Changing Risk Aversion, Unexpected Inflation, and Term Structure Variation*

Kevin Wang

October 1998

Abstract

This paper extends a consumption-based model of Campbell and Cochrane (1995) by including news on real returns as a source of variation in risk aversion and taking into account inflation-output correlation. An econometric specification using unexpected inflation as a proxy for the news is estimated with monthly data on consumption, inflation, and Treasury bonds. The results show that the changing risk aversion, which ranges between reasonable values and relates to observed business conditions, gives rise to an interesting explanation of the cyclical movements of the term structure. Unlike the standard utility model, this model generates positive ex-ante term premiums and can account for empirical rejections of the expectations hypothesis. Consistent with historical data, the model implies that the Fama and Bliss (1987) regressions have missed significant information about the ex ante term premium variation.

*Comments welcome. I thank Yacine Ait-Sahalia, John Cochrane, George Constantinides, Kent Daniel, Rajnish Mehra, and participants at the finance workshop at University of Chicago for helpful comments and suggestions.

1 Introduction

Campbell and Cochrane (1995) have recently proposed a consumption-based explanation of aggregate stock market behavior, aimed at a growing list of empirical puzzles against the standard power utility model. Their model has the power utility with an i.i.d. consumption growth driving process, but it is modified to include an external habit process. They demonstrate that the modification yields several remarkable results. Most notably, the model can explain with reasonable parameters the celebrated "equity premium puzzle" (Mehra and Prescott 1985, Hansen and Jagannathan 1992) and "risk free rate puzzle" (Weil 1989). Other results include that the model explains the long-horizon predictability of excess stock returns and predicts many of the difficulties surrounding the standard utility model.

This paper explores a simple extension of the Campbell-Cochrane model, motivated by the following considerations. First, central to the successes of the model is the changing risk aversion. Assuming a habit formation process in their framework is equivalent to specifying a process for variation in risk aversion. That is, the degree of risk aversion changes through time, driven by shocks to consumption growth. In this perspective, however, it is surely hard to defend that consumption growth in each period is the only source of news that affects risk aversion. Secondly, interdependencies between real and nominal processes may be important in the determination of asset returns. Campbell and Cochrane emphasize that to use their model to understand the term structure of nominal interest rates, one should tackle the problem of modeling inflation and then will be able to address the correlation between the stock returns and inflation. Thirdly, the local risk aversion coefficients implied by their chosen parameter values are still very large, which makes it of natural interest to investigate if an extension can retain the successes without invoking unreasonably high risk aversion.

The basic step of the extension is to include news on future real investment opportunities as an additional source of variation in risk aversion. This, of course, significantly alters evolution through time of the state dependent utilities. Consumption growth and inflation are assumed to be exogenously determined variables which follow a vector autoregression process. At least for the period from 1969 to 1993, a VAR(2) appears to be a reasonable model of the driving process. To render a tractable econometric analysis, I use the unexpected real return on the one-period nominal discount bond (or equivalently, unexpected inflation) as a proxy for news about real investment opportunity set. These lead to a complete model of the term structure of nominal interest rates in terms of the ultimate objectives of economic agents and the stochastic properties of observed forcing variables.

I estimate the extended model with monthly data on consumption, inflation, and Treasury bills, using a simulated least squares method. The results support the inclusion of the unexpected real returns as an additional source of variation in risk aversion. The estimates imply a very reasonable range of the changing risk aversion coefficient, suggesting that high risk aversion is not a necessary aspect of Campbell-Cochrane theory. More importantly, the implied variation in risk aversion is related to observed fluctuations in business conditions. The local risk aversion coefficient rises during the recessions of 1973-75 and 1979-82 but tends to be nonincreasing or decrease for the other subperiods after 1970.

When viewed as a measure or index of business conditions, the changing risk aversion gives rise to an interesting account of the cyclical movements of the term structure. The coefficient has significant correlations with a variety of term structure variables such as yields, yield spreads, expected holding period returns, and ex-ante term premiums. Most of the equilibrium relations implied by the correlations are consistent with empirical observations as well as simple intuitions of risk aversion and interactions between market forces (demand and supply).

The standard utility model with constant relative risk aversion, in addition to its poor fit with stock data, also has difficulties in explaining the movements of the term structure.

Backus, Gregory and Zin (1989) show with an artificial economy approach that the model cannot account for the empirical rejections of the expectations hypothesis. Term premium regressions produced in their economy with the standard utility invariably accept the expectations hypothesis, even for implausible parameter values. They also find that the model cannot produce positive average term premiums, in contrast to the fact that the averages observed at the short end of the term structure are reliably positive. Their finding is confirmed by Boudoukh (1993) who adds stochastic volatility of inflation and estimates the model with long-term Treasury bill data.

Unlike the constant risk aversion model, the extended Campbell-Cochrane model produces positive ex ante premiums and can account for empirical rejections of the expectations hypothesis with Treasury bill data. The Fama and Bliss (1987) term premium regressions produced by the model are very similar as those obtained from the data. In particular, the slope coefficients are close to one and all above two standard errors from zero. Moreover, the model indicates that the forward-spot spreads are not close to the ex ante term premiums on the long maturity bonds. This is consistent with the finding that expected excess returns on a variety of securities including long term Treasury bills tend to follow a countercyclical pattern (Fama and French 1989, and Fama 1990), and the fact that the forward-spot spreads tend to be procyclical.

Shiller (1990) observes that theoretical work on the term structure, while it has offered many insights, still does not allow us to say much about the term structure we observe. As a matter of fact, certain important questions such as what explains the cyclical pattern of the yield curve (level and shape) and the rejections of the expectations hypothesis are largely unanswered by theoretical models of the term structure. The exercise here suggests that to relax the constant risk aversion assumption may be a useful clue to develop more interesting models.

2 The Model

A basic hypothesis underlying the extension is that news about future economic conditions affects risk aversion. More explicitly, consumers' degrees of risk aversion depend on their views about future states of the economy, but the rational consumers adjust the views through time according to the news that pops up as the economy evolves. The forward-looking consumers tend to become more risk averse if new information indicates that future outlook of the economy is worse than previously expected. Conversely, the consumers tend to become less risk averse when good news arrives making them more optimistic about future economic conditions.

Details of the model begin with a representative consumer's utility function

$$\sum_{t=0}^{\infty} \delta^t u(C_t - H_t).$$

C_t is consumption, H_t is habit level, and δ is subjective discount factor. For simplicity, the logarithm function $u(x) = \ln x$ is used below. The level of habit H_t is assumed to be a function of aggregate or "the Joneses' consumption series, so that H_t is not a choice variable. If H_t is held fixed as C_t varies, the local coefficient of relative risk aversion is

$$RRA_t \equiv e^{\gamma_t} = -\frac{C_t u_{cc}}{u_c} = \frac{C_t}{C_t - H_t}. \tag{1}$$

The description of preferences may be completed by specification of a process for either log local risk aversion γ_t or habit H_t. To accommodate the hypothesis on variation in risk aversion, I assume that γ_t follows a mean-reverting AR(1) process

$$\gamma_t = \mu(1 - \phi) + \phi\gamma_{t-1} - e_t, \quad 0 < \phi < 1. \tag{2}$$

The innovation e_t summarizes or quantifies the effect (on risk aversion) of time t new information about future conditions of the economy. It moves up with arrival of good news and goes down when bad news is received.

The difference equation (2) provides an intuitive mechanism to generate time-varying risk aversion. The next step is to formulate specific contents of the information variable e_t. Not all the news that can forecast something in the future is of interest to the utility-maximizing consumer. A piece of new information can be so important as to affect the consumer's risk aversion only if it alters the value of his expected utility at the time. Otherwise, whatever is revealed by the information is irrelevant to the single-minded consumer.

In a multiperiod world, a consumer cares about not only future output levels (endowments) but also future real investment opportunities (distributions of one-period real returns on available assets), since asset investments are effective means to achieve an optimal allocation of consumption through time. Given a predicted path of future output levels and a conditionally fixed preference ordering, different views (due to different news) of asset real returns in future periods give rise to different outcomes of the consumer's optimization problem. For example, if new information tells that future real returns are higher than previously expected (and if everything else is as expected), the consumer will perceive that the future is better in terms of his own welfare since the maximum conditional expected utility based on the information is larger than he thought before. In short, news about future prospects of real investment opportunity set, like news about future outputs, is meaningful to the consumer since it provides new information relevant to his goal.

I assume that the innovation e_t is driven by the two sources of news as follows:

$$e_t = \lambda(\gamma_{t-1})\varepsilon^g_t + \theta(\gamma_{t-1})\varepsilon^r_t. \qquad (3)$$

Here $\varepsilon^g_t = g_t - E_{t-1}g_t$ and $\varepsilon^r_t = r_t - E_{t-1}r_t$. g_t is consumption growth rate (i.e., $g_t = \ln C_t - \ln C_{t-1}$). r_t is the real rate of return on a one-period instrument which pays one nominal dollar at time t. The unexpected change ε^g_t proxies for new information about future outputs and the shock ε^r_t serves as a proxy for news on real investment opportunities. $\lambda(\gamma_{t-1})$ and $\theta(\gamma_{t-1})$ are so-called sensitivity functions. The expectations are conditioned upon time

$t-1$ information set of the consumer. The model of Campbell-Cochrane (1995) corresponds to the special case that $\theta(\gamma_{t-1}) \equiv 0$, so that local risk aversion in their model responds only to news in consumption growth.

The assumption that news on real investment opportunity set affects the consumer's preferences (by changing his degree of risk aversion) is not in contradiction with the existing literature on multiperiod consumption-investment theories. For example, Fama (1970) considers uncertainty about investment prospects as one possible source to induce state dependent preferences. He states that 'the investment opportunities available in any given future periods may depend on events occurring in preceding periods, and such uncertainty about investment prospects induces state dependent utilities'. The view is apparently shared by Merton (1971,73) and other authors. As for the empirical ground, real returns in the data have significantly positive autocorrelations which decline very slowly at higher-order lags (e.g., Fama and Gibbons 1984)[1], and hence there is a basis to assume that an unexpected change in current real returns reveals new information about the future returns.

The use of the simplest investment instrument as a proxy variable facilitate an econometric analysis of the model. This is due to a relation implied by the Fisher's equality that unexpected real return on the instrument is just the negative of unexpected inflation[2]. That is, the Fisher's identity yields

$$\varepsilon^r_t = -\varepsilon^i_t. \qquad (4)$$

Here $\varepsilon^i_t = i_t - E_{t-1} i_t$ and i_t is inflation rate ($i_t = \ln p_t - \ln p_{t-1}$ where p_t is price level measure). Thus, the choice of the proxy enables one to extract the two shocks $(\varepsilon^g_t, \varepsilon^r_t)$ from

[1] This is also true with recent data. From April 69 to December 93, e.g., the autocorrelation of the real rate of return on the 1-month bill declines from 0.48 (at lag 1) to 0.18 at lag 36.

[2] The Fisher's identity is $R_{t-1} = r_t + i_t$, where R_{t-1} and r_t are nominal and real rates of return on the one-period instument, respectively; i_t is one period inflation rate. The identity implies that $R_{t-1} = E_{t-1} r_t + E_{t-1} i_t$ (provided that R_{t-1} is observed by the consumer at $t-1$) and hence

$$r_t - E_{t-1} r_t = -(i_t - E_{t-1} i_t).$$

an exogenously specified driving process of consumption growth and inflation rate. This leads to significant simplification for the purpose to obtain a tractable empirical analysis. Otherwise, one has to tackle either an endogenously determined variable (say, D/P) or a driving process of a higher dimension. To investigate the feasibility of other proxy choices is left for future work.

While (3) and (4) imply that variation in risk aversion is related to unexpected inflation, it is not necessarily the case that price level per se matters to the consumer. Here, unexpected inflation is merely an informational proxy. The model is consistent with a rational expectations view in which consumers care about only real variables, but price level is informative since markets set prices on the basis of rational forecasts of future real activity[3]. Fama (1981,82) argues with evidence that current inflation is caused by anticipated future real activity and the real activity is fundamental in the determination of asset real returns. His findings suggest there is nothing to preclude that unexpected inflation contains news about future real returns, even if real returns are determined within the real sector.

I specify the sensitivity functions as

$$\lambda(\gamma_{t-1}) = e^{\gamma_{t-1}} - 1; \quad \theta(\gamma_{t-1}) = \theta. \tag{5}$$

Campbell and Cochrane choose $\lambda(\gamma_{t-1})$ to satisfy three conditions: (i) the log risk free real rate is a linear function of γ_t; (ii) habit is predetermined at steady state $\gamma = \bar{\gamma}$; (iii) habit is predetermined near the steady state. It is unclear, however, if their choice of $\lambda(\gamma_{t-1})$ is differentiable over parameter space or not[4]. The sensitivity function $\lambda(\gamma_{t-1}) = e^{\gamma_{t-1}} - 1$,

[3] It appears possible, however, that the model can also be consistent with a world in which consumers have certain `nominal or monetary illusions'.

[4] In the notations here, their sensitivity function is defined as

$$\lambda(\gamma_{t-1}) = e^{\bar{\gamma}}\sqrt{1 + 2(\gamma_t - \bar{\gamma})} - 1;$$

if $\gamma_{t-1} \geq \gamma_{min} \equiv \bar{\gamma} + \frac{1}{2}(e^{-2\bar{\gamma}} - 1)$; $\lambda(\gamma_{t-1}) = 0$ if $\gamma_{t-1} < \gamma_{min}$. In a given sample, γ_{t-1} is a function of parameters, so is γ_{min}. The sample dependence makes it unclear if the sensitivity function is differentiable or not in the set of parameters which equate γ_{t-1} to γ_{min}. If not, challenging problems may arise in estimation and inference.

which is an approximation of the one they use (around the steady state), avoids the linearity restriction (i) and hence the possible nonsmoothness problem, but it still satisfies (ii) and (iii). To set $\theta(\gamma_{t-1})$ to be a constant is mainly motivated by the simplicity.

The external habit formation process implicitly defined by (1) and (2) may be viewed as follows. From (1), $H_t/C_t = 1 - e^{-\gamma_t}$. Hence the hypothesis (2) implies that the `keeping up with the Joneses' process is not isolated from the views about future prospects of the economy. The changing views have a negative effect on the habit process, consistent with the following story. When tough times are expected to come, the consumers are more conservative in consumption activities, so that a larger portion of their consumption is devoted to `stick to the Joneses'. If good times are predicted, the more optimistic consumers tend to behave more diversely. They care less about catching up with others, but are more interested in special activities that bring them utility (say, surpassing the Joneses).

With the additional assumptions (3) and (5), the habit process is approximated around the steady state by

$$\ln H_t \approx (1-\phi)h + E_{t-1}g_t + \phi \ln H_{t-1} + (1-\phi)\ln C_{t-1} - \theta_\lambda(\mu)\varepsilon^r_t;$$

where $h = \ln(1 - e^{-\mu})$ is the steady state value of $\ln H_t - \ln C_t$. The choice of $f_\lambda(\gamma_{t-1})$ ensures that the habit level H_t does not respond to contemporaneous consumption C_t around the steady state. Instead, H_t depends on the past consumptions. The real return shock affects the habit. If θ is positive, positive unexpected real return as good news has a negative effect on the habit level.

The real price of an period real riskless pure discount bond in the model is the conditional expectation of the pricing kernel $E_t m(t, t+n)$. Here the pricing kernel is

$$m(t, t+n) = \delta^n \exp\{\gamma_{t+n} - \gamma_t - (g_{t+1} + \dots + g_{t+n})\}.$$

Thus, time t equilibrium price of an period nominal discount bond with \$1 face value is $E_t M(t, t+m)$, with the nominal pricing kernel (or nominal marginal rate of substitution between time t and t+n) defined as[5]

$$M(t,t+n) = \delta^n \exp\{\gamma_{t+n} - \gamma_t - (g_{t+1} + \cdots + g_{t+n}) - (i_{t+1} + \cdots + i_{t+n})\}. \qquad (6)$$

The term $\gamma_{t+n} - \gamma_t$ marks the difference between the model and the standard utility model with constant relative risk aversion. The pricing kernels reduce to the common ones if $\gamma_{t+n} - \gamma_t = 0$ for any t.

To be able to price nominal bonds in the above framework, one need to make explicit assumptions on the joint process for consumption growth g_t and inflation rate i_t. Following Hansen and Singleton (1983) and Labadie (1989), I simply assume that g_t and i_t follow a vector autoregressive (VAR) process, which allows for codependencies of the processes and their innovations. The importance of interdependencies between real and nominal processes in determining interest rates is emphasized in a number of recent empirical works (e.g., Boudoukh 1993, Pennacchi 1991). The simple approach avoids thorny issues in endogenizing the joint process of inflation and consumption growth, but still takes into account effect of the codependencies on nominal bond prices.

More precisely, I assume that time t information set of the representative consumer is generated by exogenously determined variables $g_t, g_{t-1}, g_{t-2}, \ldots$ and $i_t, i_{t-1}, i_{t-2}, \ldots$. Denote $x_t = (g_t, i_t)'$. The process x_t is assumed to follow a stationary VAR(L) process

$$x_t = a + \sum_{l=1}^{L} \Phi_l x_{t-l} + \epsilon_t. \qquad (7)$$

Φ_l is a 2×2 matrix and a a 2×1 vector. $\epsilon_t = (\epsilon_t^g, \epsilon_t^i)'$ and ϵ_t is jointly normally distributed with covariance matrix Σ. With these assumptions on the driving process, we now have a complete model of the dynamics of the term structure of nominal interest rates.

[5] Derivation of these pricing relations is standard. See Campbell and Cochrane, Constantinides 1988, or Boudoukh 1993.

3 Estimation of the Model

3.1 The Estimation Method

I use a simulated least squares method (Wang 1995) to estimate the model. To illustrate the basic idea behind the method, I assume for the moment that the joint process of consumption growth and inflation defined in (7) is known. That is, the parameters of the VAR(L) process and L are known. In this case, one need to estimate only $\beta_0 \equiv (\mu;\nu;\phi;\theta;\delta)'$.

Let $M(t,t+n;\beta)$ denote the pricing kernel defined in (6). Given a sample of prices $\{Q_t^n\}_{t=1}^T$ of the nominal discount bond with n months to maturity, a natural thought is to apply the least squares method to estimate β_0. The least squares estimate of β_0 is the vector that minimizes $\sum_{t=1}^T q_t(\beta)$. Here

$$q_t(\beta) = (Q_t^n - E_t M(t,t+n;\beta))^2 :$$

The application of the least squares procedure is, however, frustrated by high dimensional numerical integrations since the model does not yield price functions of close form. For example, the price function $E_t M(t,t+n;\beta)$ for the 5-year bond ($n = 60$) is a 120-dimensional integral which cannot be expressed in close form. It is known that no conventional numerical methods (including the so-called 'brute force' simulation approach) are practical to evaluate a large number of integrals of such high dimensions throughout an iterative procedure used to solve for the estimates (e.g., McFadden 1989, Danielsson 1994).

The simulation method, which is aimed to circumvent the problem, proceeds as follows. For a given month t in the sample, obtain 2τ simulated pricing kernels $\{\hat M_i(t,t+n;\beta)\}_{i=1}^{2\tau}$ which are independent and identical in distribution to $M(t,t+n;\beta)$, conditional on all the observed information up to t. How to generate such a simulated kernel? First compute γ_t inductively by (2) with an assumed initial value of γ_1 and shocks (ε^g and ε^i) extracted from the known process of g_t and i_t. (γ_t is in the time t information set.) Given (g_t, i_t, γ_t) and their lagged values, generate a simulated n-month path $\{\hat g_{t+s};\hat i_{t+s};\hat\gamma_{t+s}\}_{s=1}^n$ by the transition equations (2) and (7), using a set of n random draws from $N(0;\Sigma)$ – the joint distribution of the two shocks. This produces a simulated pricing kernel

$$\hat{M}_1(t, t+n; \beta) = \delta^n \exp\{\hat{\gamma}_{t+n} - \gamma_t - (\hat{g}_{t+1} + \dots + \hat{g}_{t+n}) - (\hat{\imath}_{t+1} + \dots + \hat{\imath}_{t+n})\}:$$

The simulated least squares estimate $\hat{\beta}$ of β_o is then obtained as the minimizer of $\sum_{t=1}^{T} \hat{q}_t(\beta)$. Here

$$\hat{q}_t(\beta) = (Q_t^n - \tau^{-1} \sum_{i=1}^{\tau} \hat{M}_i(t, t+n; \beta))(Q_t^n - \tau^{-1} \sum_{i=\tau+1}^{2\tau} \hat{M}_i(t, t+n; \beta)):$$

Why use the product? The key is that the two items in the product are independent and identical in distribution, conditional on all the observed up to time t. In particular, the two have the same conditional mean which is equal to $Q_t^n - E_t M(t, t+n; \beta)$. Hence by the law of iterated expectations,

$$E\hat{q}_t(\beta) = Eq_t(\beta); \quad \forall t:$$

The unbiasedness enables a law of large numbers operating across the sample to control the joint effect of approximation errors $\hat{q}_t(\beta) - q_t(\beta)$ $(1 \le t \le T)$. More precisely, the second term on the right hand side of the identity below

$$T^{-1} \sum_{t=1}^{T} \hat{q}_t(\beta) = T^{-1} \sum_{t=1}^{T} q_t(\beta) + T^{-1} \sum_{t=1}^{T} (\hat{q}_t(\beta) - q_t(\beta))$$

converges to zero as the sample size increases. As a result, the simulation estimator $\hat{\beta}$ and the least squares estimator converge to the same limit for any given size τ.

The joint process of consumption growth and inflation rate is, of course, not known. The estimation procedure then consists of two stages. In the first stage, I estimates the vector autoregression process g_t and i_t by the maximum likelihood method. In the second stage, I obtain the simulated least squares estimate of β_o by using the fitted joint process as if the parameters of the process were estimated without error. The simulation estimator still has the standard large sample properties (consistency and asymptotic normality) without any restriction on simulation size (see Wang 1995 for further details and proofs).

A couple remarks are in order. First of all, why not specify $\lambda(\gamma_{t-1})$ as a constant to yield close form price functions? The model with constant sensitivity functions does look more elegant as it avoids the integration problem, but it has serious drawbacks. It can be shown that the constant sensitivity functions can not produce any time-variation in expected excess returns (e.g., Campbell and Cochrane). Moreover, I estimate the model by the least squares method but the estimates imply that the ex ante premiums in the model are negative. In the data, however, the observed average term premiums on bills of shorter-maturities are reliably positive (e.g., Fama 1984). This indicates that the constant sensitivity model is not a useful improvement over the standard utility models (see Backus et al. 1989).

What is wrong with a conventional approximation method? In a conventional approach, one simply replaces the price functions by some approximations. However, to ensure approximation errors within a reasonable range is often too costly or simply impractical. It is difficult, if possible at all, to gauge the magnitude of estimates' biases induced by the errors through an iterative numerical procedure. In general, the induced biases do not diminish as the sample size gets large, producing inconsistent estimates and tests of wrong sizes. These problems may be particularly serious in estimation of a term structure model, since changes in bond prices are so subtle (small relative to the prices) that moderate approximation errors may create large biases in estimates.

The simulation method utilizes the sample size to control the joint effect of approximation errors in estimation, while using the simulation size (τ) to reduce individual errors. As an econometric approach, it needs no restriction on the size of simulation to yield estimates with standard large sample properties. The flexibility enables one to select an affordable size in practice, taking into consideration the limits on memory capacity of available computers or the costs in terms of computing time.

3.2 The Consumption-Inflation Process

Monthly data from February 1959 to December 1993 are obtained from the CITIBASE database. Consumption is measured by real monthly expenditure on nondurables and services per capita, that is, $C_t = (GMCNQ + GMCSQ)/POP$. Following the common practice (e.g., Boudoukh 1993, Marshell 1992), the associated implicit price deflator is calculated and used as price level. Consumption growth (g_t) and inflation rate (i_t) are then obtained from the consumption and the price level.

The major concern here is if the interactions through the period between consumption growth and inflation rate can be well represented by a low order vector autoregression model. A simple VAR process certainly has its limitations. It can be a poor or misleading model of the forcing process, for example, if a significant structural change of the dynamic relation between consumption and inflation occurs in the middle of the period.

To check if a VAR(L) model can capture the consumption-inflation relation through the time interval, a simple approach is to compare subperiod estimation results. If the model is a reasonable one, the implied relations in subperiods should be similar. I divide the period into three: from February 1959 to January 1969, from February 1969 to September 1979, and from October 1979 to December 1993. The choice of October 1979 as division point is motivated by the well-known announced change in Federal Reserve policy at the time. The first and second periods are divided just to make each of them span a decade or so.

Table 1(a) reports maximum likelihood estimation results of VAR(2) processes for g_t and i_t for the subperiods and the overall period. VAR(1) processes are first fitted. The residuals are strongly autocorrelated, however, as indicated by the sample cross autocorrelation matrices (not shown here). For the VAR(2) processes, the fits appear reasonably clean and certainly a significant improvement over the VAR(1) cases. Most elements in the sample residual cross autocorrelation matrices of lags from one to twelve are within two standard errors from zero. There are a few outsiders, which are all within or around three approximated standard errors from zero. I have also tried VAR(3) processes, but there is little improvement.

Table 1(b) presents correlations between the actual processes g_t and i_t, between the fitted consumption growth and inflation, and between the two regression errors (shocks). The fitted values and the errors are calculated from the VAR(2) processes. For the last two subperiods, the consumption-inflation codependencies appear to be fairly stable. The correlations are all negative (more than or around two standard errors below zero), and of similar magnitude between the two periods. The first decade, however, exhibits a very different dynamic relation. The two extracted shocks as well as the actual processes are significantly positively correlated, and the correlation between the fitted processes is trivial.

The correlation patterns in the subperiods suggest there may be a significant structural change in the joint process of inflation and consumption growth during late 60's or early 70's. No matter whether such a change exists or not, the subperiod differences surely make it hard to defend that a simple VAR(2) process is adequate for the period from February 1959 to December 1993. For example, the fitted VAR(2) process for the overall period implies a correlation of -0.05 between the two shocks, which captures neither the positive correlation in the first decade nor the magnitude of negative correlations during the other two subperiods.

No attempt is made here to tackle this challenge. Instead, I simply focus the following analysis on the period from February 1969 to December 1993, during which it appears reasonable to assume that consumption growth and inflation follow a VAR(2) process.

3.3 Estimation of the Preference Parameters

Monthly price data on U.S. Treasury bonds with one-, three-, six-, twelve-, thirty-six-, and sixty-month to maturity are obtained from CRSP. The prices for the one-year, three-year, and five-year bonds are from Fama-Bliss discount bond file (FAMABLISPRI.DAT). The prices of the one-month, three-month, and six-month bills are derived from the yields in the six month yield file (FYLDAVE6.DAT).

With the six price series, a least squares estimate of β_0 (vector of the preference parameters) is the minimizer of the sum

$$\sum_{t=1}^{T}\sum_{j=1}^{6}\left(Q_t^{n_j} - E_t M(t, t+n_j; \beta)\right)^2; \tag{8}$$

provided that the joint process of consumption growth and inflation is known. Here n_j is the number of months to maturity for the n_j-month bond. While there is no particular efficiency property associated with it, the nonlinear least squares estimate seems intuitively appealing: it is the vector that minimizes the sum of squared pricing errors through time and across maturity.

Given a fitted VAR(2) process of g_t and i_t, I compute the simulated least squares estimates of the preference parameters. (The simulation approach is described in section 3.1.) A simulation size of two hundred is used [6]. That is, for each month (t) in the sample, I generate two hundred conditionally independent paths of future consumption growth, inflation rate, and log local risk aversion for the next five years ($\{(\hat{g}_{t+s}; \hat{i}_{t+s}; \hat{\alpha}_{t+s})\}_{s=1}^{60}$); each simulated path is based on a set of 120 random draws from the standard normal distribution [7]. Computation of the estimates encounters no problem when all six bonds or only long-term bonds are used. When using only short-term bills, the e04ucf optimizer (NAG library) fails to locate a

[6] This seemingly small size already consumes 30-50 percents of the memory capacities of the gsbphd machines. For the e04ucf optimizer (FORTRAN NAG library) I use, it takes about 2 or 3 hours (in a lucky day) to locate a solution. Alternatively, one may regenerate the same draws at each iterative step to avoid occupying a huge memory space; but such a procedure is too slow: it takes several days to find a minimum for the same simulation size.

[7] The draws are made independently across paths and across months.

minimum for some subperiods. I note however that when a stronger stationarity constraint, i.e., $0 \leq \phi \leq 0.99$, is imposed, the optimizer is able to find a solution for every case under consideration.

Table 2 reports the simulated least squares estimation results. The standard errors are adjusted for autocorrelation and heteroskedasticity of regression residuals, by the Newey-West method with 24 lags[8]. The estimates are economically plausible. The steady state coefficients of relative risk aversion implied by the estimated γ's range between 1.6 and 2.5. The estimates of θ are all positive, indicating that the consumer treats positive unexpected real returns or negative unexpected inflation as good news. Estimates of the subjective discount factor δ are, as expected, very close to one.

The estimates support the hypothesis that the consumer takes seriously the news contained in unexpected inflation about real returns. The estimates of the sensitivity parameter θ are above four standard errors from zero for the periods 1969-93 and 1959-69. Table 2 (a) displays a large difference in estimates of θ between the first decade and the period thereafter, which is not surprising since the dynamic relations in the two periods between consumption growth and inflation are quite different. The constrained estimates for the cases associated with the long term bonds are fairly stable. For the short end, the estimated θ for the period 1979-93 is statistically insignificant, which is the only exception in the constrained cases.

[8] I also tried 12 lags and 36 lags. There is little difference.

3.4 The Changing Risk Aversion and the Habit

I use the simulation estimates of the preference parameters (reported in Table 2 (a)(i)) to obtain a look at the variation in risk aversion and the external habit process during the period from April 1969 to December 1993. Descriptive statistics for local risk aversion ($RRA_t = e^{\gamma_t}$), proportion of consumption devoted to habit ($100 \times H_t/C_t$), and change in habit relative to consumption ($100 \times (H_t - H_{t-1})/C_{t-1}$) are presented in Table 3.

The local risk aversion coefficient varys through time and ranges between reasonable values (1.30 and 7.74). As expected, the information variable e_t is positively correlated with the unexpected consumption growth, but negatively correlated with the unexpected inflation and the change in risk aversion. The differences in magnitudes of the correlations suggest that news on real returns (or unexpected inflation) is the more dominant of the two shocks that move the degree of risk aversion. Combined with the estimation results, we see that the inclusion of the additional source of news is both statistically and economically significant.

Driven by the composition of the two shocks, the local risk aversion RRA_t has a novel behavior. Its correlations with consumption growth and inflation are much smaller in absolute value than those of ΔRRA_t with the variables. Interestingly, the correlations of RRA_t and ΔRRA_t with the real returns are of similar magnitude, but of different signs. This says that high risk aversion tends to be associated with high realized real returns, but high returns tend to drive down the risk aversion.

The fraction of consumption used to `keeping up with the Joneses' is neither trivial nor extremely high. The proportion H_t/C_t, as it is equal to $1 - e^{-\gamma_t}$, varies in the same direction as the local risk aversion. (The computed correlation between the two is 0.90 for the period.) The ratio $(H_t - H_{t-1})/C_{t-1}$ is used to measure if changes in habit are drastic or significant to the representative consumer. The table shows that the changes during the period are small in terms of the consumption level.

4 The Term Structure Variation

This section presents term structure relations implied by the estimated Campbell-Cochrane model. The first step is to compute the implied term structure variables. Since prices of the bills (and hence returns etc.) in the model do not possess close form expressions, some numerical approximation procedure has to be used.

I compute the model prices and returns via simulation. Given the parameter estimates, one can now afford a large simulation size in computation of the prices and returns as one need only to evaluate the high dimensional integrals once. The prices are calculated as simple averages of simulated nominal pricing kernels. A simulation size of 50,000 is used. That is, I use fifty thousand simulated paths generated in exactly the same way as in estimation to obtain fifty thousand simulated kernels and then take the average.

Following Backus, Gregory, and Zin (1989), I use expected holding period returns and expected term premiums defined in terms of simple returns. For example, the time t expected m-month excess return on a bill with $m + n$ months to maturity is

$$E_t h(m + n; m : t + m) - 1/Q^m_t;$$

where $h(m + n; m : t + m) = Q^n_{t+m}/Q^{m+n}_t$, and Q^m_t is the time t price of a m-month bill. As in Fama and Bliss (1987) and Shiller (1993), the expected excess return may be referred to as the expected (holding-period) term premium. The use of simple returns instead of continuously compounded ones is purely due to a technical difficulty. To calculate expected continuously compounded returns, one needs to evaluate $E_t \ln Q^n_{t+m}$, which nests a conditional expectation (i.e., $Q^n_{t+m} = E_{t+m} M(t + m; t + m + n)$) nonlinearly in another expectation operator. There seems, however, no practical way to obtain reliable approximation of such nonlinearly nested expectations.

Computation of expected returns and term premiums defined in terms of simple returns does not encounter the difficulty since $E_t Q^n_{t+m}$ can be simplified by the law of iterated expectations: $E_t Q^n_{t+m} = E_t E_{t+m} M(t+m; t+m+n) = E_t M(t+m; t+m+n)$. To compute $E_t Q^n_{t+m}$, I use each of the fifty thousand simulated five-year paths to yield a simulated pricing kernel which is identical in time-t conditional distribution to $M(t+m; t+m+n)$. $E_t Q^n_{t+m}$ is then obtained as the average of the fifty thousand simulated kernels.

For the following analysis, all the model variables are obtained from the simulation estimates of the preference parameters (Table 2 (a)(i)) and the maximum likelihood estimates of the VAR(2) process of consumption growth and inflation for the period from February 1969 to December 1993.

4.1 The Term Structure and the Changing Risk Aversion

Yields and Yield Spreads

Table 4(a) presents descriptive statistics of yields in the model. The unconditional term structure is nearly flat and the standard deviation of the implied yields decreases with maturity. The correlations among the yields are very close to one (not shown here). The yield curve is positively correlated with the representative consumer's varying degree of relative risk aversion and inflation, but has negative correlation with consumption growth. The correlation of the yields with RRA_t increases with maturity; but the correlations with g_t and i_t decrease in absolute value with maturity.

Similar correlation patterns are observed in the market yields. In particular, Table 4(b) shows that correlations between the actual yields and the local risk aversion coefficient are impressive, rising from 0.76 to 0.90. The spreads between the 5-year and the 1-year bonds also exhibit (more or less) similar patterns of correlations as in the model. Correlations between the market and the model yields rise with maturity, ranging from 0.73 to 0.90, which appears to suggest that variation in longer-maturity yields is better captured by the model.

The strong positive correlation between the yield curve and the local risk aversion is consistent with a simple explanation. During periods of high or increasing risk aversion, larger expected real returns are needed to induce consumers to hold the bonds. Given their predictions of future inflation rates, yields have to move up in order to increase the expected real returns. Conversely, when the degree of risk aversion remains low or dropping, lower yields (and hence smaller expected real returns) are sufficient to attract enough demand for the bonds.

The changing risk aversion is negatively correlated with the slope of the yield curve (spread between annualized yields). The negative correlation is consistent with the intuition that when risk aversion is high, consumers may prefer to settle with the less uncertain long-term real return by buying a long-term bond rather than take the more risky alternative of rolling over a shorter-term instrument. If a sufficient number of people behave this way, it is plausible that yields on the longer-maturity bonds tend to be depressed. To make it clearer, consider an imaginary case in which inflation is perfectly predictable or deterministic. In such a world, the bonds are risk free, but rolling a bond over involves risks due to the uncertain future prices of the bond.

Expected Term Premiums

Table 5(a) presents descriptive statistics of implied ex ante term premiums. Both means and standard deviations of the premiums rise with maturity. Interestingly, the premiums are always positive through the sample period from April 1969 to December 1993. This implication of the Campbell-Cochrane model is novel since recent works show that the standard power utility model cannot produce positive average term premiums observed at the short end of the term structure (Backus, Gregory, and Zin 1989, Boudoukh 1993)[9]. The means of the implied premiums, however, are still small. In the data, for example, the sample mean of ex post three month excess returns on the six month bill is around 0.5 percent (annu-

[9] Backus et. al. point out that introducing durable goods as in Dunn and Singleton (1986) is not a solution since their model produces positive, not negative, correlations of marginal rates of substitution.

alized simple return), which is about five times of the corresponding average in the model. Variability of the implied ex ante premiums is very small, in contrast to several empirical studies (e.g., Backus et. al. 1989, Fama and Bliss 1987, Startz 1982) which suggest that ex ante premiums in the bills vary considerably over time.

The expected term premiums are positively correlated with the local coefficient of risk aversion and inflation, and the correlations are stronger for the one year horizon than for the holding period of three months. This pattern seems to suggest that the changing degree of risk aversion matters more when more substantial risk is involved. If expected term premiums are viewed as compensation for risk, the positive correlations between the premiums and the local risk aversion are easy to understand. When the degree of risk aversion increases, higher premiums are required. Conversely, smaller premiums are good enough in periods of lower risk aversion. The positive correlations are, however, far from being close to one, which suggests there exist other determinants[10] of the term premium variation.

A technical issue deserves some attention. Since the ex ante premiums and their variability are small, one may naturally question how reliable the computation of the premiums is. The same question applies to the differences between expected holding period returns reported in Table 5(c). The issue is left for future consideration.

Expected Holding Period Return

Table 5(b) and 5(c) report descriptive statistics for implied expected holding period returns and their differences. The unconditional means and standard deviations of the expected returns are fairly similar across maturity. The conditional term structure varies through time and is always upward-sloping. As shown in panel (c), the conditional expected returns always rise with maturity but differentials between the returns remain very small through time. The expected returns are positively correlated with the changing degree of risk

[10] A obvious candidate is time-varying risk, as measured by conditional variances of excess returns. However, computation of the conditionals is frustrated by nonlinearly nested expectations.

aversion and inflation, but negatively correlated with consumption growth. In particular, the correlations with the local risk aversion coefficient are impressively high: 0.91 for the holding period of three months and 0.97 for the one year period. The correlation patterns are not hard to understand since expected returns are sums of yields and expected term premiums. Interestingly, the differences among the expected returns also have significantly positive correlations with the local risk aversion coefficient (and the correlations are larger for the 1-year holding period than for the 3-month horizon) but have little to do with the processes of inflation and consumption growth.

Relation to Business Conditions

It seems intuitive to think there is a relation between changing economic conditions and variation risk aversion. In particular, periods of strong economic conditions may be associated with low or decreasing risk aversion and recessions should be associated with high or increasing risk aversion. In the Campbell-Cochrane economy, variation in risk aversion is driven by changes in the state variables (inflation and consumption growth) and hence the local risk aversion coefficient may be regarded as a reasonable measure of the conditions of the economy. Interestingly, the implied coefficient is related to actual business conditions. Figure 4.1 plots the time path of the coefficient for the period from April 69 to December 93. The risk aversion rises during the recessions of 1973-75 and 1979-82. For periods associated with stable or strong economic activity, it tends to decrease slightly (1969-73, 1986-93), or stay around (1975-79), or drop sharply (1983-86).

With the above relation, the earlier correlation results imply that the yields and the expected term premiums (and hence the expected returns) tend to be low or decreasing when economic conditions are strong, but high or increasing during recessions. The prediction of term premium variation is consistent with the findings of Fama (1990) and Fama and French (1989). Variation of the market yields since 1969 seems also consistent with the model prediction. The yields are high or increasing during the recessions of 1973-75 and 1979-1982, while they tend to be low or decreasing in periods of strong economic activity, for example, 1970-72, 1975-78, 1983-85 (see Figure 4.1, or Figure 1 of Fama and Bliss).

4.2 Rejections of the Expectations Hypothesis

Fama-Bliss Regressions

An empirically important issue is what observables track time-variation in expected term premiums. Much of the evidence in the literature about variation in ex ante term premiums or rejections of the expectations hypothesis are based on forward rates; an example is the well-known work of Fama and Bliss (1987). Fama and Bliss regress one-year excess returns of two- to five-year bills on corresponding forward-spot spreads. The estimated slope coefficients are all positive and close to one. Their conclusion is that the forward-spot spreads track variation through time in one-year expected excess returns.

Table 6(a) reports regression results in the model and in the data. As for the dependent variables, I use not only continuously compounded excess returns but simple ones as well, since the computed ex ante premiums are in terms of simple returns. The regression results in the model are much like those obtained from the data. In particular, the estimates of the slope coefficients range between 1.0 and 2.0 and all are above or around two standard errors from zero. The standard errors are obtained by Newey-West method with lag length eleven, i.e., they are adjusted for autocorrelation and heteroskedasticity. Apparently, the model can account for the term premium regression tests which reject the expectations hypothesis.

The standard utility model with constant risk aversion, on the contrary, is not able to explain the rejections. Backus et al. (1989) use an artificial economy to show that term premium regressions with simulated data, produced by their economy with the standard utility, invariably accept the expectations hypothesis. This happens even when they use implausible parameter values.

The Forward-Spot Spreads and the Term Premiums

The forward-spot spreads in the model have strong negative correlation with the local risk aversion and the yields. Table 6(b) shows that the correlations range from -0.42 to -0.87. In the Campbell-Cochrane economy, the ex ante term premiums are by construct positively correlated with the risk aversion coefficient. These indicate that the forward-spot spreads must differ considerably from the ex ante premiums. In other words, the spreads must have missed significant information about the term premium variation. If not so, the spreads imply that the ex ante premiums are procyclical (negatively correlated with the coefficient RRA_t).

In the data, the term premium regressions also give rise to such an implication. Fama and Bliss find that at least after 1970, there seems to be a relation between the sign of the forward-spot spread and the business cycle: positive forward-spot spreads tend to be associated with periods of strong business activity and negative spreads tend to occur during recessions. If the forward-spot spreads really capture the time-varying ex ante premiums, the relation implies that the expected excess returns tend to be high when economic conditions are strong and low during recessions. This is, however, not consistent with the counter-cyclical pattern identified by Fama (1990) and Fama and French (1989).

The Term Premium Regressions with Yields

The implication that the term premium regressions are unreliable to track term premium variation raises an interesting question: is this really the case in the data? I use a set of simple tests to check if the actual data supports the implication. First, I regress the residuals in the Fama-Bliss term premium regressions on (annualized) yields of the one to four year bills. The idea is familiar: the residuals should not contain any significant and currently available information if the regressions capture the expected premiums. The choice to use the yields is motivated by the facts that yields have strong negative correlations with the forward-spot spreads and yields are certainly observable information to market participants.

The residual regression results are presented in Table 7(a). The estimated slopes of the yields are all positive. The t-statistics of the slopes based on the Newey-West standard errors (with lag length eleven) are 2.1, 2.4, 2.3, and 1.9 in the four regressions. Alternatively, the yields may be put directly into the term premium regressions:

$$p(n;m : t+12) = a + b_1 f(n;m : t) + b_2 y(n - m : t) + \epsilon(t+12):$$

Here $p(n;m : t+12)$ and $f(n;m : t)$ are (ex post) term premiums and forward-spot spreads defined exactly as in Fama and Bliss. $y(n - m : t)$ is the yield on the bond with $(n - m)$ month to maturity. Table 7(b) shows that the slope estimates for the yields are slightly larger than in the residual regressions. The t-values obtained from the Newey-West standard errors are 2.4, 2.8, 2.7, and 2.1, respectively. The R^2's in the four regressions all increase substantially after the yields are included (compare with Table 6(a)). Since the forward-spot spreads and the yields have strong negative correlations (in levels as well as in changes, Table 7(c)), the results of the two sets of regressions clearly support the implication that the Fama-Bliss regressions must have missed significant information about variation in the expected excess returns for the period from 1969 to 1993.

5 Tests and Price Errors

5.1 A Set of GMM Tests

I use a sequential or two-stage generalized method of moment (GMM) procedure to obtain Hansen's J test statistics. In the first stage, the maximum likelihood estimates of the consumption-inflation process are calculated. The GMM estimate of the preference parameter vector β_0 is obtained in the second stage by minimizing the distance of

$$w_t \otimes (M(t,t+p;\beta;\hat{\xi})R^n_{t+p} - 1)$$

to the null vector. Here $R^n_{t+p} = Q^{n-p}_{t+p}/Q^n_t$, and Q^x_s is the price of the x-month bill at time s. $M(t,t+p;\beta;\hat{\xi})$ is the pricing kernel and $\hat{\xi}$ denotes the ML estimate of the parameter vector of the driving process. w_t is the instrument vector. (For the choice of the distance matrix and other details of the sequential procedure, see Ogaki 1993).

The procedure is applied to each of six returns listed in Table 8, using the vector of instruments $w^0_t = (g_t, i_t, g_{t-1}, i_{t-1}, g_{t-2}, i_{t-2}, y_t)$. Here y_t is the yield on the one month bill. To choose only one yield is due to the fact that yields are highly correlated, so that adding another yield is not effective as the two moment restrictions (with the two yields as instrument variables) are too similar. In addition, the yields on long-term bonds may not be valid instrumental variables, since they are contaminated with possibly substantial measurement errors.

The estimates obtained from the one-year returns on the two- to five-year bonds are similar to the simulated least squares estimates; in particular, the estimates of the sensitivity parameter θ are all above three standard errors from zero. The corresponding chi-square statistics are quite low, only in the neighborhood of one. The implied ranges of local risk aversion coefficient for all the six cases are very reasonable. The estimate of θ implied by the 3-month return on the six-month bill is much smaller and the associated J test rejects the model. The χ^2 statistic from the six-month return on the one-year bond lies in the middle, not large enough to reject the model.

This set of tests appears to support the model, at least for the longer-maturity bonds. However, the small χ^2 statistics obtained in the long-term bond cases may be because the two-step estimates are too noisy or because the test procedure in these cases is lack of power. Tests with multiple returns are to be conducted. Given the assumed driving process (or explicitly specified information sets), however, we have a direct way to assess the model: check the price errors.

5.2 Implied Price Errors

A price error is defined as difference between the market and the model prices per $100 face value. Table 9 reports the price errors implied by the simulated least squares estimates of the preference parameters and the fitted VAR(2) process used in section 4.

The price errors are highly correlated through time and across maturity. The large correlations may be partly due to model misspecification and partly due to measurement errors. As in any model, it is likely (if not certainly) that misspecification of some type exists. For example, the AR(1) process of γ_t, choices of the proxy variables, the sensitivity functions, or the assumed driving process, all are possible sources of misspecification. On the other hand, mismeasurement of consumption, inflation, and prices of the bills may also be a significant factor to produce correlated price errors.

Another pattern of the errors is that the magnitude increases with maturity. The errors are large in economic sense. Nonetheless, the price errors for the longer-maturities are similar in size to those of a two-factor CIR model estimated by Pearson and Sun (1994), who use monthly data in a shorter sample period (from December 1971 to December 1986). For example, the standard deviations of price errors in the estimated CIR model (p.1295 Pearson and Sun) for the 3-year, 4-year, and 5-year bills are 1.95, 2.95, and 3.95, with maximums 5.08, 8.76, and 9.62, respectively.

6 References

1. Abel, Andrew B., 1990, "Asset Prices under Habit Formation and Catching Up with the Joneses", American Economic Review Papers and Proceedings 80, 38-42.

2. Backus, David K., Allan W. Gregory and Stanley E. Zin, 1989, "Risk Premiums in the Term Structure: Evidence from Artificial Economies", Journal of Monetary Economics 24, 371-99.

3. Boudoukh, Jacob, 1993, "An Equilibrium Model of Nominal Bond Prices with Inflation-Output correlation and Stochastic Volatility", Journal of Money, Credit, and Banking 25, 637-665.

4. Campbell, John Y. and John H. Cochrane, 1995, "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior", NBER Working Paper No. 4995.

5. Cochrane, John H. and Lars Peter Hansen, 1992, "Asset pricing Lessons for Macroeconomics" in Oliver Blanchard and Stanley Fischer, Eds. 1992 NBER Macroeconomics Annual.

6. Constantinides, George, 1988, "Theory of Valuation: Overview and Recent Developments", in Bhattacharia, Sudipto and George Constantinides, Eds. Theory of Valuation, Frontiers of Modern Financial Theory, Vol. 1, 1-23.

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

8. Danielsson, Jon, 1994, "Stochastic Volatility in Asset Prices: Estimation with Simulated Likelihood", Journal of Econometrics 64, 375-400.

9. Dunn, Kenneth B. and Kenneth J. Singleton, 1986, "Modelling the Term Structure of Interest Rates under Non-separable Utility and Durability of Goods", Journal of Financial Economics 17, 27-55.

10. Fama, Eugene F., 1970, "Multiperiod Consumption-Investment Decisions", American Economic Review 60, 163-174.

11. Fama, Eugene F., 1981, "Stock Returns, Real Activity, Inflation, and Money", American Economic Review 71, 545-565.

12. Fama, Eugene F., 1982, "Inflation, Output, and Money", Journal of Business 55, 201-231.

13. Fama, Eugene F., 1990, "Term Structure Forecasts of Interest Rates, Inflation and Real Returns", Journal of Monetary Economics 25, 59-76.

14. Fama, Eugene F. and Robert R. Bliss, 1987, "The Information in Long-Maturity Forward Rates", American Economic Review 77, 680-692.

15. Fama, Eugene F. and Kenneth R. French, 1989, "Business Conditions and Expected Returns on Stocks and Bonds", Journal of Financial Economics 25, 23-49.

16. Fama, Eugene F. and Michael R. Gibbons, 1984, "A Comparison of Inflation Forecasts", Journal of Monetary Economics 13, 327-348.

17. Hansen, Lars Peter and Ravi Jagannathan, 1992, "Restrictions on Intertemporal Marginal Rates of Substitution Implied by Asset Returns", Journal of Political Economy 99, 225-262.

18. Hansen, Lars Peter and Kenneth J. Singleton, 1983, "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns", Journal of Political Economy 91, 249-268.

19. Labadie, Pamela, 1989, "Stochastic Inflation and the Equity Premium", Journal of Monetary Economics 24, 277-298.

20. Lucas, Robert E., Jr., 1982, "Interest Rates and Currency Prices in a Two-Country World", Journal of Monetary Economics 10, 335-360.

21. Marshall, David, 1992, "Inflation and Asset Returns in a Monetary Economy", Journal of Finance 47, 1315-1342.

22. McFadden, D., 1989, "A Method of Simulated Moments for Estimation of Discrete Response Models without Numerical Integrations", Econometrica 57, 995-1026.

23. Mehra, Rajnish and Edward Prescott, 1985, "The Equity Premium Puzzle", Journal of Monetary Economics 15, 145-161.

24. Merton, Robert C., 1971, "Optimum Consumption and Portfolio Rules in a Continuous-time Model", Journal of Economic Theory 3, 373-413.

25. Merton, Robert C., 1973, "An Intertemporal Capital Asset Pricing Model", Econometrica 41, 867-887.

26. Ogaki, Masao, 1993, "Generalized Method of Moments: Econometric Applications", in G. S. Maddala, C. R. Rao and H. D. Vinod, Eds. Handbook of Statistics, Vol. 11., 455-488.

27. Pearson, Neil D. and Tong-sheng Sun, 1994, "Exploiting the Conditional Density in Estimating the Term Structure: An Application to the Cox, Ingersoll, and Ross Model", Journal of Finance 50, 1279-1304.

28. Pennacchi, George G., 1991, "Identifying the Dynamics of Real Interest Rates and Inflation: Evidence Using Survey Data", Review of Financial Studies 4, 53-86.

29. Shiller, Robert J., 1990, "The Term Structure of Interest Rates", chap. 13 in Handbook of Monetary Economics, edited by Frank Hahn and Benjamin Friedman, vol. 1.

30. Startz, Richard, 1982, "Do Forecast Errors or Term Premia Really Make the Difference between Long and Short Rates?", Journal of Financial Economics 10, 323-329.

31. Wang, Qing, 1995, "Simulated Least Squares and Likelihood Estimation of Markov Models of Bond Prices".

32. Weil, Philippe, 1989, "The Equity Premium Puzzle and the Risk-Free Rate Puzzle", Journal of Monetary Economics 24, 401-421.

Table 1: Fitted VAR(2) Processes for Consumption Growth and Inflation

(a) Parameter estimates of the autoregressions (std. err. in parentheses):

Subperiod 1: Feb. 59 – Jan. 69

	intercept	g_{t-1}	i_{t-1}	g_{t-2}	i_{t-2}
consumption growth g_t	3.11e−3	−0.277	0.110	−0.056	−0.229
	(0.64e−3)	(0.090)	(0.178)	(0.090)	(0.179)
inflation rate i_t	1.90e−3	−0.025	−0.051	−0.027	0.103
	(0.33e−3)	(0.046)	(0.092)	(0.047)	(0.092)

Subperiod 2: Feb. 69 – Sept. 79

	intercept	g_{t-1}	i_{t-1}	g_{t-2}	i_{t-2}
consumption growth g_t	6.33e−3	−0.342	−0.207	−0.185	−0.474
	(1.07e−3)	(0.088)	(0.153)	(0.089)	(0.154)
inflation rate i_t	2.12e−3	−0.012	0.323	−0.048	0.324
	(0.60e−3)	(0.049)	(0.086)	(0.050)	(0.087)

Subperiod 3: Oct. 79 – Dec. 93

	intercept	g_{t-1}	i_{t-1}	g_{t-2}	i_{t-2}
consumption growth g_t	2.68e−3	−0.337	−0.144	0.040	−0.163
	(0.61e−3)	(0.077)	(0.110)	(0.076)	(0.109)
inflation rate i_t	1.76e−3	0.026	0.310	0.055	0.232
	(0.42e−3)	(0.053)	(0.075)	(0.052)	(0.074)

The joint period of the last two: Feb. 69 – Dec. 93

	intercept	g_{t-1}	i_{t-1}	g_{t-2}	i_{t-2}
consumption growth g_t	3.38e−3	−0.300	−0.112	−0.020	−0.227
	(0.54e−3)	(0.058)	(0.090)	(0.058)	(0.091)
inflation rate i_t	1.66e−3	0.020	0.346	0.018	0.291
	(0.34e−3)	(0.036)	(0.057)	(0.036)	(0.057)

The overall period: Feb. 59 – Dec. 93

	intercept	g_{t-1}	i_{t-1}	g_{t-2}	i_{t-2}
consumption growth g_t	3.38e−3	−0.283	−0.063	−0.033	−0.267
	(0.39e−3)	(0.048)	(0.075)	(0.048)	(0.075)
inflation rate i_t	1.37e−3	−0.021	0.326	−0.014	0.340
	(0.24e−3)	(0.030)	(0.046)	(0.030)	(0.045)

Table 1 (continued)

(b) Correlations between consumption growth and inflation, between the fitted values, between the unexpected changes:

periods	actual processes	fitted processes	implied shocks	(std. errors)
Feb 59 − Jan 69	0.18	−0.01	0.23	(0.09)
Feb 69 − Sep 79	−0.27	−0.53	−0.19	(0.09)
Oct 79 − Dec 93	−0.17	−0.35	−0.16	(0.08)
Feb 69 − Dec 93	−0.19	−0.47	−0.14	(0.06)
Feb 59 − Dec 93	−0.13	−0.46	−0.05	(0.05)

Note:

The data are monthly observations from CITIBASE database. Consumption is real monthly expenditure on nondurables and services divided by population. Price level is the associated implicit deflator of the consumption. Consumption growth and inflation rate are obtained from the consumption and the price level. The VAR(2) processes are estimated by the maximum likelihood method. The 'actual processes' denote the two time series g_t and i_t. The fitted processes and implied shocks are fitted regression values and regression errors, respectively, obtained from the estimated VAR(2) processes.

Table 2: Simulation Estimates of the Preference Parameters

(a) Estimates obtained from the six bonds:

(i) February 69 $-$ December 93

parameter	μ	ϕ	θ	δ
estimate	0.687	0.996	22.881	0.9995
(std. err.)	(0.766)	(0.001)	(4.206)	(0.0027)

(ii) February 59 $-$ January 69

parameter	μ	ϕ	θ	δ
estimate	0.527	0.926	3.466	1.0003
(std. err.)	(0.072)	(0.159)	(2.711)	(0.0007)

(iii) February 59 $-$ December 93

parameter	μ	ϕ	θ	δ
estimate	0.698	0.983	5.614	0.9993
(std. err.)	(0.085)	(0.004)	(1.281)	(0.0005)

(b) Estimates under the constraint $0.00 \leq \phi \leq 0.99$:

February 69 $-$ September 79:

parameter	6 bonds		1m-3m-6m		1y-3y-5y	
	estimate	(std. error)	estimate	(std. error)	estimate	(std. error)
μ	0.860	(0.102)	0.741	(0.131)	0.862	(0.125)
ϕ	0.990	(—)	0.990	(—)	0.990	(—)
θ	9.110	(2.937)	11.910	(3.873)	9.055	(3.467)
δ	1.0001	(0.0007)	1.0006	(0.0008)	1.0001	(0.0006)

Table 2 (continued)

October 79 ¡ December 93

parameter	6 bonds		1m-3m-6m		1y-3y-5y	
	estimate	(std. error)	estimate	(std. error)	estimate	(std. error)
½	0 884	(0 188)	0 .702	(0 200)	0 874	(0 174)
Á	0 990	(| {)	0 990	(| {)	0 990	(| {)
µ	7.784	(2 260)	3 010	(2 931)	7.786	(2 977)
±	1 0001	(0 0011)	1 0022	(0 0017)	1 0001	(0 0009)

February 69 ¡ December 93

parameter	6 bonds		1m-3m-6m		1y-3y-5y	
	estimate	(std. error)	estimate	(std. error)	estimate	(std. error)
½	0 868	(0 079)	0 .777	(0 083)	0 869	(0 082)
Á	0 990	(| {)	0 990	(| {)	0 990	(| {)
µ	11 147	(1 898)	9 813	(2 011)	11 196	(1 934)
±	0 9993	(0 0006)	1 0002	(0 0003)	0 9993	(0 0006)

Note: The notation `1m-3m-6m' (or `1y-3y-5y') means that the prices of the 1-month, 3-month, and 6-month (or 1-year, 3-year, and 5-year) discount bonds are used to obtain the estimates below. `6 bonds' denote the case in which the prices of all the six bonds are used. The price data are from CRSP. In each of the cases, the simulation estimate of the parameter vector is the vector which minimizes a simulation-based approximation of the sum of squared pricing errors through time and across the bonds being used. The size of simulation is 200. (An illustration of the method and further details of the estimation are provided in subsections 3 1 and 3 3.)

Table 3: The Changing Risk Aversion and the Habit

(April 69 – December 93)

Descriptive statistics:

	RRA_t	e_t	$HB1_t$	$HB2_t$
mean	3.15	0.00	61.0	0.03
std. dev.	1.63	0.05	15.5	1.12
min	1.30	-0.28	23.1	-7.54
max	7.74	0.29	87.1	5.86
1st autocor.	0.99	-0.03	0.98	-0.08

Cross correlations:

	ΔRRA_t	e_t	ε^g_t	ε^i_t	g_t	i_t	r_t	r^a_t
RRA_t	0.07	-0.11	-0.02	0.10	-0.06	0.23	0.45	0.40
ΔRRA_t		-0.87	-0.30	0.82	-0.28	0.70	-0.63	-0.58
e_t			0.27	-0.98	0.26	-0.83	0.72	0.67

Notations: $RRA_t = e^{\gamma_t}$, $e_t = \lambda(\gamma_t)\varepsilon^g_t + \theta\varepsilon^i_t$, $HB1_t = 100(\frac{H_t}{C_t})$, $HB2_t = 100(\frac{H_t - H_{t-1}}{C_{t-1}})$, and $\Delta RRA_t = RRA_t - RRA_{t-1}$. r_t and r^a_t are the real rates of return on the one-month bill in the model and in the actual data, respectively. All the variables in the model are calculated using the simulation estimates (Table 2 (a)(i)) of the preference parameters and the fitted VAR(2) process of consumption growth and inflation for the period from February 1969 to December 1993.

Table 4: Yields and Yield Spreads

(April 69 – December 93)

Descriptive statistics:

(a) Implied yields

	1-m	3-m	6-m	9-m	1-y	3-y	5-y	sprd 1	sprd 2
mean	7.99	8.02	8.06	8.09	8.11	8.18	8.23	33.04	0.12
std. dev.	2.90	2.49	2.33	2.24	2.19	2.01	1.90	7.35	0.37
1st autocor.	0.66	0.91	0.94	0.96	0.97	0.98	0.98	0.99	0.74
corr. with									
RRA_t	0.79	0.90	0.93	0.95	0.95	0.97	0.97	0.97	-0.65
g_t	-0.53	-0.23	-0.16	-0.13	-0.11	-0.07	-0.06	-0.05	0.31
i_t	0.56	0.55	0.48	0.43	0.40	0.32	0.30	0.27	-0.81

(b) Market yields

	1-m	3-m	6-m	1-y	3-y	5-y	sprd 1	sprd 2
mean	6.69	7.20	7.43	7.68	8.12	8.34	34.03	0.66
std. dev.	2.66	2.77	2.81	2.65	2.37	2.23	8.70	0.95
1st autocor.	0.94	0.97	0.95	0.96	0.97	0.97	0.97	0.93
corr. with								
RRA_t	0.76	0.79	0.80	0.82	0.88	0.90	0.90	-0.18
g_t	-0.14	-0.13	-0.11	-0.11	-0.08	-0.07	-0.06	0.14
i_t	0.40	0.37	0.37	0.33	0.23	0.17	0.12	-0.52
the model yields	0.73	0.81	0.81	0.83	0.88	0.90	0.90	0.50

Note:

Yields reported here are continuously compounded. They are annualized and multiplied by 100; i.e., they are percents per year. `sprd 1' denotes the spread (difference) between unannualized yield on the 5-year bill and that on the 1-year bond. `sprd 2' is the spread between annualized yields on the 5-year and the 1-year bonds.

Table 5: Expected Term Premiums and Holding-Period Returns

(April 69 – December 93)

(a) Expected term premiums

	3-month holding period			1-year holding period			
	6m	9m	12m	2y	3y	4y	5y
mean	0.103	0.168	0.213	0.192	0.336	0.470	0.600
std.dev.	0.016	0.021	0.027	0.019	0.027	0.035	0.041
min	0.053	0.100	0.130	0.147	0.280	0.402	0.512
max	0.149	0.227	0.285	0.253	0.423	0.590	0.752
1st autocor.	-0.08	-0.01	0.05	0.23	0.42	0.47	0.54
corr. with RRA_t	0.19	0.24	0.30	0.46	0.58	0.64	0.68
corr. with g_t	-0.04	0.00	-0.01	-0.05	-0.09	-0.07	-0.06
corr. with i_t	0.10	0.14	0.15	0.16	0.18	0.19	0.24

(b) Expected holding period returns

	3-month holding period			1-year holding period			
	6m	9m	12m	2y	3y	4y	5y
mean	8.21	8.28	8.32	8.66	8.80	8.94	9.07
std.dev.	2.54	2.54	2.55	2.39	2.40	2.40	2.41
min	1.56	1.61	1.64	3.88	4.04	4.14	4.26
max	14.91	15.01	15.06	14.22	14.35	14.53	14.66
1st autocor.	0.91	0.91	0.91	0.97	0.97	0.97	0.97
corr. with RRA_t	0.90	0.90	0.90	0.96	0.96	0.96	0.96
corr. with g_t	-0.23	-0.23	-0.23	-0.11	-0.11	-0.11	-0.11
corr. with i_t	0.55	0.55	0.55	0.40	0.40	0.40	0.40

Table 5 (continued)

(c) Differences in Expected holding period returns

	3-month holding period		1-year holding period		
	9m $-$ 6m	12m $-$ 9m	3y $-$ 2y	4y $-$ 3y	5y $-$ 4y
mean	0.065	0.045	0.144	0.135	0.129
std. dev.	0.014	0.016	0.017	0.018	0.019
min	0.022	0.004	0.097	0.049	0.056
max	0.114	0.097	0.206	0.188	0.191
1st autocor.	0.07	0.01	0.11	0.09	0.19
corr. with RRA_t	0.13	0.19	0.42	0.36	0.33
corr. with g_t	0.04	-0.01	-0.09	0.01	0.00
corr. with i_t	0.09	0.06	0.11	0.09	0.19

Note:

The premiums and returns are simply compounded, annualized, and multiplied by 100. (See the beginning of section 4 for further details.)

Table 6: Fama-Bliss Term Premium Regressions

(April 69 − December 92)

(a) Term premium regressions

$p(n;m : t+12) = a + bf(n;m : t) + u(n;m : t+12)$

(1) Continuously compounded excess returns as dependent variables

In the data:

	a	s(a)	b	s(b)	R^2	ρ_1
$p(24;12 : t+12)$	0.002	(0.003)	0.994	(0.250)	14.8	0.88
$p(36;12 : t+12)$	0.000	(0.006)	1.368	(0.317)	17.8	0.89
$p(48;12 : t+12)$	-0.003	(0.009)	1.721	(0.396)	20.6	0.88
$p(60;12 : t+12)$	0.001	(0.013)	1.286	(0.624)	8.1	0.89

In the model:

	a	s(a)	b	s(b)	R^2	ρ_1
$p(24;12 : t+12)$	0.001	(0.001)	1.154	(0.272)	13.4	0.77
$p(36;12 : t+12)$	0.001	(0.003)	1.467	(0.435)	9.6	0.83
$p(48;12 : t+12)$	0.001	(0.004)	1.561	(0.585)	7.0	0.86
$p(60;12 : t+12)$	0.002	(0.005)	1.512	(0.727)	5.2	0.88

(2) Simply compounded excess returns as dependent variables

In the data:

	a	s(a)	b	s(b)	R^2	ρ_1
$p(24;12 : t+12)$	0.002	(0.004)	1.084	(0.277)	14.4	0.88
$p(36;12 : t+12)$	0.001	(0.007)	1.488	(0.350)	17.1	0.89
$p(48;12 : t+12)$	-0.002	(0.010)	1.879	(0.449)	19.8	0.88
$p(60;12 : t+12)$	0.004	(0.014)	1.375	(0.710)	7.5	0.89

In the model:

	a	s(a)	b	s(b)	R^2	ρ_1
$p(24;12 : t+12)$	0.001	(0.002)	1.253	(0.300)	13.5	0.77
$p(36;12 : t+12)$	0.001	(0.003)	1.587	(0.480)	9.6	0.83
$p(48;12 : t+12)$	0.002	(0.004)	1.686	(0.645)	7.0	0.86
$p(60;12 : t+12)$	0.002	(0.005)	1.625	(0.800)	5.1	0.88

Table 6 (continued)

(b) Cross correlations in the model

	$f(24;12:t)$	$f(36;12:t)$	$f(48;12:t)$	$f(60;12:t)$
$y(12:t)$	-0.59	-0.72	-0.81	-0.87
$y(24:t)$	-0.53	-0.67	-0.77	-0.83
$y(36:t)$	-0.51	-0.65	-0.75	-0.81
$y(48:t)$	-0.50	-0.64	-0.74	-0.81
RRA_t	-0.42	-0.57	-0.67	-0.74

Notations: $p(n;12:t+12)$ is ex post 12-month excess return on an n-month bill. $f(n;12:t)$ is the corresponding forward-spot spreads (defined exactly as in Fama-Bliss). $y(n-12:t)$ is the (annualized) yield of the $(n-12)$-month bill. The spreads and yields are continuously compounded. For 'in the data' cases, the variables are for the CRSP data set. For 'in the model' cases, the variables being used are the ones calculated from the model. The standard errors are obtained by the Newey-West method with lag length 11.

Table 7: Term Premium Regressions with Yields (actual data)

(April 69 $-$ December 92)

(a) Residual regressions

$u(n;m:t+12) = a + by(n-m:t) + \epsilon(t+12)$

	a	$s(a)$	b	$s(b)$	R^2	ρ_1
$u(24;12:t+12)$	-0.019	(0.009)	0.25	(0.12)	9.9	0.90
$u(36;12:t+12)$	-0.042	(0.017)	0.52	(0.22)	12.5	0.90
$u(48;12:t+12)$	-0.064	(0.025)	0.77	(0.33)	13.4	0.89
$u(60;12:t+12)$	-0.076	(0.036)	0.90	(0.48)	8.1	0.91

(b) Term premium regressions with yields

$p(n;m:t+12) = a + b_1 f(n;m:t) + b_2 y(n-m:t) + \epsilon(t+12)$

	a	$s(a)$	b_1	$s(b_1)$	b_2	$s(b_2)$	R^2	ρ_1
$p(24;12:t+12)$	-0.024	(0.010)	1.39	(0.29)	0.30	(0.12)	25.1	0.88
$p(36;12:t+12)$	-0.055	(0.020)	1.89	(0.38)	0.63	(0.23)	30.2	0.88
$p(48;12:t+12)$	-0.080	(0.028)	2.19	(0.40)	0.87	(0.34)	32.6	0.87
$p(60;12:t+12)$	-0.089	(0.041)	1.80	(0.64)	1.01	(0.51)	18.5	0.90

(c) Correlations

$f(24;12:t)$ and $y(12:t)$	-0.43	$\Delta f(24;12:t)$ and $\Delta y(12:t)$	-0.59
$f(36;12:t)$ and $y(24:t)$	-0.42	$\Delta f(36;12:t)$ and $\Delta y(24:t)$	-0.51
$f(48;12:t)$ and $y(36:t)$	-0.34	$\Delta f(48;12:t)$ and $\Delta y(36:t)$	-0.37
$f(60;12:t)$ and $y(48:t)$	-0.34	$\Delta f(60;12:t)$ and $\Delta y(48:t)$	-0.50

Note:

All the variables are continuously compounded, calculated from CRSP Fama-Bliss discount bond file. $p(n;m;t+12)$ and $f(n;m;t)$ are (ex post) term premiums and forward-spot spreads, respectively, as in Table 7. $u(n;m;t+12)$ are residuals from the Fama-Bliss term premium regressions (Table 7). $y(n-m:t)$ is the (annualized) yield on the $(n-m)$-month bond. The standard errors are obtained by the Newey-West method with lag length 11.

Table 8: A Set of Two-stage GMM Estimates and Tests

	R(6;3)	R(12;6)	R(24;12)	R(36;12)	R(48;12)	R(60;12)
μ	0.852	1.388	1.006	1.096	1.125	1.169
	(2.228)	(3.500)	(1.898)	(2.086)	(2.245)	(2.341)
ϕ	0.967	0.966	0.967	0.970	0.972	0.973
	(0.024)	(0.011)	(0.010)	(0.010)	(0.010)	(0.010)
θ	5.360	10.543	16.249	18.223	19.550	20.392
	(3.786)	(6.006)	(4.899)	(5.076)	(5.161)	(5.190)
δ	0.999	0.996	0.998	0.998	0.997	0.997
	(0.008)	(0.015)	(0.007)	(0.008)	(0.008)	(0.009)
$\chi^2(3)$	21.68	5.92	0.92	1.06	1.18	1.27
p-value	0.00	0.12	0.78	0.76	0.75	0.74
RRA:						
mean	2.33	3.87	2.79	3.06	3.18	3.33
dev.	0.17	0.69	0.61	0.80	0.93	1.06
max	2.75	5.45	4.54	5.39	5.91	6.45
min	2.05	2.12	2.03	2.10	2.09	2.10

Note:

$R(x;y)$ is the y-month gross simple return on the x-month bill, i.e., Q^{x-y}_{t+y}/Q^{x}_{t}, where Q^{x}_{t} is the time t price of the x month bill. The vector of instruments being used is $w'_t = (g_t; i_t; g_{t-1}; i_{t-1}; g_{t-2}; i_{t-2}; y_t)$. Here g_t, i_t, and y_t are consumption growth, inflation rate, and the yield on the one month bill, respectively. In the first step, the maximum likelihood estimates of the parameters of the joint process of $\{g_t$ and $i_t\}$ are calculated. Then the GMM procedure in Ogaki (1993) is applied. The sample period is from April 69 to December 93.

Table 9: Implied Price Errors

(April 69 – December 93)

Descriptive statistics:

	1-month	3-month	6-month	1-year	3-year	5-year
mean	0.11	0.20	0.31	0.41	0.22	-0.27
std.dev.	0.17	0.40	0.79	1.35	2.57	3.13
min	-0.42	-0.94	-1.78	-3.18	-6.63	-9.31
max	0.53	1.20	4.03	3.21	5.70	6.79
1st autocor.	0.29	0.73	0.76	0.85	0.86	0.86
2nd autocor.	0.32	0.61	0.65	0.73	0.75	0.75
3rd autocor.	0.28	0.52	0.56	0.63	0.66	0.67
6th autocor.	0.13	0.36	0.39	0.43	0.45	0.47
9th autocor.	0.15	0.31	0.33	0.35	0.34	0.35
12th autocor.	-0.03	0.13	0.17	0.20	0.17	0.20

Crosscorrelation:

	1-month	3-month	6-month	1-year	3-year	5-year
1-month	1.00					
3-month	0.84	1.00				
6-month	0.74	0.94	1.00			
1-year	0.68	0.92	0.94	1.00		
3-year	0.60	0.83	0.84	0.94	1.00	
5-year	0.55	0.76	0.75	0.85	0.97	1.00

Note:

A price error is defined as difference between the market and the model price per $100 face value. The model prices are calculated via simulation, using the simulation estimates (Table 2 (a)(i)) of the preference parameters and the fitted VAR(2) process of consumption growth and inflation for the period from February 1969 to December 1993. The size of simulation is fifty thousand.