

NBER WORKING PAPER SERIES

EQUILIBRIUM YIELD CURVES

Monika Piazzesi
Martin Schneider

Working Paper 12609
<http://www.nber.org/papers/w12609>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2006

Email addresses: piazzesi@uchicago.edu, martin.schneider@nyu.edu. We thank Pierpaolo Benigno and John Campbell for helpful discussions. We also thank Andy Atkeson, David Backus, Frederico Belo, John Cochrane, Lars Hansen, Anil Kashyap, Patrick Kehoe, Narayana Kocherlakota, Ricardo Mayer, Ellen McGrattan, Lubos Pastor, Chris Sims, Harald Uhlig, Michael Woodford and seminar participants at UCLA, the Minneapolis Fed and the University of Chicago for their comments. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2006 by Monika Piazzesi and Martin Schneider. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Equilibrium Yield Curves
Monika Piazzesi and Martin Schneider
NBER Working Paper No. 12609
October 2006
JEL No. E0,E3,E4,G0,G12

ABSTRACT

This paper considers how the role of inflation as a leading business-cycle indicator affects the pricing of nominal bonds. We examine a representative agent asset pricing model with recursive utility preferences and exogenous consumption growth and inflation. We solve for yields under various assumptions on the evolution of investor beliefs. If inflation is bad news for consumption growth, the nominal yield curve slopes up. Moreover, the level of nominal interest rates and term spreads are high in times when inflation news are harder to interpret. This is relevant for periods such as the early 1980s, when the joint dynamics of inflation and growth was not well understood.

Monika Piazzesi
Graduate School of Business
University of Chicago
5807 S. Woodlawn
Chicago, IL 60637
and NBER
piazzesi@uchicago.edu

Martin Schneider
Department of Economics
New York University
269 Mercer Street, 7th Floor
New York, NY 10003
ms1927@nyu.edu

I Introduction

The main theme of this paper is that investors dislike surprise inflation not only because it lowers the payoff on nominal bonds, but also because it is bad news for future consumption growth. The fact that nominal bonds pay off little precisely when the outlook on the future worsens makes them unattractive assets to hold. The premium that risk averse investors seek as compensation for inflation risk should thus depend on the extent to which inflation is perceived as a carrier of bad news.

One implication is that the nominal yield curve slopes upward: long bonds pay off even less than short bonds when inflation, and hence bad news, arrives. Therefore, long bonds command a term spread over short bonds. Moreover, the level of interest rates and term spreads should increase in times when inflation news are harder to interpret. This is relevant for periods such as the early 1980s, when the joint dynamics of inflation and growth had just become less well understood.

We study the effect of inflation as bad news in a simple representative agent asset pricing model with two key ingredients. First, investor preferences are described by recursive utility. One attractive feature of this preference specification is that – in contrast to the standard time-separable expected utility model – it does not imply indifference to the temporal distribution of risk. In particular, it allows investors to prefer a less persistent consumption stream to a more persistent stream, even if overall risk of the two streams is the same. In our context, aversion to persistence generates a heightened concern with news about the future and makes investors particularly dislike assets that pay off little when bad news arrives.

The second ingredient of the model is a description of how investor beliefs about consumption and inflation evolve over time. Investor beliefs determine to what extent inflation is perceived to carry bad news at a particular point in time. We consider various specifications, some of which take into account structural change in the relationship between consumption growth and inflation over the postwar period in the United States. Given investor beliefs about these two fundamentals, we determine interest rates implied by the model from the intertemporal Euler equation.

We perform two broad classes of model exercises. First, we consider stationary rational expectations versions of the model. Here we begin by estimating a stochastic process for U.S. consumption growth and inflation over the entire postwar period. We assume that investor beliefs are the condition-

als of this process, and derive the properties of the model-implied yield curve. The estimated process in this benchmark exercise has constant conditional variances. As a result, all asset price volatility derives from changes in investors' conditional expectations. In particular, the dynamics of yields is entirely driven by movements in expected consumption growth and inflation.

The benchmark model captures a number of features of observed yields. Both model implied and observed yields contain a sizeable low frequency component (period > 8 years) that is strongly correlated with inflation. At business cycle frequencies (between 1.5 and 8 years), both the short rate and the term spread are driven by the business cycle component of inflation, which covaries positively with the former and negatively with the latter. Both a high short rate and a low term spread forecast recessions, that is, times of low consumption growth. Finally, average yields are increasing, and yield volatility is decreasing, in the maturity of the bond.

The fact that the model implies an upward-sloping nominal yield curve depends critically on both preferences and the distribution of fundamentals. In the standard expected utility case, an asset commands a premium over another asset only when its payoff covaries more with consumption growth. Persistence of consumption growth and inflation then implies a downward sloping yield curve. When investors exhibit aversion to persistence, an assets commands a premium also when its payoff covaries more with news about future consumption growth. The estimated process implies that inflation brings bad news. The implied correlation between growth and inflation is critical; if inflation and consumption growth were independent, the yield curve would slope downward even if investors are averse to persistence.

The role of inflation as bad news suggests that other indicators of future growth might matter for term premia. Moreover, one might expect the arrival of other news about growth or inflation to make yields more volatile than they are in our benchmark model. In a second exercise, we maintain the rational expectations assumption, but model investors' information set more explicitly by exploiting information contained in yields themselves. In particular, we begin by estimating an unrestricted stochastic process for consumption growth, inflation, the short rate, and the term spread. We then derive model-implied yields given the information set described by this stochastic process.

The resulting model-implied yields are very similar to those from our benchmark. It follows that, viewed through the lens of our consumption-based asset pricing model, inflation itself is the key

predictor of future consumption, inflation, and yields that generates interest rate volatility. Conditional on our model, we can rule out the possibility that other variables – such as investors’ perception of a long run inflation target, or information inferred from other asset prices – generates volatility in yields. Indeed, if observed yields had been generated by a version of our model in which investors price bonds using better information than we modelers have, our exercise would have recovered that information from yields.

We also explore the role of inflation as bad news in a class of models that accommodate investor concern with structural change. Here we construct investor beliefs by sequentially estimating the stochastic process for fundamentals. We use a constant gain adaptive learning scheme where the estimation for date t places higher weight on more recent observations. The investor belief for date t is taken to be the conditional of the process estimated with data up to date t . We then compute a sample of model-implied yields from the Euler equations, using a different investor belief for each date. We apply this model to consider changes in yield curve dynamics, especially around the monetary policy experiment.

It has been suggested that long interest rates were high in the early 1980s because investors at the time were only slowly adjusting their inflation expectations downward. In the context of our model, this is not a plausible story. Indeed, it is hard to write down a sensible adaptive learning scheme in which the best forecast of future inflation is not close to current inflation. Since inflation fell much more quickly in the early 1980s than nominal interest rates, our learning schemes do not generate much inertia in inflation expectations. At the same time, survey expectations of inflation also fell relatively quickly in the early 1980s, along with actual inflation and the forecasts in our model.

We conclude that learning can help understand changes in the yield curve only if it entails changes in subjective uncertainty that have first order effects on asset prices. In a final exercise, we explore one scenario where this happens. In addition to sequential estimation, we introduce parameter uncertainty which implies that investors cannot easily distinguish permanent and transitory movements in inflation. With patient investors who are averse to persistence, changes in uncertainty then have large effects on interest rates and term spreads. In particular, the uncertainty generated by the monetary policy experiment leads to sluggish behavior in interest rates, especially at the long end of the yield curve, in the early 1980s.

A by-product of our analysis is a decomposition into real and nominal interest rates, where the former are driven by expected consumption growth, whereas the latter also move with changes in expected inflation. Importantly, inflation as an indicator of future growth affects both nominal and real interest rates. Loosely speaking, our model says that yields in the 1970s and early 1980s were driven by nominal shocks – inflation surprises – that affect nominal and real rates in opposite directions. Here an inflation surprise lowers real rates because it is bad news for future consumption growth. In contrast, prior to the 1970s, and again more recently, there were more real shocks – surprises in consumption growth – that make nominal and real interest rates move together.

Our model also predicts a downward sloping real yield curve. In contrast to long nominal bonds, long indexed bonds pay off when future real interest rates – and hence future expected consumption growth – are low, thus providing insurance against bad times. Coupled with persistence in growth, this generates a downward sloping real yield curve in an expected utility model. The effect is reinforced when investors are averse to persistence. Unfortunately, the available data series on U.S. indexed bonds, which is short and comes from a period of relatively low interest rates, makes it difficult to accurately measure average long indexed yields. However, evidence from the United Kingdom suggests that average term spreads are positive for nominal, but negative for indexed bonds.

The paper is organized as follows. Section II presents the model, motivates our use of recursive utility and outlines the yield computations. Section III reports results from the benchmark rational expectations version of the model. Section IV maintains the rational expectations assumption, but allows for more conditioning information. Section V introduces learning. Section VI reviews related literature. Appendix A collects our estimation results. Appendix B presents summary statistics about real rate data from the US and the UK. Appendix C contains results with alternative data definitions and estimation strategies.

II Model

We consider an endowment economy with a representative investor. The endowment – denoted $\{C_t\}$ since it is calibrated to aggregate consumption – and inflation $\{\pi_t\}$ are given exogenously. Equilibrium prices adjust such that the agent is happy to consume the endowment. In the remainder of this section, we define preferences and explain how yields are computed.

A. Preferences

We describe preferences using the recursive utility model proposed by Epstein and Zin (1989) and Weil (1989), which allows for a constant coefficient of relative risk aversion that can differ from the reciprocal of the intertemporal elasticity of substitution (IES). This class of preferences is now common in the consumption-based asset pricing literature. Campbell (1993, 1996) derives approximate loglinear pricing formulas (that are exact if the IES is one) to characterize premia and the price volatility of equity and real bonds. Duffie, Schroeder, and Skiadas (1997) derive closed-form solutions for bond prices in a continuous time version of the model. Restoy and Weil (1998) show how to interpret the pricing kernel in terms of a concern with news about future consumption. For our computations, we assume a unitary IES and homoskedastic lognormal shocks, which allows us to use a linear recursion for utility derived by Hansen, Heaton, and Li (2005).

We fix a finite horizon T and a discount factor $\beta > 0$. The time t utility V_t of a consumption stream $\{C_t\}$ is defined recursively by

$$(1) \quad V_t = C_t^{1-\alpha_t} \text{CE}_t(V_{t+1})^{\alpha_t},$$

with $V_{T+1} = 0$. Here the certainty equivalent CE_t imposes constant relative risk aversion with coefficient γ ,

$$\text{CE}_t(V_{t+1}) = E_t \left(V_{t+1}^{1-\gamma} \right)^{1/(1-\gamma)},$$

and the sequence of weights α_t is given by

$$(2) \quad \alpha_t := \frac{\sum_{j=1}^{T-t} \beta^j}{\sum_{j=0}^{T-t} \beta^j}.$$

If $\beta < 1$, the weight α_t on continuation utility converges to β as the horizon becomes large. If $\gamma = 1$, the model reduces to standard logarithmic utility. More generally, the risk aversion coefficient can be larger or smaller than one, the (inverse of the) intertemporal elasticity of substitution.

Discussion

Recursive preferences avoid the implication of the time-separable expected utility model that decision makers are indifferent to the temporal distribution of risk. A standard example, reviewed by Duffie and Epstein (1992), considers a choice at some date zero between two risky consumption plans A and B. Both plans promise contingent consumption for the next 100 periods. Under both plans, consumption in a given period can be either high or low, with the outcome determined by the toss of a fair coin. However, the consumption stream promised by plan A is determined by *repeated* coin tosses: if the toss in period t is heads, consumption in t is high, otherwise consumption in t is low. In contrast, the consumption stream promised by plan B is determined by a *once and for all* coin toss at date 1: if this toss is heads, consumption is high for the next 100 periods, otherwise, consumption is low for the next 100 periods.

Intuitively, plan A looks less risky than plan B. Under plan B, all eggs are in one basket, whereas plan A is more diversified. If all payoffs were realized at the same time, risk aversion would imply a preference for plan A. However, if the payoffs arrive at different dates, the standard time-separable expected utility model implies indifference between A and B. This holds regardless of risk aversion and of how little time elapses between the different dates. The reason is that the time-separable model evaluates risks at different dates in isolation. From the perspective of time zero, random consumption at any given date – viewed in isolation – does have the same risk (measured, for example, by the variance.) What the standard model misses is that the risk is distributed differently over time for the two plans: plan A looks less risky since the consumption stream it promises is less persistent.

According to the preferences (1), the plans A and B are ranked differently if the coefficient of relative risk aversion γ is not equal to one. In particular, $\gamma > 1$ implies that the agent is averse to the persistence induced by the initial shock that characterizes plan B and therefore prefers A. This is the case we consider in this paper. When $\gamma < 1$, the agent likes the persistence and prefers B.

Another attractive property of the utility specification (1) is that the motives that govern consumption smoothing over different states of nature and consumption smoothing over time are allowed to differ. For example, an agent with recursive utility and $\gamma > 1$ would not prefer an erratic deterministic consumption stream A to a constant stream B. Indeed, there is no reason to assume why the two smoothing motives should be tied together like in the power utility case, where the risk aversion

coefficient γ is the reciprocal of the elasticity of intertemporal substitution. After all, the notion of smoothing over different states even makes sense in a static economy with uncertainty, while smoothing over time is well defined in a dynamic but deterministic economy.

We specify a (long) finite horizon T because we want to allow for high discount factors, $\beta > 1$. There is no a priori reason to rule out this case. The usual justification for low discount factors is introspection: when faced with a constant consumption stream, many people would prefer to shift some consumption into the present. While this introspective argument makes sense in the stochastic environment in which we actually live – where we may die before we get to consume, and so we want to consume while we still can – it is not clear whether the argument should apply to discounting in a deterministic environment with some known horizon (which is the case for which the discount factor β is designed.)

Pricing kernel

We divide equation (1) by current consumption to get

$$\frac{V_t}{C_t} = \text{CE}_t \left(\frac{V_{t+1}}{C_{t+1}} \frac{C_{t+1}}{C_t} \right)^{\alpha_t}.$$

Taking logarithms, denoted throughout by small letters, we obtain the recursion

$$v_t - c_t = \alpha_t \ln \text{CE}_t [\exp(v_{t+1} - c_{t+1} + \Delta c_{t+1})].$$

Assuming that the variables are conditionally normal, we get

$$(3) \quad v_t - c_t = \alpha_t E_t(v_{t+1} - c_{t+1} + \Delta c_{t+1}) + \alpha_t \frac{1}{2} (1 - \gamma) \text{var}_t(v_{t+1}).$$

Solving the recursion forward and using our assumption that the agent's beliefs are homoskedastic, we can express the log ratio of continuation utility to consumption as an infinite sum of expected discounted future consumption growth,

$$(4) \quad v_t - c_t = \sum_{i=0}^{T-t} \alpha_{t,1+i} E_t(\Delta c_{t+1+i}) + \text{constant}.$$

For $\beta < 1$ and $T = \infty$, the weights on expected future consumption growth are simply $\alpha_{t,i} = \beta^i$. Even

for large finite T , equation (4) can be viewed as a sum of expected consumption growth with weights that are independent of the forecasting horizon $1 + i$.

For finite T , the weights $\alpha_{t,i}$ are given by

$$\alpha_{t,i} := \sum_{j=i}^{T-t} \beta^j \bigg/ \sum_{j=0}^{T-t} \beta^j ,$$

so that $\alpha_{t,1} = \alpha_t$. For $\beta > 1$, the weights on expected future consumption growth are decreasing and concave in the forecast horizon i . For large T , they remain equal to one for many periods. If consumption growth reverts to its mean – that is, $E_t(\Delta c_{t+1+i})$ converges to the unconditional mean of consumption growth as i becomes large – then the log ratio of continuation utility is approximately given by the infinite-horizon undiscounted sum of expected consumption growth.

Payoffs denominated in units of consumption are valued by the real pricing kernel

$$(5) \quad M_{t+1} = \beta \left(\frac{C_{t+1}}{C_t} \right)^{-1} \left(\frac{V_{t+1}}{CE_t(V_{t+1})} \right)^{1-\gamma} .$$

The random variable M_{t+1} represents the date t prices of contingent claims that pay off in $t + 1$. In particular, the price of a contingent claim that pays off one unit if some event in $t + 1$ occurs is equal to the expected value of the pricing kernel conditional on the event, multiplied by the probability of the event. In a representative agent model, the pricing kernel is large over events in which the agent will feel bad: claims written on such events are particularly expensive.

Again using normality, we obtain the log real pricing kernel

$$(6) \quad \begin{aligned} m_{t+1} &= \ln \beta - \Delta c_{t+1} - (\gamma - 1)(v_{t+1} - E_t(v_{t+1})) \\ &\quad - \frac{1}{2}(1 - \gamma)^2 \text{var}_t(v_{t+1}) \\ &= \ln \beta - \Delta c_{t+1} - (\gamma - 1) \sum_{i=0}^{T-t-1} \alpha_{t+1,i} (E_{t+1} - E_t) \Delta c_{t+1+i} \\ &\quad - \frac{1}{2}(\gamma - 1)^2 \text{var}_t \left(\sum_{i=0}^{T-t-1} \alpha_{t+1,i} E_{t+1}(\Delta c_{t+1+i}) \right) . \end{aligned}$$

The logarithmic expected utility model (the case $\gamma = 1$) describes “bad events” in terms of future *realized* consumption growth – the agent feels bad when consumption growth is low. This effect is

represented by the first term in the pricing kernel. Recursive utility introduces a new term that reflects a concern with the temporal distribution of risk. In the case we consider, $\gamma > 1$, the agent fears downward revisions in consumption expectations. More generally, a source of risk is not only reflected in asset prices if it makes consumption more volatile, as in the standard model, but it can also affect prices if it affects only the temporal distribution of risk, for example if it makes consumption growth more persistent.

Finally, we define the log nominal pricing kernel, that we use below to value payoffs denominated in dollars:

$$(7) \quad m_{t+1}^{\$} = m_{t+1} - \pi_{t+1}.$$

B. Nominal and Real Yield Curves

The agent's Euler equation for a real bond that pays 1 unit of consumption n periods later determines its time- t price $P_t^{(n)}$ as the expected value of its payoff tomorrow weighted by the real pricing kernel:

$$(8) \quad P_t^{(n)} = E_t \left(P_{t+1}^{(n-1)} M_{t+1} \right) = E_t \left(\prod_{i=1}^n M_{t+i} \right).$$

This recursion starts with the one-period bond at $P_t^{(1)} = E_t [M_{t+1}]$. Under normality, we get in logs

$$(9) \quad \begin{aligned} p_t^{(n)} &= E_t \left(p_{t+1}^{(n-1)} + m_{t+1} \right) + \frac{1}{2} \text{var}_t \left(p_{t+1}^{(n-1)} + m_{t+1} \right) \\ &= E_t \left(\sum_{i=1}^n m_{t+i} \right) + \frac{1}{2} \text{var}_t \left(\sum_{i=1}^n m_{t+i} \right). \end{aligned}$$

The n -period real yield is defined from the relation

$$(10) \quad y_t^{(n)} = -\frac{1}{n} p_t^{(n)} = -\frac{1}{n} E_t \left(\sum_{i=1}^n m_{t+i} \right) - \frac{1}{n} \frac{1}{2} \text{var}_t \left(\sum_{i=1}^n m_{t+i} \right).$$

For a fixed date t , the *real yield curve* maps the maturity n of a bond to its real yield $y_t^{(n)}$. Throughout this paper, we assume that the agent's beliefs are homoskedastic. To the extent that we observe heteroskedasticity of yields in the data, we will attribute it to the effect of learning about the dynamics of fundamentals.

Analogously, the price of a nominal bond $P_t^{(n)\$}$ satisfies the Euler equation (8) with dollar signs attached. From equations (9) and (10), we can write the nominal yield as

$$(11) \quad y_t^{(n)\$} = -\frac{1}{n} p_t^{(n)\$} = -\frac{1}{n} E_t \left(\sum_{i=1}^n m_{t+i}^{\$} \right) - \frac{1}{n} \frac{1}{2} \text{var}_t \left(\sum_{i=1}^n m_{t+i}^{\$} \right).$$

By fixing the date t , we get the *nominal yield curve* as the function that maps maturity n to the nominal yield $y_t^{(n)\$}$ of a bond.

Equations (9) and (10) show that log prices and yields of real bonds in this economy are determined by expected future marginal utility. The log prices and yields of nominal bonds additionally depend on expected inflation. To understand the behavior of yields, it is useful to decompose yields into their unconditional mean and deviations of yields from the mean. Below, we will see that while the implications for average yields will depend on whether we assume recursive or expected (log) preferences, the dynamics of yields – and thus volatility – will be the same for both preference specifications.

The dynamics of real yields can be derived from the conditional expectation of the real pricing kernel (6) together with the yield equation (10). Specifically, we can write the deviations of real yields $y_t^{(n)}$ from their mean $\mu^{(n)}$ as

$$(12) \quad y_t^{(n)} - \mu^{(n)} = \frac{1}{n} E_t \sum_{i=1}^n (\Delta c_{t+i} - \mu_c),$$

where μ_c denotes the mean consumption growth rate. This equation shows that the dynamics of real yields are driven by changes in expected future consumption growth. Importantly, these dynamics do *not* depend on any preference parameters. In particular, the equation (12) is identical for recursive utility and expected log utility. Of course, equation (12) does depend on the elasticity of intertemporal substitution, which we have set equal to one.

Similarly, the dynamics of nominal yields can be derived from the conditional expectation of the nominal pricing kernel (7) together with the yield equation (11). As a result, we can show that de-measured nominal yields are expected nominal growth rates over the lifetime of the bond

$$(13) \quad y_t^{(n)\$} - \mu^{(n)\$} = \frac{1}{n} E_t \sum_{i=1}^n (\Delta c_{t+i} - \mu_c + \pi_{t+i} - \mu_\pi).$$

The dynamics of real and nominal yields in equations (12) and (13) show that changes in the difference

between nominal and real yields represent changes in expected future inflation.

The unconditional mean of the one-period real rate is

$$(14) \quad \mu^{(1)} = -\ln \beta + \mu_c - \frac{1}{2} \text{var}_t(\Delta c_{t+1}) - (\gamma - 1) \text{cov}_t \left(\Delta c_{t+1}, \sum_{i=0}^{T-t-1} \alpha_{t+1,i} (E_{t+1} - E_t) \Delta c_{t+1+i} \right).$$

The first three terms represent the mean real short rate in the log utility case. The latter is high when β is low, which means that the agent is impatient and does not want to save. An intertemporal smoothing motive increases the real rate when the mean consumption growth rate μ_c is high. Finally, the precautionary savings motive lowers the real rate when the variance of consumption growth is high. With $\gamma > 1$, an additional precautionary savings motive is captured by the covariance term. It not only lowers interest rates when realized consumption growth is more volatile, but also when it covaries more with expected consumption growth, that is, when consumption growth is more persistent.

The mean of the nominal short rate is

$$(15) \quad \begin{aligned} \mu^{(1)\$} &= \mu^{(1)} + \mu_\pi - \frac{1}{2} \text{var}_t(\pi_{t+1}) - \text{cov}_t(\pi_{t+1}, \Delta c_{t+1}) \\ &\quad - (\gamma - 1) \text{cov}_t \left(\pi_{t+1}, \sum_{i=0}^{T-t-1} \alpha_{t+1,i} (E_{t+1} - E_t) \Delta c_{t+1+i} \right). \end{aligned}$$

There are several reasons for why the Fisher relation fails or, put differently, for why the short rate is not simply equal to the real rate plus expected inflation. First, the variance of inflation enters due to Jensen's inequality. Second, the covariance of consumption growth and inflation represents an inflation risk premium. Intuitively, nominal bonds – including those with short maturity — are risky assets. The real payoff from nominal bonds is low in times of surprise inflation. If the covariance between inflation and consumption is negative, nominal bonds are unattractive assets, because they have low real payoffs in bad times. In other words, nominal bonds do not provide a hedge against times of low consumption growth. Investors thus demand higher nominal yields as compensation for holding nominal bonds. Recursive utility introduces an additional reason why nominal bonds may be unattractive for investors: their payoffs are low in times with bad news about future consumption growth. These bonds may thus not provide a hedge against times with bad news about the future.

We define $rx_{t+1}^{(n)} = p_{t+1}^{(n-1)} - p_t^{(n)} - y_t^{(1)}$ as the return on buying an n -period real bond at time t for $p_t^{(n)}$ and selling it at time $t + 1$ for $p_{t+1}^{(n-1)}$ in excess of the short rate. Based on equation (9), the

expected excess return is

$$(16) \quad E_t \left(r x_{t+1}^{(n)} \right) = -\text{cov}_t \left(m_{t+1}, E_{t+1} \sum_{i=1}^{n-1} m_{t+1+i} \right) - \frac{1}{2} \text{var}_t \left(p_{t+1}^{(n-1)} \right).$$

The covariance term on the right-hand side is the risk premium, while the variance term is due to Jensen's inequality. Expected excess returns are constant whenever conditional variances are constant, as in our benchmark belief specification. With learning, however, the conditional probabilities that are used to evaluate the conditional covariances in equation (16) will be derived from different beliefs each period. As a result, expected excess returns will vary over time.

The risk premium on real bonds is positive when the pricing kernel and long bond prices are negatively correlated. This correlation is determined by the autocorrelation of marginal utility. The risk premium is positive if marginal utility is negatively correlated with expected changes in future marginal utility. In this case, long bonds are less attractive than short bonds, because their payoffs tend to be low in bad times (when marginal utility is high). The same equation also holds for nominal bonds after we attach dollar signs everywhere. Here, the sign of the risk premium also depends on the correlation between (nominal) bond prices and inflation. Over long enough samples, the average excess return on an n -period bond is approximately equal to the average spread between the n -period yield and the short rate.¹ This means that the yield curve is on average upward sloping if the right-hand side of equation (16) is positive on average.

In our model, expected changes in marginal utility depend on expected future consumption growth. The expected excess return (16) can therefore be rewritten as

$$(17) \quad E_t \left(r x_{t+1}^{(n)} \right) = \text{cov}_t \left(m_{t+1}, E_{t+1} \sum_{i=1}^{n-1} \Delta c_{t+1+i} \right) - \frac{1}{2} \text{var}_t \left(p_{t+1}^{(n-1)} \right).$$

Real term premia are thus driven by the covariance of marginal utility with expected consumption

¹To see this, we can write the excess return as

$$\begin{aligned} p_{t+1}^{(n-1)} - p_t^{(n)} - y_t^{(1)} &= n y_t^{(n)} - (n-1) y_{t+1}^{(n-1)} - y_t^{(1)} \\ &= y_t^{(n)} - y_t^{(1)} - (n-1) \left(y_{t+1}^{(n-1)} - y_t^{(n)} \right) \end{aligned}$$

For large n and a long enough sample, the difference between the average $(n-1)$ -period yield and the average n -period yield is zero.

growth. The expected excess return equation (16) for an n -period nominal bond becomes

$$(18) \quad E_t \left(r x_{t+1}^{(n)\$} \right) = \text{cov}_t \left(m_{t+1}^{\$}, E_{t+1} \sum_{i=1}^{n-1} \Delta c_{t+1+i} + \pi_{t+1+i} \right) - \frac{1}{2} \text{var}_t \left(p_{t+1}^{(n-1)\$} \right).$$

This equation shows that nominal term premia are driven by the covariance of the nominal pricing kernel with expected nominal growth.

III Benchmark

In this section, we derive investor beliefs from a state space system for consumption growth and inflation that is estimated with data from the entire postwar sample. The conditional probabilities that we use to evaluate the agent's Euler equation, and thus to compute yields, come from this estimated system.

Data

We measure aggregate consumption growth with quarterly NIPA data on nondurables and services and construct the corresponding price index to measure inflation. We assume that population growth is constant. The data on bond yields with maturities one year and longer are from the CRSP Fama-Bliss discount bond files. These files are available for the sample 1952:2-2005:4. The short (1-quarter) yield is from the CRSP Fama riskfree rate file. These data, our MATLAB programs, and Appendix C which contains additional results based on alternative inflation and population series can be downloaded from our websites.

Beliefs about Fundamentals

The vector of consumption growth and inflation $z_{t+1} = (\Delta c_{t+1}, \pi_{t+1})^\top$ has the state-space representation

$$(19) \quad \begin{aligned} z_{t+1} &= \mu_z + x_t + e_{t+1} \\ x_{t+1} &= \phi_x x_t + \phi_x K e_{t+1} \end{aligned}$$

where $e_{t+1} \sim N(0, \Omega)$, the state vector x_{t+1} is 2-dimensional and contains expected consumption and inflation, ϕ_x is the 2×2 autoregressive matrix, and K is the 2×2 gain matrix. Our benchmark model

assumes that the agent’s beliefs about future growth and inflation are described by this state space system evaluated at the point estimates. Based on these beliefs, the time- t conditional expected values in the yield equations (12) and (13) are simply linear functions of the state variables x_t . We estimate this system with data on consumption growth and inflation using maximum likelihood. Table A.1 in Appendix A reports parameter estimates.

The state space system (19) nests a first-order Vector-Autoregression. To see this, start from the VAR $z_{t+1} = \mu_z + \phi z_t + e_{t+1}$ and set $x_t = \phi(z_t - \mu_z)$. This will result in a system like (19) but with $K = I$ (and $\phi_x = \phi$). Since K is a 2×2 matrix, setting $K = I$ imposes four parameter restrictions, which we can test with a likelihood ratio test. The restrictions are strongly rejected based on the usual likelihood ratio statistic $2 \times [\mathcal{L}(\theta_{\text{unrestricted}}) - \mathcal{L}(\theta_{\text{restricted}})] = 34.3$, which is greater than the 5 percent and 1 percent critical $\chi^2(4)$ values of 9.5 and 13.3, respectively.

The reason for this rejection is that the state space system does a better job at capturing the dynamics of inflation than the first-order VAR. Indeed, quarterly inflation has a very persistent component, but also a large transitory component, which leads to downward biased estimates of higher order autocorrelations in the VAR. For example, the n th-order empirical autocorrelations of inflation are .84 for $n = 1$, .80 for $n = 2$, .66 for $n = 5$ and .52 for $n = 10$. While the state space system matches these autocorrelations almost exactly (as we will see in Figure 1 below), the VAR only matches the first autocorrelation and understates the others: the numbers are .84 for $n = 1$, .72 for $n = 2$, .43 for $n = 5$ and .19 for $n = 10$.

For our purposes, high-order autocorrelations are important, because they determine long-horizon forecasts of inflation and thus nominal yields through equation (13). By contrast, this issue is not important for matching the long-horizon forecasts of consumption growth and thus real yields in equation (12). The autocorrelation function of consumption growth data starts low at .36 for $n = 1$, .18 for $n = 2$ and is essentially equal to zero thereafter. This function can be matched well with a first-order VAR in consumption growth and inflation.

To better understand the properties of the estimated dynamics, we report covariance functions which completely characterize the linear Gaussian system (19). Figure 1 plots covariance functions computed from the model and from the raw data. At 0 quarters, these lines represent variances and contemporaneous covariances. The black lines from the model match the gray lines in the data




Figure 1: Covariance functions computed from the estimated benchmark model and from the raw data. Dotted lines indicate $2 \times$ standard errors bounds around the covariance function from the data computed with GMM. For example, the graph titled "consumption, lagged consumption" shows the covariance of current consumption growth with consumption growth lagged x quarters, where x is measured on the horizontal axis.

quite well. The dotted lines in Figure 1 are $2 \times$ standard error bounds around the covariance function estimated with raw data. These standard error bounds are not based on the model; they are computed with GMM. (For more details, see Appendix A.) To interpret the units, consider the upper left panel. The variance of consumption growth is .22 in model and data, which amounts to $\sqrt{.22 \times 4^2} = 1.88$ percent volatility. Figure 1 shows that consumption growth is weakly positively autocorrelated. For example, the covariance $\text{cov}(\Delta c_t, \Delta c_{t-1}) = \rho \text{var}(\Delta c_t) = \rho \times .22 = 0.08$ in model and data which implies that the first-order autocorrelation is $\rho = .36$. Inflation is clearly more persistent, with an autocorrelation of 84%.

An important feature of the data is that consumption growth and inflation are negatively correlated contemporaneously and forecast each other with a *negative* sign. For example, the upper right panel in Figure 1 shows that high inflation is a leading recession indicator. Higher inflation in quarter t predicts lower consumption growth in quarter $t + n$ even $n = 6$ quarters ahead of time. The lower left

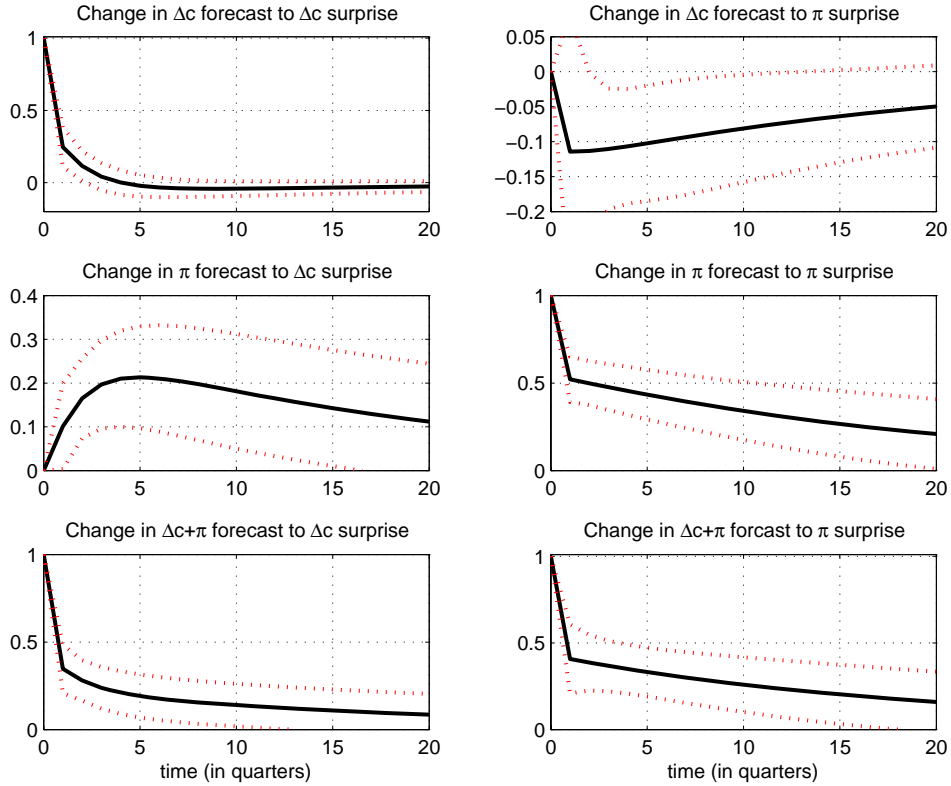


Figure 2: Impulse responses to 1-percentage point surprises e_{t+1} in consumption growth and inflation. The responses are measured in percent. Dotted lines are $2 \times$ standard error bounds based on maximum likelihood.

panel shows that high consumption also forecasts low inflation, but with a shorter lead time. These cross-predictability patterns will be important for determining longer yields.

From equations (12) and (13) we know that the dynamics of equilibrium interest rates are driven by forecasts of growth and inflation. Real yield movements are generated by changes in growth forecasts over the lifetime of the bond, while nominal yield movements are generated by changing nominal growth forecasts. To understand the conditional dynamics of these forecasts better – as opposed to the unconditional covariances and thus univariate regression forecasts from Figure 1 – we plot impulse responses in Figure 2. These responses represent the change in forecasts following a 1-percent shock e_{t+1} . The signs of the own-shock responses are not surprising in light of the unconditional covariances; they are positive and decay over time. This decay is slower for inflation, where a 1-percent surprise increases inflation forecasts by 40 basis points even two years down the road. However, the cross-shock responses reveal some interesting patterns. The middle left plot shows that a 1-percent growth surprise predicts inflation to be *higher* by roughly 20 basis points over the next 2-3 years. The top right plot

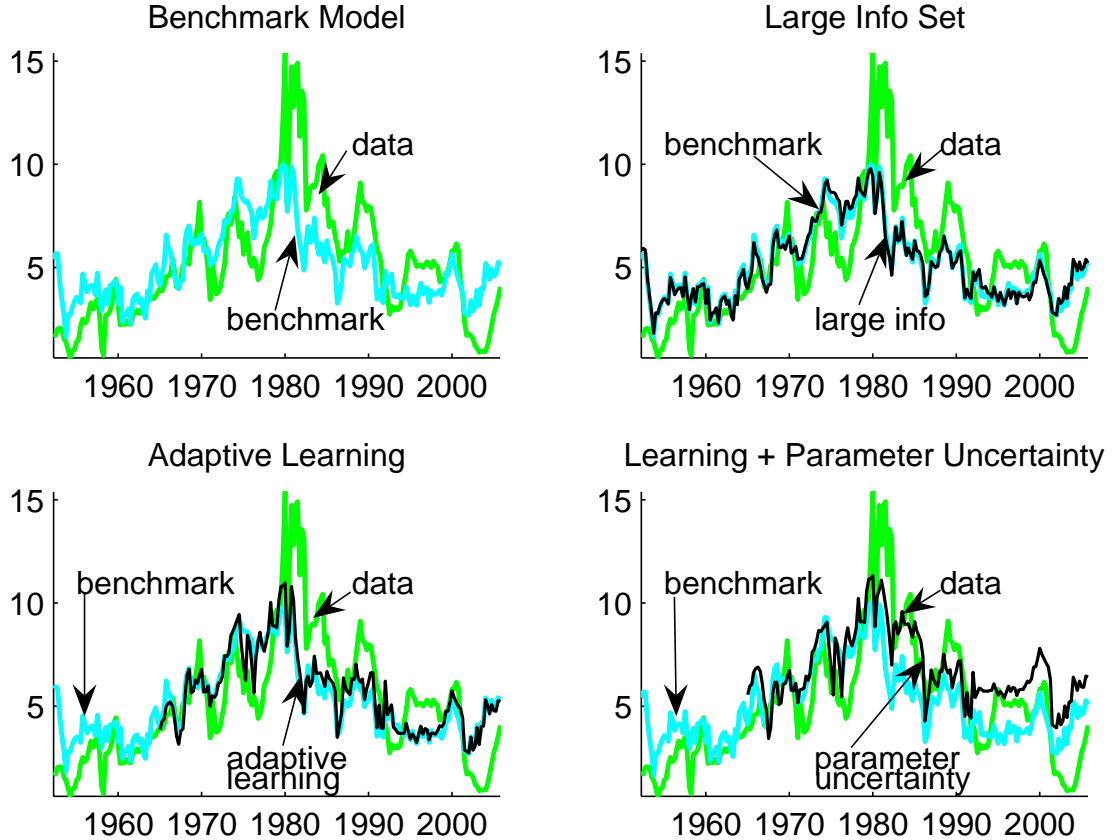


Figure 3: The top left panel plots the nominal short rate in the data and the nominal rate implied by the benchmark model. For comparison, these two lines are also included in the other plots. The additional line in the other plots is the short rate implied by the model indicated in the title.

shows that a 1-percent inflation surprise lowers growth forecasts over the next year by roughly 10 bp.

While we can read off the impulse responses of real rates directly from the top row of plots in Figure 2, we need to combine the responses from the top two rows of plots to get the response of nominal growth or, equivalently, nominal interest rates. This is done in the bottom row of plots in Figure 2. Here, inflation and growth surprises both lead to higher nominal growth forecasts – even over longer horizons. From the previous discussion, we know that this effect is entirely due to the long-lasting effect of both types of shocks on inflation. These findings imply that growth surprises and inflation surprises move short-maturity real rates in opposite directions, but won't affect long-maturity real rates much. In contrast, growth and inflation surprises affect even longer-maturity nominal rates, because they have long-lasting effects on inflation forecasts. In particular, these shocks move nominal rates in the same direction.

An inspection of the surprises e_{t+1} in equation (19) reveals that the historical experience in the U.S. is characterized by a concentration of large nominal shocks in the 1970s and early 1980s. (We do not include a plot for space reasons.) Outside this period, inflation shocks occurred rarely and were relatively small. By contrast, real surprises happened throughout the sample and their average size did not change much over time. As a consequence, our benchmark model says that yields in the 1970s and early 1980s were mainly driven by nominal shocks – inflation surprises – that affect nominal and real rates in opposite directions. Here an inflation surprise lowers real rates because it is bad news for future consumption growth. In contrast, prior to the 1970s, and again more recently, there were more real shocks – surprises in consumption growth – that make nominal and real interest rates move together.

Preference Parameters and Equilibrium Yields

The model’s predictions for yields are entirely determined by the agent’s beliefs about fundamentals and two preference parameters, the discount factor β and the coefficient of relative risk aversion γ . We select values for the preference parameters to match the average short and long end of the nominal yield curve. For our benchmark, those values are $\beta = 1.005$ and $\gamma = 59$. These numbers indicate that the agent does not discount the future and is highly risk averse. The nominal short rate implied by the benchmark model is shown in the top left plot in Figure 3. The benchmark model produces many of the movements that we observe in the data. For example, higher nominal growth expectations in the mid 1970s and early 1980s make the nominal rate rise sharply.

Average Nominal Yields

Panel A in Table 1 compares the properties of average nominal yields produced by the model with those in the data. Interestingly, the model with recursive utility produces, on average, an upward sloping nominal yield curve – a robust stylized fact in the data. The average difference between the 5-year yield and the 3-month yield in the data is roughly 1 percentage point, or 100 basis points (bp). This difference is statistically significant; it is measured with a 13 bp standard error. By contrast, the average level of the nominal yield curve is not measured precisely. The standard errors around the 5.15 percent average short end and the 6.14 percent average long end of the curve are roughly 40 bp. The intuitive explanation behind the positive slope is that high inflation means bad news about future consumption. During times of high inflation, nominal bonds have low payoffs. Since inflation

affects the payoffs of long bonds more than those of short bonds, agents requires a premium, or high yields, to hold them.

TABLE 1: AVERAGE YIELD CURVES (IN % PER YEAR)

PANEL A: AVERAGE NOMINAL YIELD CURVE						
	1 quarter	1 year	2 year	3 year	4 year	5 year
Data	5.15	5.56	5.76	5.93	6.06	6.14
SE	(0.43)	(0.43)	(0.43)	(0.42)	(0.41)	(0.41)
Benchmark Model	5.15	5.33	5.56	5.78	5.97	6.14
Benchmark Model 1-5	5.43	5.56	5.73	5.88	6.02	6.14
Expected (Log) Utility	4.92	4.92	4.91	4.90	4.89	4.88
Large Info Set with same β, γ	5.06	5.14	5.29	5.44	5.60	5.74
Large Info Set	5.15	5.28	5.48	5.71	5.93	6.14
SE Spreads	5-year minus 1 quarter yield (0.13)		5-year minus 2-year yield (0.07)			
PANEL B: AVERAGE REAL YIELD CURVE						
Benchmark Model	0.84	0.64	0.49	0.38	0.30	0.23
Expected (Log) Utility	1.22	1.21	1.21	1.21	1.21	1.21
Large Info Set with same β, γ	0.84	0.63	0.47	0.38	0.31	0.26
Large Info Set	0.70	0.40	0.17	0.04	-0.06	-0.14

Note: Panel A reports annualized means of nominal yields in the 1952:2-2005:4 quarterly data sample and the various models indicated. "SE" represent standard errors computed with GMM based on 4 Newey-West lags. "SE Spreads" represent standard errors around the average spreads between the indicated yields. For example, the 0.99 percentage point average spread between the 5-year yield and the 1-quarter yield has a standard error of 0.13 percentage points.

Panel A in Table 1 also shows that the average nominal yield curve in the data has more curvature than the curve predicted by the model. A closer look reveals that the curvature in the data comes mostly from the steep incline from the 3-month maturity to the 1-year maturity. If we leave out the extreme short end of the curve, the model is better able to replicate its average shape.² This idea is explored in the line "Benchmark Model 1-5 year" where we select parameter values to match the average 1-year and 5-year yields. The resulting parameter values are $\beta = 1.004$ and $\gamma = 43$.

²We are grateful to John Campbell for this suggestion.




Figure 4: Risk premia in the expected utility model with coefficient of relative risk aversion γ (in percent per year). The plot shows the contribution of the individual terms on the right-hand side of the expected excess return equation (20) as a function of γ .

A potential explanation for the steep incline in the data are liquidity issues that may depress short T-bills relative to other bonds. These liquidity issues are not present in our model.

In contrast, the expected utility model generates average nominal yield curves that are downward sloping. For the case with expected log utility, the negative slope is apparent from line 3 in Panel A. To see what happens in the more general case with coefficient of relative risk aversion γ , we need to re-derive the equation for expected excess returns (18). The equation becomes

$$(20) \quad E_t \left(r x_{t+1}^{(n)\$} \right) = - \text{cov}_t \left(\gamma \Delta c_{t+1} + \pi_{t+1}, E_{t+1} \sum_{i=1}^{n-1} \gamma \Delta c_{t+1+i} + \pi_{t+1+i} \right) - \frac{1}{2} \text{var}_t \left(p_{t+1}^{(n-1)\$} \right).$$

Figure 4 plots the individual terms that appear on the right-hand side of this equation as a function of γ . Most terms have negative signs and thus do not help to generate a positive slope. The only candidate involves the covariance between inflation and expected future consumption growth, $\text{cov}_t \left(\pi_{t+1}, E_{t+1} \sum_{i=1}^{n-1} \gamma \Delta c_{t+1+i} \right)$. This term is positive, because of the minus sign in equation (20)




Figure 5: The top left panel plots the nominal term spread (5-year minus 3-month) in the data and the spread implied by the benchmark model. For comparison, these two lines are also included in the other plots. The additional line in the other plots is the spread implied by the model indicated in the title.

and the fact that positive inflation surprises forecast lower future consumption growth. With a higher γ , the importance of this term goes up. However, as we increase γ , the persistence of consumption growth becomes more and more important, and the real yield curve becomes steeply downward sloping. Since this effect is quadratic in γ , it even leads to a downward-sloping nominal curve. The intuitive explanation is that long real bonds have high payoffs precisely when current and future expected consumption growth is low. This makes them attractive assets to hold and so the real yield curve slopes down. When γ is high, this effect dominates also for nominal bonds.

Average Real Yields

At the preference parameters we report, the benchmark model also produces a downward sloping real yield curve. The short real rate is already low, .84 percent, while long real rates are an additional .60 percentage point lower. It is difficult to assess the plausibility of this property of the model without a long sample on real yields for the United States. In the United Kingdom, where indexed bonds have been trading for a long time, the real yield curve seems to be downward sloping. Table B.3 reports statistics for these bonds. For the early sample (January 1983 – November 1995), these numbers are taken from Table 1 in Evans (1998). For the period after that (December 1995 – March 2006), we use data from the Bank of England website. Relatedly, Table 1 in Barr and Campbell (1997) documents that average excess returns on real bonds in the U.K. are negative.

In the United States, indexed bonds, so-called TIPS, have started trading only recently, in 1997. During this time period, the TIPS curve has been mostly upward sloping. For example, mutual funds that hold TIPS – such as the Vanguard Inflation-Protected Securities Fund – have earned substantial returns, especially during the early years. Based on the raw TIPS data, J. Huston McCulloch has constructed real yield curves. Table B.4 in Appendix B documents that the average real yield curve in these data is upward sloping. The average real short rate is .8 percent, while the average 5-year yield is 2.2 percent.

These statistics have to be interpreted with appropriate caution. First, the short sample for which we have TIPS data and, more importantly, the low risk of inflation during this short sample make it difficult to estimate averages. Second, TIPS are indexed to lagged CPI levels, so that additional assumptions are needed to compute ex ante real rates from these data. Third, there have been only few issues of TIPS, so that the data are sparse across the maturity spectrum. Finally, TIPS were highly illiquid at the beginning. The high returns on TIPS during these first years of trading may reflect liquidity premia instead of signaling positive real slopes.

Volatility of Real and Nominal Yields

Table 2 reports the volatility of real and nominal yields across the maturity spectrum. We only report one row for the benchmark recursive utility model and the (log) expected utility model, because the two models imply the same yield dynamics in equations (12) and (13). Panel A shows that the benchmark model produces a substantial amount of volatility for the nominal short rate. According to the estimated state space model (19), changes in expected fundamentals – consumption growth and inflation – are able to account for 1.8 percent volatility in the short rate. This number is lower than the 2.9 percent volatility in the data, but the model is two-thirds there. In contrast, the model predicts a smooth real short rate. This effect is due to the low persistence of consumption growth.

TABLE 2: VOLATILITY OF YIELDS (IN % PER YEAR)

	1 quarter	1 year	2 year	3 year	4 year	5 year
Data	2.91	2.92	2.88	2.81	2.78	2.74
SE	(0.36)	(0.33)	(0.32)	(0.32)	(0.31)	(0.30)
Benchmark Model + Exp. (Log) U	1.80	1.64	1.47	1.34	1.22	1.12
Large Info Set	1.81	1.68	1.54	1.43	1.34	1.25
PANEL B: REAL YIELDS						
Benchmark Model + Exp. (Log) U	0.75	0.55	0.46	0.41	0.38	0.34
Large Info Set	0.83	0.62	0.49	0.42	0.36	0.32

Panel A also reveals that the model predicts much less volatility for long yields relative to short yields. For example, the model-implied 5-year yield has a volatility of 1.1 percent, while the 5-year yield in the data has a volatility of 2.7 percent. While the volatility curve in the data is also downward sloping, the slope of this curve is less pronounced than in the model. This relationship between the volatility of long yields *relative* to the volatility of short yields is the excess volatility puzzle. This puzzle goes back to Shiller (1979) who documents that long yields derived from the expectations hypothesis are not volatile enough. According to the expectations hypothesis, long yields are conditional expected values of future short rates. It turns out that the persistence of the short rate is not high enough to generate enough volatility for long yields. Shiller’s argument applies to our benchmark specification,

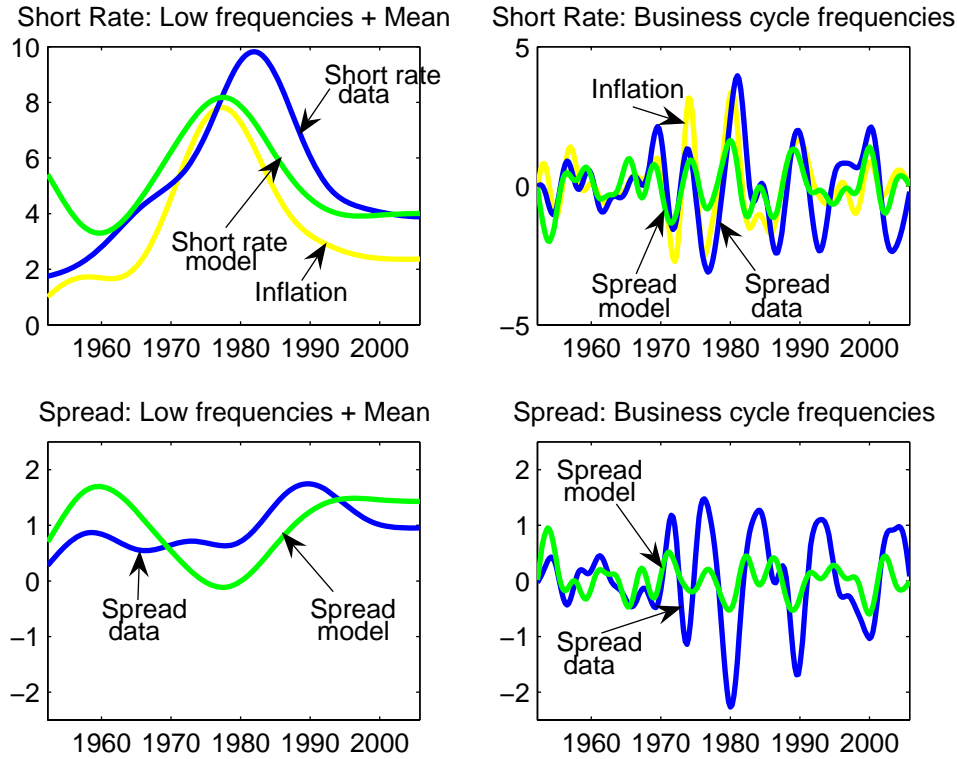


Figure 6: Low frequency components and business cycle components of nominal yields and spreads. Top row of panels: nominal short rate in the data and the benchmark model together with inflation. Bottom row of panels: nominal spread in the data and the benchmark model.

because risk premia in equation (17) are constant, and the expectations hypothesis holds. Below, we will show that our specification with learning produces more volatility for long yields.

Panel B shows that the volatility curve of real bonds also slopes down. Tables B.3 and B.4 in Appendix B show that this feature is also present in the U.K. indexed yield data and the McCulloch real yields for the U.S.

Frequency Decompositions and the Monetary Experiment

To better understand the properties of the model, we use a band-pass filter to estimate trend and cyclical components of yields. The filters isolate business-cycle fluctuations in yields that persist for periods between 1.5 and 8 years from those that persist longer than 8 years. Figure 6 plots the various estimated components. The top left panel shows the low frequency components of the model-implied short rate as well as the observed short rate and inflation. The plots shows that the model captures the fact that the low frequency component in nominal yields is strongly correlated with inflation. At

these low frequencies, the main difference between model and data is the experience of the mid 1980s. When inflation started to come down at the end of the 1970s, nominal yields stayed high well into the 1980s. According to benchmark beliefs – which are estimated over the whole data sample and which ignore parameter uncertainty – inflation forecasts came down as soon as inflation started to decline. The basic mechanism behind these changes in inflation expectations is persistence; since inflation is close to a random walk, inflation forecasts for next quarter are close to this quarter’s value of inflation. As a consequence, inflation forecasts in the early 1980s fell dramatically, right after inflation went down. In the model, changes in the nominal short rate during this period are driven by changes in inflation expectations, and so the short rate falls as well. Below, we will explore how these findings are affected by learning.

The top right panel in Figure 6 shows the business cycle movements of the same three series: nominal rate in data and model together with inflation. At this frequency, the short rate is driven by the business cycle movements in inflation. The model captures this effect, but does not generate the amplitude of these swings in the data. The bottom right panel in Figure 6 shows the business cycle movements in data on the spread and consumption growth together with those in the model. The plot reveals that the three series are strongly correlated at this frequency. In contrast, the bottom left panel shows that the series do not have clear low-frequency components.

Autocorrelation of Yields

Another feature of the benchmark model is that it does a good job in matching the high autocorrelation of short and long yields. The autocorrelation in the nominal short rate is 93.6 percent, while the model produces 93.4 percent. For the 5-year nominal yield, the autocorrelation in the model is 94.8 percent and only slightly underpredicts the autocorrelation in the data, which is 96.5 percent. These discrepancies are well within standard error bounds. As in the data, long yields in the model are more persistent than short yields. These findings are quite remarkable, because we did not use any information from nominal yields to fit the dynamics of the state space system.

TABLE 3: AUTOCORRELATION OF YIELDS

PANEL A: NOMINAL YIELDS						
	1 quarter	1 year	2 year	3 year	4 year	5 year
Data	0.936	0.943	0.953	0.958	0.962	0.965
SE	(0.031)	(0.030)	(0.028)	(0.027)	(0.027)	(0.025)
Benchmark Model + Exp. (Log) U	0.934	0.942	0.945	0.947	0.947	0.948
Large Info Set	0.946	0.954	0.959	0.961	0.962	0.962

PANEL B: REAL YIELDS						
Benchmark Model + Exp. (Log) U	0.733	0.851	0.922	0.944	0.951	0.954
Large Info Set	0.768	0.846	0.898	0.919	0.929	0.935

IV The Role of Investor Information

In the benchmark exercise of the previous section, the fundamentals – inflation and consumption growth – play two roles. On the one hand, they determine the pricing kernel: all relevant asset prices can be written in terms of their conditional moments. On the other hand, they represent investors’ information set: all conditional moments are computed given the past record of consumption growth and inflation, and nothing else. This is not an innocuous assumption. It is plausible that investors use other macroeconomic variables in order to forecast consumption growth and inflation. Moreover, investors typically rely on sources of information that do not come readily packaged as statistics, such their knowledge of institutional changes or future monetary policy.

In this section, we extend the model to accommodate a larger investor information set. In particular, we use yields themselves to model agents’ information. We proceed in two steps. First, we estimate an unrestricted state space system of the type (19) that contains not only consumption growth and inflation, but also the short rate and the yield spread. At this stage, we ignore the fact that the model itself places restrictions on the joint dynamics of these variables – the only purpose of the estimation is to construct agents’ information set. The second step of the exercise is then the same as in the benchmark case: we compute model-implied yields and compare them to the yields in the data.

The motivation for this particular way of modelling investor information comes from the theoretical model itself. If the data were in fact generated by a model economy in which yields are equal to

investors' expectations of consumption growth and inflation, our approach would perfectly recover all investor information relevant for the analysis of the yield curve. To illustrate, suppose that the short rate is given by

$$y_t^{(1)\$} = E_t [\Delta c_{t+1} + \pi_{t+1} | \mathcal{I}_t] + \text{constant},$$

where \mathcal{I}_t is the investor information set, which contains past consumption growth, inflation, and yields, but perhaps also other variables that we do not know about.

Suppose further that our unrestricted estimation delivers the true joint distribution of Δc_{t+1} , π_{t+1} , $y_t^{(1)\$}$ and $y_t^{(20)\$}$. The sequence of model-implied short rates computed in the second step of our exercise, is then, up to a constant,

$$E_t \left[\Delta c_{t+1} + \pi_{t+1} | \left(\Delta c_\tau, \pi_\tau, y_\tau^{(1)\$}, y_\tau^{(20)\$} \right)_{\tau=1}^t \right].$$

The law of iterated expectations implies that this sequence should exactly recover the data $y_t^{(1)\$}$. A similar argument holds for the yield spread. The series of model-implied yield changes would thus be identical to yield changes in the data. In other words, if the benchmark model replicates observed yield changes for *some* information structure under rational expectations, then it will generate observed yield changes also under the particular information structure we consider here.

The joint model of fundamentals and yields takes the same general form as the system (19), except that it allows for four state variables and four observables, which implies that 42 parameters must be estimated. Table A.2 in Appendix contains these parameter estimates. Figure 7 compares the autocovariance functions of the four observables in the data and for the estimated model. A first order state space structure appears to do a reasonable job in capturing the joint dynamics of fundamentals and yields. According to these estimated dynamics, low short rates and high spreads predict lower consumption growth. Moreover, high short rates and low spreads predict high inflation rates. To key question for our model is whether these real and nominal growth predictions arise from additional information contained in yields.

When we compute the model-implied short rate and term spread with a “Large Info Set”, they look very much like those from the benchmark. The top right hand panel of Figures 3 and 5 plots these series, together with the data and the benchmark results. Summary statistics on model-implied

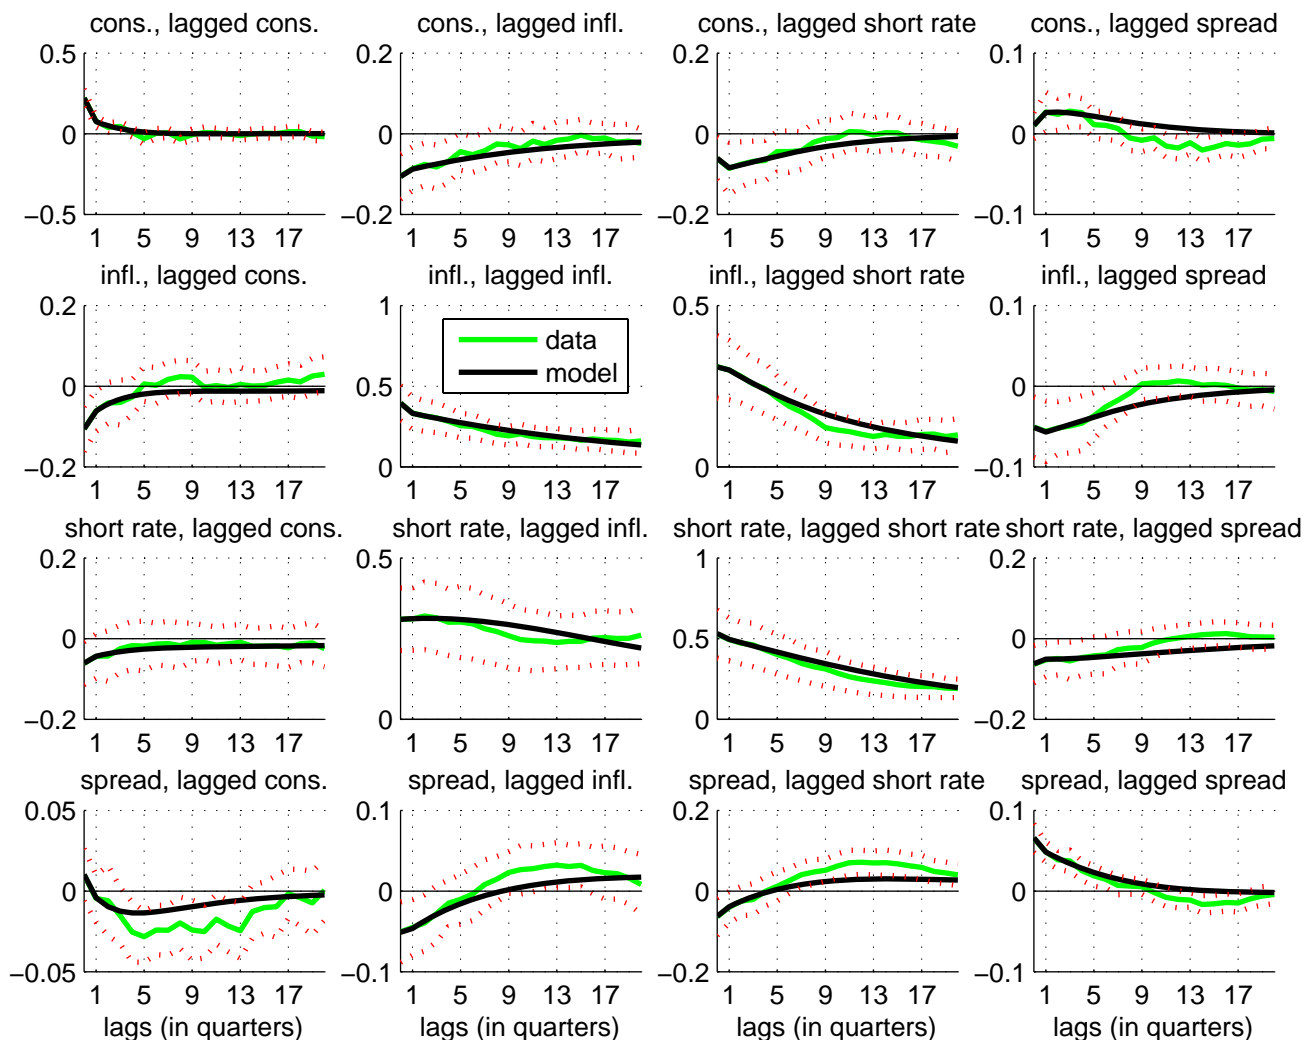


Figure 7: Covariance functions from the state space system based on a “Large Info Set” – consumption growth, inflation, the short rate, and the spread. Dotted lines indicate $2 \times$ standard errors bounds around the covariance function from the data computed with GMM.

yields from this “Large Info Set” model are also included in Tables 1 and 2. Interestingly, average nominal yields in Table 1 based on a “Large Info Set” are somewhat *lower* than in the benchmark, when we evaluate the two models at the same preference parameter values. The intuitive explanation is that more information lowers risk in the model. Line 5 of Table 1 re-phrases this finding: if we want to match the average slope of the nominal yield curve with a “Large Info Set”, we need to rely on more risk aversion, $\gamma = 85$ instead of the benchmark value of $\gamma = 59$, and a similar discount factor $\beta = 1.005$. Nevertheless, the results are overall very similar to the benchmark case. We conclude that not much is lost by restricting the investor information set to contain only past inflation and consumption growth.

The key point from this exercise is that the short rate and the yield spread do not contain much more information about future consumption growth and inflation than is already contained current and past consumption growth and inflation. Another way to see this is to run regressions of future real and nominal growth rates on current values of the four variables Δc_t , π_t , $y_t^{(1)\$}$ and $y_t^{(20)\$}$. In the one-step ahead real growth regression, the coefficient on consumption growth is .26 with a t-statistic of 4.2 and the coefficient on inflation is $-.11$ with a t-statistic of -1.85 . (These t-statistics are based on Newey-West standard errors.) The coefficients on yields are not significant and also economically tiny, around 0.0015. The R^2 in this regression is 16%. In 4-step ahead and 8-step ahead growth regressions, inflation becomes more important, but yields remain insignificant. In the one-step ahead nominal growth regression, we find the same pattern. The coefficient on consumption is .21 with a t-stat of 2.5, the coefficient of inflation is .58 with a t-stat of 5.1, and yields do not enter significantly. The R^2 of this regression is 31%. In the 4-step ahead and 8-step ahead nominal growth regressions, we get the same patterns. We can conclude that the bivariate autocovariances between, say, current consumption growth and lagged spreads in Figure 7 do not survive in multivariate regressions.

Our results may appear surprising in light of the observed volatility in yields. On the one hand, one might have expected that it is always easy to back out a latent factor from observed yields that generates a lot of volatility in model-implied yields as well.³ On the other hand, it would seem easy to change the information structure of the model in order to have information released earlier, again making conditional expectations, and hence yields, more volatile. However, an important feature of the exercise here is that we not only compute model-implied yields from an Euler equation, but also check the correlation of model implied and observed yields.

To see the difference between our exercise and other ways of dealing with information unknown to the modeler, consider the following stylized example. Assume that the true data generating process for consumption growth is constant, while inflation and the short rate are both iid with unit variance, but independent of each other. If we had performed our benchmark exercise on these data, we would have found an iid inflation process. With constant consumption growth and iid inflation, computing the short rate from the Euler equation would have delivered a constant model-implied nominal short

³Indeed, the quarterly variation in bond yields is well explained using a statistical factor model with only two latent factors, or principal components. Intuitively, the lion share of the movements in nominal yields are up/down movements across the curve. The first principal component of yields captures these so-called “level” movements which explain 98.22% of the total variation in yields. An additional 1.58% of the movements in yields is captured by the second principal component, which represents movements in the slope of the curve. Together, “level” and “slope” explain almost all, 99.80%, of the variation in yields.

rate, which is much less volatile than the observed short rate.

Now consider two alternative exercises. Exercise A assumes that investors' expected inflation is driven by a perceived "inflation target", which is backed out from the short rate (for simplicity, suppose it is set equal to the short rate). Exercise B assumes that investors' expected inflation is driven by a perceived inflation target that is equal to next period's realized inflation. This exercise may be motivated by the fact that investors read the newspaper and know more than past published numbers at the time they trade bonds. Suppose further that both exercises maintain the assumption that the Euler equation holds: model-implied short rates are computed as investors' subjective expected inflation. Both exercises then generate model-implied short rates that – when viewed in isolation – have exactly the same distribution as observed short rates.

In spite of their success in generating volatility, both exercises miss key aspects of the joint distribution of inflation and the short rate. In Exercise A, model-implied expected inflation is independent of actual inflation one period ahead, which is inconsistent with rational expectations. This happens because the inflation target is backed out from the short rate, which here moves in the data for reasons that have nothing to do with inflation or inflation expectations. In Exercise B, the model implied short rate is perfectly correlated with inflation one period ahead, while these variables are independent in the data.

The exercise of this section avoids the problems of either Exercise A or B. If the first step estimation had been done using the example data, we would have found independence of inflation and the short rate. As a result, the model-implied short rate based on the estimated information set would be exactly the same as in the benchmark case. The model would thus again imply constant short rates. We would thus have correctly inferred that yields do not contain information about future inflation and consumption growth, than is contained in the fundamentals themselves. As a result, any model economy where the Euler equation holds and beliefs are formed via rational expectations produces model-implied yields that are less volatile than observed yields.

V Learning

In the benchmark exercise of Section III, investor beliefs about fundamentals are assumed to be conditional probabilities of a process that was estimated using all data through 2005. This approach has three a priori unattractive properties. First, it ignores the fact that investors in, say, 1980 only had access to data up to 1980. Second, it assumes that agents believed in the same stationary model throughout the postwar period. This is problematic given that the 1970s are often viewed as a period of structural change. Indeed, the decade witnessed the first ever peacetime inflation in the US, the breakdown of leading macroeconomic models, as well as significant innovation in bond markets. Third, the benchmark beliefs were based on point estimates of the forcing process, ignoring the fact that the parameters of the process itself are not estimated with perfect precision, and investors know this.

In this section, we construct a sequence of investor beliefs that do not suffer from the above drawbacks. We maintain the hypothesis that, at every date t , investors form beliefs based on a state space system of the form (19). However, we reestimate the system for every date t using only data up to date t . To accommodate investor concern with structural change, we maximize a modified likelihood function that puts more weight on more recent observations. To model investor concern with parameter uncertainty, we combine the state space dynamics with a Bayesian learning scheme about mean fundamentals.

A. Beliefs

Formally, beliefs for date t are constructed in three steps. We first remove the mean from the fundamentals $z_t = (\Delta c_t, \pi_t)^\top$. Let $v \in (0, 1)$ denote a “forget factor” that defines a sequence of geometrically declining sample weights. The weighted sample mean for date t is

$$(21) \quad \hat{\mu}_z(t) = \left(\sum_{i=0}^{t-1} v^i \right)^{-1} \sum_{i=0}^{t-1} v^i z_{t-i}.$$

The sequence of estimated means for consumption growth and inflation is plotted in Figure 8. It essentially picks up the low frequency components in fundamentals.

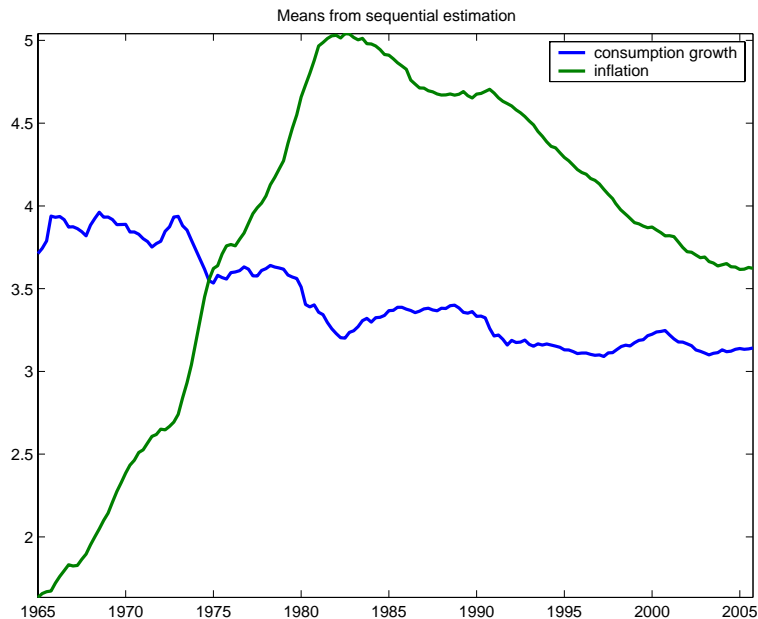


Figure 8: Sequentially estimated means of consumption growth and inflation (in percent per year). This estimation uses $v = .99$.

Adaptive Learning

In a second step, we estimate the state space system (19) using data up to date t by minimizing the criterion

$$(22) \quad -\frac{1}{2} \sum_{i=0}^{t-1} v^i \left[\log \det \Omega + (z_{t-i} - \hat{\mu}_z(t) - x_{t-1-i})^\top \Omega^{-1} (z_{t-i} - \hat{\mu}_z(t) - x_{t-1-i}) \right]$$

starting at $x_0 = 0$. Maximum likelihood estimation amounts to the special case $v = 1$; it minimizes the equally weighted sum of squared in-sample forecast errors. In contrast, the criterion (22) penalizes recent forecast errors more heavily than those in the distant past. Ljung and Soderstrom (1987) and Sargent (1993) advocate this approach to adaptive learning in situations where the dynamics of a process may change over time.

The forget factor v determines how quickly past data are downweighted. For most of our results, we use $v = .99$, which implies that the data point from 17 years ago receives about one half the weight of the most recent data point. To allow an initial sample for the estimation, the first belief is constructed for 1965:1. The analysis of yields in this section will thus be restricted to the period since 1965. As in the benchmark case, the estimation step not only delivers estimates for the matrices ϕ_x , K and Ω , but also estimates for the sequence of states $(x_\tau)_{\tau=1}^t$, starting from $x_0 = 0$. In particular,

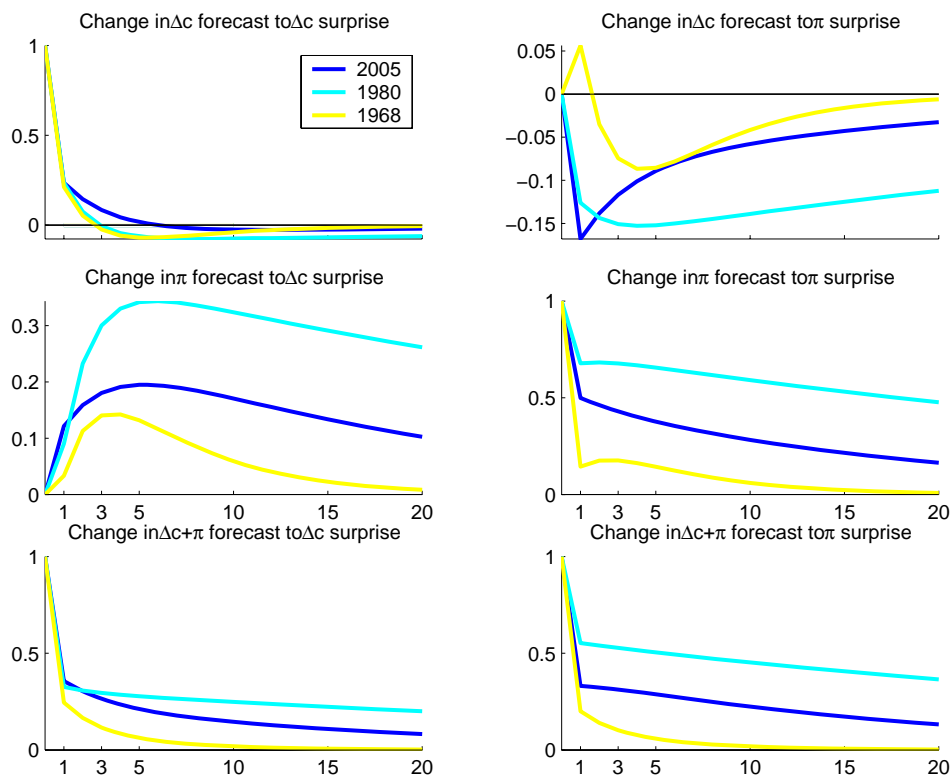


Figure 9: Impulse responses to 1-percent consumption growth and inflation surprises, in % per year, for real consumption growth, inflation, and nominal consumption growth. Time is measured in quarters along the horizontal axis.

we obtain an estimate of the current state x_t that can be taken as the basis for forecasting future fundamentals under the system estimated with data up to date t .

Figure 9 illustrates how the dynamics of consumption growth and inflation has changed over time. In each panel, we plot estimated impulse responses to consumption growth and inflation surprises, given data up to the first quarter of 1968, 1980 and 2005. In a rough sense, the three selected years represent “extreme points” in the evolution of the dynamics: impulse responses for years between 1968 and 1980 would for the most part lie in between the lines for these two years, and similarly for the period 1980-2005. The response of real growth to a growth surprise has not changed much over the years. In contrast, an inflation surprise led to a much larger revision of inflation forecasts – at all horizons – in 1980 than in 1968; the effect has diminished again since then.

Growth surprises also had a larger (positive) effect on inflation forecasts in 1980 than either before or after. While this is again true for all forecast horizons, the effect of inflation surprises on growth forecasts changed differently by horizon. For short horizons, it has decreased over time; only for longer

horizons is it largest in 1980. The bottom line is that both the persistence of inflation and its role as an indicator of bad times became temporarily *stronger* during the great inflation of the 1970s.

Performing the estimation step for every date t delivers not only sequences of parameter estimates, but also estimates of the current state x_t . Computing conditional distributions given x_t date by date produces a sequence of investor beliefs. The subjective belief at date t determines investors' evaluation of future utility and asset payoffs at date t . We thus use this belief below to calculate expectations of the pricing kernel, that is, yields, for date t . In contrast to the benchmark approach, the exercise of this section does not impose any direct restriction on beliefs across different dates; for example, it does not require that all beliefs are conditionals of the same probability over sequences of data. The updating of beliefs is thus implicit in the sequential estimation.

The model also does not impose a direct link between investor beliefs and some "true data generating process", as the benchmark approach does by imposing rational expectations. The belief at date t captures investors' subjective distribution over fundamentals at date t . It is constrained only by past observations (via the estimation step), and not by our (the modelers') knowledge of what happened later. At the same time, our approach does impose structural knowledge on the part of investors: their theory of asset prices is based on the representative agent preferences that we use.

Parameter Uncertainty

The third step in our construction of beliefs introduces parameter uncertainty. Here we focus exclusively on uncertainty about the estimated means. Our goal is to capture the intuition that, in times of structural change, it becomes more difficult to distinguish permanent and transitory changes in the economy. We thus assume that, as of date t , the investor views both the true mean μ_z and the current persistent (but transitory) component x_t as random. The distribution of z_t can be represented by a system with four state variables:

$$(23) \quad \begin{aligned} z_{\tau+1} &= \mu_z + x_\tau + e_{\tau+1}, \\ \begin{pmatrix} \mu_z \\ x_{\tau+1} \end{pmatrix} &= \begin{pmatrix} I_2 & 0 \\ 0 & \phi_x \end{pmatrix} \begin{pmatrix} \mu_z \\ x_\tau \end{pmatrix} + \begin{pmatrix} 0 \\ \phi_x K e_{\tau+1} \end{pmatrix}. \end{aligned}$$

The matrices ϕ_x , K and Ω are assumed to be known and are taken from the date t estimation step.

In order to describe investors' perception of risk, it is helpful to rewrite (23) so that investors'

conditional expectations – rather than the unobservables μ_z and x – are the state variables. Let $\hat{\mu}_z(\tau)$ and \hat{x}_τ denote investors' expectations of μ_z and x_τ , respectively, given their initial knowledge at date t as well as data up to date τ . We can rewrite (23) as

$$(24) \quad \begin{aligned} z_{\tau+1} &= \hat{\mu}_z(\tau) + \hat{x}_\tau + \hat{e}_{\tau+1}, \\ \begin{pmatrix} \hat{\mu}_z(\tau+1) \\ \hat{x}_{\tau+1} \end{pmatrix} &= \begin{pmatrix} I_2 & 0 \\ 0 & \phi_x \end{pmatrix} \begin{pmatrix} \hat{\mu}_z(\tau) \\ \hat{x}_\tau \end{pmatrix} + \begin{pmatrix} K_\mu(\tau+1) \\ \phi_x K_z(\tau+1) \end{pmatrix} \hat{e}_{\tau+1}, \end{aligned}$$

where $\hat{e}_{\tau+1}$ is investors' one step ahead forecast error of the data $z_{\tau+1}$. The matrices $K_\mu(\tau+1)$ and $K_z(\tau+1)$ can be derived by applying Bayes' Rule. They vary over time, because the learning process is nonstationary. Early on, the investor expects to adjust his estimate of, say, mean inflation, a lot in response to an inflation shock. As time goes by, the estimate of the mean converges, and the matrix K_μ converges to zero, while the matrix K_z reverts to the matrix K from (19).

To complete the description of investors' belief, it remains to specify the initial distribution of μ_z and x_t at date t . We assume that these variables are jointly normally distributed, with the mean of μ_z given by the point estimate (21) and the mean of x_t given by its point estimate from the date t estimation step. To specify the variance, we first compute the weighted sum of squares

$$(25) \quad \Sigma_z(t) = \left(\sum_{i=0}^{t-1} v^i \right)^{-1} \sum_{i=0}^{t-1} v^i (z_{t-i} - \hat{\mu}_z(t))^\top (z_{t-i} - \hat{\mu}_z(t)).$$

This provides a measure of overall uncertainty that the investor has recently experienced. We then compute the variance of the estimates $(\hat{\mu}_z(t), \hat{x}_t)$ under the assumption that the system (24) was initialized at some date $t-n$, at a variance of $\Sigma_z(t)$ for $\mu_z(t-n)$ and a variance of zero for x_{t-n} .

The idea here is to have investors' relative date t uncertainty about μ_z and x depend not only on the total variance in recent history, captured by $\Sigma_z(t)$, but also by the nature of recent dynamics, captured by the estimation step. For example, it should have been easier to disentangle temporary and permanent movements in inflation from the data if inflation has been less persistent recently. The above procedure captures such effects. Indeed, the main source of variation in investor beliefs for this exercise comes from the way the estimated dynamics of Figure 9 change the probability that an inflation surprise signals a permanent change in inflation. The patterns for yields we report below remain essentially intact if we initialize beliefs at the same variance Σ_z for all periods t . Similarly, the

results are not particularly sensitive to the choice of n . For the results below, we use $n = 25$ years.

The presence of parameter uncertainty adds permanent components to the impulse responses of growth and inflation surprises. This is because a surprise \hat{e} changes the estimate of the unconditional mean, which is relevant for forecasting at any horizon. The direction of change is given by the coefficients in the K_μ matrices. In particular, the matrix $K_\mu(t)$ will determine investors' subjective covariances between forecasts of growth and inflation in period $t + 1$ – the key determinants of risk premia in the model. For the typical date t , the coefficients in $K_\mu(t)$ reflect similar correlation patterns as the impulse responses in Figure 9. Growth surprises increase the estimates of both mean growth and mean inflation. Inflation surprises affect mean inflation positively, and mean growth negatively.

B. Yields

To compute yields, we evaluate equation (11), where all conditional means and variances for date t are evaluated under the date t subjective distribution. The results are contained in Table 4 and the bottom panels of Figures 3 and 5, which show realized yields predicted by the model. We report two types of results. The results in Table 4 and the bottom left panels of Figures 3 and 5 allow only for adaptive learning, without parameter uncertainty. For this case, we select the preference parameters so that the model matches the mean short rate and term spread, as for the previous exercises. Model-implied yields from an example with parameter uncertainty are included in the bottom right panels of Figures 3 and 5.

Implementing the case of parameter uncertainty for patient investors ($\beta \geq 1$) requires us to choose a third parameter, the planning horizon T . To see why, consider how continuation utility (4) enters the pricing kernel (6). Utility next quarter depends on next quarter's forecasts of future consumption growth, up to the planning horizon. As discussed above, the case of parameter uncertainty adds a permanent component to the impulse response of, say, an inflation surprise: an inflation surprise next quarter will lower expected consumption growth for all quarters up to the planning horizon. The “utility surprise” $v_{t+1} - E_t v_{t+1}$ therefore depends on the length of the planning horizon. Intuitively, an investor who lives longer and cares more strongly about the future, is more affected by the outcomes of future learning.⁴

⁴This effect is not present without parameter uncertainty, because the random component of future consumption growth forecasts then converges to zero with the forecast horizon. Therefore, as long as the planning horizon is long

TABLE 4: RESULTS WITH ADAPTIVE LEARNING

PANEL A: NOMINAL YIELD CURVE						
	1 quarter	1 year	2 year	3 year	4 year	5 year
Data starting 1965:1						
Mean	5.95	6.39	6.63	6.80	6.94	7.02
Volatility	2.84	2.80	2.73	2.64	2.58	2.52
Adaptive Learning Model						
Mean	5.95	6.14	6.39	6.61	6.82	7.02
Volatility	2.10	2.24	2.46	2.67	2.85	3.01

PANEL B: REAL YIELD CURVE						
Adaptive Learning Model						
Mean	1.27	1.16	1.05	0.97	0.89	0.82
Volatility	0.72	0.60	0.60	0.65	0.71	0.77

Note: The implications of the learning models can only be studied from 1965:1 onwards, because we need some initial observations to start the algorithms.

It follows that, for patient investors with a long planning horizon, the effect of risk on utility can be as large (or larger) as the effect of mean consumption growth and inflation. Since parameter uncertainty becomes the main driver of risk premia in this case, the planning horizon and the risk aversion coefficient have similar effects on the model results. For the results below, we use $T = 25000$ years and $\gamma = 4$, together with $\beta = 1$. At these parameter values, the model has interesting implications for the behavior of the short rate and spread during the monetary experiment.

Adaptive Learning

The short rate in the economy with adaptive learning behaves similarly to that in the benchmark model as long as there is little turbulence – the 1960s and early 1970s, and the 1990s. However, the model generates significantly higher short rates during the monetary experiment and also somewhat higher rates during the mid 1980s. The new movements are brought about by changes in the dynamics. In particular, the investor’s subjective covariance between inflation and future expected consumption increased a lot around 1980. This development was not just due to inflation volatility: the correlation between inflation and future consumption also increased. As the stagflation experience of the 1970s made its way into the beliefs of adaptive learners, our basic “inflation as bad news” mechanism was thus reinforced.

enough, it does not matter for the utility surprise even if $\beta > 1$.

Since inflation became such an important carrier of bad news, the 1980s not only increased the inflation premium on short bonds in the adaptive learning economy, but also introduced large spikes in the term spread. In the data, the high short rates of 1980 were accompanied by historically low term spreads. In contrast, the adaptive learning model generates a large term spread, for the same reason as it generates high short rates. Apart from this outlier, the model economy does exhibit a low frequency trend in the spread, with higher spreads after the 1980s than before.

Model implied yields from the adaptive learning economy are remarkably similar to the benchmark model immediately after the monetary experiment ended. The reason is that inflation forecasts from both models drop immediately as inflation itself comes down. This result is quite robust to alternative specifications of the learning scheme, obtained for example by changing the forget rate or switching from geometric downweighting to a rolling window approach. We conclude that learning does not induce inertia in inflation forecasts that can explain why interest rates remained high in the early 1980s.

Parameter Uncertainty

The results with parameter uncertainty also look very different in the early 1980s compared to other years. The short rate tracks the benchmark until the late 1970s. However, it then peaks at a higher rate in 1981 and it remains high thereafter. Parameter uncertainty thus generates the sluggish adjustment of yields at the end of the monetary experiment. The economy with parameter uncertainty also exhibits a transition of the spread from negative values in the late 1970s to historically high values throughout the first half of the 1980s. A similar transition took place in the data. Towards the end of the sample yields and spreads come down again, especially for the latter, the decline is more pronounced than in the data.⁵

Importantly, this is not due to sluggish inflation expectations: by design, inflation forecasts are the same in the adaptive learning and the parameter uncertainty exercises. Instead, the role of inflation as bad news is here enhanced by the difficulty investors face in disentangling permanent from transitory moves in inflation. The increase in parameter uncertainty through the 1970s implies that, in the early 1980s, there is a greater chance that an inflation surprise signals a permanent shift in inflation

⁵The parameter uncertainty model also generates low spreads at the beginning of the sample. As for the adaptive learning model, the behavior in this period is driven in part by the fact that the samples used in the sequential estimation are as yet rather short.

that would generate bad news. Since the (subjective) means of inflation and consumption growth are also negatively correlated, the inflation surprise would generate permanent bad news. For a patient investor, we obtain large movements in risk premia.

VI Related Literature

The literature on the term structure of interest rates is vast. In addition to a substantial body of work that documents the behavior of short and long interest rates and summarizes it using statistical and arbitrage-free models, there are literatures on consumption based asset pricing models, as well as models of monetary policy and the business cycle that have implications for yields. There is also a growing set of papers that documents the importance of structural change in the behavior of interest rates and the macroeconomy. We discuss these groups of papers in turn.

Statistical and Arbitrage-Free Models

Average nominal yields are increasing and concave in maturity. Excess returns on nominal bonds are positive on average and also increasing in maturity. They are also predictable using interest rate information (Fama and Bliss 1987, Campbell and Shiller 1991). The latter fact contradicts the expectations hypothesis, which says that long rates are simply averages of expected future short rates, up to a constant. The expectations hypothesis also leads to an “excess volatility puzzle” for long bond prices, which is similar to the excess volatility of stock prices: under rational expectations, one cannot reconcile the high volatility of nominal rates with observed persistence in short rates (Shiller 1979). A related literature documents “excess sensitivity” of long rates to particular shocks, such as macroeconomic announcements (Gurkaynak, Sack, and Swanson 2005).

Another stylized fact is that nominal yields of all maturities are highly correlated. Litterman and Scheinkman (1991) have shown that a few principal components explain much of the variation in yields. For example, in our quarterly postwar panel data, 99.8% of the variation is explained by the first and second principal components. Here the elephant in the room is the first component, which alone captures 98.2% of this variation and stands for the “level” of the yield curve. The second component represents changes the “slope” of the curve, while the third principal component represents “curvature” changes.

This fact has motivated a large literature on arbitrage-free models of the term structure. The goal here is to summarize the dynamics of the entire yield curve using a few unobservable factors. Recent work in this area explores the statistical relationship between term structure factors and macroeconomic variables. For example, the arbitrage-free model by Ang, Piazzesi, and Wei (2006) captures the role of the term spread as a leading indicator documented by the predictive regressions surveyed in Stock and Watson (1999). In this work, the only cross-equation restrictions on the joint distribution of macro variables and yields come from the absence of arbitrage.

In the present paper, our focus is on cross-equation restrictions induced by Euler equations, which directly link yields to conditional moments of macroeconomic variables. In particular, we focus on properties of the short rate and a single yield spread and use these to link the level and slope of the yield curve to inflation and the business cycle. The rational expectations exercises in Sections III and IV also impose the expectations hypothesis through our assumptions on preferences and the distribution of shocks. While this implies that the model economies do not exhibit predictability and excess volatility of long yields, they are useful for understanding the macro underpinnings of average yields as well as the volatility of the level factor, which accounts in turn for the lion's share of yield volatility. The learning exercises in Section V do generate predictability in yields because of time variation in perceived risk.

Consumption-based asset pricing models

The representative agent asset pricing approach we follow in this paper takes the distribution of consumption growth and inflation as exogenous and then derives yields from Euler equations. Early applications assumed power utility. Campbell (1986) shows analytically that positive serial correlation in consumption growth and inflation leads to downward sloping yield curves. In particular, term spreads on long indexed bonds are negative because such bonds provide insurance against times of low expected consumption growth. Backus, Gregory, and Zin (1989) document a “bond premium puzzle”: average returns of long bonds in excess of the short rate are negative and small for coefficients of relative risk aversion below 10. Boudoukh (1993) considers a model with power utility where the joint distribution of consumption growth and inflation is driven by a heteroskedastic VAR. Again, term premia are small and negative. The latter two papers also show that heteroskedasticity in consumption growth and inflation, respectively, is not strong enough to generate as much predictability in excess bond returns as is present in the data. Chapman (1997) documents that ex-post real rates and

consumption growth are highly correlated, at least outside the monetary policy experiment.

Our results show that the standard result of negative nominal term spreads is overturned with recursive utility if inflation brings bad news. The form of recursive utility preferences proposed by Epstein and Zin (1989) and Weil (1989) has become a common tool for describing investors' attitudes towards risk and intertemporal substitution. Campbell (1999) provides a textbook exposition. An attractive feature of these preferences is that they produce plausible quantity implications in business cycle models even for high values of the coefficient of relative risk aversion, as demonstrated in Tallarini (2000). Bansal and Yaron (2004) show that a model with recursive utility can also generate a high equity premium and a low riskfree rate if consumption growth contains a small, but highly persistent, component. They argue that, even though empirical autocovariances of consumption growth do not reveal such a component, it is hard to refute its presence given the large transitory movements in consumption growth.

Our benchmark rational expectations exercise postulates a consumption process parametrized by our maximum likelihood point estimates. As a result, the autocovariances of consumption growth in our model are close to their empirical counterparts. The effects we derive are mostly due to the forecastability of consumption growth by inflation, again suggested by our point estimates. Our learning exercise with parameter uncertainty plays off the fact that permanent and persistent transitory components can be hard to distinguish.

The literature has also considered utility specifications in which current marginal utility depends on a mean-reverting state variable. In habit formation models as well as in Abel's (1990) model of "catching up with the Joneses", marginal utility not only depends on current consumption but also on consumption growth which is mean reverting. The presence of a mean-reverting state variable in marginal utility tends to generate an upward sloping yield curve: it implies that bond prices (expected changes in marginal utility) are negatively correlated with marginal utility. Since bonds thus pay off little precisely in times of need, they command a premium. Quantitative analysis of models of habit formation and catching up with the Joneses showed that short real interest rates become very volatile when the models are calibrated to match the equity premium.

Campbell and Cochrane (1995) introduce a model in which marginal utility is driven by a weighted average of past innovations to aggregate consumption, where the weight on each new innovation is

positively related to the level of the marginal utility. With this feature, low current marginal utility need not imply extremely high bond prices, since the anticipation of less volatile innovations in the future discourages precautionary savings and lowers bond prices. In their quantitative application, Campbell and Cochrane focus on equity and short bonds, and pick the weight function so that the real riskless rate is constant and the term structure is flat. Wachter (2006) instead picks the weight function to match features of the short rate dynamics. In a model driven by iid consumption and an estimated inflation process, she shows that this approach accounts for several aspects of yield behavior, while retaining the results for equity from the Campbell-Cochrane model.⁶

Monetary and business cycle models

The consumption based asset pricing approach we follow in this paper assumes a stochastic trend in consumption. In contrast, studies in the business cycle literature often detrend real variables, including consumption, in a first step and then compare detrended data to model equilibria in which the level of consumption is stationary. This distinction is important for the analysis of interest rates, since the pricing kernel (6), derived from the Euler equation, behaves very differently if consumption is stationary in levels (Labadie 1994).⁷ Alvarez and Jermann (2005) have shown that a permanent component must account for a large fraction of the variability of state prices if there are assets that have large premia over long term bonds, as is the case in the data. A stochastic trend in consumption directly induces a large permanent component in real state prices.

Recently various authors have examined the term-structure implications of New Keynesian models. The “macro side” of these models restricts the joint distribution of output, inflation and the short nominal interest rate through an Euler equation – typically allowing for an effect of past output on current marginal utility as well as a taste shock – , a Phillips curve and a policy reaction function for the central bank. Longer yields are then linked to the short rate via an exogenous pricing kernel (Rudebusch and Wu 2005, Beechey 2005) or directly through the pricing kernel implied by the Euler equation (Bekaert, Cho and Moreno 2005, Hordahl, Tristani, and Vestin 2005, Ravenna and Seppala 2005). Our model differs from these studies in that it does not put theoretical restrictions on the

⁶The New Keynesian model of Bekaert, Cho and Moreno (2005) assumes “catching up with the Joneses” together with a taste shock to marginal utility. This is another way to reconcile the behavior of yields with a habit formation model.

⁷In particular, if consumption reverts to its mean, “good” shocks that increase consumption lead to lower expected consumption growth and hence lower real interest rates and higher real bond prices. This is exactly the opposite of the effect discussed in Section II, where “good” shocks that increase consumption growth leads to higher expected consumption growth and hence higher real interest rates and lower bond prices.

distribution of the macro variables and does not allow for taste shocks.

Our model assumes frictionless goods and asset markets. In particular, there are no frictions associated with the exchange of goods for assets, which can help generate an upward sloping yield curve. For example, Bansal and Coleman (1996) derive a liquidity premium on long bonds in a model where short bonds are easier to use for transactions purposes. Alvarez, Atkeson, and Kehoe (1999) show that money injections contribute to an upward sloping real yield curve in a limited participation model of money. This is because money injections generate mean reversion in the level of consumption of bond market participants. Seppala (2004) studies the real yield curve in a model with heterogeneous agents and limited commitment. He shows that incomplete risk sharing helps to avoid a bond premium puzzle.

Learning

Our learning exercise builds on a growing literature that employs adaptive learning algorithms to describe agent beliefs. This literature is surveyed by Evans and Honkapohja (2001). Empirical applications to the joint dynamics of inflation and real variables include Sargent (1999) and Marcet and Nicolini (2003). Carceles-Poveda and Giannitsarou (2006) consider a Lucas asset pricing model where agents learn adaptively about aspects of the price function. In these studies, learning often concerns structural parameters that affect the determination of endogenous variables. In our setup, investors learn only about the (reduced form) dynamics of exogenous fundamentals; they have full structural knowledge of the price function. Another feature of many adaptive learning applications is that standard errors on the reestimated parameters are not taken into account by agents. In our model, standard errors are used to construct subjective variances around the parameters and investors' anticipation of future learning is an important determinant of risk premia.

Learning has been applied to the analysis of the term structure by Fuhrer (1996), Kozicki and Tinsley (2001), and Cogley (2005). In these papers, the expectations hypothesis holds under investors' subjective belief, as it does in our model. Fuhrer's work is closest to ours in that he also considers the relationship between macrovariables and yields, using an adaptive learning scheme. However, the link between yields and macroeconomic variables in his model is given by a policy reaction function with changing coefficients, rather than by an Euler equation as in our setting. The paper shows that changing policy coefficients induce expectations about short rates that generate inertia in long rates

in the 1980s. In other words, inertia is due to changing conditional means. This is different from our results, where interest rates are tied to expected consumption growth and inflation. This is why, in the context of our model, changes in conditional variances are more important.

Kozicki and Tinsley (2001) and Cogley (2005) use different learning models to show that the expectations hypothesis may seem to fail in the data even if it holds under investors's subjective belief. Kozicki and Tinsley consider an adaptive learning scheme, while Cogley derives beliefs from a Bayesian VAR with time-varying parameters for yields, imposing the expectations hypothesis. Regime-switching models of interest rates deal with some of the same stylized facts on structural change as learning models. (For a survey, see Singleton 2006.) A key property is that they allow for time variation in conditional variances. Since this is helpful to capture the joint movements of inflation and the short rate, regime switching is a prominent feature of statistical models that construct ex ante real rates from inflation and nominal yield data. Veronesi and Yared (2001) consider an equilibrium model of the term structure with regime switching and power utility.

VII Conclusion

We see at least two interesting tasks for future research. The first is to understand better the sources of yield volatility at business cycle frequencies. While some of the models presented in this paper exhibit substantial volatility, and do quite well on low frequency movements in interest rate levels, none of them exhibits as much volatility at business cycle frequencies as we find in the data, especially for the yield spread. One natural extension of our benchmark rational expectations model is to capture nonlinear features of the inflation process through regime switching or other devices that allow conditional heteroskedasticity. In addition to generating more volatility, this might have interesting implications for the predictability of excess long bond returns. To evaluate rational expectations models, the analysis in Section 4 – where we capture investors' information using asset prices in a first step before computing model implied yields – provides a way to evaluate many different information structures at the same time.

A second task is to develop further models in which changes in uncertainty have first order effects on interest rates. We have provided one example of such a model and have shown that it holds some promise for understanding why interest rates were high in the 1980s, although inflation expectations

were low. However, more work is needed to reconcile the learning process with interest rates during other periods, and to integrate it more tightly with survey expectations. To this end, the tractable approach to learning that we consider in Section 5 – combining adaptive learning and parameter uncertainty – is less involved than a full Bayesian learning setup, but can nevertheless capture both agents’ understanding of the future dynamics of fundamentals and agents’ confidence in how well they understand this dynamics.

References

- Abel, A.B. 1999. "Risk Premia and Term Premia in General Equilibrium." *Journal of Monetary Economics* 43, pp. 3-33.
- Alvarez, F. and U. Jermann 2005. "Using Asset Prices to Measure the Persistence of the Marginal Utility of Wealth." *Econometrica* 73(6), pp. 1977-2016.
- Alvarez, F., A. Atkeson, and P.J. Kehoe 2005. "Time-Varying Risk, Interest Rates and Exchange Rates in General Equilibrium." *Federal Reserve Bank of Minneapolis Working Paper 627*.
- Alvarez, F., A. Atkeson, and P.J. Kehoe 1999. "Money and Interest Rates with Endogeneously Segmented Markets." *NBER Working Paper 7060*.
- Ang, A., G. Bekaert, and M. Wei 2006. "Do Macro Variables, Asset Markets, or Surveys Forecast Inflation Better?" *NBER Working Paper 11538*.
- Ang, A., M. Piazzesi, and M. Wei 2006. "What Does the Yield Curve tell us about GDP Growth?" *Journal of Econometrics* 131, pp. 359-403.
- Backus, D., A. Gregory, and S. Zin 1989. "Risk Premiums in the Term Structure: Evidence from Artificial Economies." *Journal of Monetary Economics* 24, pp. 371-399.
- Bansal, R. and W.J. Coleman II 1996. "A Monetary Explanation of the Equity Premium, Term Premium, and Riskfree Rate Puzzles." *Journal of Political Economy* 104, 1135-1171.
- Bansal, R. and A. Yaron 2004. "Risks for the long run: A potential resolution of asset pricing puzzles." *Journal of Finance* 59, pp. 1481-1509
- Barr, D.G. and J.Y. Campbell 1997. "Inflation, Real Interest Rates, and the Bond Market: A Study of UK Nominal and Index-Linked Government Bond Prices." *Journal of Monetary Economics* 39, pp. 361-383.
- Beechey, M. 2005. "Excess Sensitivity and Volatility of Long Interest Rates: The Role of Limited Information in Bond Markets." Working Paper, UC Berkeley.
- Bekaert G., S. Cho, and A. Moreno 2005. "New-Keynsian Macroeconomics and the Term Structure." *NBER Working Paper 11340*.
- Boudoukh J. 1993. "An Equilibrium Model of Nominal Bond Prices with Inflation-Output Correlation and Stochastic Volatility." *Journal of Money, Credit, and Banking* 25, pp. 636-665.
- Brandt, M.W. and K.Q. Wang 2003. "Time-Varying Risk Aversion and Unexpected Inflation." *Journal of Monetary Economics* 50, 1457-1498.
- Campbell, J.Y. 1986. "Bond and Stock Returns in a Simple Exchange Model." *Quarterly Journal of Economics* 101(4), pp. 785-803.
- Campbell, J.Y. 1993. "Intertemporal Asset Pricing without Consumption Data." *American Economic Review* 83, pp. 487-512.
- Campbell, J.Y. 1996. "Understanding Risk and Return." *Journal of Political Economy* 104, pp. 298-345.

- Campbell, J.Y. 1999. "Asset Prices, Consumption, and the Business Cycle." in: J. Taylor and M. Woodford, eds, *Handbook of Macroeconomics* 1 (19), pp. 1231-1303.
- Campbell, J.Y. and J.H. Cochrane 1995. "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior." *NBER Working paper* 4995.
- Campbell, J.Y. and R.J. Shiller 1991. "Yield Spreads and Interest Rate Movements: A Bird's Eye View." *Review of Economic Studies* 58, pp. 495-514.
- Campbell, J.Y. and L.M. Viceira 2001. "Who should buy long-term bonds?" *American Economic Review* 91(1), p. 99-127.
- Carceles-Poveda, E. and C. Giannitsarou 2006. "Asset Pricing with Adaptive Learning." Working Paper, University of Cambridge.
- Chapman, D. 1997. "The Cyclical Properties of Consumption Growth and the Real Term Structure." *Journal of Monetary Economics* 39, 145-172.
- Cogley, T. 2005. "Changing Beliefs and the Term Structure of Interest Rates: Cross-Equation Restrictions with Drifting Parameters." *Review of Economic Dynamics* 8, pp. 420-451.
- Cogley, T. and T.J. Sargent 2001. "Evolving Post-World War II U.S. Inflation Dynamics" *NBER Macro Annual* 2001.
- Cox, J.C., J.E. Ingersoll, and S.A. Ross 1985. "A theory of the term structure of interest rates." *Econometrica* 53, 385-407.
- Duffie, D. and L. Epstein 1992. "Stochastic Differential Utility." *Econometrica* 60(2), pp. 353-394.
- Duffie, D., M. Schroeder, and C. Skiadas 1997. "A Term Structure Model with Preferences for the Timing of Resolution of Uncertainty." *Economic Theory* 9, pp. 3-22.
- Epstein, L. and S. Zin 1989. "Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework." *Econometrica* 57, pp. 937-69.
- Evans, M.D. 1998. "Real Rates, Expected Inflation, and Inflation Risk Premia." *Journal of Finance* 53(1), pp. 187-218.
- Evans, C.L. and D. Marshall 1998. "Monetary policy and the term structure of nominal interest rates: Evidence and theory." *Carnegie-Rochester Conference Series on Public Policy* 49, 53-111.
- Fama, E.F. and R.R. Bliss 1987. "The Information in Long-Maturity Forward Rates." *American Economic Review* 77, pp. 680-92.
- Fuhrer, J. 1996. "Monetary Policy Shifts and Long-Term Interest Rates." *Quarterly Journal of Economics* 111(4), pp. 1183-1209.
- Gurkaynak, R., B. Sack, and E. Swanson 2005. "The Sensitivity of Long-Term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models." *American Economic Review* 95(1), March 2005, 425-36.
- Hansen, L.P., J.C. Heaton, and N. Li 2005. "Consumption strikes back? Measuring Long Run Risk." Working Paper, University of Chicago.
- Hordahl, P., O. Tristani, and D. Vestin 2005. "The Yield Curve and Macroeconomic Dynamics." Working Paper, European Central Bank.

- Kozicki, S. and P.A. Tinsley 2001. "Shifting Endpoints in the Term Structure of Interest Rates." *Journal of Monetary Economics* 47, pp. 613-652.
- Litterman, R. and J. Scheinkman 1991. "Common factors affecting bond returns." *Journal of Fixed Income* 1, 54-61.
- Ljung, L. and T. Soderstrom 1987. "*Theory and Practice of Recursive Identification.*" MIT Press.
- Marcet, A. and J.P. Nicolini 2003. "Recurrent Hyperinflations and Learning." *American Economic Review* 93(5), 1476-1498.
- Piazzesi, M. and M. Schneider 2006. "Inflation and the Price of Real Assets." Working Paper, University of Chicago and FRB Minneapolis.
- Piazzesi, M., M. Schneider, and S. Tuzel 2006. "Housing, Consumption, and Asset Pricing." Forthcoming *Journal of Financial Economics*.
- Restoy, F. and P. Weil 1998. "Approximate Equilibrium Asset Prices." *NBER Working Paper* 6611.
- Rudebusch, G. and T. Wu 2005. "A Macro-Finance Model of the Term Structure, Monetary Policy and the Economy." Working Paper, Federal Reserve Bank of San Francisco.
- Sargent, T.J. 1993. "*Bounded Rationality in Macroeconomics*" Oxford University Press.
- Sargent, T.J. 1999. "*The Conquest of American Inflation.*" Princeton University Press, Princeton, NJ.
- Seppala, J. 2004. "The Term Structure of Real Interest Rates: Theory and Evidence from UK index-linked Bonds." *Journal of Monetary Economics* 51, pp. 1509-1549.
- Shiller, R. 1979. "The Volatility of Long-Term Interest Rates and Expectations Models of the Term Structure." *Journal of Political Economy* 87(6) pp. 1190-1219.
- Sims, C. and T. Zha. "Were there Regime Switches in U.S. Monetary Policy?" Working paper, Princeton University.
- Singleton, K.J. 2006. "*Empirical Dynamic Asset Pricing.*" Princeton University Press.
- Stock, J.H. and M. Watson 1999. "Business Cycle Fluctuations in U.S. Macroeconomic Time Series." in: J. Taylor and M. Woodford, eds, *Handbook of Macroeconomics* 1, pp. 3-61.
- Tallarini, T.D. 2000. "Risk-sensitive Real Business Cycles." *Journal of Monetary Economics* 45, pp. 507-532.
- Uhlig, H. 2006. "Asset Pricing with Epstein-Zin Preferences." Working Paper, Humboldt University, Berlin.
- Veronesi, P. and F. Yared 2001. "Short and Long Horizon Term and Inflation Risk Premia in the U.S. Term Structure." Working Paper, University of Chicago.
- Wachter, J. 2006. "A Consumption-Based Model of the Term Structure of Interest Rates." *Journal of Financial Economics* 79, pp. 365-399.
- Weil, P. 1989. "The Equity Premium Puzzle and the Risk-Free Rate Puzzle." *Journal of Monetary Economics* 24, pp. 401-421.

Appendix

A Estimation of the State Space System

Given the normality assumption on the disturbance vector e_{t+1} , the log likelihood function of the vector z_{t+1} is easily derived as the sum of log Gaussian conditional densities. In setting up these conditional densities, we compute the state vector x_t recursively as $x_t = \phi_x x_{t-1} + \phi_x K (z_t - x_{t-1})$ starting with $x_0 = 0$. The resulting parameter estimates are reported in Tables A1 and A2. The data are in percent and sampled at a quarterly frequency, 1952:2-2005:4. For example, this means that $\mu_c = 0.823$ represents a mean annualized consumption growth rate of $0.823 \times 4 = 3.292$ percent. We de-mean the series for the estimation, which is why we do not report standard errors for the means.

The dotted lines in Figure 1 are $2 \times$ standard error bounds computed using GMM. We use these bounds to answer the question whether the point estimate of the covariance function from the model is within standard error bounds computed from the data, without imposing the structure from the model. For each element of the covariance function, we estimate a separate GMM objective function. For example, we use moments of the type $h(t, \theta) = (\Delta c_t - \mu_c)(\Delta c_{t-1} - \mu_c) - \theta$ or $h(t, \theta') = (\Delta c_t - \mu_c)(\pi_{t-1} - \mu_\pi) - \theta'$. We compute the GMM weighting matrix with 4 Newey-West lags.

TABLE A.1: MAXIMUM LIKELIHOOD FOR BENCHMARK

	μ_z	chol(Ω)		ϕ_x		$\phi_x K$	
Δc	0.823	0.432	0	0.544	-0.099	0.242	-0.117
	-	(0.021)	-	(0.170)	(0.054)	(0.074)	(0.097)
π	0.927	-0.092	0.293	0.280	1.019	0.089	0.526
	-	(0.021)	(0.014)	(0.118)	(0.037)	(0.050)	(0.067)

Note: This table contains the parameter estimates for the "Benchmark" system

$$\begin{aligned} z_{t+1} &= \mu_z + x_t + e_{t+1} \\ x_{t+1} &= \phi_x x_t + \phi_x K e_{t+1} \end{aligned}$$

where $z_{t+1} = (\Delta c_{t+1}, \pi_{t+1})^\top$. The system starts at $x_0 = 0$. "chol(Ω)" is the Cholesky decomposition of $\text{var}(e_{t+1}) = \Omega$. Brackets indicate maximum-likelihood asymptotic standard errors computed from the Hessian.

TABLE A.2: MAXIMUM LIKELIHOOD FOR LARGE INFO SET MODEL

	μ_z				$\text{chol}(\Omega)$				
Δc	0.823				0.422	0	0	0	
	–				(0.021)	–	–	–	
π	0.927				–0.082	0.288	0	0	
	–				(0.020)	(0.014)	–	–	
$y^{(1)\$}$	1.287				0.031	0.045	0.234	0	
	–				(0.016)	(0.016)	(0.011)	–	
$y^{(20)\$} - y^{(1)\$}$	0.248				–0.013	–0.017	–0.112	0.119	
	–				(0.011)	(0.011)	(0.010)	(0.006)	
		ϕ_x				$\phi_x K$			
Δc	0.604	0.256	0.139	–0.096	0.243	0.070	0.119	–0.088	
	(0.156)	(0.109)	(0.096)	(0.073)	(0.083)	(0.052)	(0.041)	(0.029)	
π	–0.057	1.042	0.126	–0.036	–0.075	0.440	0.098	–0.098	
	(0.070)	(0.048)	(0.043)	(0.028)	(0.107)	(0.076)	(0.056)	(0.039)	
$y^{(1)\$}$	–0.008	–0.027	0.906	0.023	–0.239	0.142	0.7701	0.043	
	(0.047)	(0.032)	(0.030)	(0.019)	(0.192)	(0.113)	(0.093)	(0.064)	
$y^{(20)\$} - y^{(1)\$}$	0.151	–0.030	–0.022	0.883	0.090	–0.195	0.286	0.548	
	(0.115)	(0.081)	(0.074)	(0.049)	(0.246)	(0.165)	(0.137)	(0.101)	

Note: This table contains the parameter estimates for the "Large Info Set" system

$$\begin{aligned} z_{t+1} &= \mu_z + x_t + e_{t+1} \\ x_{t+1} &= \phi_x x_t + \phi_x K e_{t+1} \end{aligned}$$

where $z_{t+1} = (\Delta c_{t+1}, \pi_{t+1}, y_{t+1}^{(1)}, y_{t+1}^{(20)\$} - y_{t+1}^{(1)\$})^\top$. The system starts at $x_0 = 0$. "chol(Ω)" is the Cholesky decomposition of $\text{var}(e_{t+1}) = \Omega$. The data are in percent and sampled quarterly, 1952:2 to 2005:4. Standard errors are computed from the Hessian.

B U.K. and U.S. Evidence on Real Bonds

Table 1 in Evans (1998) reports means, volatilities and autocorrelations for U.K. indexed yields for the monthly sample January 1983 – November 1995. The Bank of England posts its own interpolated real yield curves from U.K. indexed yields. The sample of these data starts later and has many missing values for the early years, especially for short bonds. Panel A in Table B.3. therefore reproduces the statistics from Table 1 in Evans (1998) for the early sample. Panel B in Table B.3 reports statistics based on the data from the Bank of England starting in December 1995.

The data from the Bank of England can be downloaded in various files from the website <http://www.bankofengland.co.uk/statistics/yieldcurve/index.htm>. The data are daily observations.

To construct a monthly sample, we take the last observation from each month. The shortest maturity for which data are available consistently is 2 1/2 years. There are a few observations on individual maturities missing. We extrapolate these observations from observations on yields with similar maturities.

TABLE B.3: U.K. INDEXED BONDS

PANEL A: JANUARY 1983 - NOVEMBER 1995

	2 years	3 years	4 years	5 years	10 years
mean	6.12	5.29	4.62	4.34	4.12
volatility	1.83	1.17	0.70	0.53	0.45
autocorrelation	0.63	0.66	0.71	0.77	0.85

PANEL B: DECEMBER 1995 - MARCH 2006

	2 1/2 yr.	3 years	4 years	5 years	10 years	15 years	20 years
mean	2.59	2.56	2.51	2.48	2.41	2.38	2.33
volatility	0.86	0.78	0.70	0.67	0.66	0.69	0.74
autocorrelation	0.98	0.97	0.97	0.97	0.98	0.98	0.99

J. Huston McCulloch has constructed interpolated real yield curves from TIPS data. His website <http://www.econ.ohio-state.edu/jhm/ts/ts.html> has monthly data that start in January 1997. Table B.4 reports the properties of these real yields together with the McCulloch nominal yields from January 2000 until January 2006.

TABLE B.4: McCULLOGH DATA

PANEL A: REAL YIELD CURVE

	1 quarter	1 year	2 year	3 year	4 year	5 year
mean	0.79	1.06	1.39	1.69	1.95	2.16
volatility	1.86	1.61	1.37	1.23	1.15	1.09
autocorrelation	.847	.872	.908	.935	.947	.951

PANEL B: NOMINAL YIELD CURVE

mean	2.92	3.14	3.42	3.69	3.93	4.14
volatility	1.84	1.69	1.51	1.36	1.22	1.10
autocorrelation	.963	.960	.954	.945	.935	.923

C Alternative Data Definitions and Estimation Strategies

Section C.1 investigates the impact of using alternative inflation measures for the estimation. We show that the choice of inflation series matters for the model’s predictions regarding the level of yields in the 1970s and the volatility of yields. These alternative measures, especially CPI inflation, are, however, not appropriate because they do not correspond to our measure of consumption. Section C.2 investigates the impact of using alternative assumptions about population growth. We document that the standard population growth series look strange and different from each other. We estimate the model based on these alternatives and find that the numbers change, but not in ways that affect the story we tell in the paper. Section C.3 investigates the impact of estimating the initial state x_0 . Again, the numbers change, but not by much.

C.1 Inflation Data

We measure inflation with the NIPA price index that corresponds to precisely our measure of aggregate consumption. This measure is the most appropriate for matching theory to data, and therefore we use it in our empirical work. Our measure of inflation differs from other conventional inflation measures, such as the Consumer Price Index. Here, we check the impact of using a different inflation measure in our estimation.

Figure 10 plots our inflation measure together with CPI inflation in the top panel and inflation measured by the GNP deflator in the lower panel. It is clear from the plot that our measure of inflation has less high-frequency noise and is lower in the 1970s than CPI inflation. The measure is more similar to the GDP deflator and other inflation measures (such as the personal consumption deflator) that are computed from time-averaged data.

To check how these differences affect the implications of the benchmark model, we re-estimated the state space system (19) with CPI inflation. The CPI-based system improves upon Figure 3, because it generates higher nominal yields in the 1970s. It also improves upon Table 2 by producing more volatility for nominal yields. For example, the volatilities of the $n = 1, 4$, and 20 maturity CPI-implied yields are 2.14, 1.90, and 1.31 percent as opposed to the 1.80, 1.64, and 1.12 percent in Table 2. However, a frequency decomposition as in Figure 6 reveals that the volatility in the CPI




Figure 10: Comparison of our inflation measure (black line) with the alternative inflation measure (green line) indicated in the title of the subplot.

numbers generates volatility at the “wrong frequency” — volatility at frequencies higher than the business cycle. Other popular inflation measures (such as the GNP Deflator or Personal Consumption Deflator) generate results that are more similar to our benchmark results.

C.2 Population Data

As we explain on page 14, the results in the paper are based on the assumption that population growth is constant. Under this assumption, per capita consumption growth is equal to the growth rate of the raw consumption NIPA data minus a constant. The size of the constant only affects the interpretation of the discount factor β in the model. The results in the paper are based on a constant which is equal to zero.

The advantage of working with the assumption of constant population growth is that we do not have to take a stance on which population series we want to use. There are surprisingly large differences in the standard population series available from various data sources. Not only do these series look




Figure 11: Data on annualized population growth from two different sources. The top panel plots NIPA population growth. The bottom panel plots the original BLS series. The growth rates are computed from the population series N_t as $400 \times (\log N_t - \log N_{t-1})$.

strange, they also differ widely from each other.

Figure 11 plots two standard population growth rates. The two panels in the figure use quarterly population numbers N_t and compute the annualized growth rate as $400 \times (\log N_t - \log N_{t-1})$. The top panel plots the growth rate from the NIPA quarterly tables. The bottom panel plots the original BLS numbers, where N_t is the quarterly average of the monthly data. The BLS data is downloadable from <http://www.bls.gov/cps>.

Both series look strange, in that they exhibit large spikes. The spikes often happen at the beginning of decades, which indicates that the construction of these population series suffers from interpolation issues between each census. The other issue is that the two series look nothing alike. For example, the NIPA population growth rate exhibits high-frequency fluctuations that are clearly not present in the BLS data. We checked into various other databases (such as the population numbers posted on FRED), and found that they also look strange and different from the two series in Figure 11.

We re-estimated the model with per-capita consumption based on the BLS data. Figure 12 in this

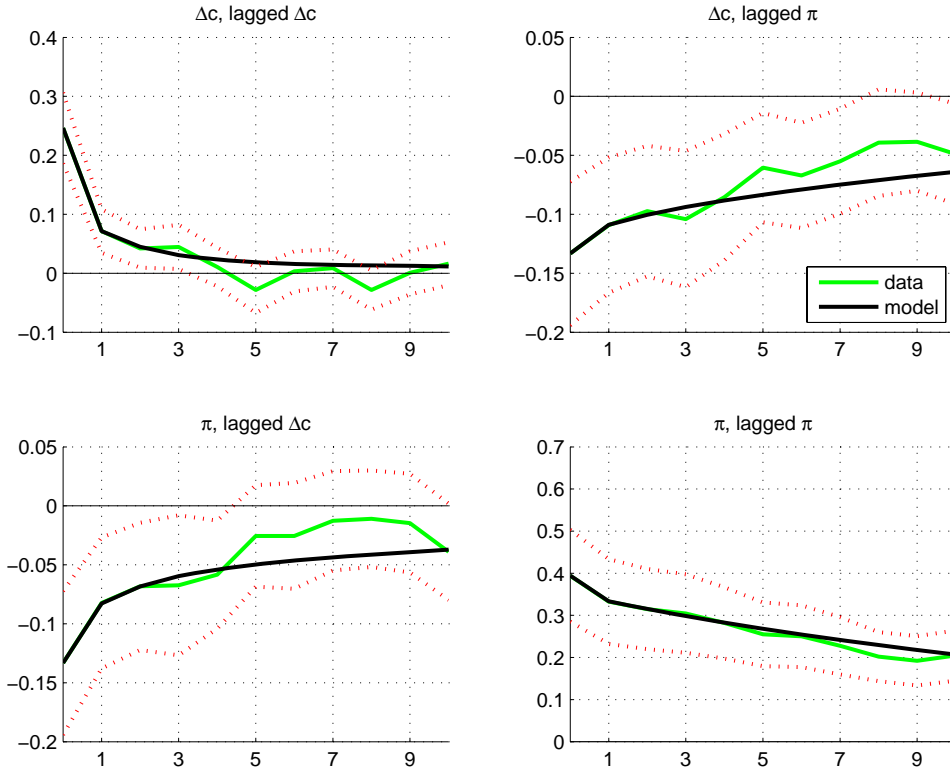


Figure 12: Results with BLS population growth.

appendix is the analogue of Figure 1. We can see that the precise numbers differ, but the findings are similar. Not surprisingly, the benchmark model now needs a lower discount factor, $\beta = 1.0016$ (as opposed to the $\beta = 1.005$ number we report on page 19) in order to match the sample average of the shortest and longest yield in the dataset. The model also needs a lower coefficient of relative risk aversion $\gamma = 37$ (as opposed to the $\gamma = 59$ in the paper).

We also re-estimated the model with per-capita consumption based on the NIPA data. Again, Figure 13 is the analogue of Figure 1. Again, the precise numbers differ, but the findings are similar. Again, not surprisingly, the benchmark model now needs a lower discount factor, $\beta = 1.0014$ (as opposed to the $\beta = 1.005$ number we report on page 19) in order to match the sample average of the shortest and longest yield in the dataset. Differently from the BLS data, the model now needs a higher coefficient of relative risk aversion $\gamma = 210$ (as opposed to the $\gamma = 59$ in the paper).




Figure 13: Results with NIPA population growth.

C.3 Initial State x_0

To construct benchmark beliefs in Section 3 of the paper and beliefs with a large information set in Section 4, we estimate the parameters μ_z , ϕ_x , K , and $\text{chol}(\Omega)$ of the state space model (19) using maximum likelihood. We form the log likelihood function as the sum of log conditional Gaussian densities of the observed vector z_{t+1} conditional on x_t for $t = 1, \dots, T$. Appendix A reports results based on the assumption that the density for the first observation z_1 conditions on the value $x_0 = 0$ for the initial state, its unconditional mean. An alternative approach is to treat the initial state x_0 as parameter which we estimate along with μ_z etc. Here, we report results based on this alternative approach.

Tables C.1 and C.2 are the analogues of Tables A.1 and A.2 in Appendix A, but now we also report estimates for the parameter vector x_0 . The parameter estimates for the benchmark model are essentially the same. The results for the large infoset model are also similar.

TABLE C.1: MAXIMUM LIKELIHOOD FOR BENCHMARK

	μ_z	chol(Ω)		ϕ_x		$\phi_x K$		x_0
Δc	0.823	0.424	0	0.514	-0.105	0.246	-0.114	1.070
	-	(0.021)	-	(0.149)	(0.052)	(0.068)	(0.098)	(0.417)
π	0.927	-0.085	0.291	0.255	1.013	0.101	0.522	-0.738
	-	(0.020)	(0.014)	(0.105)	(0.034)	(0.049)	(0.065)	(0.281)

TABLE C.2: MAXIMUM LIKELIHOOD FOR LARGE INFO SET MODEL

	μ_z	x_0		chol(Ω)				
Δc	0.823	0.620		0.420	0	0	0	
	-	(0.570)		(0.022)	-	-	-	
π	0.927	-0.903		-0.071	0.281	0	0	
	-	(0.252)		(0.020)	(0.014)	-	-	
$y^{(1)\$}$	1.287	-0.962		0.043	0.035	0.226	0	
	-	(0.257)		(0.016)	(0.016)	(0.011)	-	
$y^{(20)\$} - y^{(1)\$}$	0.248	-0.167		-0.011	-0.019	-0.119	0.110	
	-	(0.168)		(0.011)	(0.012)	(0.010)	(0.005)	
		ϕ_x		$\phi_x K$				
Δc	0.659	0.333	0.132	-0.066	0.181	0.039	0.122	-0.098
	(0.127)	(0.113)	(0.099)	(0.074)	(0.095)	(0.047)	(0.044)	(0.030)
π	-0.057	1.068	0.136	-0.036	-0.163	0.365	0.064	-0.098
	(0.058)	(0.045)	(0.041)	(0.030)	(0.091)	(0.064)	(0.060)	(0.043)
$y^{(1)\$}$	-0.005	-0.031	0.901	0.027	-0.398	0.060	0.684	0.052
	(0.041)	(0.034)	(0.028)	(0.019)	(0.182)	(0.129)	(0.104)	(0.075)
$y^{(20)\$} - y^{(1)\$}$	0.139	-0.028	-0.004	0.869	-0.203	-0.354	0.161	0.558
	(0.103)	(0.086)	(0.069)	(0.049)	(0.248)	(0.196)	(0.149)	(0.115)

Based on the new set of parameter values in Tables C.1 and C.2, we now select values for the preference parameters to match the average short and long end of the nominal yield curve. For our benchmark, those values are $\beta = 1.005$ and $\gamma = 57$. Tables C.3 and C.4 report results based on these parameters.

TABLE C.3: AVERAGE YIELD CURVES (% PER YEAR)

PANEL A: AVERAGE NOMINAL YIELD CURVE

	1 quarter	1 year	2 year	3 year	4 year	5 year
Data	5.15	5.56	5.76	5.93	6.06	6.14
SE	(0.43)	(0.43)	(0.43)	(0.42)	(0.41)	(0.41)
Benchmark Model	5.15	5.33	5.56	5.78	5.97	6.14
Expected (Log) Utility	4.92	4.92	4.91	4.90	4.89	4.88
Large Info Set with same β, γ	4.94	5.03	5.16	5.29	5.41	5.52
Large Info Set	5.15	5.31	5.53	5.74	5.95	6.14
Benchmark Model 1-5 year	5.43	5.56	5.73	5.88	6.02	6.14
SE Spreads	5-year minus 1 quarter yield (0.13)			5-year minus 2-year yield (0.07)		

PANEL B: AVERAGE REAL YIELD CURVE

Benchmark Model	0.84	0.64	0.49	0.38	0.30	0.23
Expected (Log) Utility	1.22	1.21	1.21	1.21	1.21	1.21
Large Info Set	0.70	0.40	0.17	0.04	-0.06	-0.14
Benchmark Model 1-5 year	1.32	1.19	1.08	1.01	0.95	0.91

TABLE C.4: VOLATILITY OF YIELDS (% PER YEAR)

PANEL A: NOMINAL YIELDS

	1 quarter	1 year	2 year	3 year	4 year	5 year
Data	2.92	2.92	2.88	2.81	2.78	2.74
SE	(0.36)	(0.33)	(0.32)	(0.32)	(0.31)	(0.30)
Benchmark Model + Exp. (Log) U	1.80	1.64	1.48	1.34	1.23	1.13
Large Info Set	1.80	1.72	1.63	1.56	1.50	1.44

PANEL B: REAL YIELDS

Benchmark Model + Exp. (Log) U	0.76	0.54	0.45	0.41	0.37	0.34
Large Info Set	0.83	0.63	0.50	0.42	0.37	0.32