Learning and the Long Rate

Emi Nakamura and Jón Steinsson*

Harvard University

March 2004

Abstract

In this paper we seek to explain the behavior of the long term interest rate in the U.S. over the last 50 years using the idea that investors had only limited information about the process generating short term interest rates. Investors are assumed to behave according to the expectations hypothesis of the term structure but use econometric methods to learn about the appropriate model for short term interest rates and form expectations about future short rates. In particular, they use methods for identifying structural breaks. Although the approach seems reasonably "rational" from an econometric perspective, it differs in important ways from traditional formulations of Bayesian learning. Our approach yields considerably more successful predictions about the behavior of the long term interest rate than existing rational expectations based explanations. In particular, our model provides an explanation for why the long term interest rate has tended to fall when the yield spread is high.

Keywords: Expectation Hypothesis of the Term Structure, Structural Breaks, Learning.

JEL Classification: E43, E47, D83.

*We would like to thank Attila Ambrus, Assaf Ben-Shoham, Daniel Benjamin, John Campbell, Gary Chamberlain, David Laibson, Emily Oster, Jesse Shapiro and seminar participants at Harvard University for valuable comments and discussions. Date of last revision: April 11, 2004.

1 Introduction

The expectatons hypothesis of the term structure has long been the dominant idea employed by economists to think about the relationship between yields on bonds of different maturities. The basic idea underlying the expectations hypothesis is that the expected holding period return on bonds of different maturities should be equal (or possibly differ by a constant risk premium). The theory has strong implications about movements of the term structure that have been subject to countless empirical tests. Its empirical record is mixed overall but quite poor when it comes to explaining the movements of long term interest rates.

The expectation hypothesis implies that the yield on a long term bond should be the average of expected future short term interest rates over the life-time of the long bond. One approach to assessing the expectations hypothesis is to combine it with a model of the short term interest rate and see whether the expected future interest rates implied by that model explain the evolution of the long term interest rate. While early research in this vein such as Modigliani and Shiller (1973) provided support of the expectations hypothesis, later work has been less supportive.

By far the most widely used models of the short term interest rate in finance and macroeconomics are autoregressive models. Kozicki and Tinsley (2001) show that if a model of the short rate is estimated in which the short rate is assumed to be stationary and follow an autoregressive process the implied long term interest rate is much less volatile than the actual long term interest rate (see figure 1.1).[1] They also show that an estimated model in which the short rate is assumed to follow an autoregressive process with a unit root implies that the long rate should track the short rate much more closely than it actually does (see figure 1.2).

Kozicki and Tinsley point out that in both cases the behavior of the implied long rate is heavily influenced by the long run mean of the short rate process. If the short rate is stationary in levels, its long run mean is constant and the implied long rate deviates too little from this constant level to explain the actual behavior of the long rate. If the short rate is stationary in first differences, its long run mean level at each point in time is equal to a moving average of current and past short rates. Estimates of this process put large weight on the current short rate and therefore imply that the long rate should track the short rate much more closely than it actually does.

While this type of analysis is quite informative, it does not constitute a direct test of the expectations hypothesis. The poor results need not reflect a failure of the expectations hypothesis. They could entirely be due to the fact that these models of the short rate are misspecified. In order

[1] The data we use in this study are quarterly observations on nominal zero-coupon yields with maturity three months and ten years. The sample period is 1952:1 through 1999:2. This data was constructed by McCulloch and Kwon (1993) and updated by Kamakura Risk Information Services. This is the same data as Campbell and Viceira (2002) use.

to get away from this problem, direct tests of the expectations hypothesis are needed.

Applied to returns over one period, the expectations hypothesis implies that the one-period return on a one-period bond (its current yield) should equal the one-period expected return on an n-period bond. This implication of the expectation hypothesis should hold regardless of the process followed by the short rate. Furthermore, as long as market participants have rational expectations, this implication of the expectation hypothesis can be tested using data on actual bond market yields.

It is useful to develop some notation. The fact that the one-period return on a one-period bond should equal the expected one-period return on an n-period bond may be stated more succinctly as $y_{1,t} = E_t r_{n,t+1}$, where $y_{1,t}$ denotes the yield of a one-period bond at time t, $r_{n,t+1}$ denoted the one-period return of an n-period bond from period t to period $t+1$ and E_t denotes the expectations opperator conditional on information at time t. We can now use the fact that $r_{n,t+1} = ny_{n,t} - (n-1)y_{n-1,t+1}$, to rewrite this equation as

$$E_t[y_{n-1,t+1} - y_{n,t}] = \frac{1}{n-1} s_{n,t}, \tag{1}$$

where $s_{n,t} = y_{n,t} - y_{1,t}$ denotes the yield spread between an n-period bond and a one-period bond.[2] This equation tells us that when the yield spread is high the long rate is expected to rise. Under the maintained hypothesis of rational expectation, this implication of the expectations hypothesis of the term structure may be tested by running the regression

$$y_{n-1,t+1} - y_{n,t} = \alpha_n + \beta_n \left(\frac{s_{n,t}}{n-1} \right) + \epsilon_{n,t}. \tag{2}$$

The expectations hypothesis implies that $\beta_n = 1$ for all n. Campbell et al. (1997, ch. 10) report the results of this test for maturities from 2 months to 120 months. They find that β_n is negative for all maturities. In all cases β_n is significantly less than one and for maturities from 12 months to 120 months it is significantly less than zero. This means that when the yield spread is large the long rate has tended to fall rather than rise as the expectation hypothesis implies that it should. This type of relationship implies large predictable excess returns on the long bond.[3]

There are a limited number of possible explanations for why this test could fail. First, it could fail because of measurement error in the long term interest rate. Since the long term interest rate appears in the dependent variable with a negative sign and in the regressor with a positive sign, measurement error in the long rate will bias the estimated coefficient downward and may even

[2] $r_{n,t+1} = p_{n-1,t+1} - p_{n,t} = ny_{n,t} - (n-1)y_{n-1,t+1}$ where $p_{n,t}$ is the log-price of an n-period bond at time t. See chapter 10 of Campbell et al. (1997) for a detailed presentation of basic concepts related to fixed income securities.

[3] A legion of other papers has also performed this test of the expectation hypothesis and reached similar conclusions. Notable early work in this vein includes Fama and Bliss (1987) and Campbell and Shiller (1991).

result in a negative estimate. Campbell and Shiller (1991) use instrumental variable techniques to assess whether measurement error is the cause of the negative estimates but find that the negative coefficients are quite robust.

A second explanation is that the last fifty years of data may have been unusual. If this period does not constitute a representative sample, it will not reveal the true underlying relationship between long rates and short rates. This kind of "Peso problem" explanation can of course always be invoked to explain the failure of tests of rational expectations. Rational expectations tests are often tests of correlation—and one can always invoke the argument that, for some reason, the sample correlation in a particular data set differs considerably from the "true" population correlation. For the case of term structure anomalies the "Peso problem" explanation has been explored by Bekaert et al. (2001).

A third explanation is that the expectations hypothesis doesn't hold because risk premia vary over time. However, for time varying risk premia to explain a finding of $\beta = 0$ in the regression above, all variation in the slope of the yield curve must be due to variation in risk.[4] Variation in risk premia must be even larger to explain negative estimates of β. This view is rather unsatisfying unless it is combined with a theory that explains the large variation in risk premia. To date the literature has only made limited progress at formulating such a theory.

If the reader is not satisfied with the explanations above, he has no choice but to believe that the failure of the expectations hypothesis is due to a failure of some aspect of the rational expectations hypothesis. The results in Froot (1989) support this view. Froot uses survey data of investor's expectations to separately test the expectations hypothesis and the assumptions of rational expectations. For longer term bonds, Froot is unable to reject the expectation hypothesis. He finds that expectation errors rather than deviations from the expectations hypothesis explain the bulk of the deviation of the estimated β from one.

Formally, the rational expectations hypothesis states that investors' expectations conditional on available information coincide with the true conditional expectation in the model. Investors can fail to have rational expectations because they don't optimize or don't effectively use the information that they have. But investors can also fail to have rational expectations simply because they don't have as much information as rational expectations models assume that they do. Rational expectations models make the strong assumption that the market participants know the true underlying model of the the short rate. If the investors are learning about the parameters of the true model, their expectations may be systematically biased. This may occur even if they know the "true"

[4]If $\beta = 0$, the change in the yield on the long bond is unrelated to the spread. So, if the spread is 3% then that means that the expected excess return on the long bond is 3% (since its yield is not expected to change). In other words, a 3% spread corresponds to a 3% risk premium.

distribution of these parameters.[5] This type of learning may go on for a long time if information about the parameters is revealed slowly. For example, information about the characteristics of different regimes may not be revealed until the regimes actually occur.

In this paper, we take the view that, indeed, investors in the bond market have very limited information about the "true" process generating short term interest rates. But given their limited information, we assume that the investors use sophisticated econometric methods to form exptectations, in particular, methods for identifying structural breaks. The only information investors are assumed to have is past data on bond yields and other potentially relevant variables. We then ask how the long term interest rate would have behaved over the last 50 years if investors behaved according to the expectation hypothesis—i.e., bought bonds that had high expected returns and sold bonds that had low expected returns until the expected return on all bonds were equal—but formed their expectations using econometric methods. In this way, we attempt to differentiate "true irrationality" on the part of market participants from what may look like irrationality but is more appropriately associated with limited information.

The econometric approach we assume combines classical and Bayesian econometric techniques. The agents are assumed to first select a broad class of models using a standard econometric model selection criterion, and then use Bayesian techniques to form expectations conditional on this class of models. The main advantage of this technique relative to a fully Bayesian approach is that it is considerably more computationally tractable (both for us, and perhaps for market participants as well). Indeed, virtually all empiricial work in the econometrics literature uses formal or informal model selection (rather than model averaging) techniques, at least to some extent. A second advantage of using a model selection criterion is that it provides a transparent framework for analyzing the general issues involved in inference in the context of an unknown model, and does not rely on particular characteristics of a prior over different possible models.

When we apply this framework to the behavior of the long term interest rate over the last 50 years, we find that we are able to provide a remarkably good explanation of the movements in the long rate. We find that allowing investors to consider models of the short term interest rate that include structural breaks greatly increases our ability to explain movements in the long rate. In particular, our model provides an explanation for why the long rate has tended to fall when the yield spread is large rather then rise as rational expectations models imply it should. This occured because learning about the current interest rate regime tended to yield larger movements in the long rate than the forces embodied in rational expectations versions of the expectations hypothesis.

[5] Formally, uncertainty about fixed parameters of the model is inconsistent with the ergodicity assumption inherent in rational expectations tests.

We use 10 year rolling regressions to provide evidence that the relationship between short term changes in the long rate and the yield spread has not been stable over the last 50 years. In the 1950's and through most of the 1960's the behavior of the long rate was in line with the expectations hypothesis. The anomalous relationship between the long rate and the yield spread was consentrated mostly in the 1970's and 1980's. It peaked in the early 1980's and has been diminishing steadily since then. We show that the long rate produced by our model matches these broad patterns.

There has been quite a large literature within macroeconomics, under the heading of learning, that investigates the idea that economic agents act like statisticians or econometricians in making inference about the economy. A prominent example from this literature is Sargent's (2002) work on least-squares learning, where agents form expectations using OLS regressions. Evans and Honkapohja (2001) present a thorough overview of the learning literature in macroeconomics. There is also a large literature on learning in the game theory literature. Fudenberg and Levine (1999) present an overview of this literature. The approach we describe can also be interpreted as a complex learning rule. However, this paper diverges from most of the learning literature in that we allow for considerably more flexibility in how the agents learn about the world. In particular, the agents in our framework are allowed to consider models with structural breaks. In contrast, agents engaged in "fictitious play" or "least squares learning" continue to use the same model of the world even if circumstances change drastically.[6]

In this respect, our approach is more related to the work in the artificial intelligence literature in computer science which attempt to codify rules of optimal inference that can be programmed into computers.[7] Like our approach, this literature has placed an emphasis on allowing for a very broad specification of possible true models of the world.

The paper proceeds as follows. In section 2, we show that the behavior of the long term interest rate can be described by an extremly simple function of past short rates. More precisely, we show that a simple exponentially weighted moving average of past short rates tracks the long rate remarkably well. In section 3, we present a theoretical framework for learning in the context of an unknown model. This framework is based primarily on insights from the econometrics and statistics literatures. In section, 4 we apply this framework to the term structure of interest rates.

[6]In fictitious play, the probabilities of various outcomes equal the historical frequency of their occurence. The implicit assumption is, thus, that conditioning variables are not useful in forming expectations.

[7]A prominent example in the computer science literature is Q-learning, e.g., Kearns and Vazirani (1994).

2 Fitting the Long Rate

Before seeking a behavioral explanation of the long rate, let us note some striking empirical regularities about the long rate. In figure 2.1 we plot the yield on a ten year zero coupon bond over the time period 1952-1999 at a quarterly frequency. We also plot an exponentially weighted moving average (EWMA) of past 3 month yields given by

$$C + (1-\lambda)\sum_{j=0}^{\infty} \lambda^j y_{1,t-j}. \tag{3}$$

It is evident from the figure that this simple function of past short rates fits the long rate very well. Actually, it fits the long rate slightly better than a totally unrestricted distributed lag with 16 lags of the short rate. The R^2 of the EWMA with $\lambda = 0.15$ and $C = 0.013$ is equal to 0.94 while the R^2 of the unrestricted distributed lag is 0.93. Perhaps surprisingly, the long-range history dependence of the EWMA makes up, in terms of fit, for the much larger number of free parameters in the unrestricted 16 variable distributed lag. A slightly better fit can be achieved if a distributed lag is used that puts more weight on the current short rate and less on past short rates than the EWMA does. But the improvement is marginal.[8]

Moreover, the EWMA can account for a sizable fraction of the variation in the spread between the long term interest rate and the short term interest rate. Figure 2.2 plots the actual spread and the spread between the EWMA and the short term interest rate. The figure clearly shows that the EWMA is able to account for a very substantial fraction of the variation in the spread. To be precise, the R^2 of the EWMA prediction for the spread is 0.70.

These findings are related to the literature on the term structure of interest rates from the 1960's and 70's. That literature found that distributed lags of past short rates fit the long rate in the 1950's and 60's quite well (see, e.g., Modigliani and Sutch, 1966; Modigliani and Shiller, 1973 and Dobson et al., 1976). Figures 2.1 and 2.2 show that the same kind of relationship that held between the long rate and past short rates in the 1950's and 60's continued to hold through the 1970's, 80's and 90's. Furthermore, the figures show that the rather complicated distributed lag specifications used in the earlier literature can be replaces by a simple EWMA.

It is actually rather remarkable that a simple EWMA of past short rates should track the long rate as well as it has for as long as it has. Over the last 50 years we have seen great variation in the behavior of the short term interest rate, from the low and relatively stable rates of the 1950's and 60's, to the ever rising and more volatile rates of the 1970's and early 80's and back to a period of low and stable rates over the last 20 years. One might have expected this to mean that the

[8] Exponentially weighted moving avegares are sometimes also referred to as geometric distributed lags.

relationship between the long rate and past short rates would change much more than it actually has.

Finding such a simple way to characterize the relationship between the long term interest rate and past short term interest rates is useful in that it can help us think about which kinds of behavioral relationships are likely to be able to explain the term structure of interest rates. In the literature of the 1960's and 70's, distributed lags of past short rates were often ment to proxy for the market's perception of a "normal" level of interest rates to which current rates were expected to return.[9] The stability of the relationship between the long rate and an EWMA of past short rates suggests that this notion may play an important role in investor's long term expectations about nominal interest rates. In the remainder of this paper we explore behavioral explanations of long term interest rates that may explain why this is the case.

3 Econometric Learning

Consider a model in which investors are optimizing agents who seek to maximize the expected returns on their portfolios. The investors, however, don't know the data generating process for the returns on assets. The only knowledge they have is a dataset X_{t-1} that includes past returns and also past observations of other potentially relevant variables.

The first order conditions of these investors say that the expected return on all assets should be equal for all holding periods. For example, $\hat{E}_t R_{n,t+j} = \hat{E}_t R_{m,t+j}$, where $R_{n,t+j}$ denotes the one-period return on an n-period bond at time $t+j$ and $\hat{E}_t$ denotes the investor's expectation conditional on his information set at time t. Unlike in the case of rational expectations, we do not assume that the agent's expectations coincide with the "true" expectations in the model.

A log-linear approximation of the agent's first order conditions yields the expectations hypothesis of the term structure. In the empirical work that we present in section 4, we focus on one particular implication of the expectations hypothesis, namely the fact that the yield on a long term zero-coupon bond should equal a simple average of the short term interest rate over the life-time of the long bond:

$$y_{n,t} = \hat{E}_t \left[\frac{1}{n} \sum_{j=0}^{n-1} y_{1,t+j} \right] \tag{4}$$

The main problem faced by the investors in this model is how to form expectations. We assume that the investors consider a variety of possible predictive models for short term interest rates, $\mathcal{P} = \{P_\theta : \theta \in \Theta\}$, where for concreteness we assume that the models may be represented by

[9] The idea that the market expects the interest rate to regress toward a "normal" level based on past experience was put forth by Keynes (1936). See also Modigliani and Sutch (1966).

probability distributions indexed by a parameter θ. These models differ in their implications regarding expected future short rates. The investors must choose one of these expectations or combine them in some way to form a single expectation.

The fully Bayesian decision theoretic approach to the investors' problem in forming expectations would be to assume that the investors have a prior $\Pi(\theta)$ over all the possible models. Given this prior they could use the likelihood $f_\theta(X_{t-1})$ of each of the models on the past data to form a posterior $Q(\theta|X_{t-1}) \propto f_\theta(X_{t-1})\Pi(\theta)$. They could then form expectations as a weighted average of the expectation for each of the models,

$$\hat{E}_t y_{i,t+j} = \int_\Theta E_t^\theta y_{i,t+j} dQ(\theta|X_{t-1}), \tag{5}$$

where $E_t^\theta y_{i,t+j}$ denotes the conditional expectation of $y_{i,t+j}$ according to model θ and the weights are provided by the posterior density. This approach is theoretically appealing. However, if the set $\mathcal{P}$ is large, it can be difficult to implement computationally. Furthermore, this approach leaves as a black box the source of the subjective beliefs embedded in the prior. Surely, some priors are more "rational" than others.

The investor's problem of forming expectations is of course exactly the type of problem econometrics seeks to answer. It is a problem of selecting and estimating a time series model. Most applied economists do not however use the fully Bayesian approach. Rather, they use formal or informal model choice methods to choose a particular model class, such as a linear regression model, and then use methods such as maximum likelihood or GMM to estimate the free parameters within this model class. Bayesian analyses of parameter uncertainty are generally only carried out on reasonably tightly parameterized models if they are carried out at all.[10] Computational costs can be thought of as one reason why the agents in our models, as well as empirical researchers, might use these types of methods as opposed to a fully Bayesian approach.[11] Since the goal of econometrics is to provide methods for carrying out "good" inference, it seems reasonable to assume that investors form their expectations in a way that is broadly consistent with modern econometric analysis. This is the approach we take.

An advantage of this assumption is that it imposes a set of well-known rules on the expectation formation mechanism. A primary criticism of the literature that has tried to provide alternatives to the idea that "agents know the model" has been its arbitrariness. If not rational expectation then what? Sims (1980) cites the "wilderness" of alternatives once one deviates from rational

[10]For examples of such Bayesian analysis, see Barberis (2000) and Knox (2002).

[11]Indeed, a number of papers in the behavioral economics literature argue that model selection, as opposed to model averaging, is a good way of modeling human behavior. See, e.g., Mullainathan (2002) and Hong and Stein (2003).

expectations. However, though methods of empirical analysis are by no means formulaic, there are nevertheless clear principles that guide scientific empirical analysis. If the agents in the model are assumed to follow these principles, then the "econometric learning" approach we propose is no more arbitrary than rational expectations approaches in which the appliead economist uses these same principles (perhaps implicitly) in specifying the "true" empirical process.

In our empirical analysis, we assume that investors use model choice methods (described below) to select a subset of $\mathcal{P}$. They then calculate parameter uncertainty within this narrower class of models. This approach combines a model selection approach with a Bayesian model averaging approach. An advantage of this approach is that it is considerably more computationally tractable than a fully Bayesian approach.

Most model choice methods use a model's fit in past data as a major determinant when selecting among many models. However, if the set $\mathcal{P}$ is large enough, model choice methods based solely on fit, such as maximum likelihood, become dysfunctional. For example, among nested linear models, the model with the largest number of variables will always have the highest likelihood. In order to mitigate this problem, econometric model selection techniques are generally increasing in the likelihood of the model but also include a penalty for increasing the complexity of the model. The Akaike Information Criterion (AIC) and the Bayesian Information Criterion (BIC) are examples of model selection criteria that have been shown to have desirable properties in selecting among linear models. These criteria have the form,

$$IC = \log f(m) + p(m), \tag{6}$$

where $f(m)$ is the maximum attainable level of likelihood given the complexity level m and $p(m)$ is a penalty for complexity. The "complexity" m generally reflects the number of effective free parameters. Intuitively, the role of the complexity penalty $p(m)$ is to prevent overfitting associated with complex models. The BIC and AIC have been shown to asymptotically select the correct model in various model choice problems, including the one we consider in Section 4.

More generally, a large literature in econometrics, computer science, physics and philosophy provides various arguments for the importance of fit and complexity in model choice.[12] The idea that, all else equal, the simplest model should be selected is sometimes referred to as the "principle of parsimony" in the statistics literature. The Occam's razor principle that the simplest explanation is best may also be thought of as reflecting this principle.[13]

[12]For example, Rissanen (1987)'s concept of "minimum description length" (MDL) provides a general framework for representing the tradeoff between complexity and fit inherent in most scientific model selection criteria. According to MDL, one should select the model which "encodes" the data most efficiently, where the description of the data consists of two parts: the model itself and deviations from the model (the error term). See Hansen and Yu (2003).

[13]See Zellner et al. (2001) for a discussion of simplicity in scientific inference.

The model choice framework allows us to provide the following stylized representation of the agent's approach to inference. Let us define the function $V(f_\theta(X_{t-1}), m(\theta))$ over the space of models where V is increasing in the model's likelihood on past data $f_\theta(X_{t-1})$ and decreasing in the number of effective free parameters $m(\theta)$. In this case, an agent who is learning according to the "econometric learning" framework described above can be viewed as if he is selecting models to maximize the function $V(f_\theta(X_{t-1}), m(\theta))$ in much the same way that standard utility maximizing agents choose bundles of goods to maximize their utility function U. The idea that the investor selects models of the short rate to maximize V means that given equal fit of two possible models of the short-term interest rate, the investor chooses the simpler one. This assumption is essentially a mathematical representation of the "principle of parsimony" discussed above. Similar tradeoffs between complexity and fit would, of course, arise in particular applications of Bayesian model averaging, though in a somewhat less transparent form.

The tradeoff between parsimony and fit does not arise in standard models of learning such as least squares learning and fictitious play. The reason is that in these learning frameworks, the agent is assumed to know or at least believe in a class of approximating models with a fixed number of effective parameters. The agent is assumed to continue using this class even if the data are considerably better described by a more complicated model, or if certain variables in the model could be discarded with minimal loss of fit.

Given the notation we have adopted, one way of understanding the goal of the econometric forecasting literature as well as the artificial intelligence literature in computer science is as an attempt to construct a V that is optimal over the many different prediction problems. However, it is important to note that even if one takes the optimistic view that econometric methods have achieved this goal, rational expectations tests will still fail when evaluated for a particular time series if an agent employs these methods.

The problem is that, while the agent's expectation errors are in expectation zero if one averages over both "prediction problems" and time, the agent could be wrong on average for any particular prediction problem—such as predicting short term interest rates. For example, suppose the agent's utility depends on predicting the future values of several different time series $z_{i,t}$ where i indexes the time series and t indexes time. One might then assume (optimistically) that econometric forecasting techniques have been optimized to solve the problem,

$$\min_p E[(z_{i,t} - p(X_{t-1}))^2], \tag{7}$$

where p is some function of past data X_{t-1} and the expectation runs over both realizations of $z_{i,t}$ and prediction problems.

It is important to note that we have implicitly assumed that the prediction function p only depends indirectly on i through the data X_{t-1}. This seems appropriate for the type of situation that we have envisioned, in which the agent has little direct information about the data generating function. Thus, the prediction algorithm is designed to work well for a variety of potential prediction problems. Practically, artificial intelligence inference rules and off-the-shelf econometric techniques are often designed with this type of flexibility in mind. Indeed, forecasting algorithms are often evaluated according to their success on a wide variety of macroeconomic time series. (see Stock and Watson paper).

Time series rational expectations tests are based on comparing the agent's subjective expectation of a variable z, $E^*(z|X)$ with the true conditional expectation $E(z|X)$ using data on the time series z. However, the solution to the minimization problem (7) does not generally yield $p_i(X_{t-1}) = E[(z_{1,t}|(X_{t-1})]$ (where the expectation now runs only over realizations of $z_{1,t}$), even if the agent's forecasting rule solves the minimization problem (7). The problem is that the forecasting rule p_i is chosen taking into consideration the agent's uncertainty over which prediction problem he faces. However, the rational expectations test averages only over realizations of $z_{1,t}$, rather than over prediction problems as well.

In our particular empirical application,the tradeoff between parsimony and fit arises at a number of points in the investor's estimation problem. First, we use a Bayesian information criterion to select the number of breaks in the short-term interest rate. We also make a number of more general modelling decisions– such as restricting the number of variables in the model– on the basis of providing a parsimonious representation of the data. We discuss these specific issues in more detail in the next section.

4 A Learning Model of Long Term Interest Rates

In section 2, we presented a striking empirical regularity for long-term interest rates. Namely, long-term interest rates are well represented as an exponentially weighted moving average of lagged short-term interest rates. The objective of this section is to argue that this empirical regularity can be "rationalized" by investors' concerns that the best model for short-term interest rates might involve structural breaks.

A number of recent studies have noted the possibility that structural breaks may be important in characterizing the behavior of short-term interest rates and the inflation rate. In particular, favorable empirical evidence for the structural break view has been obtained by Altissimo and Corradi (2003) for short-term nominal interest rates. More generally, a large literature has begun to develop about the fundamental difficulty of distinguishing between data generated by a unit root process versus a process with structural breaks (see, e.g., Perron, 1989; Stock, 1994; Granger and Hyung, 1999; Bos et al., 1999; Diebold and Inoue, 2001).

4.1 A Simple Model with Constant Probability of a Break

Let us begin with a simple model of breaks that provides considerable intuition about how a model with breaks is able to account for the reduced form results discussed above. Consider an investor trying to predict long-term interest rates using an autoregressive model. Suppose he assumes at any given point in time that there has been a structural break in nominal interest rates some time in the past 5 years. However, for simplicity, assume that beyond this he has no idea of when the break occurred. Moreover, he also has no idea of the correct parameters of the model.

According to the expectation hypothesis of the term structure the yield on a 10-year zero coupon bond depend on expectations of short- term interest rates far into the future. This means that if the short rate is stationary the long rate is primarily determined by the mean of the short rate process. Indeed, if the first order autocorrelation of the short rate in quarterly data is less than 0.9, the investor's prediction about the 10-year rate is almost *entirely* determined by his assessment of the mean of the short rate process.

A natural estimator for the mean of the short rate, conditional on a particular break date $t-j$, is the sample mean since that date

$$s_{t-j,t} = \frac{1}{j}\sum_{i=0}^{j-1} y_{1,t-i}. \tag{8}$$

A Bayesian approach to estimating the mean of the short rate, within this class of models, would be to then construct a weighted average of the sample means for different break dates:

$$\hat{\mu}_t = \sum_{j=1}^{J} p_{t-j,t} s_{t-j,t}, \tag{9}$$

where $\hat{\mu}_t$ denotes the estimate of the mean of the short rate, $p_{t-j,t}$ is the probability conditional on information at time t that the break occured at time $t-j$ and J is the maximum number of periods since the break. Our simplifying assumptions that the investor knows a break occured in the last 5 years but has no idea when it occured withing that time period translate into $J = 20$ and $p_{t-j,t} = 0.05$.

Using equation (8) to substitute for $s_{t-j,t}$ in equation (9) we get an expression for $\hat{\mu}_t$ as a function of past short-term interest rates. This substitution yields a set of declining weights on past short-term interest rates, as in the case of the exponentially weighted moving average. There are two reasons why the weights on past interest rates decline as the interest rates move further into the past. First, a past interest rate only provides information about the current mean when no break has occurred since that date—an event which is less likely for an observation farther into the past. Second, interest rates further in the past appear only in sample means $s_{t-j,t}$ with larger numbers of terms—and thus contribute less to the estimated current mean $\hat{\mu}_t$.

Figure 4.1 plots the investor's real-time estimates of the current mean of the short rate using the approach described above.[14] It also plots the yield on a 10-year zero coupon bond. Evidently, this simple approach provides a good fit to the long-term interest rate, producing very similar predictions to the exponentially weighted moving average.

4.2 A Model with Estimated Break Dates

Although the approach discussed above comes considerably closer to providing a behavioral interpretation of the data than the exponentially weighted moving average, a number of unanswered questions remain. Why should the investor choose a 5 year window? Why only consider one break? Why not use the data to provide information on the *location* of the breaks as well as their magnitude? And so on.

We address these issues by taking a more serious econometric approach to modelling the investor's learning process. The approach described below retains the intuition of the simple break model described above, but endogenizes both the number of breaks and the probability that a break has occurred at a given point in time. In the spirit of the discussion in section 3, our approach closely resembles the approach that some applied economists have taken in modelling interest rates in empirical work. In particular, our approach builds on Wang and Zivot's (2000) Bayesian model of multiple structural breaks.

Suppose the investor models the short term interest rate as an autoregressive process with occasional structural breaks in its mean,

$$y_{1,t} = \mu_t + \alpha y_{1,t-1} + \epsilon_t, \tag{10}$$

where α is an autoregressive parameter, ϵ_t is an error term and μ_t is a parameter that changes only occasionally. These breaks occur at unknown times $\tau_1, \tau_2, ... < t$. The parameters of the model are also unknown to the investor.

The investor is interested in determining what the appropriate level of the long-term interest rate is according to the expectation hypothesis of the term structure. As we discussed above, if the short-term interest rate follows a process in the class described by equation (10), the appropriate

[14]We add a constant equal to 0.013 to this series produced by the procedure described above.

level of the long term interest rate will depend almost entirely on the current mean of the short rate process provided that the autoregressive parameter is not very close to one.[15]

Suppose the investor takes the following partially Bayesian approach to estimating the mean of the short rate process. First, he uses a model selection criterion to select the number of structural breaks that have occured in the mean of the short rate in the time-span over which he has data. He then uses Bayesian methods to estimate the autoregressive parameter as well as a probability distribution for the location of each of the structural breaks. As we discuss in section 3, this type of partially Bayesian approach is probably more common in real-life empirical work than the fully Bayesian approach. (For example, see Wang and Zivot, 2000). Thus, this approach seems to constitute a reasonable "scientific" approach to the problem.

A common model selection criterion in the structural breaks literature is the Bayesian Information Criterion,

$$BIC = \log \hat{\sigma}^2 + \kappa A \frac{\log T}{T}, \tag{11}$$

where $\hat{\sigma}^2$ is the sum of the squared residuals from the model (the estimate of the error variance), κ is the number of "effective parameters" and $A \log(T)/T$ is the penalty associated with increasing the number of parameters. The number of effective parameters here is $2k+2$, where k is the number of breaks.[16]

Yao (1988) shows that selecting the model with the lowest value of the BIC defined above yields a consistent estimate of the number of structural breaks if the model is correctly specified. More generally, the BIC embodies the fundamental tradeoff between complexity and fit descussed in Section 3. The fit of the model inevitably improves as the number of breaks is increased; however, increasing numbers of breaks are penalized to avoid "overfitting" the data.

The parameter A determines the extent to which the investor penalizes additional complexity in the interest rate model. This parameter is typically set equal to 1 in empirical applications. However, from an asymptotic point of view, the value of the parameter does not matter since the consistency proof does not depend on the value of A. In finite sample situations, the "optimal" penalty for complexity from a behavioral perspective would depend on how often the investor encounters structural breaks. (For example, if he never encounts time series with structural breaks, the optimal value of A would be infinity.) We "calibrate" the parameter A as 0.5.[17]

[15] Here we are implicitely assuming that the agents expectation of the long run mean of the short rate in the future is equal to the current long run mean of the short rate. The rationale for this assumption is discussed in more detail in section 4.3.

[16] This definition explicitly accounts for the fact that the location of the breaks must be estimated in addition to the value of the constant term between breaks.

[17] $A = 1$ yields no breaks in the nominal interest rate despite breaks detected in both the real interest rate and the rate of inflation.

The first step in the investor's estimation strategy at any given point in time is to select the number of structural breaks he thinks have occurred in short-term interest rates so far. In order to do this, he estimates the model with different numbers of breaks and chooses the model with the lowest value of the BIC. The investor uses the methodology of Bai and Perron (1998a,b) to choose the location of the breaks for the purpose of this calculation.[18]

Once the investor has chosen the number of breaks in the constant term, he uses a Bayesian approach to estimate the autoregressive parameter and a probability distribution for the location of each break date. We use the Markov chain Monte Carlo (MCMC) approach described in Wang and Zivot (2000) to compute numerically the probabilities assigned to the various break dates.[19] The priors we adopt for this estimation are those used in Wang and Zivot (2000).[20] We then estimate the parameters of the current regime, μ_t and α from equation (10), as the probability weighted average of the OLS estimates for these parameters conditional on each breakdate. Finally, these estimates of parameters of the current regime are used to calculate forcasts of the short rate over the next ten years. The investor's estimate of the appropriate level for the long rate is the average of these forcasted short rates.

To summarize, the investor is assumed to first estimate the number of breaks using a model selection criterion that trades off complexity and fit. The investor then estimates the parameters of the current regime as weighted averages of OLS estimates over possible past break dates. Finally, the investor calculates his forecast for the long-term interest rate using equation (4), taking into consideration the persistence in the interest rate process implied by his estimate of the autoregressive parameter. The investor carries out this procedure repeatedly using real-time data at each point in time.

In both the calculation of the number of breaks and the MCMC procedure, we assume that a minimum period must elapse between break dates, as is standard in the literature. We set the minimum time between breaks at 12 quarters. We also restrict the date of the last break to be at least 12 quarters prior to the date of estimation. The minimum time between breaks is chosen, roughly speaking, to match the model's predictions to the long-term interest rate. However, our results are reasonably insensitive to the choice of this parameter—increasing or decreasing the

[18]We use slightly modified versions of Gauss programs written by Bai and Perron to perform this calculation.

[19]We use a slightly modified version of Gauss programs written by Wang and Zivot to perform this calculation. Given that we use an MCMC procedure, the issue of convergence arises. However, we find that our results are essentially invariant to increasing the number of iterations in the procedure. The current graph is created using a 10,000 iteration burn- in and 10,000 subsequent iterations. [CHECK THIS] [MCMC REGULARITY CONDITIONS?]

[20]Wang and Zivot (2000) assume that priors for all the parameters are mutually independent. The prior for the location of the breaks is a discrete uniform distribution. The proirs for the autoregressive parameter and the constant term for each regime are normal distributions with large variances. The prior for the variance of ϵ_t for each regime is the natural conjugate inverted gamma prior with parameters chosen so that this prior is also relatively non-infomative.

minimum time between breaks by a year has little effect on the results.

It is important to note that the minimum duration between breaks is the *only* parameter that is chosen to fit the long-term interst rate, aside from a constant risk-premium and the weight on complexity in the BIC. The remainder of the estimated parameters are estimated solely off of the short-term interest rate. In contrast, most behavioral models of the long-term interest rate include a great number of parameters estimated off the long-term interest rate. (See, e.g., Campbell and Shiller, 1987; Kozicki and Tinsley, 2001; Fuhrer, 1996)

The results from our multiple break model are represented graphically in figures 4.2 and 4.3.[21] Figure 4.2 shows that the multiple break model yields a fit that is nearly as good as the simpler, less behavioral models. The only substantial deterioration in fit is that the model predicts a higher long rate for several periods around the time interest rates peaked in the early 1980's. Furthermore, figure 4.3 shows that the multiple break model doesn't merely fit the long rate to the extent that the short rate fits the long rate. Rather, the multible break model actually fits most of the variation in the spread between the short rate and the long rate.

The intuition for the fit is the same as in the simple break model described in section 4.1. Namely, since more recent observations are more likely to be informative about the current interest rate regime, the investor places more weight on these observations in calculating his estimate of the current mean. However, unlike in the simple break model, the investor now also uses the data to estimate how many breaks have occured as well as how recently the last break occured. For example, during the rapid rise and fall of interest rates in the 1970's and 1980's the investor places much more weight on a recent structural break than during more tranquil periods such as the 1960's and 1990's.

While the graphical fit we are able to obtain is informative, in order to properly asses our model we need to know whether the long term rate implied by our model has certain key properties that actual long rates have. One key property about the behavior of the long rate is that in our sample it has tended to fall when the yield spread is high. As we discussed in section 1, this empirical regularity is inconsistent with the joint hypothesis of rational expectations and the expectations hypothesis. In panel A of table 1 we show that the long rates implied by our models, both the equal probability model and the multiple break model, match this property of actual long rates. More, precicely, we run the regression in equation (2) using different long rate series. The first three rows show results using the actual long rate for three different sample periods. They replicate the standard result that the estimates of β_n are negavite and significantly less than zero. We then run this same regression using the long rate implied by our equal probability model and our multiple

[21] We again add a constant equal to 0.013 to the series produced by the procedure described above.

break model. The estimates of β_n in these regressions also turn out to be negative and significantly less than zero. Our model is therefore able to explain the fact that long rates fall when the yield spread is high. The explanation is that movements in the long rate due to learning about the mean of the current short rate regime are larger than movements that have to do with the forces embodied in rational expectations models.

Another key property of the yield curve is that the short rate tends to rise when the yield spread is high. Unlike the property discussed above, this property is consistent with rational expectations and the expectations hypothesis. In panel B of table 1 we present estimates of the coefficient β_n in the regression

$$s^*_{n,t} = \alpha_n + \beta_n s_{n,t} + \epsilon_{n,t} \tag{12}$$

where

$$s^*_{n,t} \equiv \sum_{i=1}^{n-1} (1 - i/n) \Delta y_{i,t+i}.$$

Again, the first three rows present results using actual long rates for three different sample periods. The results for the 1952:1-1991:1 sample period line up well with the results presented in Campbell et al. (1997) for this same sample period. However, the results in rows 2 and 3 show that this relationship is not quite as strong for the more recent sample periods. The estimates of β_n fall from 1.482 to around 0.7 when the 1950's are dropped from the sample and the 1990's are added. However, the estimates of β_n are still significantly greater than zero. Rows 4 and 5 in panel B show that the estimates of β_n using the long rate from our equal probability model and our multiple break model line up quite will with the estimates for the same sample periods using the actual long rate. This provides further support for our claim that the long rates produced by our models closely replicate the actual behavior of long term interest rates.

It may seem strange that our model, which is meant to describe a process of "rational" learning, yields predictability of excess returns on long term bonds. However, it turns out that the "predictability" of excess returns is actually quite unstable over time. Moreover, it is heavily consentrated around 1980. Figure 4.4 shows the evolution of the β_n coefficient in equation (2) when estimated on a rolling 10 year sample.[22] From this figure we can see that the behavior of the long rate was in line with the expectations hypothesis in the 1950's and through most of the 1960's. The behavior of the long rate starts to deviate from the expectations hypothesis in the late 1960's. These deviations persist through the 1970's and seem to peak in the early 1980's. However, since the mid-1980's, the deviations from the expectations hypothesis have sharply diminished. Thoughout most of the 1990's it is no longer possible to reject the expectations hypothesis using a ten year

[22] More precisely, each point in figure 4.4 is the β_n coefficient in equation (2) when estimated on the last ten years of data before that date. Standard error bands for two standard errors above and below are also reported.

sample.

Figure 4.4 shows that the negative coefficients in panel A of table 1 can not be interpreted as a stable behavioral bias that sophisticated investors should be able to take advantage of. Rather, it suggests that these negative coefficients are due to the unusual and unpredictable behavior of interest rates in the 1970's and 1980's. Figure 4.4 suggests that if the behavior of interest rates continues to be as stable as it has been in the 1990's the bond market will learn to accurately predict interest rate movements and the negative coefficient may disappear.

Figure 4.5 shows that the long rate implied by our multiple break model produces a very simmilar pattern of β_n coefficients as the actual long rate series. The β_n's are larger in the 1980's than in the 1970's and have sharply diminished in size through the 1990's. However, the β_n's produced by our model are consistently more negative than the β_n's produced by the actual long rate. This is due to the fact that changes in the actual long rate tend to lead changes in the long rate implied by our model. This temporal mismatch is small enough that it is hard to detect in figure 4.3. But the regression in eqation (2) is quite sensitive to this. It should be noted that if no breaks are detected in the short rate the β_n's coefficients produced by our model will converge to one over time.

One noticable feature of the results presented in figures 4.2 and 4.3 is that the fit of our model deteriorates markedly after 1995. One possible explanation of this is increased credibility of the Federal Reserve. Our model of expectation formation is based solely on past data on interest rates. A central bank that has built up credibility by "calling its shots" for some time will eventually induce investors to include its announcements in their forecasting models. A central bank in this position will have the option of moving the market with words rather than actions. In line with this view, many authors and commentators in recent years have noted that the bond market has become increasingly sensitive to Fed announcements.

An interesting observation about this exercise is that, according to our behavioral model, the estimated persistence in the interest rate has almost *no impact* on the investor's forecast of the long-term interest rate. The reason is that the investor's estimates of the persistence of the interest rate are much lower given that he allows for breaks—generally less than 0.9. Thus, virtually all of the variation in the investor's forecast of the long-term interest rate is associated with his reassessment of the current mean.

The estimated numbers of breaks is not monotonically increasing. If the investor observes a rapid run-up in interest rates, he may at first estimate a structural break, but then revise his opinion about whether this break occurred if the run-up is followed soon after by a decline in interest rates–making the rapid increase look more like a temporary blip than a permanent increase. However, the estimated numbers of breaks are approximately as follows: 0 breaks from 1952 to 1962, 1 break from 1962 to 1968, between 1 and 3 from 1968 to 1979, between 2 and 3 breaks from 1979 to 1983, 3 breaks from 1983 to 1992 and 4 breaks thereafter.

4.3 Discussion of Alternative Modelling Strategies

The point of this section is to argue that while investors were not rational in the rational expectations sense, their predictions regarding interest rates can be interpreted as conforming to the expectations hypothesis of the term structure given a reasonable empirical specification of short-term interst rates. In particular, given the large changes in interest rates over this period, it was reasonable for market participants to think that there might have been structural changes in the short rate process—leading to changes in its long-term mean.

There are three main ways in which the tradeoff between fit and complexity, discussed in Section 3, comes into play in our analysis of short term interest rates, in addition to the use of the BIC described above. First, a more familiar way of modelling the changes in the mean of the short rate from a theoretical perspective might be as part of a regime shift model with a finite number of regimes. However, a regime change model would require nearly as many regimes as there are breaks in the structural break model to fit the data equally well, and would additionally require estimating the transition probabilities between regimes. Since there is no clear evidence of a finite number of alternating regimes for the case of short-term interest rates, we argue that a modelling strategy based on trading off complexity and fit as in Section 3, might reasonably prefer the structural break model over a regime change model. More generally, in the literature on macroeconomic forecasting structural breaks play an important role while regime-switching models do not (see, e.g., Clements and Hendry, 1998, 1999).

It is also worth mentioning that regime change models cannot explain the data on the long-term interest rate. Driffill (1992) shows that Hamilton's (1989) regime-switching model applied to the short-term interest rate cannot explain the behavior of the long-term interest rate–despite the large number of parameters that are estimated based on the long-term interest rate in this model. Driffill's work considers a model with only two interest rate regimes. However, we find in unreported work that regime switching models with multiple regimes also fail to explain the behavior of long-term interest rates without adding a large number of parameters to the model.

The second way in which the tradeoff between fit and complexity comes into play is in motivating our focus on univariate models of the short rate. The existing empirical evidence suggests that additional variables aside from lagged values of the short-rate do not contribute significantly to explaining interest rates. Indeed, this has been traditionally true even of Fed announcements. [ADD

REFERENCES] Again, although we do not carry out any formal analysis, it seems quite reasonable that a parsimonious modelling strategy would condition solely on past values of the short-term interest rate, at least before the late 1990's.[23]

Third, our procedure for forming expectations makes the implicit assumption of "no expected change" regarding the evolution of the long run mean of the short rate. Notice, that this does not mean that agents believe there will not be any further structural breaks in the long run mean of the short rate. It merely means that the agents don't know in which direction future breaks will be. This assumption is analogous to the ubiquitous practice in the rational expectations literature of estimating models based on "modern" data to make inferences about today and the future. For example, monetary policy models are often estimated only using data since 1982 since monetary policy is thought to have been considerably different before this time (see, e.g., Taylor, 1999). Yet, predictions for the future are made from these models without taking into consideration the possibility of another regime change. Again, we argue for the "no change" prediction as the simplest available prediction, though arbitrary patterns of *future* expected interest rates would be equally consistent with the past data in real-time.[24] The situation would, of course, be different if a clear and recognizable pattern were present in the mean of the short term interest rate.

One might argue that the investor should consider a variety of other types of models aside from the linear structural break and regime switching models discussed above—i.e. models with more autoregressive terms and latent variable models. However, a tractable model of learning must always restrict the model space in some way. From the perspective of forecasting long-term interest rates, models with additional lagged terms are likely to have reasonably similar implications to our AR(1) model, given how unimportant short-term interest rate dynamics are for the long-run predictions of the model in the case of stationary interest rate processes. Furthermore, an analogous set of issues relating to the inclusion of structural breaks are likely to arise in more complicated AR models and latent variable models.

Finally, it is worth noting once again that more complicated autoregressive models (without structural breaks) also cannot explain the behavior of the long-rate even when these models are estimated in real-time. Not surprisingly, more complex AR models behave essentially like the simple AR model presented in the previous section; while Kalman filter models including both unit root and stationary components have similar implications to the unit root model discussed in section 1 for the long-term interest rate.

[23]Formal econometric model selection criteria are often used to select the number of variables in a model. To our knowledge, however, the econometric properties of model selection criteria for simultaneously selecting the number of breaks and the number of variables have not been investigated.

[24]This issue is often refered to as the grue problem in the philosophy literature.

5 Conclusion

Much of what makes decision-making hard in everyday life, as well as in economic policymaking, is developing the appropriate model—rather than optimizing once the model has been agreed upon. It has been long noted in the economics literature that the rational expectations assumption that people "know the true model" seems inappropriate when it is clear that people have little direct information about the true data generating process. Yet models of how people form expectations in an unknown model are difficult to investigate empirically since expectations are rarely directly observable.

The term structure of interest rates provides an unusual window onto peoples' expectations. For this reason, the term structure of interest rates has been a focal point for debates on theories of expectation formation for at least 50 years. In this paper, we show that an econometric approach to modelling investors' expectations has realistic implications for the behavior of long-term interest rates. Thus, it seems possible to explain failures of rational expectations tests by limited information, rather irrationality.

There are large potential gains from developing effective models of expectation formation. For example, a central bank that has a realistic model of how people form expectation is likely to be able to much more effectively "manage" the inflationary expectations of economic agents. The would be particularly important when the central bank would like to transition from one regime to another (e.g. disinflate). The present results suggest that rapid changes in interest rates have the advantage over more gradual shifts in that they quickly alert market participants that a policy change has occurred. We should note, however, that these insights are most relevant to a central bank that has little credibility and therefore cannot use announcements to sway investors expectations—a description that perhaps applies less to the Federal Reserve today than it did in the late 1970's.

Similarly, a crucial determinant of workers' incentives to invest effort and human capital in a firm have to do with what they think the firm's future hiring (and firing) policy will be. Thus, an important question in the literature on organizational economics is how firms can effectively manage their employees' expectations about their future labor management policies. The analysis of investors' expectations in the laboratory of the term structure of interest rates has the potential to shed some light on these questions. These are facinating areas for future research.

References

ALTISSIMO, F., AND V. CORRADI (2003): "Strong Rules for Detecting the Number of Breaks in a Time Series," *Journal of Econometrics*, 117, 207–244.

BAI, J., AND P. PERRON (1998a): "Estimating and Testing Models with Multiple Structural Changes," *Econometrica*, 66(1), 47–78.

——— (1998b): "Computation and Analysis of Multiple Structural Change Models," manuscript, Boston University.

BARBERIS, N. C. (2000): "Investing for the Long Run when Returns are Predictable," *Journal of Finance*, 55, 225–264.

BEKAERT, G., R. J. HODRICK, AND D. MARSHALL (2001): "Peso Problem Explanations for Term Structure Anomalies," *Journal of Monetary Economics*, 48(2), 241–270.

BOS, C. S., P. H. FRANSES, AND M. OOMS (1999): "Long Memory and Level Shifts: Re-Analyzing Inflation Rates," *Empirical Economics*, 24, 427–449.

CAMPBELL, J. Y., A. W. LO, AND A. C. MACKINLAY (1997): *The Econometrics of Financial Markets*. Princeton University Press, Princeton, New Jersey.

CAMPBELL, J. Y., AND R. SHILLER (1987): "Cointigration and Tests of Present Value Models," *Journal of Political Economics*, 95(5), 1062–1088.

——— (1991): "Yield Spreads and Interest Rate Movements: A Bird's Eye View," *Review of Economic Studies*, 58(3), 495–514.

CAMPBELL, J. Y., AND L. M. VICEIRA (2002): *Strategic Asset Alloccation: Portfolio Choice for Long-Term Investors*. Oxford University Press, Oxford, UK.

CLEMENTS, M. P., AND D. F. HENDRY (1998): *Forcasting Economic Time Series*. Cambridge University Press, Cambridge, UK.

——— (1999): *Forcasting Non-Stationary Economic Time Series*. MIT Press, Cambridge, Mass.

DIEBOLD, F. X., AND A. INOUE (2001): "Long Memory and Regime Switching," *Journal of Econometrics*, 105, 131–159.

DOBSON, S. W., R. C. SUTCH, AND D. E. VANDERFORD (1976): "An Evaluation of Alternative Empirical Models of the Term Structure of Interest Rates," *Journal of Finance*, 31(4), 1035–1065.

DRIFFILL, J. (1992): "Changes in Regime and the Term Structure: A Note," *Journal of Economic Dynamics and Control*, 16, 165–173.

EVANS, G. W., AND S. HONKAPOHJA (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton, New Jersey.

FAMA, E. F., AND R. R. BLISS (1987): "The Information in Long-Maturity Forward Rates," *American Economic Review*, 77(4), 680–692.

Froot, K. A. (1989): "New Hope for the Expectations Hypothesis of the Term Structure of Interest Rates," *Journal of Finance*, 44(2), 283–305.

Fudenberg, D., and D. K. Levine (1999): *The Theory of Learning in Games*. MIT Press, Cambridge, Mass.

Fuhrer, J. C. (1996): "Monetary Policy Shifts and Long-Term Interest Rates," *Quarterly Journal of Economics*, 111(4), 1183–1209.

Granger, C. W. J., and N. Hyung (1999): "Occasional Structural Breaks and Long Memory," Discussion Paper, University og California, San Diego.

Hamilton, J. D. (1989): "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle," *Econometrica*, 57, 357–384.

Hansen, M. H., and B. Yu (2003): "Model Selection and the Principle of Minimum Description Length," Working Paper, University of California at Berkeley.

Hong, H., and J. C. Stein (2003): "Simple Forecasting and Paradigm Shifts," Working Paper, Princeton University.

Kearns, M. J., and U. V. Vazirani (1994): *An Introduction to Computational Learning Theory*. MIT Press, Cambridge, Mass.

Keynes, J. M. (1936): *The General Theory of Employment, Interest and Money*. Harcourt, Brace and Co., London, UK.

Knox, T. A. (2002): "Learning How to Invest when Returns are Uncertain," Working Paper, Harvard University.

Kozicki, S., and P. Tinsley (2001): "Shifting endpoints in the term structure of interest rates," *Journal of Monetary Economics*, 47, 613–652.

McCulloch, H. J., and H.-C. Kwon (1993): "US Term Structure Data 1947-1991," Working Paper No 93-6, Ohio State University.

Modigliani, F., and R. Shiller (1973): "Inflation, Rational Expectations and the Term Structure of Interest Rates," *Economica*, 40, 12–43.

Modigliani, F., and R. C. Sutch (1966): "Innovations in Interest Rate Policy," *American Economic Review*, 56(1/2), 178–197.

Mullainathan, S. (2002): "Thinking Through Catagories," Working Paper, MIT.

Perron, P. (1989): "The Great Crash, the Oil Price Shock and the Unit Root Hypothesis," *Econometrica*, 57, 1361–1401.

Rissanen, J. (1987): "Stochastic Complexity," *Journal of the Royal Statistical Society. Series B (Methodological)*, 49(3), 223–239.

Sargent, T. J. (2002): *The Conquestion of American Inflation*. Princeton University Press, Princeton, New Jersey.

Sims, C. A. (1980): "Macroeconomics and Reality," *Econometrica*, 48(1), 1–48.

Stock, J. H. (1994): "Unit Roots, Structural Breaks, and Trends," in *Handbook of Econometrics*, ed. by R. Engle, and D. McFadden, vol. IV, pp. 2740–2843. Elsevier, Amsterdam, Holland.

Taylor, J. B. (1999): "A Historical Analysis of Monetary Policy Rules," in *Monetary Policy Rules*, ed. by J. B. Taylor, pp. 319–341. University of Chicago Press, Chicago, Ill.

Wang, J., and E. Zivot (2000): "A Bayesian Time Series Model of Multiple Structural Changes in Level, Trend and Variance," *Journal of Business and Economic Statistics*, 18(3), 374–386.

Yao, Y.-C. (1988): "Estimating the Number of Change-Points via Schwarz' Criterion," *Statistics & Probability Letters*, 6, 181–189.

Zellner, A., H. A. Keuzenkamp, and M. McAleer (eds.) (2001): *Simplicity, Inference and Modelling*. Cambridge University Press, Cambridge, UK.

Table 1: Regression Estimates

Data	Sample Period	β_n (S.E.)
Panel A: Long-yield changes		
Actual Long Rate	1952:1-1991:1	-4.235
		(1.431)
Actual Long Rate	1957:1-1999:3	-3.682
		(1.202)
Actual Long Rate	1962:3-1999:3	-3.908
		(1.218)
Long Rate from Equal Probability Model	1957:1-1999:3	-5.893
		(0.399)
Long Rate from Multiple Break Model	1962:3-1999:3	-10.417
		(1.692)
Panel B: Short-rate changes		
Actual Long Rate	1952:1-1991:1	1.482
		(0.109)
Actual Long Rate	1957:1-1999:3	0.711
		(0.165)
Actual Long Rate	1962:3-1999:3	0.663
		(0.163)
Long Rate from Equal Probability Model	1957:1-1999:3	0.674
		(0.143)
Long Rate from Multiple Break Model	1962:3-1999:3	0.290
		(0.138)

Panel A presents estimates of the parameter β in equation (2), with White standard errors. Panel B presents estimates of equation (12), with Newey-West standard errors. All the regressions are run using historical data on the 3 month rate for $y_{1,t}$. In each panel the first three rows present results using historical data on yields of 10 year zero coupon bonds for $y_{n,t}$. The fourth row in each panel uses the implied long rate from our equal probability of break model $y_{n,t}$. The fifth row in each panel uses the implied long rate from our multiple break model for $y_{n,t}$. In panel A, $y_{n-1,t+1}$ is approximated by $y_{n,t+1}$. In all cases n = 40 quarters.

Figure 1.1: Stationary Short Rate

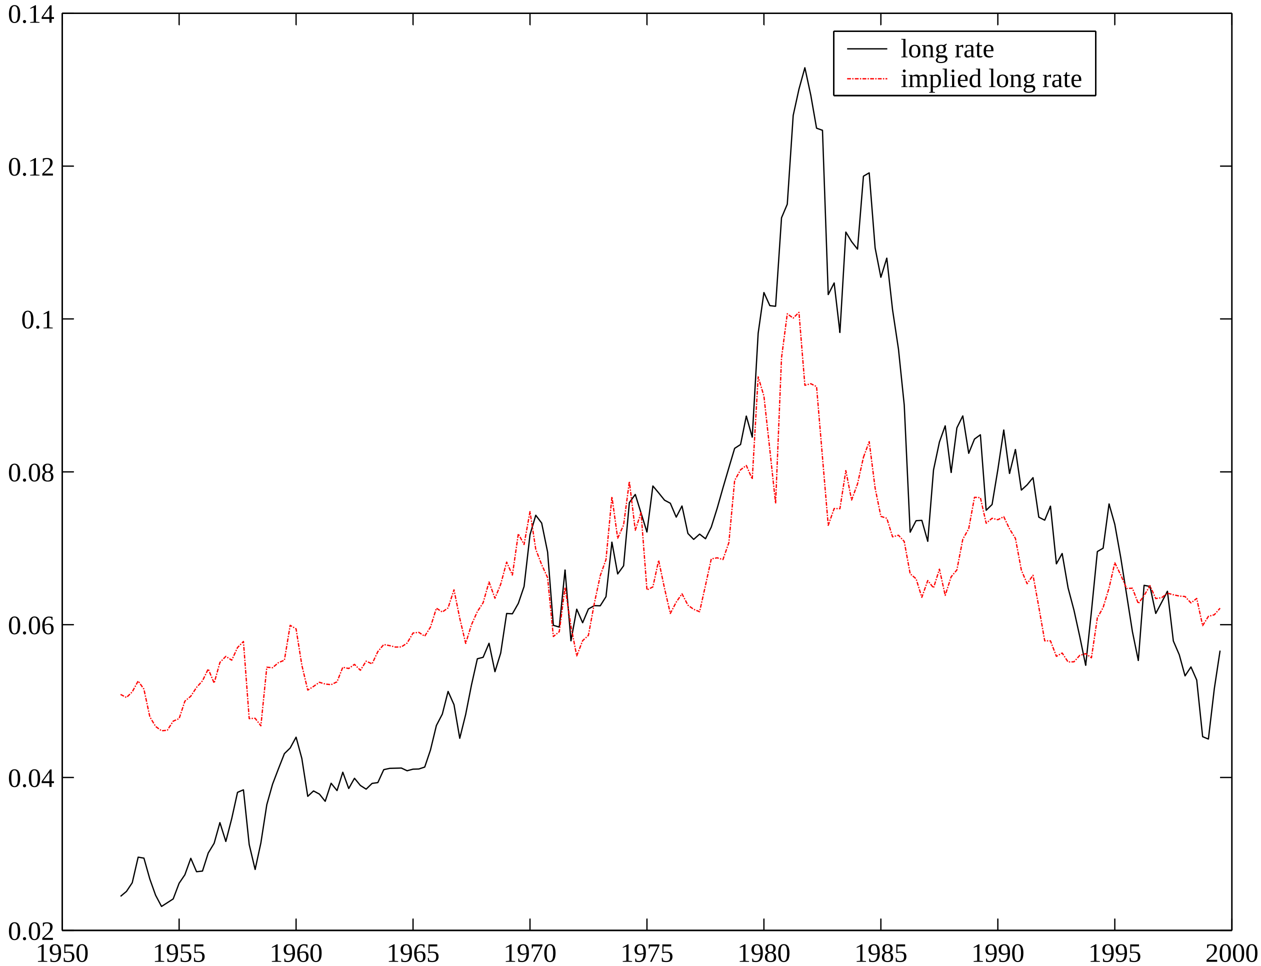

Student Version of MATLAB

Figure 1.2: Short Rate with a Unit Root

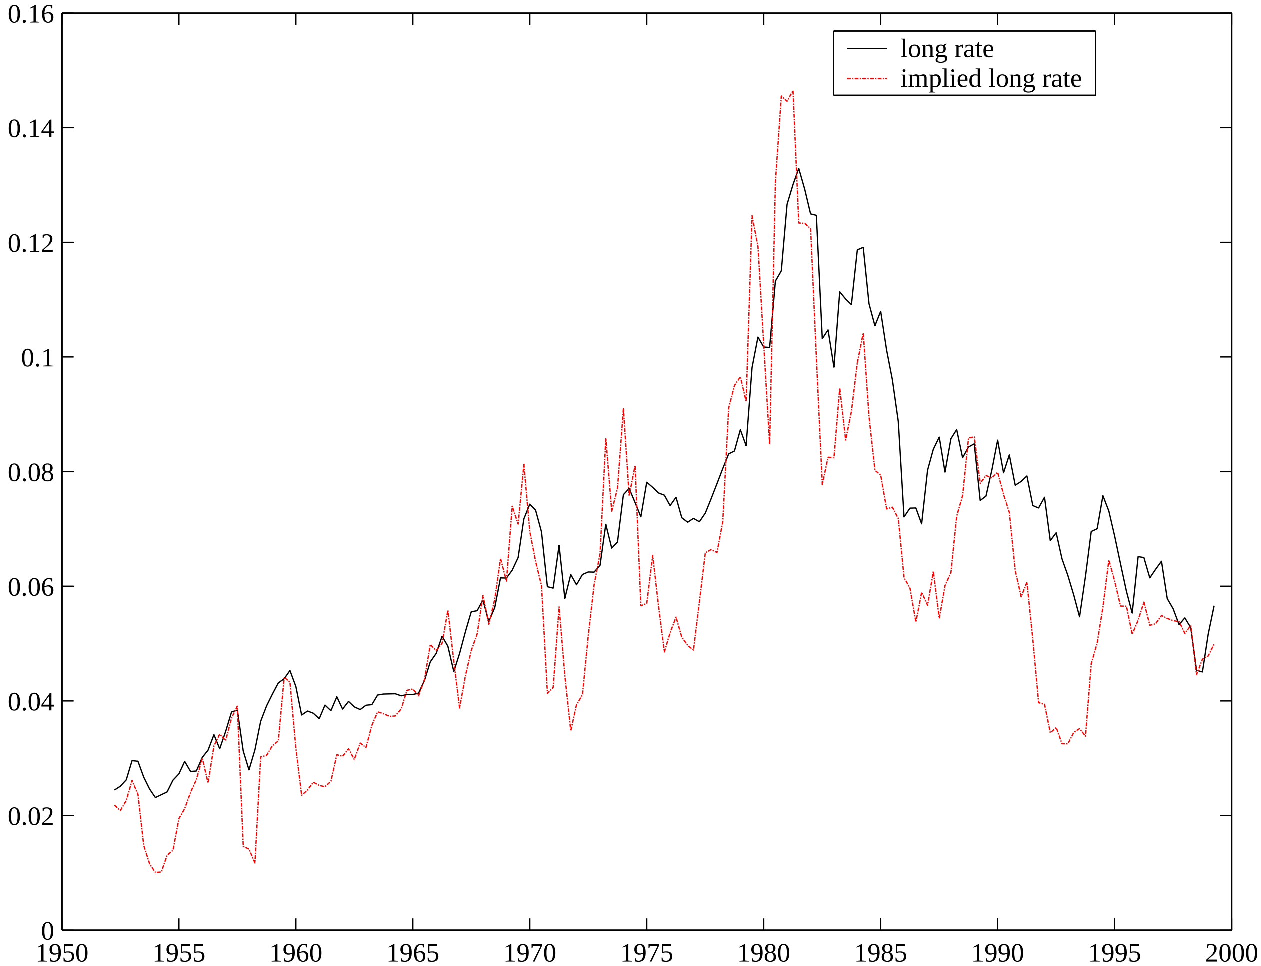

Figure 2.1: Long Rate and EWMA

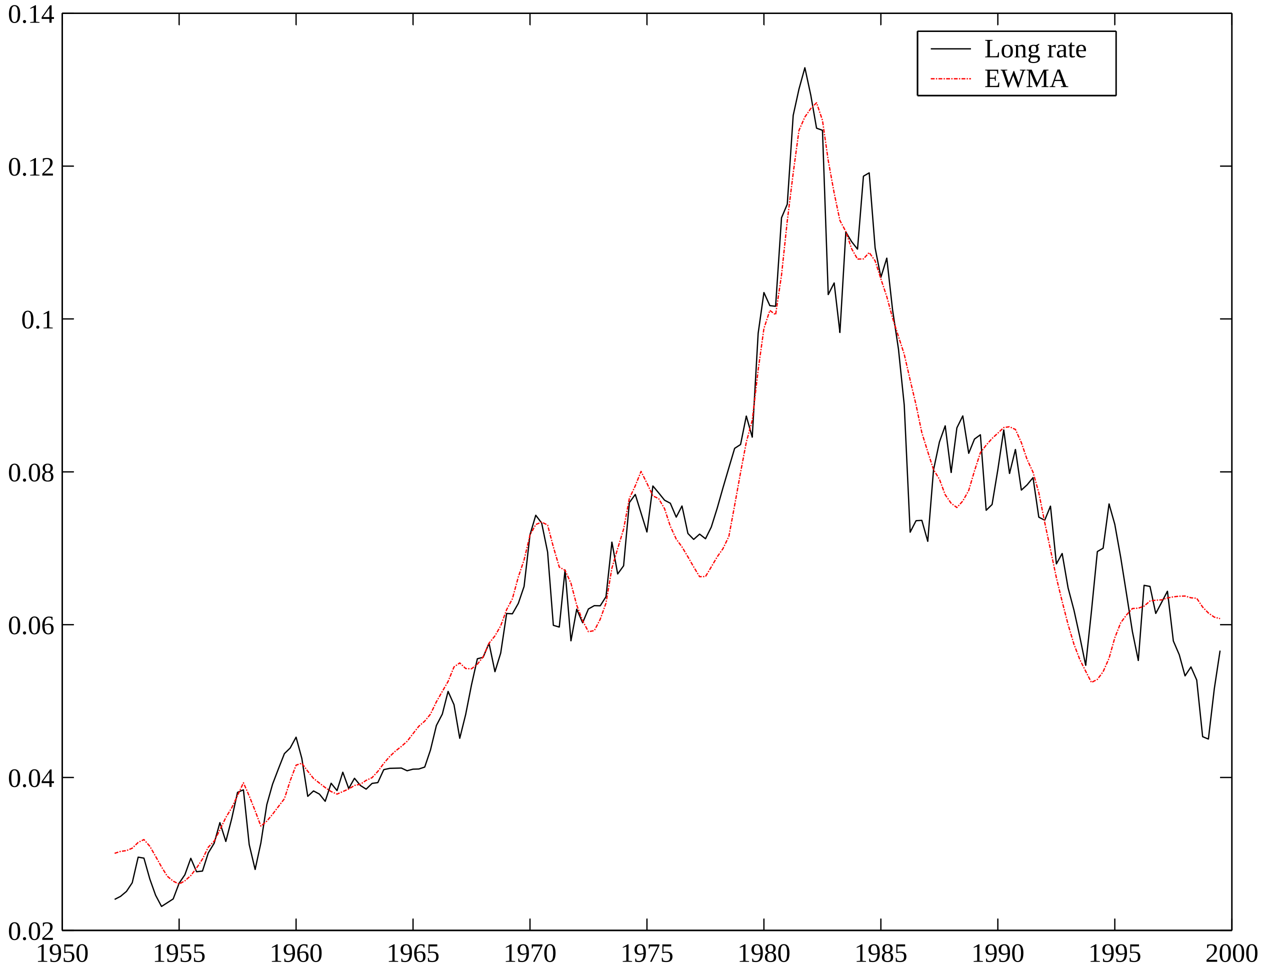

Figure 2.2: Spread and Implied Spread using an EWMA

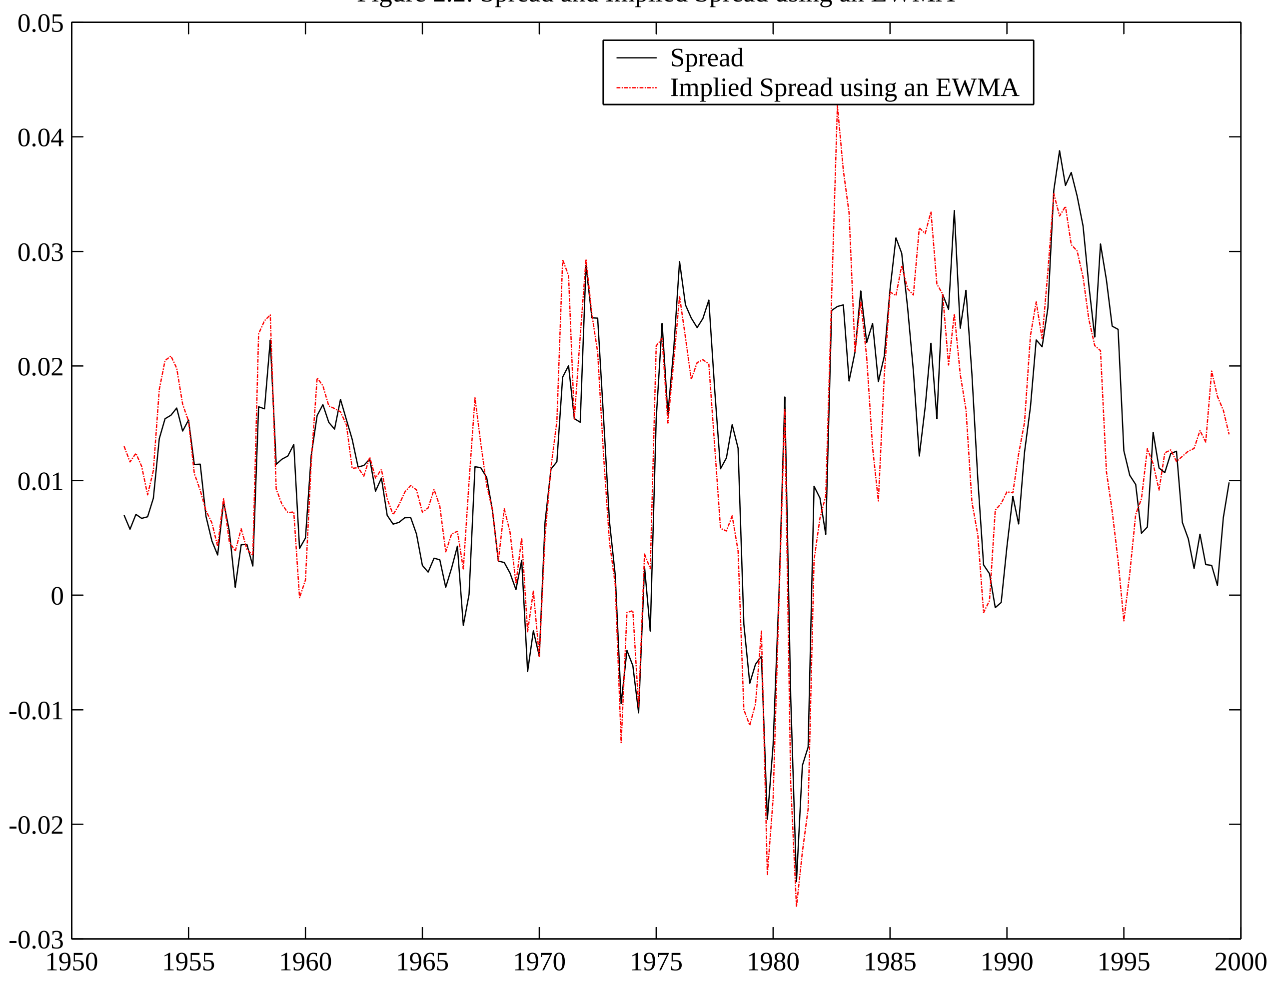

Student Version of MATLAB

Figure 4.1: Long Rate from the Constant Probability of Break Model

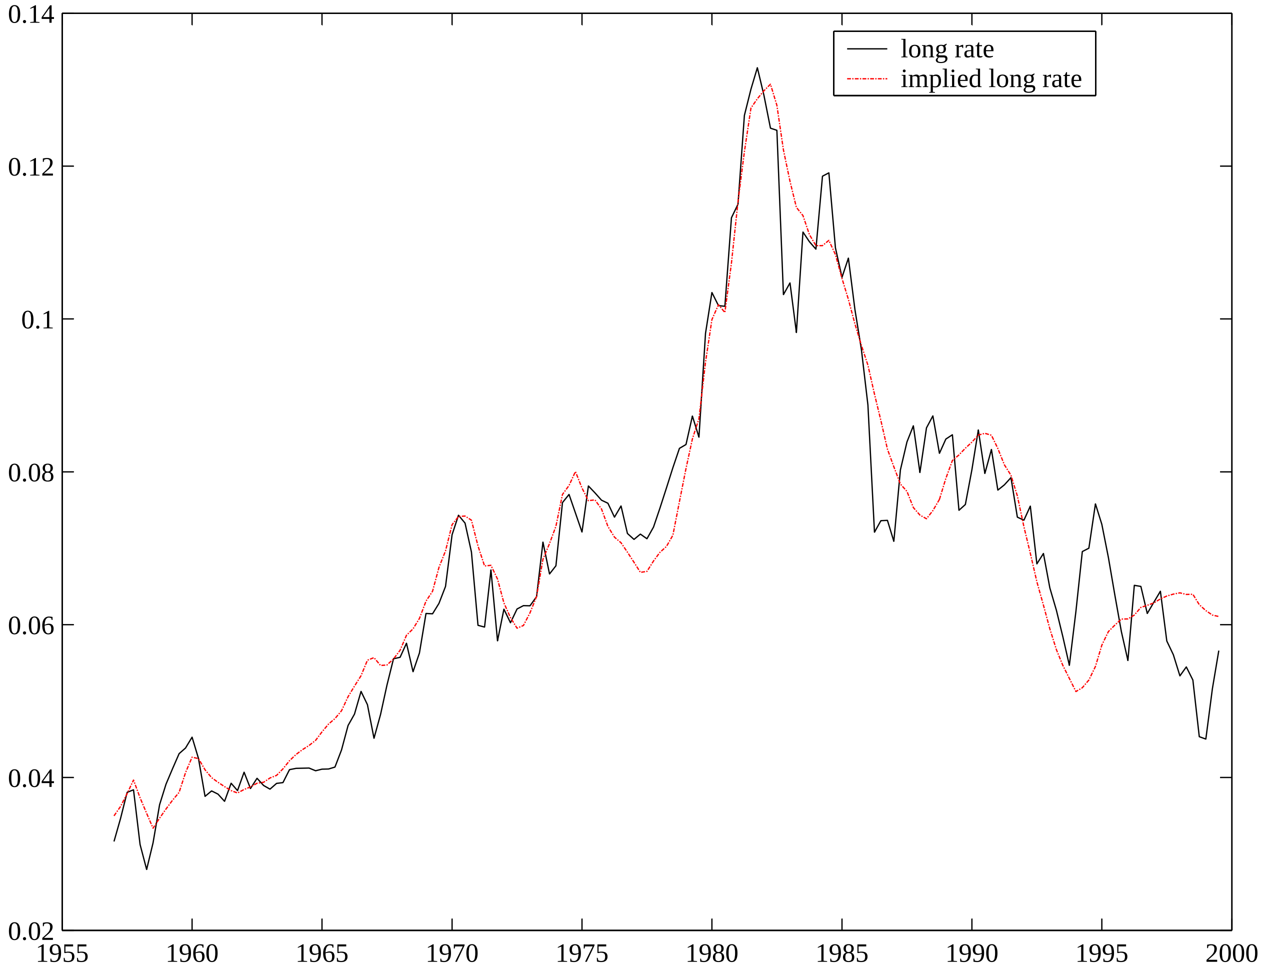

Figure 4.2: Long Rate from the Multiple Break Model

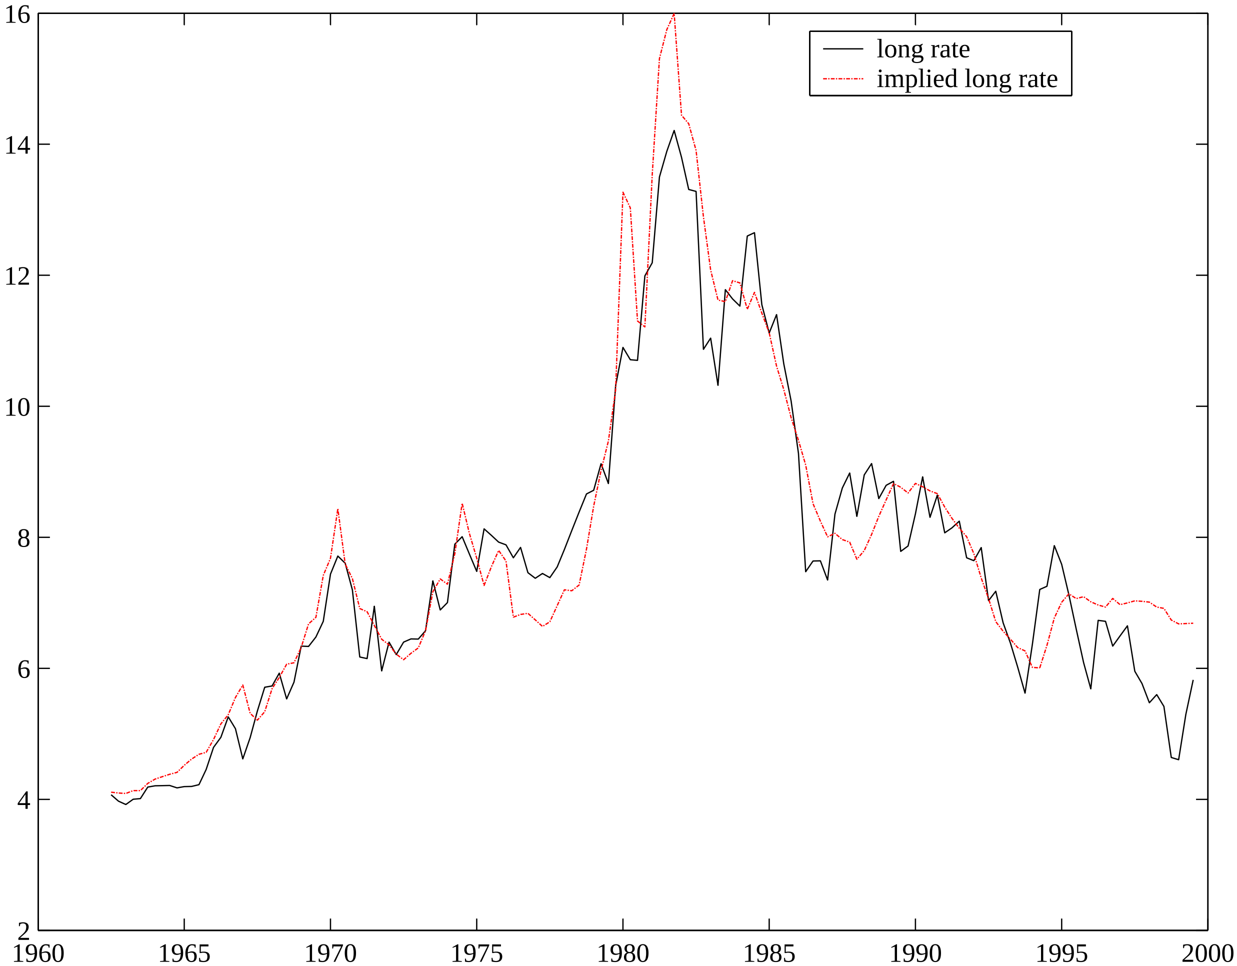

Figure 4.3: Yield Spread from the Multiple Break Model

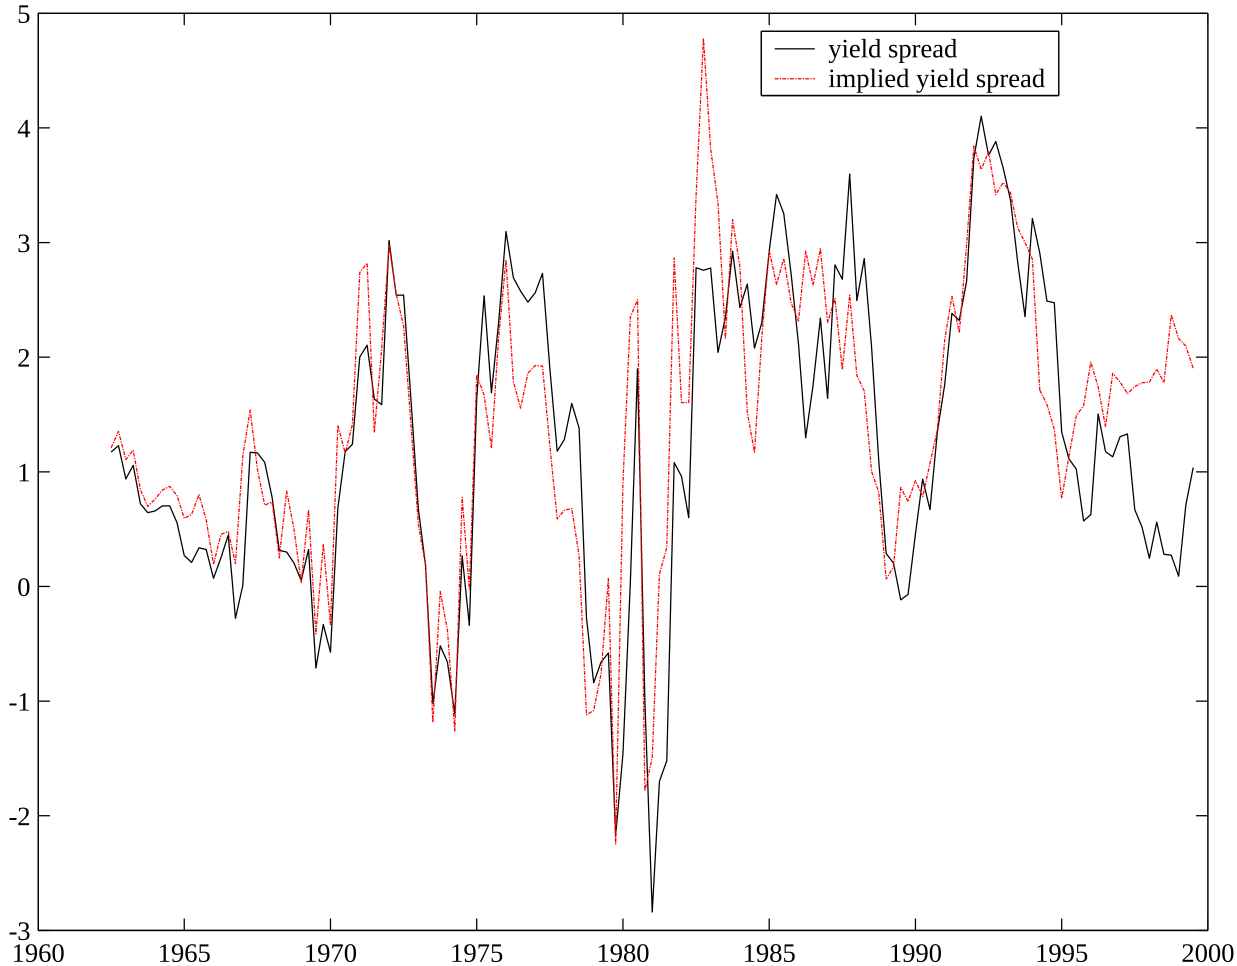

Figure 4.4: Estimates of beta with a rolling 10 year sample using the actual long rate

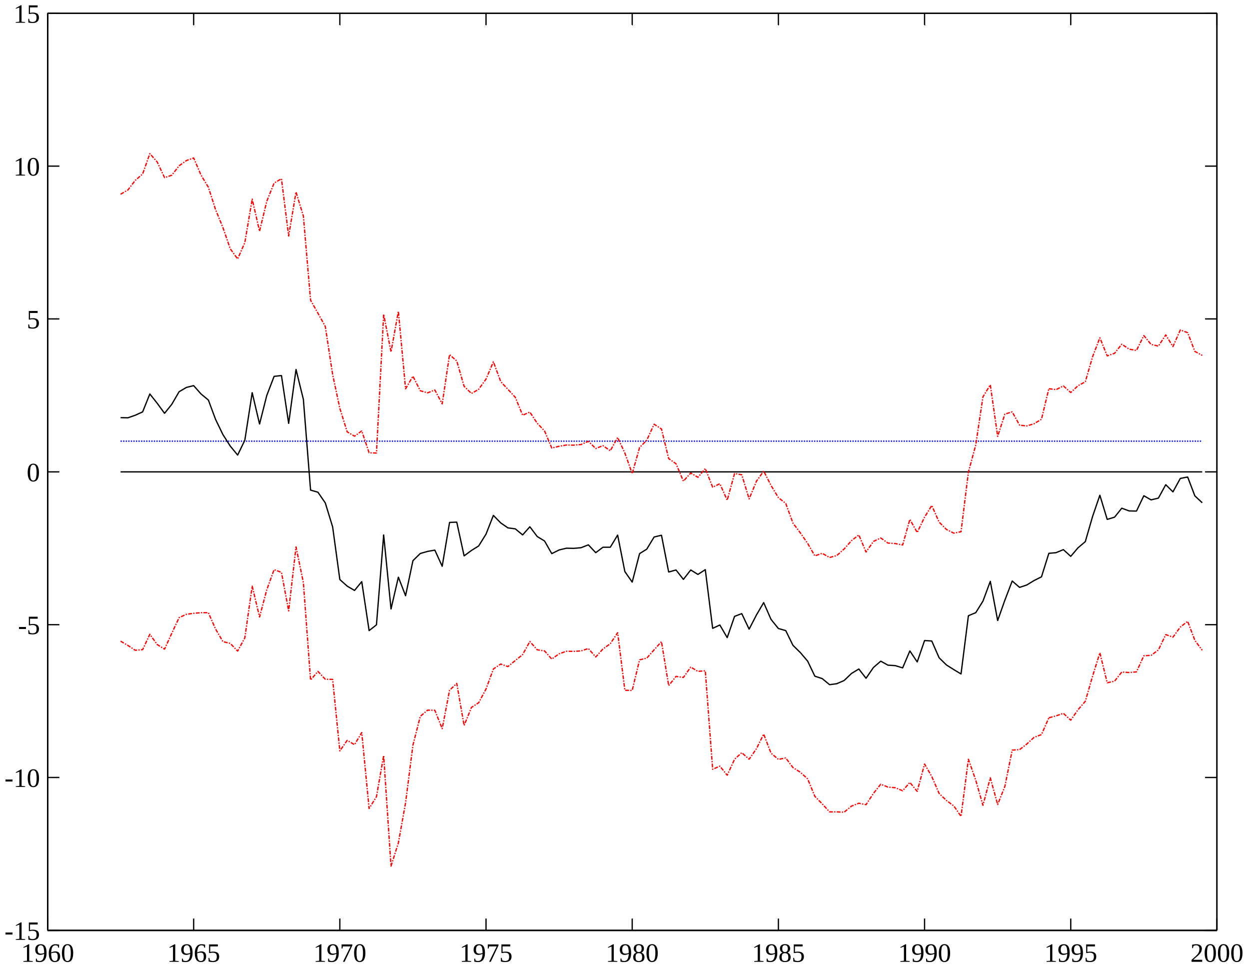

Figure 4.5: Estimates of beta with a rolling 10 year sample using the multible break model