilities are complex and interesting. Business cycle turning points are characterized by high volatility and strongly time-varying term premiums. These implications are then investigated using data on bond returns. Nonparametric estimation results are broadly consistent with the model. Using the slope of the term structure as a conditioning variable, we can identify periods with negative term premiums and volatile returns.

1 Introduction

There is increasing evidence in the literature that the U.S. economy is characterized not by a single stationary process for real economic variables such as GNP, industrial production and consumption, but by periodic shifts between distinct regimes.[1] These regime-shift models have the following empirical properties: (i) different mean growth rates across the regimes describing the aggregate economic series (e.g., expansions versus contractions), (ii) persistence of the regimes, and (iii) minimal persistence in the growth rates of the aggregate series within regimes. Taking this description of the aggregate economy as given, what are the implications for the expected returns and the volatilities of returns on financial assets?

This question is particularly important because most rational, equilibrium, asset pricing models have had only limited success at linking the prices of financial assets to the underlying aggregate economic series.[2] Empirically, movements in the conditional means and volatilities of asset returns seem too large to be explained by movements in the aggregate economy. In particular, the lack of persistence in aggregate economic growth rates, coupled with reasonably chosen risk aversion levels, dooms the model, preventing it from matching key characteristics of the data. Therefore, it may seem unlikely that a regime-shift driven economy, with i.i.d. (or near i.i.d) growth rates, would add much to the discussion.

This paper takes a first, and we believe promising, look at the relation between bond prices and a regime-shift driven aggregate economy. Specifically, we examine the implications of a regime-shift economy for the the mean and volatility of bond returns across different maturities and generate several thought-provoking results.

First, and most important, the conditional mean and volatility of bond returns (across

[1]The notion that the business cycle consists of distinct phases goes at least as far back as Burns and Mitchell (1946). See Hamilton (1989, 1994) and Filardo (1994) for recent examples of regime-shift models applied to aggregate economic series, and see Diebold and Rudebusch (1996) for an excellent survey of the literature.

[2]The more successful of these models have relied on either market imperfections (see Polkovnichenko (1998) and Guo (1999) for recent examples) or non-standard preferences (see Campbell and Cochrane (1999) for a recent example).

maturities) swing wildly within this economy, even if economic growth rates are i.i.d. (within regimes) and risk aversion levels are *low*. This occurs because the regimes offer different economic environments and the regimes themselves are persistent. These results contrast markedly with those from a standard linear time series model of the aggregate economy, which generates little variation in either term premiums or volatilities.

Second, the model provides definitive implications for when these swings in the mean and volatility of returns take place. In particular, we find that, theoretically, the largest changes should occur around turning points in the economy, i.e., in states of the world with high probabilities of regime shifts. If these probabilities are empirically related to the peaks and troughs of business cycles, as some previous authors have found, then this would suggest a relation between the mean and volatility of returns and the aggregate economy.

Third, these model implications are examined using data on U.S. government bonds and aggregate U.S. economic series. We find strong evidence that changes in both the conditional mean and volatility of returns on bonds, as well as the term structure of these conditional moments, coincide with the theory. From an empirical point of view, in other words, aggregate economic variables (albeit via a complex functional relation) have good explanatory power for the conditional moments of returns.

While the literature on regime shifts focuses primarily on business cycle modeling and understanding the functioning of the real economy, this paper shows that regime shifts also have strong implications for the pricing of financial assets. Specifically, when the phases of the business cycle are viewed as persistent shifts in the underlying economy rather than unexpected shocks around a long-run stationary process, investors' rational expectations are affected accordingly. Transitions between expansions and contractions are seen as having potentially large and relatively long-lived effects on asset payoffs and discount rates. In terms of the ICAPM, switches between regimes have dramatic effects on the investment opportunity set. Consequently, assets are priced based partly on how they covary with these changes in investment opportunities. In other words, multiple regimes introduce an additional state variable to the investor's optimization problem – what regime is the economy in and what regime is it likely to be in, in the future? However, this variable also enters in a highly nonlinear way, leading to complex asset return dynamics.

The remainder of the paper is organized as follows. In Section 2, we provide the asset pricing framework and the relevant economic intuition associated with bond pricing in a regime-shift driven economy. In Section 3, a two-regime model is estimated via maximum likelihood, using aggregate data. In brief, the model produces dramatic time-variation in the maturity structure of term premiums, ranging from positive and upward sloping to humped to negative and downward sloping. Moreover, variation in term premiums is not related directly to the conditional volatility of returns, yet varies systematically over the business cycle, with the largest changes occurring at cycle turning points. In Section 4, the paper examines these implications using return data on U.S. Treasury securities. Several measures of business cycle turning points are employed, including the NBER business cycle dating scheme and the slope of the term structure of interest rates. A simple nonparametric analysis shows that these variables identify periods of large swings in both the mean and the volatility of the term structure of bond returns. An interesting side result is that the standard link between the risk and return of financial assets is broken as volatility is no longer a sufficient statistic for the asset's priced risk. Instead, the peaks and troughs of the business cycle exhibit quite different behavior in the term structures of the means and volatilities of bond returns. As such, the two-regime model of the aggregate economy provides new, and important, insights into the behavior of bond returns.

2 The Theory

Consider the standard, no-arbitrage, complete markets model of Harrison and Kreps (1979), in which all assets can be priced via a unique stochastic discount factor, i.e,

$$p_t = \mathrm{E}_t \left[m_{t,t+1} (p_{t+1} + y_{t+1}) \right],$$

where p_t is the price of the real asset, y_{t+1} is the asset's payout of the real good, and $m_{t,t+1}$ is the real stochastic discount factor or pricing kernel. For example, in the context of a representative agent model, this discount factor is the representative agent's marginal rate of substitution. The above equation can be rewritten in nominal terms as

$$p_t^* = \mathrm{E}_t \left[m_{t,t+1}(p_{t+1}^* + y_{t+1}^*)\pi_{t,t+1}{}^{-1} \right], \tag{1}$$

where $\pi_{t,t+1}$ is the inflation rate on the real good, and the * denotes the nominal price.

Therefore, the resulting price at time t of a τ-period, zero-coupon bond that pays \$1 is

$$q_t^*(\tau) = \mathrm{E}_t[m_{t,t+\tau}\pi_{t,t+\tau}{}^{-1}] \tag{2}$$

where

$$m_{t,t+\tau} = m_{t,t+1}m_{t+1,t+2}\cdots m_{t+\tau-1,t+\tau}. \tag{3}$$

The 1-period, gross return (t to $t+1$) on this τ-period bond is

$$r_{t,t+1}^*(\tau) = \frac{q_{t+1}^*(\tau-1)}{q_t^*(\tau)} \tag{4}$$

$$= \frac{\mathrm{E}_{t+1}[m_{t+1,t+\tau}\pi_{t+1,t+\tau}{}^{-1}]}{\mathrm{E}_t[m_{t,t+\tau}\pi_{t,t+\tau}{}^{-1}]} \tag{5}$$

If we further assume that the stochastic discount factor, $m_{t+1,t+\tau}$, and the inflation rate, $\pi_{t,t+\tau}$, are uncorrelated, then it is possible to derive an expression for expected real returns on nominal bonds in terms of only real variables, i.e., inflation plays no role.[3] Specifically, using this assumption, equation (1) and equation (5), and substituting for the expectation of the product of the discount factor and the asset return and applying some algebra, the expected real return on a τ-period nominal bond can be written as

$$\mathrm{E}_t[r_{t,t+1}^*(\tau)\pi_{t,t+1}^{-1}] = r_{ft} - r_{ft}\,\mathrm{Cov}_t[m_{t,t+1}, r_{t,t+1}(\tau)] \tag{6}$$

$$= r_{ft} - \frac{r_{ft}}{q_t(\tau)}\mathrm{Cov}_t[m_{t,t+1}, q_{t+1}(\tau-1)] \tag{7}$$

$$= r_{ft} - \frac{r_{ft}}{q_t(\tau)}\mathrm{Cov}_t[m_{t,t+1}, \mathrm{E}_{t+1}[m_{t+1,t+\tau}]], \tag{8}$$

[3]The assumption that inflation is statistically independent of the real discount factor is strong. Though it is a common assumption in the bond pricing literature (e.g., see Cox, Ingersoll and Ross (1985) and Gibbons and Ramaswamy (1993)), we make the assumption here to isolate the real component of bond pricing. Extensions of this assumption are discussed in Section 5 of the paper.

where r_{ft} and $q_t(\tau)$ are the real rates and prices on indexed (i.e., real) bonds.[4] The expected excess real return on the bond is proportional to the conditional covariance between the discount factor next period and the expectation next period of the discount factor further in the future. If the discount factor is positively (negatively) autocorrelated, then the covariance is positive (negative) and the term premium (expected excess return) is negative (positive).

To put an economic interpretation on this relation, consider an economy with a representative agent with CRRA utility. In this case, the discount factor is the marginal rate of substitution:

$$m_{t,t+1} = \beta \left(\frac{F_{t+1}}{F_t} \right)^{-\alpha}, \tag{9}$$

where β is the subjective time discount factor, F_t is the real level of the aggregate economic factor such as aggregate consumption or production (see Lucas (1978) or Cochrane (1992)), and α is the coefficient of relative risk aversion.

For this specification, the discount factor essentially inherits the autocorrelation properties of economic growth. If economic growth is positively autocorrelated, then the term premiums on bonds are negative. Investors are willing to buy securities that have expected returns lower than the risk-free rate because these securities provide a hedge against economic risk. A negative shock to economic growth next period means that growth rates will also be lower in the future. Times will be bad and payoffs in these states of the world are extremely valuable. However, bad times (i.e., low growth) also mean low future interest rates. Investors want to save for the low economic growth periods to come; therefore, bond prices are high. A decrease in interest rates generates high returns on bonds, which hedge the negative shock to economic growth.

If the logarithm of economic growth is modeled as a stationary ARMA process (as is often

[4]Note that the real, risk-free rate, i.e., the certain 1-period real return on a 1-period bond that pays in goods, is

$$r_{ft} \equiv r_{t,t+1}(1) = \frac{1}{q_t(1)} = \frac{1}{\mathrm{E}_t[m_{t,t+1}]}.$$

The result in (7) follows from this equation.

the case), then the scaled covariance in equation (9) is approximately constant over time, i.e., term premiums do not vary. In other words the conditional relation between next period's economic growth and economic growth further in the future is fixed by the parameters of the model.

In contrast, a model of regime shifts with time-varying transition probabilities has the potential to break this link. The autocorrelation of economic growth will depend on both the process within each regime and the probability of switching regimes. For example, assume that economic growth is positively autocorrelated within an expansionary (high mean economic growth) regime, i.e., conditional on remaining in the regime a positive shock to economic growth next period implies high economic growth further in the future. However, assume further that the probability of moving to the contraction (the low mean economic growth regime) is positively related to the level of economic growth. Depending on the exact parameterization, a positive shock to the economy may now be associated with lower expectations of economic growth in the future due to the associated increase in the probability of entering a contraction.

The implications of regime shifts for asset prices and returns have begun to be explored in the financial economics literature. For example, Gray (1996) successfully models interest rates as following a two-regime process with time-varying transition probabilities. While the paper does not directly address the link between interest rates and the underlying economy, it does suggest such an interpretation. Cecchetti, Lam, and Mark (1990) use a simple two-regime model of consumption and dividends to assess the implications of regime shifts for equity returns. Specifically, they show that even if neither regime is mean-reverting, the switches between regimes lead to serial correlation in equity returns. In other words, the two-regime model produces time-varying expected returns in a rational expectations setting. Whitelaw (1998) employs a more complex specification with time-varying transition probabilities to investigate the puzzling weak empirical relation between expected returns and volatility in the stock market. The effects of regime switches on volatility can generate this result in equilibrium.

A natural complement to the work of Cecchetti et al (1990) and Whitelaw (1998) is to look at the implications of regime shifts in the real economy for bond returns. Specifically, since multi-regime models have their strongest implications for the conditional moments of returns and their time-variation, this paper investigates term premiums (expected excess bond returns) and their volatilities in a two-regime economy. In a representative agent economy, the logarithm of the growth rate of the aggregate economic factor is modeled as a two-regime i.i.d. switching process, where the transition probabilities between the regimes also depend on the level of the economy. Bonds are then priced using the first order condition of a representative agent with CRRA utility. As such, the model is an extension of the log-linear model that has been used extensively in the literature. (See, for example, Hansen and Singleton (1983) and Breeden (1985). Cox, Ingersoll, and Ross (1985) provide the continuous-time counterpart.) One implication of the log-linear model is that the term premium is constant (for a given maturity); therefore, extensions have been investigated in the literature (e.g., stochastic volatility (Boudoukh (1993)). The extension to multiple regimes in this paper also produces time-varying term premiums, in addition to generating interesting predictions about the relation between term premiums and the conditional volatilities of bond returns.

3 A Two-Regime Model

It is well documented that, from time to time, aggregate economic series go through apparently dramatic distributional changes, which some researchers treat as structural breaks in the data. Intuitively, the onset of war, changes in government policy, technological shifts, and systemic financial collapse are examples of episodes which have profound effects on the economy. One convenient way of modeling these structural breaks, and their tendency to re-occur, is via a regime-switching process. For example, in an early, seminal study of this issue, Hamilton (1989) models GNP as a simple, two-regime process with constant transition probabilities. While both the time series processes within the regimes and the regime switch probabilities are modeled in a parsimonious manner, this model provides a better fit to the data than a more complex, single-regime ARMA model. Moreover, the parameters of the two regimes lead to a natural interpretation of the regimes as the expansionary and contractionary phases of the business cycle.

Subsequent research has extended the multi-regime model to other economic series and more general specification.[5] One noteworthy extension is provided by Filardo (1994), who models industrial production as a two-regime process with time-varying transition probabilities between regimes. Specifically, the probability of switching from an expansion to a contraction, or vice versa, depends on the level of various (exogenous) state variables. Not only does this generalization provide a superior fit, but it also generates more complex and interesting dynamic behavior of the real economy. As a result, the model is able, ex ante, to identify the turning points of the business cycle, as labeled, ex post, by the NBER.[6]

3.1 Model Specification and Estimation Methodology

The proposed model for aggregate economic growth, following Filardo (1994), is a two-regime i.i.d. model with time-varying transition probabilities.[7] Specifically, define the natural logarithm of the economic growth factor,

$$g_{t+1} \equiv \ln(F_{t+1}/F_t) \tag{10}$$

then the model is specified as

$$g_{t+1} = \begin{cases} a_1 + \epsilon_{1t+1} & \epsilon_{1t+1} \sim N(0, \sigma^2) \quad \text{for } I_{t+1} = 1 \\ a_2 + \epsilon_{2t+1} & \epsilon_{2t+1} \sim N(0, \sigma^2) \quad \text{for } I_{t+1} = 2 \end{cases} \tag{11}$$

[5]See Hamilton (1994, Chapter 22) and Diebold and Rudebusch (1996) for surveys of the literature.

[6]There is some debate in the literature as to whether it is possible to statistically reject ARMA models of aggregate economic series in favor of regime-shift models. However, recent evidence from papers that consider multivariate specifications in the context of a dynamic factor model and also permit duration dependence in regimes is supportive of the regime-shift specification (see, for example, Kim and Nelson (1998a,b)).

[7]Extensions to more elaborate specifications, such as ARMA processes for the growth rate as well as the transition probabilities, are discussed in Section 5.

where I_{t+1} indexes the regime. The transition probabilities between regimes are parameterized as

$$
\begin{aligned}
P_{t+1}(1,1) &\equiv \Pr[I_{t+1}=1|I_t=1,g_t] = \tfrac{\exp(p_0+p_1 g_t)}{1+\exp(p_0+p_1 g_t)} \\
P_{t+1}(1,2) &\equiv \Pr[I_{t+1}=2|I_t=1,g_t] = 1-P_{t+1}(1,1) \\
P_{t+1}(2,2) &\equiv \Pr[I_{t+1}=2|I_t=2,g_t] = \tfrac{\exp(q_0+q_1 g_t)}{1+\exp(q_0+q_1 g_t)} \\
P_{t+1}(2,1) &\equiv \Pr[I_{t+1}=1|I_t=2,g_t] = 1-P_{t+1}(2,2)
\end{aligned}
\tag{12}
$$

Contingent on the data, the regimes can be thought of as the expansionary and contractionary phases of the real business cycle. Within each regime log economic growth follows an i.i.d process with normally distributed innovations. The means of the innovations, however, are all allowed to vary across regimes. The probability of switching between regimes in any period is determined by the level of economic growth. The functional form ensures that these probabilities are bounded between zero and one. The parameters p_1 and q_1 govern the time-variation in these transition probabilities. If the parameters are positive (negative) then the probability of a regime switch decreases (increases) as economic growth increases.

At first glance, the two-regime model specified above appears difficult to estimate due to the importance of the unobserved (by the econometrician) state variable – the regime at any point in time. Conditional on this state variable, economic growth follows an i.i.d process, but unconditionally the distribution is much more complex, especially because of the transition probabilities that depend on the level of the growth rate. However, Gray (1996) shows that by reparameterizing the model in terms of the probability of being in a given regime at each date rather than the transition probability, it is possible to construct the likelihood function in a recursive manner similar to that of a GARCH model. In the two-regime model, the unobserved regime state variable is analogous to the unobserved volatility parameter in GARCH models. The details are provided in Appendix A.

3.2 Data and Estimation Results

The model in Section 3.1 can be estimated using a variety of economic series as proxies for the aggregate economy and the resulting stochastic discount factor. From a practical perspective, all of the candidate data series suffer from a number of problems that have been discussed in the literature, including measurement errors and time-aggregation bias. We focus primarily on real GDP – a series that clearly reflects the aggregate fluctuations inherent in the phases of the business cycle. However, we also examine industrial production and consumption of nondurables and services. The GDP and consumption series are reported quarterly and span the period from the first quarter of 1953 to the last quarter of 1998. The industrial production data are monthly from January 1954 to December 1998. Earlier data are excluded due to measurement issues arising from changes in calculation methodologies. All the series are from the Basic Economics database.

Table 1, Panel A provides descriptive statistics for the log growth rates of the three series. In each case, real growth averages between 3% and 4% on an annualized basis. However, GDP growth is approximately twice as volatile as consumption growth, although it is less volatile then production growth. All three series exhibit small and positive first-order autocorrelations. This positive autocorrelation suggests positive autocorrelations in the stochastic discount factor and negative term premiums (see Section 2).

Table 1, Panel B presents the parameter estimates for the two-regime model described in Section 3.1 (with standard errors in parentheses), using the three series. Note that the regimes are denoted as “Expansion” and “Contraction” to reflect the relative levels of mean growth within the regimes. In each case, mean growth in the expansion greatly exceeds that in the contraction. Moreover, these parameters are estimated quite precisely, so it is easy to reject the hypothesis that growth is the same in both regimes. The estimated within-regime conditional volatilities are also estimated well. They are lower than the corresponding unconditional volatilities, as expected, and the relative magnitudes are preserved.

The transition probability parameters are estimated with somewhat lower precision. The coefficients of greatest interest are those on lagged growth since they determine the magnitude and direction of time-variation in transition probabilities. For all three series, the point estimates indicate that the probability of a transition from an expansion to a contraction is decreasing in the level of real growth. In other words, high growth is associated with a low probability of entering the contraction regime. This effect will induce positive autocorrelations in the growth series, and it is strongest for consumption and weakest for industrial production. For consumption and GDP, the coefficient in the contraction has the same sign, although it is smaller in magnitude. The implication is that high growth in a contraction is associated with a lower probability of switching to the expansion, resulting in negative autocorrelations in the growth series. For industrial production, the coefficient is negative; therefore, conditional autocorrelations should be positive in both regimes.

Figure 1 shows the estimated probability of switching regimes conditional on the current regime and the level of GDP growth. For the expansion (solid line), the large magnitude of the coefficient p_1 implies high sensitivity to the level of lagged growth. As a result, the switch probability rises dramatically from close to zero to almost 100% as growth drops below 0%. In contrast, q_1 is much smaller yet still positive,; therefore, the probability of switching from the contraction to the expansion (dashed line) rises less quickly as growth falls.

While the parameter values discussed above are broadly consistent with our intuition regarding business cycle effects, it is worthwhile to look more closely at how well the 2-regime model captures aggregate fluctuations. To that end, the top graph of Figure 2 presents the estimated ex ante probabilities, using GDP data, of being in regime 1 (the expansion) against the NBER business cycle peaks and troughs.[8] Clearly, the model does an excellent job of identifying expansions and contractions. The ex ante probabilities are most often near 1 or 0, indicating that the model identifies the regime at any point in time with a good deal of precision. Moreover, the model provides ex ante dating of the cycle in contrast to the ex post dating employed by the NBER. The reason for the good performance of the model is apparent in the bottom graph, which shows a 5-quarter moving average of GDP growth. Although growth is quite variable, recessions are characterized by a sustained period of negative growth.

[8]Graphs based on consumption and industrial production data are qualitatively similar.

3.3 Implications for Term Premiums

In order to extract the implications of the model for the term premiums, it is necessary to price bonds in this model. For single-regime, log-linear models, closed form solutions exist; however, the time-varying transition probabilities in this model preclude such solutions. Instead we turn to a discrete state space approximation technique (see Tauchen and Hussey (1991)). Log growth is approximated by a variable that can only take on a discrete number of values. The dynamics of the economy are then described by the transition probabilities between these discrete states. As the number of discrete states increases, the discrete approximation comes arbitrarily close to the true continuous state space model. The advantage of the discrete approximation is that integration becomes summation, and the solutions to the pricing equations are relatively simple matrix equations. The details of the approximation scheme and the corresponding pricing equations are provided in Appendix B. For the purposes of this analysis, it is sufficient to note that the discrete approximation we use of 11 states per regime is effectively perfect. None of the moments of interest change as the number of states is increased further, and the discrete approximation has no effects on the inferences drawn from the model.

Given the estimated parameters for the GDP series (see Table 1), Table 2 presents the resulting annualized term premiums and volatilities for 2 maturities – 1 year and 5 years.[9] The term premium is defined as the conditional expected 1-period (one quarter in the case of the quarterly GDP series) return on a zero coupon bond in excess of the 1-period risk-free rate, i.e., $\mathrm{E}_t[r_{t,t+1}(\tau) - r_{ft}]$. The volatility is the conditional volatility of the same 1-period bond return. The moments are conditional on the current state of the world, i.e., the current regime and level of GDP growth. Consequently there are 22 conditional term premiums and volatilities for each maturity – 11 values of growth for each of 2 regimes. Table 1 also reports the level of growth in each state and the unconditional probability of each state. Note that the more extreme states have a negligible unconditional probability of occurring, and

[9]All results are based on a subjective time discount factor (β) of 1 and a coefficient of relative risk aversion (α) of 2.

therefore they are of little practical interest. Nevertheless, they do provide extreme examples that serve to illustrate the intuition behind some of the results.

These same term premiums and volatilities are also shown in graphical form in Figure 3. Each line represents the term premium (top) or volatility (bottom) of a particular maturity for the 11 states within a given regime. Specifically, the solid and dashed lines represent the 1-year and 5-year maturities, respectively, in the expansion. The dotted and short-dashed lines represent the same maturities in the contraction.

The most striking result is that term premiums and volatilities vary dramatically over time, depending on the phase of the business cycle. For example, volatility varies from less than 4% to more than 10% for the 5-year bond. For the same bond, term premiums vary from -0.94% to 0.15%. Term premiums are negative both for high growth states in expansions and low growth states in contractions. Interestingly, these are the same states of the world for which volatility is high. Neither term premiums nor volatilities are monotonic in growth in either regime. Note that the patterns are the same for the 1-year and 5-year bonds. In fact, the same effects are present throughout the term structure. The major effect of increasing maturity is to increase volatility and to increase the magnitude (either positive or negative) of the term premium.

What is the intuition behind these results? Essentially all the interesting features are being driven by the time-varying transition probabilities between regimes. Within each regime, growth is i.i.d. Consequently, if the probability of a regime switch is zero in every state, then term premiums are identically zero and volatilities are constant within a regime. The zero term premium is a result of zero autocorrelation in the stochastic discount factor, and the constant volatility is a result of assuming constant conditional volatility within each regime.

Introducing time-varying transition probabilities introduces both time-varying conditional autocorrelations and time-varying volatilities. In this specification, the probability of a regime switch depends on the level of growth. Specifically, it is decreasing in growth in both expansions and contractions. This dependence induces positive serial correlation in growth (and the stochastic discount factor) in expansions. A positive shock in an expansion increases the probability of remaining in the expansion and hence increases the probability of future high growth. Good news in the short-run is also good news in the long-run. This effect explains the negative term premiums in the high growth expansion states. However, the term premiums are positive in expansions when growth is low. What explains this phenomenon? In these states, a switch to the contraction next period is already inevitable (see Figure 1). In a contraction, a positive shock to growth increases the probability of remaining in the contraction. Good news in the short-run is bad news in the long run. This negative serial correlation produces positive term premiums.

The volatilities are also determined by the transition probabilities. Basically, if there is a lot of uncertainty about the regime next period then volatility is high. Intuitively uncertainty is maximized when the switch probability is approximately 0.5. Consequently, volatility is high in low growth contraction states. In theory, it is also high in expansion states with growth close to 0%; however, the sensitivity of the transition probability to the level of growth means that these are low probability events.

The patterns in Table 2 and Figure 3 are striking, but they are quite sensitive to some of the parameter values, although not to others. For example, α and σ affect the magnitudes of the fluctuations but not the form of the patterns themselves. Increasing either parameter increases both volatilities and term premiums. Consequently, using consumption data instead of GDP data generates the same intuition with smaller magnitudes because of lower volatility. Similarly, the qualitative results are not sensitive to the magnitude of the differences in mean growth rates across regimes. As long as the regimes are different (and persistent), then regime shifts are important.

The key parameters are those governing the regime transitions. As an illustration, Figure 4 presents the same information as Figure 3 but for the parameters estimated using industrial production data (see Table 1). In contractions, term premiums and volatilities are relatively flat across states because the transition probability is relatively insensitive to the level of growth. In other words, little time-variation in transition probabilities implies little time-variation in the conditional moments of bond returns. In marked contrast, expansion states exhibit large fluctuations for intermediate growth levels. Not surprisingly, these states correspond to the values of growth for which a transition to the contraction is possible but not certain. Finally, note that term premiums are negative in all states. Because persistence in the contraction is decreasing in growth (in contrast to the GDP and consumption estimates), both regimes exhibit positive serial correlations and negative term premiums.

Given the uncertainty in the parameter estimates and the reduced form nature of the model, can we make any predictions about bond returns? Basically, all 2-regime models with time-varying transition probabilities have two things in common. First, they exhibit large increases in bond return volatility around regime shifts, i.e., business cycle turning points. Second, term premiums also exhibit large fluctuations in these same periods. These fluctuations may be large changes in the magnitude of the premium (see Figure 4), or they may also include changes in the sign of the premium (see Figure 3). In fact, multi-regime business cycle models are perhaps the only parsimonious, equilibrium specifications that generate repeated and predictable switches in the sign of the term premium, i.e., the sign of the price of interest rate risk.

4 Empirical Evidence on Bond Returns

4.1 Data

Are the "theoretical" results in Section 3.3 consistent with the data on bond returns? To answer this question we turn to the CRSP Bond Files which provide monthly data on returns on U.S. Treasury securities grouped by maturity for the period February 1959 to December 1997. A full data series is available for 6 maturity ranges: up to 12 months, 12 to 24 months, 24 to 36 months, 36 to 48 months, 48 to 60 months, and 60 to 120 months.

Descriptive statistics for these bond returns (in excess of the 1-month, risk-free rate) are presented in Table 3. The unconditional mean return is increasing in maturity (with the exception of the 4-5 year maturities), suggesting an upward sloping term structure of term premiums on average. Volatility is also increasing in maturity, suggesting an upward sloping term structure of volatilities on average. Returns are highly correlated across maturities (as expected), and there is a small amount or return predictability using lagged returns.

Of course, the key predictions of the 2-regime model in Section 3 concern conditional term premiums and volatilities and how they vary over time, not unconditional (average) quantities. In order to estimate and model these conditional moments, we use the slope of the term structure of interests rates as the conditioning variable. Both theory and empirical evidence suggest that the slope is a useful conditioning variable for bond returns, and it has also been argued that this variable is a good business cycle predictor (see Boudoukh et al (1999) and the references therein). The slope of the term structure is defined as the 5-year yield minus the 1-month yield.

4.2 Results

Before turning to a direct estimation of the conditional moments of bond returns, it is worthwhile to take a closer look at the business cycle properties of bond returns. Figure 5 presents 13-month rolling estimates of the mean (solid line) and standard deviation (dashed line) of excess returns on the 5-10 year portfolio, with NBER business cycle peaks and troughs market by vertical solid and dashed lines, respectively. It is clear that return volatility is high around cycle turning points, particularly at the peak of the cycle. The only exception to this pattern is during the recession in the early 1990s. The evidence on mean returns is, if anything, even more striking. Returns are low (even negative) just prior to the peak of the cycle, and they then climb precipitously during the recession. Given the short length of this phase of the cycle, the evidence is remarkable.

These movements around cycle turning points can be seen equally clearly in the estimated mean and volatility of returns, conditional on the phase of the cycle, as presented in Table 4. Note that some of these estimates are conditional on ex post information (i.e., the future state of the business cycle), but they are still suggestive of variations in ex ante conditional term premiums and volatilities. The top 3 lines of the table show average returns over the

12-month period prior to a cycle peak, during the subsequent recession, and for the 12-month period after the trough. Consistent with the evidence in Figure 5, term premiums are negative, then strongly positive, before falling again as the economy proceeds into and out of a recession. The last line of the table shows the volatility in the time period spanning the recession (including 6 months before and after). For many maturities, volatility is almost twice as high during this period relative to its unconditional level.

We now turn to a direct estimation of the conditional moments of bond returns. Figure 6 shows a time series plot of the conditioning variable, the slope of the term structure of interest rates. The business cycle nature of this variable is clear. Moreover, the rapid increase in the slope in the period immediately before, during and after recessions is at least superficially similar to the pattern exhibited by bond returns (see Figure 5).

To get a feel for the predictive power of the slope of the term structure, we first conduct a simple nonparametric analysis. Specifically, we estimate both the term premium and the volatility conditional on whether the term structure is upward or downward sloping at the beginning of the month. The moments are estimated via GMM for the sample period 2/59-12/97. The results from this estimation for each of the maturity categories are presented in Table 6. Upward sloping term structures are defined as those for which the 5-year yield exceeds the 1-month yield. Months with upward sloping term structures comprise approximately 90% of the sample.

In spite of the small sample of downward sloping term structures, the results are quite striking. Estimated term premiums are negative for every maturity for downward sloping term structures, reaching a magnitude of approximately -6% (annualized) for the longest maturity bonds. Moreover, the term structure of term premiums is downward sloping. In contrast, conditioning on upward sloping term structures, all term premiums are positive with a maximum of about 2.5%. For conditional volatilities the pattern is reversed. Volatility is higher for downward sloping term structures by factor of about 2. In other words, the slope of the term structure is able to identify periods with negative term premiums and high volatility.

While we have not linked these results directly to the business cycle, Figure 6 shows that flat and inverted term structures are associated with cycle turning points from expansion to recession. As such these empirical results are consistent with the results for the estimated 2-regime model. Transitions between the phases of the cycle are associated with high volatility.

Although Table 5 provides interesting summary statistics it does not reveal much about the functional relation between the term structure slope and term premiums. To further explore this relation we estimate a nonparametric (kernel) regression of excess bond returns and volatilities on the slope. The results from these regressions are shown in Figure 7. The graphs show the conditional term premium and volatility for term structure slopes ranging from -1% to 3%. This range corresponds to the availability of reasonable data within the sample period. The moments are estimated for 3 maturity categories – 1 year (solid), 3-4 years (dashed), and 5-10 years (dotted).

Both term premiums and volatilities are generally monotonic in maturity; however, term premiums switch from monotonically upward sloping to monotonically downward sloping at approximately flat term structures. As suggested by Table 5, volatilities are high for downward sloping term structures, but they also exhibit an increase as the term structure gets steeper. The analysis in Table 5 cannot pick up this phenomenon, but it is potentially consistent with the results of Section 3. Figure 6 shows that the term structure tends to be steep at or just after the trough of the cycle. Consequently, these states may be picking up the volatility from the other turning point of the business cycle, from recession to expansion.

Overall, the nonparametric regression results suggest a complex dynamic behavior of term premiums and volatilities. Both quantities appear to exhibit significant time-variation, and price of interest rate risk appears to switch signs. The 2-regime model of Section 3 provides one possible and parsimonious explanation for these results.

5 Conclusion

Historically, business cycles have been characterized by movements about trends in aggregate economic variables following low order ARMA processes in first differences. This is a problem for rational, equilibrium, asset pricing models because the key component in these models is the covariance between the aggregate economic factor and returns on the financial asset. Consequently, asset pricing theories have had little success at explaining the variation in both the mean and the volatility of security returns via real business cycle effects. Recently, the macroeconomics literature has questioned the above view of the business cycle, developing, among other alternatives, regime-shift models. This alternative model for describing business cycles has strong implications for time-variation of the conditional distribution of asset returns. Both theoretical and empirical results in this paper suggest that these models have a promising future. The main implication, that expected returns and the volatility of returns move dramatically around regime switches, seems to be borne out in the data. Specifically, the conditional distribution of the term structure of bond returns moves most around business cycle peaks and troughs.

Perhaps most interesting is the simplicity of the model. Economic growth is i.i.d within regimes, the regime switching probabilities have especially simple forms, and inflation is ignored completely. One would expect that extensions to a more complex environment will be even more fruitful given the additional degrees of freedom available. For example, allowing economic growth to follow an ARMA process within regimes would generate more persistence in the underlying factor, which in turn yields more elaborate term structure shapes. In addition, building in an inflation process that is correlated with the stochastic discount factor would add another factor for explaining movements in the conditional distribution of bond returns. Because inflation has comparable volatility to real rates, the potential explanatory power of this extended model would increase. We hope to explore some of these issues in future research.

This paper has also documented some interesting statistical facts about the term structure of the mean and volatility of term premiums as a function of interesting economic variables such as the stage of the business cycle and the slope of the term structure. For example, when the term structure is upward (downward) sloping, the term structure of term premiums is increasing (decreasing). In contrast, the term structure of the volatility of these premiums is always increasing, irrespective of the term structure slope. Moreover, the overall magnitude of the volatilities is highest when term premiums are low, the opposite of the effect expected in a standard risk/return model. Here, the likelihood and persistence of the regimes completely drives the conditional distribution of returns. Independent of the model presented in this paper, these empirical stylized facts provide a hurdle for future financial models to explain.

A Estimating a Two-Regime Model

Consider the 2-regime model:

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

Following Gray (1996), note that the conditional density of g_{t+1} can be written as

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

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

$$
f_{it+1} \equiv f(g_{t+1}|I_{t+1} = i, g_t) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{g_{t+1} - a_i}{\sigma^2}\right) \tag{13}
$$

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 growth (denoted g_t^*), so p_{1t+1} can be rewritten as

$$
\begin{aligned}
\Pr[I_{t+1} = 1 | g_t^*] &= \sum_{i=1}^{2} \Pr[I_{t+1} = 1 | I_t = i, g_t^*] \Pr[I_t = i | g_t^*] \\
&= P_{t+1}(1,1) \Pr[I_t = 1 | g_t^*] + P_{t+1}(2,1)(1 - \Pr[I_t = 1 | g_t^*]) \qquad (14)
\end{aligned}
$$

By Bayes rule,

$$
\Pr(I_t = 1 | g_t^*) = \frac{f(g_t | I_t = 1, g_{t-1}) \Pr[I_t = 1, g_{t-1}]}{f(g_t | I_t = 1, g_{t-1}) \Pr[I_t = 1, g_{t-1}] + f(g_t | I_t = 2, g_{t-1})(1 - \Pr[I_t = 1, g_{t-1}])} \tag{15}
$$

Therefore, substituting (15) into (14),

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

where f_{it} is defined in equation (13). 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.[10]

B The Discretization Methodology

B.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.[11] 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.

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

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

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

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

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.

B.2 Pricing Assets

Given $\widehat{x}_t$ and the transition matrix Π, the solutions to certain expectation equations become relatively easy to calculate. Assume, for example, that x_t represents aggregate real growth and will be approximated by an m point discretization. Let $\widehat{l}$ denote the $m \times 1$ vector that contains the values that the MRS takes on in each of the states:[12]

$$\widehat{l} \;=\; \left\{\beta\left(x_{t+1}\right)^{-\alpha}\right\}$$

The price of a 1-period, riskless bond is the expectation of the MRS

$$q_t(1) \;=\; \mathrm{E}_t\left[\beta\left(x_{t+1}\right)^{-\alpha}\right],$$

In the discrete economy, the expectation is probability-weighted summation over the future states, i.e.,

$$\widehat{q_t(1)} = \Pi * \widehat{l}, \tag{16}$$

where $\widehat{q_t(1)}$ is an m-vector of bond prices, one for each state of the world. The price of a τ-period bond is determined recursively, utilizing the result that a τ-period bond is the expected, discounted value of a $(\tau - 1)$-period bond next period, i.e.,

$$q_t(\tau) = \mathrm{E}_t[m_{t,t+1}q_{t+1}(\tau - 1)] \tag{17}$$

In the discrete state space economy

$$\widehat{q_t(\tau)} = \Pi * (\widehat{l}. * q_{t+1}\widehat{(\tau - 1)}), \tag{18}$$

where $.*$ denotes element-by-element multiplication.

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 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.

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

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 B.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 pricing a τ-period bond, this time in an economy with two regimes. $\widehat{l}$ is the $m \times 1$ vector that contains the values of the MRS 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, $(\widehat{l}, \Pi_1)$ defines a single-regime model, and the rows of Π_1 sum to one. Similarly, let $(\widehat{l}, \Pi_2)$ 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 $\widehat{l}^*$ and Π^* as follows:

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

There are now $2m$ states of the world to reflect the fact that for each value of the MRS and 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 then used for pricing.

Panel A: Descriptive Statistics

	Mean	Std. Dev.	Autocorr.
GDP	0.775	0.993	0.338
Industrial Production	0.286	0.960	0.397
Consumption	0.800	0.483	0.346

Panel B: Two-Regime Model Estimation

	GDP			
	a_1	σ	p_0	p_1
Expansion	1.01	0.84	1.64	8.13
	(0.08)	(0.05)	(1.06)	(5.43)
	a_2	σ	q_0	q_1
Contraction	-0.34	0.84	1.44	0.98
	(0.20)	(0.05)	(0.68)	(0.69)
	Industrial Production			
	a_1	σ	p_0	p_1
Expansion	0.45	0.81	3.95	3.35
	(0.04)	(0.03)	(0.76)	(1.00)
	a_2	σ	q_0	q_1
Contraction	-1.33	0.81	0.38	-0.23
	(0.19)	(0.03)	(0.90)	(0.50)
	Consumption			
	a_1	σ	p_0	p_1
Expansion	0.92	0.41	-4.97	16.47
	(0.04)	(0.02)	(3.41)	(9.48)
	a_2	σ	q_0	q_1
Contraction	0.28	0.41	0.91	1.01
	(0.12)	(0.02)	(0.75)	(1.41)

Table 1: Descriptive Statistics and Model Estimation

Panel A presents descriptive statistics for quarterly log growth in real GDP (53Q1-98Q4) and monthly log growth in real industrial production (1/54-12/98). Panel B shows parameter estimates for the two-regime model as given in equations (11)-(12), using the same data. Standard errors are in parentheses.

		Term Premium		Volatility		
State	g_t	1-year	5-year	1-year	5-year	Prob.
			Expansion			
1	-3.35	0.12	0.15	2.85	3.92	0.00
2	-2.30	0.12	0.15	2.85	3.92	0.02
3	-1.40	0.12	0.15	2.86	3.92	0.52
4	-0.57	0.01	-0.03	4.04	5.56	5.17
5	0.23	-0.33	-0.53	5.18	7.14	18.92
6	1.01	-0.28	-0.45	4.77	6.57	28.86
7	1.79	-0.28	-0.45	4.77	6.57	18.92
8	2.59	-0.28	-0.45	4.77	6.57	5.17
9	3.42	-0.28	-0.45	4.77	6.57	0.52
10	4.32	-0.28	-0.45	4.77	6.57	0.02
11	5.37	-0.28	-0.45	4.77	6.57	0.00
			Contraction			
1	-3.35	-0.46	-0.75	6.29	8.67	0.02
2	-2.30	-0.58	-0.94	7.22	9.95	0.66
3	-1.40	-0.56	-0.92	7.40	10.21	4.05
4	-0.57	-0.40	-0.66	6.72	9.27	8.09
5	0.23	-0.20	-0.36	5.63	7.76	6.52
6	1.01	-0.05	-0.13	4.60	6.33	2.18
7	1.79	0.03	0.01	3.84	5.28	0.36
8	2.59	0.08	0.08	3.36	4.62	0.00
9	3.42	0.10	0.12	3.09	4.25	0.00
10	4.32	0.11	0.13	2.96	4.06	0.00
11	5.37	0.12	0.14	2.89	3.97	0.00

Table 2: Term Premiums and Volatilities in a Two-Regime Model

State-by-state values for the log growth rate, the term premium and volatility (annualized) for 1-year and 5-year maturities, and the unconditional state probability from a 2-regime, 22-state model. Results are based on the estimated parameter values given in Table 1 for quarterly GDP data.

		$0 < \tau \leq 1$	$1 < \tau \leq 2$	$2 < \tau \leq 3$	$3 < \tau \leq 4$	$4 < \tau \leq 5$	$5 < \tau \leq 10$
	Mean	0.066	0.101	0.117	0.126	0.112	0.138
	SD	0.284	0.703	1.054	1.294	1.492	1.836
	Auto	0.181	0.184	0.134	0.133	0.130	0.120
Corr	$0 < \tau \leq 1$	1.000	0.948	0.908	0.860	0.835	0.796
	$1 < \tau \leq 2$		1.000	0.979	0.950	0.934	0.896
	$2 < \tau \leq 3$			1.000	0.983	0.969	0.938
	$3 < \tau \leq 4$				1.000	0.988	0.965
	$4 < \tau \leq 5$					1.000	0.973
	$5 < \tau \leq 10$						1.000

Table 3: Bond Return Descriptive Statistics

Descriptive statistics for monthly returns (in percent) on maturity-sorted (in years) bond portfolios for the period 2/59-12/97.

State	%	$0 < \tau \leq 1$	$1 < \tau \leq 2$	$2 < \tau \leq 3$	$3 < \tau \leq 4$	$4 < \tau \leq 5$	$5 < \tau \leq 10$
				Mean			
1	15	-0.05	-0.19	-0.30	-0.40	-0.50	-0.56
2	13	0.22	0.39	0.50	0.55	0.58	0.63
3	15	0.05	0.08	0.08	0.06	0.00	-0.02
				Volatility			
4	31	0.45	1.06	1.54	1.81	2.06	2.48

Table 4: Conditional Means and Volatilities of Returns

Means and volatilities of returns on maturity-sorted bond portfolios (2/59-12/97) for different phases of the business cycle. The states are defined as follows: (1) the 12 months prior to the business cycle peak, (2) the contraction (i.e., the months between the peak and trough), (3) the 12 months after the business cycle trough, and (4) the period from 6 months prior to the peak to 6 months after the trough. Dates of business cycle peaks and troughs are those identified by the NBER.

	$0 < \tau \leq 1$	$1 < \tau \leq 2$	$2 < \tau \leq 3$	$3 < \tau \leq 4$	$4 < \tau \leq 5$	$5 < \tau \leq 10$
			Term Premium			
Up	0.07	0.12	0.15	0.17	0.17	0.21
Down	-0.01	-0.12	-0.21	-0.31	-0.38	-0.50
			Volatility			
Up	0.24	0.59	0.90	1.15	1.34	1.67
Down	0.55	1.34	1.94	2.19	2.46	2.89

Table 5: Conditional Term Premiums and Volatilities

Conditional monthly term premiums and volatilities on maturity-sorted bond portfolios (2/59-12/97). The moments are conditional on either upward or downward sloping term structures at the beginning of the month.

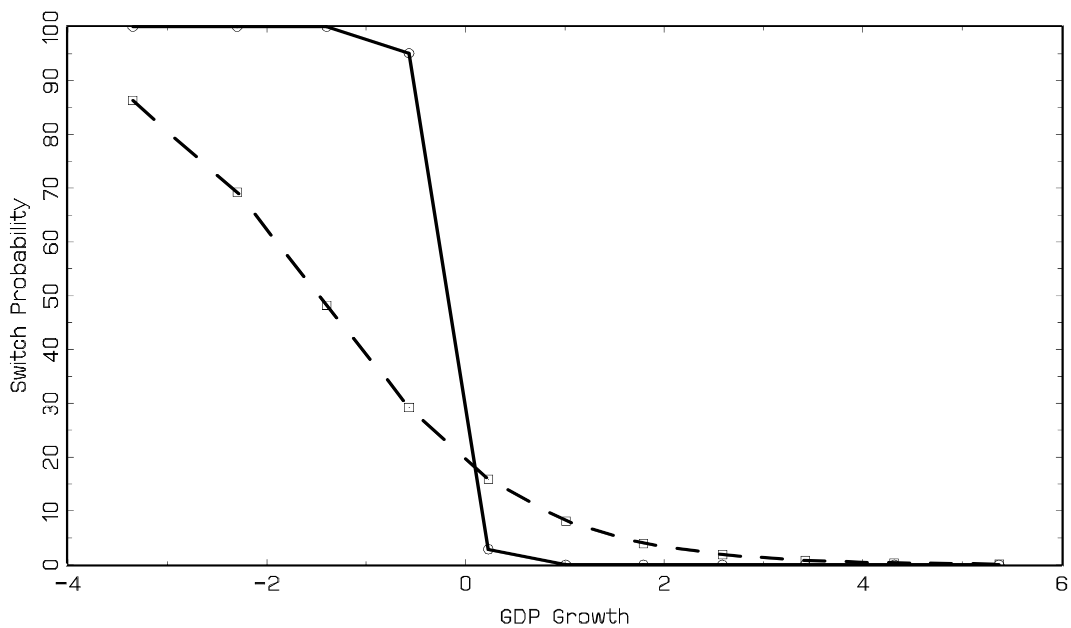

Figure 1: Regime Transition Probabilities

The probability of switching from the expansion to the contraction (solid line) and from the contraction to the expansion (dashed line) in the 2-regime, GDP model, conditional on the level of GDP growth. The model is given in equations (11)-(12) and the parameters estimates are in Table 1.

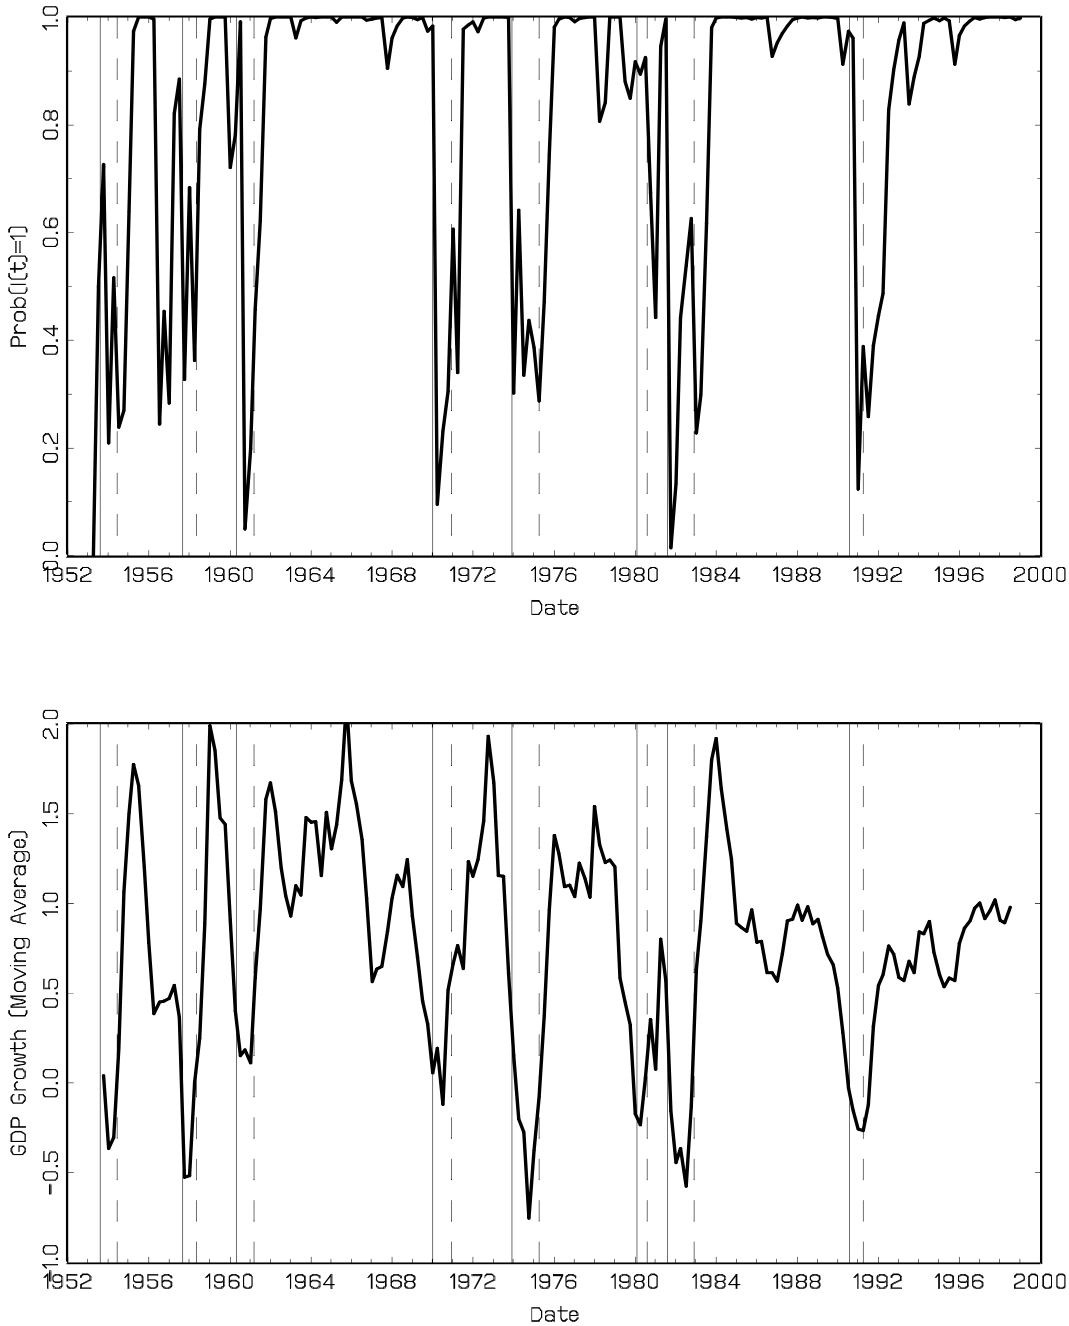

Figure 2: Regime Identification

The top graph shows the estimated ex ante probability that the economy is in state 1 (expansion) for a 2-regime model of quarterly real GDP. The model is given in equations (11)-(12) and the parameters estimates are in Table 1. The bottom graph shows a 5-quarter moving average of log growth in GDP. NBER business cycle peaks and troughs are marked by solid and dashed vertical lines, respectively.

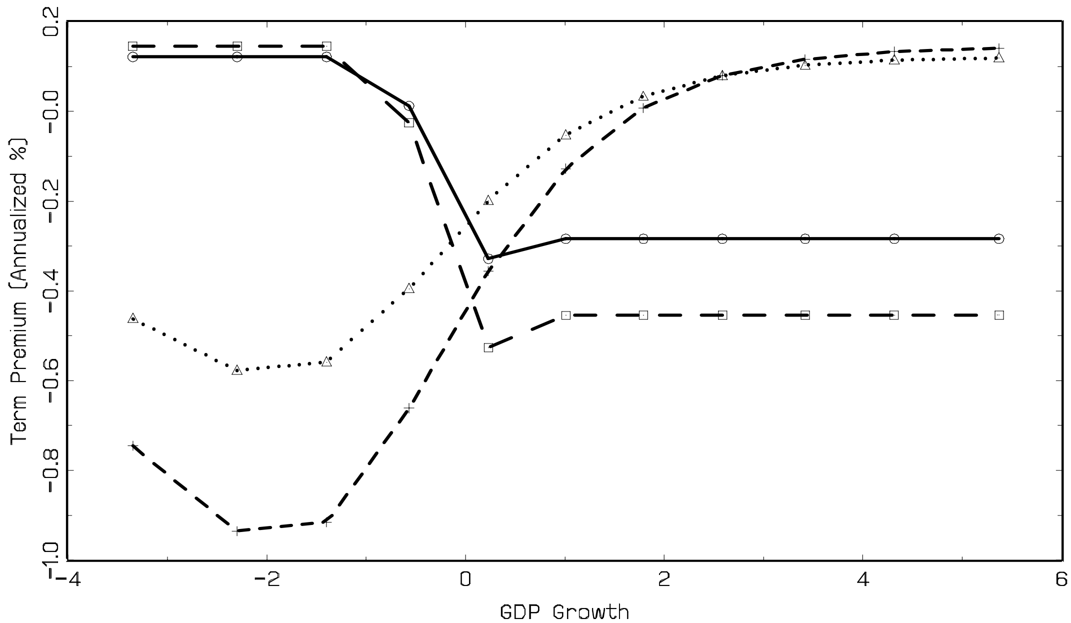

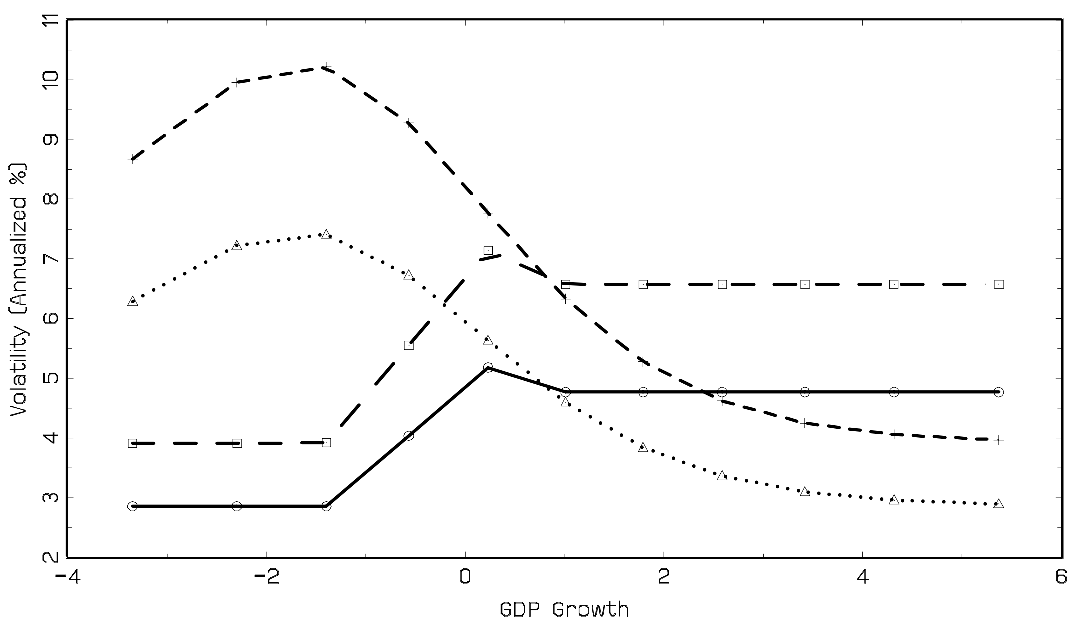

Figure 3: Term Premiums and Volatilities

Term premiums (top) and volatilities (bottom) for 1-year and 5-year maturities in the 2-regime, GDP model, conditional on the current regime and GDP growth. Solid and dashed lines represent the expansion (1-year and 5-year, respectively), and dotted and short-dashed lines represent the contraction (1-year and 5-year, respectively). The model is given in equations (11)-(12) and the parameter estimates are in Table 1.

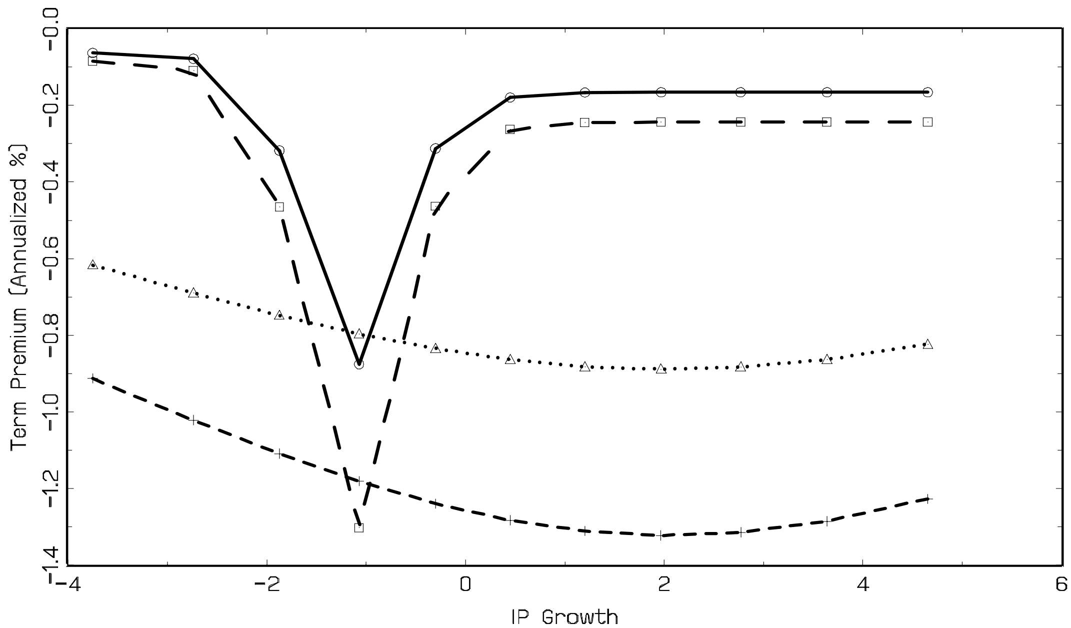

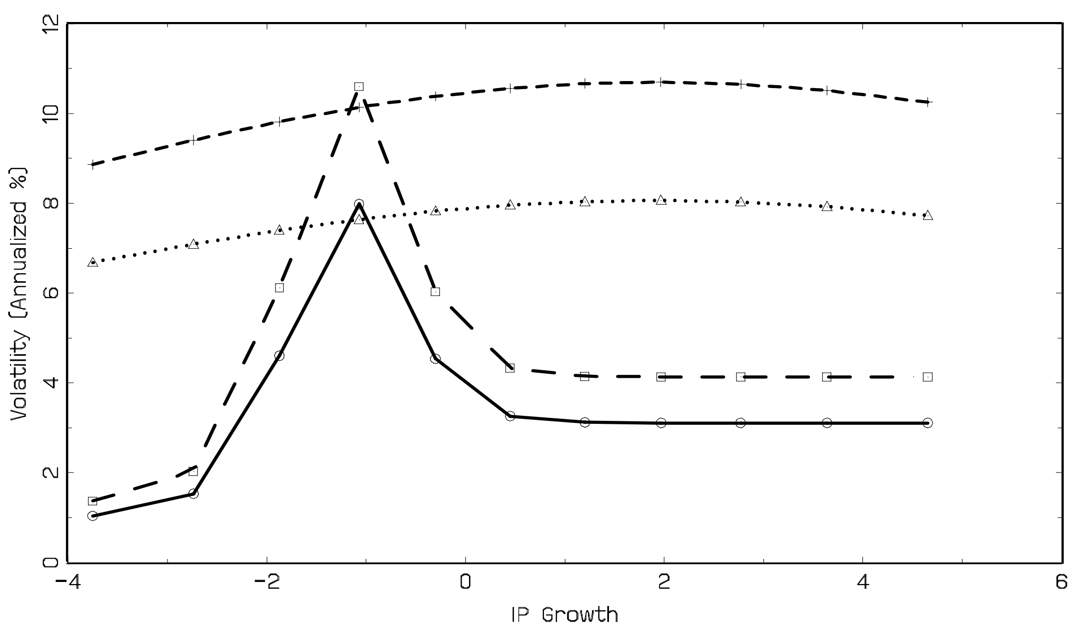

Figure 4: Term Premiums and Volatilities

Term premiums (top) and volatilities (bottom) for 1-year and 5-year maturities in the 2-regime, industrial production model, conditional on the current regime and industrial production growth. Solid and dashed lines represent the expansion (1-year and 5-year, respectively), and dotted and short-dashed lines represent the contraction (1-year and 5-year, respectively). The model is given in equations (11)-(12) and the parameter estimates are in Table 1.

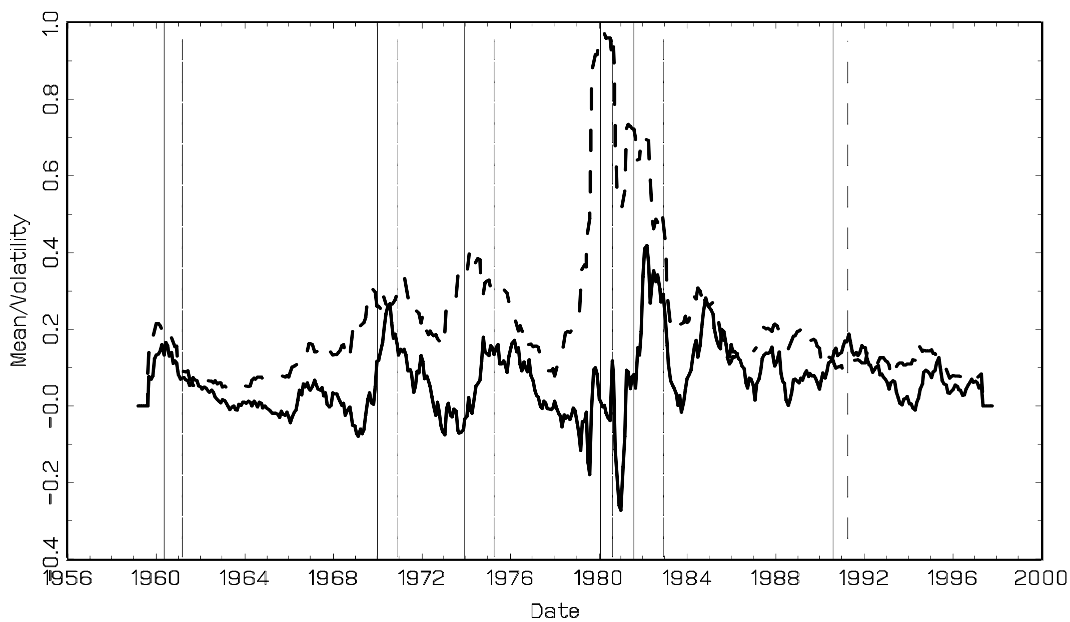

Figure 5: Rolling Means and Volatilities of Bond Returns

13-month rolling means and volatilities of the excess return on the 5- to 10-year bond portfolio. NBER business cycle peaks and troughs are marked by solid and dashed vertical lines, respectively.

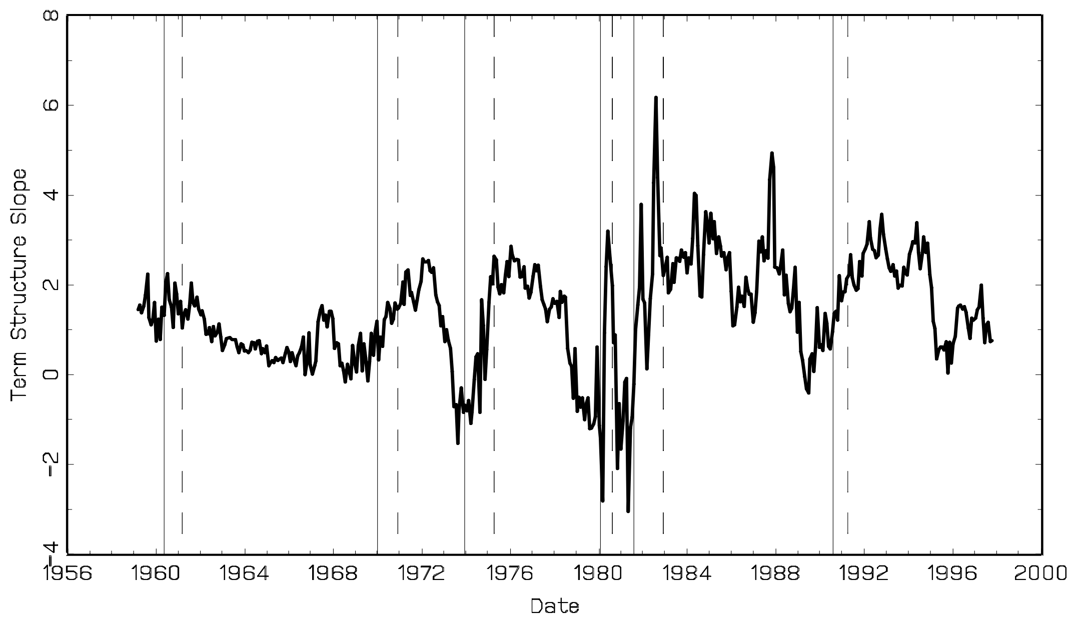

Figure 6: The Term Structure Slope

The term structure slope, defined as the difference between the yields on 5-year and 1-month securities. NBER business cycle peaks and troughs are marked by solid and dashed vertical lines, respectively.

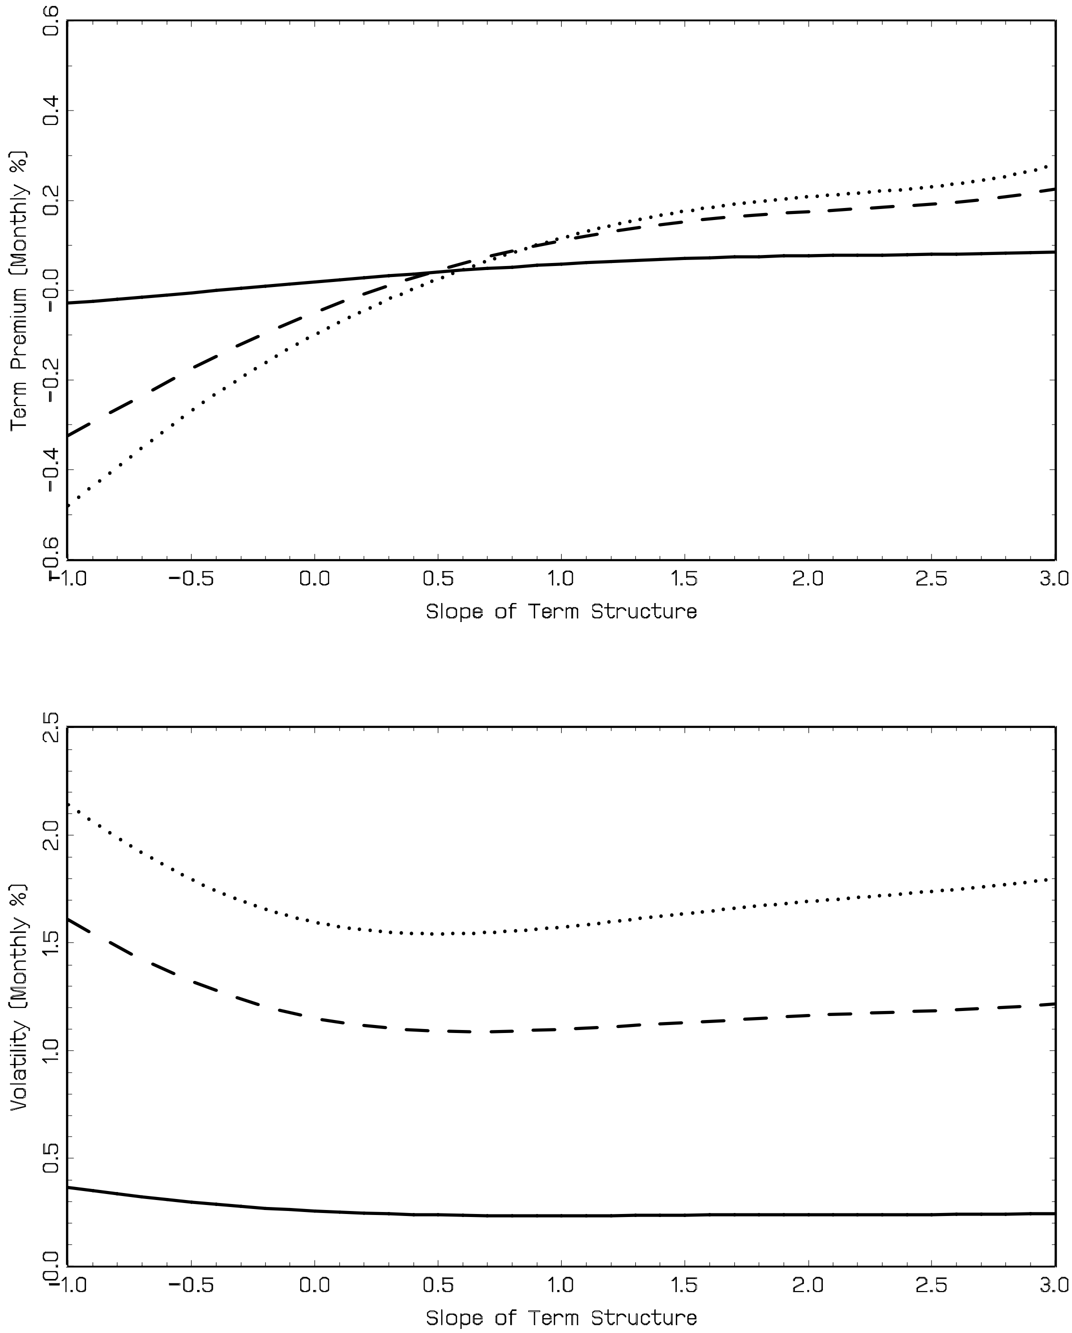

Figure 7: Nonparametric Conditional Term Premiums and Volatilities

Kernel estimation of conditional monthly term premiums and volatilities for maturities of up to 1 year (solid), 3-4 year (dashed), and 5-10 year (dotted) maturity portfolios. The conditional moments are estimated as a function of the slope of the term structure of interest rates at the beginning of the month.

References

[1] Boudoukh, Jacob, 1993, "An Equilibrium Model of Nominal Bond Prices with Inflation-Output Correlation and Stochastic Volatility," *Journal of Money, Credit, and Banking*, 25, 636-665.

[2] Boudoukh, Jacob, Matthew Richardson, Tom Smith, and Robert F. Whitelaw, 1998, "Ex Ante Bond Returns and the Liquidity Preference Hypothesis," *Journal of Finance*, 54, 1153-1167.

[3] Breeden, Douglas T., 1985, "Consumption, Production, and Interest Rates: A Synthesis," *Journal of Financial Economics*, 16, 3-39.

[4] Burns, Arthur F., and Wesley Mitchell, 1946, *Measuring Business Cycles*, National Bureau of Economic Research, New York.

[5] Campbell, John Y., and John H. Cochrane, 1999, "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107, 205-251.

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

[7] Cox, John C., Jonathan E. Ingersoll, Jr., and Stephen A. Ross, 1985, "A Theory of the Term Structure of Interest Rates," *Econometrica*, 53, 385-407.

[8] Diebold, Francis X., and Glenn D. Rudebusch, 1996, "Measuring Business Cycles: A Modern Perspective," *Review of Economics and Statistics*, 78, 67-77.

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

[10] Gray, Stephen F., 1996, "Modeling the Conditional Distribution of Interest Rates as a Regime-Switching Process," *Journal of Financial Economics*, 42, 27-62.

[11] Guo, Hui, 1999, "Business Conditions and Asset Prices in a Dynamic Economy," working paper.

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

[13] Hamilton, James D., 1994, *Time Series Analysis*, Princeton University Press, Princeton, New Jersey.

[14] Hansen, Lars P., and Kenneth J. Singleton, 1983, "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns," *Journal of Political Economy*, 91, 249-265.

[15] Harrison, M., and D. Kreps, 1979, "Martingales and Arbitrage in Multiperiod Security Markets," *Journal of Economic Theory*, 20, 381-408.

[16] Kim, Chang-Jim, and Charles R. Nelson, 1998a, "Business Cycle Turning Points, a New Coincident Index, and Tests of Duration Dependence Based on a Dynamic Factor Model with Regime Switching," *Review of Economics and Statistics*, 80, 188-201.

[17] Kim, Chang-Jim, and Charles R. Nelson, 1998b, "A Bayesian Approach to Testing for Markov Switching in Univariate and Dynamic Factor Models," working paper.

[18] Polkovnichenko, Valery, 1999, "Heterogeneity and Proprietary Income Risk: Implications for Stock Market Participation and Asset Prices," working paper.

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

[20] Whitelaw, Robert F., 1998, "Stock Market Risk and Return: An Equilibrium Approach," New York University Working Paper.