Federal Reserve Bank of New York
Staff Reports

Pricing TIPS and Treasuries with Linear Regressions

Michael Abrahams
Tobias Adrian
Richard K. Crump
Emanuel Moench

Staff Report No. 570
September 2012

This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in this paper are those of the authors and are not necessarily reflective of views at the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

Pricing TIPS and Treasuries with Linear Regressions

Michael Abrahams, Tobias Adrian, Richard K. Crump, and Emanuel Moench

Federal Reserve Bank of New York Staff Reports, no. 570
September 2012
JEL classification: G12, E44

Abstract

We present an affine term structure model for the joint pricing of Treasury Inflation-Protected Securities (TIPS) and Treasury yield curves that adjusts for TIPS' relative illiquidity. Our estimation using linear regressions is computationally very fast and can accommodate unspanned factors. The baseline specification with six principal components extracted from Treasury and TIPS yields, in combination with a liquidity factor, generates negligibly small pricing errors for both real and nominal yields. Model-implied expected inflation provides a better prediction of actual inflation than breakeven inflation. The value of the deflation floor calculated from the model is generally small in magnitude, but spiked during the recent crisis.

Key words: TIPS, inflation expectations, affine term structure models

---

Abrahams: Massachusetts Institute of Technology (e-mail: miabraha@mit.edu). Adrian, Crump, Moench: Federal Reserve Bank of New York (e-mail: tobias.adrian@ny.frb.org, richard.crump@ny.frb.org, emanuel.moench@ny.frb.org). The authors thank Benjamin Mills for excellent research assistance and colleagues from the Federal Reserve System for feedback on the modeling approach used in this study. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

1 Introduction

The evolution of inflation expectations is an important input to monetary policy decisions and is closely watched by financial market participants. Breakeven inflation - the difference between nominal yields from Treasuries and real yields from TIPS for a given maturity - reflects inflation expectations, but is subject to two important biases. First, TIPS are often perceived to be less liquid than Treasuries, especially in times of financial market stress. Second, breakeven inflation incorporates compensation for bearing inflation risk, the so-called inflation risk premium.

In this paper, we present a Gaussian affine term structure model (ATSM) for the joint pricing of the Treasury and TIPS yield curves that adjusts for the illiquidity of TIPS and generates estimates of inflation risk premiums at various maturities. Our approach has a number of advantages relative to the existing literature. First, we adjust for TIPS illiquidity in a fashion that is internally consistent within the model based on an observable illiquidity factor. A second advantage of our approach is that it allows for a large number of pricing factors to be included in the model without impairing computational feasibility. This is because our estimation approach is based on linear regressions which are easy and fast to implement. Finally, the pricing errors implied by our model are negligibly small allowing us to decompose breakeven inflation rates into its components with almost no error.

We fit our pricing model to the joint term structures of TIPS and Treasuries. ATSMs capture the $\mathbb{P}-$ and $\mathbb{Q}-$dynamics of the term structure explicitly as vector autoregressions. The traditional estimation approach for this class of models is to estimate parameters via maximum likelihood, and factors via the Kalman Filter (see Piazzesi (2003) and Singleton (2006) for surveys). While the traditional estimation approach has theoretical appeal, these models prove challenging to estimate in practice for several reasons (see Kim (2009)). First, the optimization problem is of high dimensionality, quadratic in the number of factors, making it hard to ever know for sure whether numerical solutions represent global optima. The second problem with the filtering approach is that it tends to be computationally intensive. In fact, the likelihood function often turns out to be very flat, lengthening the optimization process, a serious issue for policy applications where models need to be updated quickly.

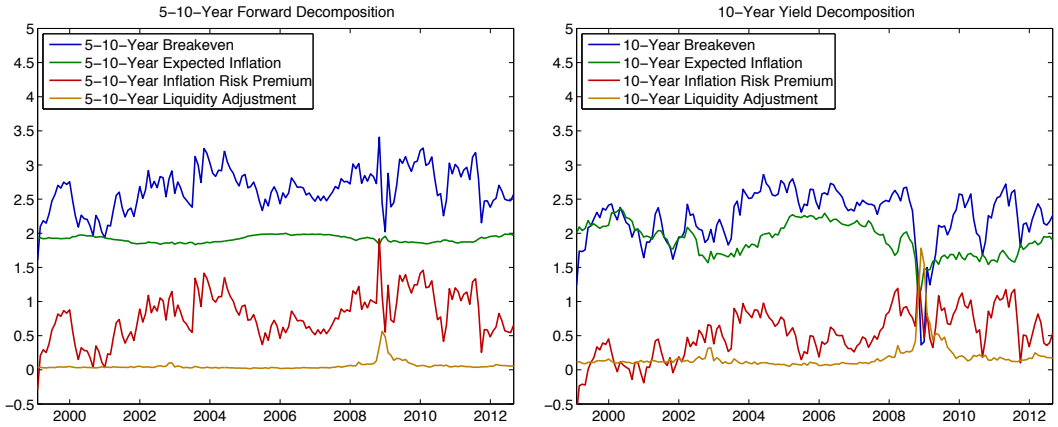

Figure 1: This figure shows the decomposition of breakeven inflation rates into the model-implied expected inflation, the inflation risk premium as well as the liquidity components. The left-hand panel shows this decomposition for 5-10 year forward breakeven inflation whereas the right-hand panel displays the decomposition for the 10 year horizon.

Third, likelihood based estimation approaches assume yield observation errors to be serially uncorrelated. However, as demonstrated in Adrian, Crump and Moench (2012), the assumption of serially uncorrelated yield pricing errors implies excess return predictability which violates that returns are arbitrage free.

In this paper, we apply a variant of the three step linear regression estimator introduced by Adrian, Crump and Moench (2012) to the joint pricing of TIPS and Treasuries. In the first step, we decompose observable pricing factors into predictable components and factor innovations by regressing factors on their lagged levels. In the second step, we estimate exposures of Treasury and TIPS excess returns with respect to lagged levels of pricing factors and contemporaneous pricing factor innovations. In the third step, we obtain the market price of risk parameters from a cross-sectional regression of the exposures of returns to the lagged pricing factors onto exposures to contemporaneous pricing factor innovations. As our pricing factors, we employ the first six principal components extracted from the joint cross-section of Treasury and TIPS yields as well as an illiquidity factor. The latter is constructed as the average absolute fitting errors from the Nelson and Siegel (1987) curve for TIPS relative to the absolute fitting error for Treasuries.

An important implication from our model relates to the decomposition of far in the future breakevens into expected inflation, the inflation risk premium, and an illiquidity component. It has long been argued that variations of 5-10 year forward rates mainly reflect changes in risk and illiquidity premia, not in inflation expectations (see Sack and Elsasser (2004) and Dudley, Roush and Ezer (2009)). Our model confirms that conjecture. The left-hand panel of Figure 1 shows that model implied expected 5-10 year forward inflation is very stable, while the variation in the forward breakeven rates mainly captures variation in the estimated inflation risk premium. Note that the picture looks very different for the decomposition of 10 year breakeven inflation (the right-hand panel). In fact, for the 10 year decomposition, both expected inflation and the inflation risk premium vary considerably. The liquidity adjustment is quantitatively important for both the 5-10 year and the 10 year maturities during the fall of 2008.

While our inflation risk premium is a linear combination of the factors used in our model, we find that it is highly correlated with a number of observable macroeconomic and financial time series such as disagreement about future inflation amongst professional forecasters, consumer confidence, as well as measures of option-implied Treasury volatility. We compare inflation forecasts from our model to the forecasts from breakeven inflation rates and show that the risk premium and illiquidity adjustment of the breakevens leads to better inflation forecasts than using unadjusted breakevens or a random walk. We also use the model to estimate the value of the deflation floor embedded in TIPS. We find that although the value of this option is generally very small, in the recent financial crisis the option value represented a significant portion of TIPS prices as market-based measures of inflation expectations dropped considerably.

The remainder of the paper is organized as follows. In Section 2 we introduce our joint model for Treasury and TIPS yields. In Section 3 we discuss the econometric procedure. Section 4 summarizes the estimation results and the model fit. In Section 5 we then use the model for a number of applications relevant to policy makers and practitioners. Section 6 reviews the related literature and Section 7 concludes. Detailed derivations are relegated to an appendix.

2 The Model

2.1 State variable dynamics and pricing kernel

We begin with a review of the distributional assumptions made under the Gaussian ATSM framework. An introduction to this class of models in continuous time is presented by Piazzesi (2003). We cover the main results succinctly to familiarize the reader with our notation. Suppose that a $K \times 1$ vector of pricing factors evolves according to the autoregression

$$X_{t+1} = \mu + \Phi X_t + \nu_{t+1} \tag{1}$$

where ν_t are *i.i.d.* Gaussian with $\mathbb{E}_t\left[\nu_{t+1}\right] = 0_{K\times 1}$ and$\mathbb{V}_t\left[\nu_{t+1}\right] = \Sigma$ satisfying $\text{rank}(\Sigma) = K$. Suppose also that assets are priced by the stochastic discount factor

$$M_t = \exp\left(-r_t - \frac{1}{2}\lambda_t'\lambda_t - \lambda_t \Sigma^{-1/2}\nu_{t+1}\right). \tag{2}$$

We follow Duffee (2002) in assuming that the $K \times 1$ price of risk vector λ_t takes the essentially affine form

$$\lambda_t = \Sigma^{-1/2}\left(\lambda_0 + \lambda_1 X_{t+1}\right). \tag{3}$$

From these assumptions it is straightforward to calculate that under the risk neutral probability measure $\mathbb{Q}$, X_t evolves according to the autoregression

$$X_{t+1} = \tilde{\mu} + \tilde{\Phi} X_t + \nu^{\star}_{t+1} \tag{4}$$

where $\tilde{\mu} = \mu - \lambda_0,\ \tilde{\Phi} = \Phi - \lambda_1,\ \nu^{\star}_{t+1} = \nu_{t+1} + \Sigma^{-1/2}\lambda_t$, and that under $\mathbb{Q}$ the innovations $\nu^{\star}_{t+1}$ are *i.i.d.* Gaussian with $\mathbb{E}^{\mathbb{Q}}_t\left[\nu^{\star}_{t+1}\right] = 0_{K\times 1}$ and $\mathbb{V}^{\mathbb{Q}}_t\left[\nu^{\star}_{t+1}\right] = \Sigma$.

2.2 No-Arbitrage Pricing

2.2.1 Treasuries

Assume that log prices of risk free discount bonds (henceforth Treasuries) take the form

$$\log P_t^{(n)} = A_n + B_n' X_t. \tag{5}$$

This assumption carries the implication that the risk free nominal short rate is affine in the state variables as well, denoted by the short rate parameters

$$r_t = \delta_0 + \delta_1 X_t. \tag{6}$$

From the absence of arbitrage it follows that the coefficients in equation (5) must be determined cross-sectionally by the following difference equations:

$$A_n = A_{n-1} + B_{n-1}'\tilde{\mu} + \frac{1}{2} B_{n-1}' \Sigma B_{n-1} - \delta_0 \tag{7}$$

$$B_n' = B_{n-1}'\tilde{\Phi} - \delta_1 \tag{8}$$

$$A_0 = 0, \;\; B_0 = 0. \tag{9}$$

2.2.2 Inflation-Indexed Securities

We expand the ordinary Gaussian ATSM framework to allow the pricing of inflation-indexed securities jointly with nominal securities so that both yield curves are affine in the state variables. A wide variety of similar models have been studied before in both continuous and discrete time, and Section 6 discusses the relationship of our model to several others.

Let Q_t be a price index at time t, and let $P_{t,R}^{(n)}$ denote the price at time t of an inflation-indexed bond with face value 1, paying out the quantity $\frac{Q_{t+n}}{Q_t}$ at time $t+n$. Per construction of $\mathbb{Q}$, the price of such a bond satisfies

$$P_{t,R}^{(n)} = \mathbb{E}_t^{\mathbb{Q}} \left[\exp\left(-r_t - ... - r_{t+n-1}\right) \frac{Q_{t+n}}{Q_t} \right]. \tag{10}$$

Denote one period log inflation by $\pi_t = \ln\left(\frac{Q_t}{Q_{t-1}}\right)$, so that

$$\frac{Q_{t+n}}{Q_t} = \exp\left(\sum_{i=1}^{n} \pi_{t+i}\right). \tag{11}$$

Assume that log prices of inflation-indexed bonds are also affine in the state variables, satisfying

$$\log P_{t,R}^{(n)} = A_{n,R} + B'_{n,R} X_t \tag{12}$$

and finally assume that inflation itself is also a linear function of the pricing factors, denoted by the inflation short rate parameters

$$\pi_t = \pi_0 + \pi_1 X_t. \tag{13}$$

In order to derive the pricing recursions for inflation-indexed bonds, we rewrite equation (10) in terms of an indexed bond purchased one period ahead. This gives

$$P_{t,R}^{(n)} = \mathbb{E}_t^{\mathbb{Q}}\left[\exp\left(-r_t + \pi_{t+1}\right) P_{t+1,R}^{(n-1)}\right]. \tag{14}$$

Under our assumptions, taking logs on both sides of (14), and making use of the fact that $\nu_{t+1}^{\star} \sim N(0, \Sigma)$ under $\mathbb{Q}$, we can calculate the expectation explicitly. After matching coefficients we find that the coefficients in equation (12) are determined by the recursive equations

$$\begin{aligned}
A_{n,R} &= A_{n-1,R} + \left(B'_{n-1,R} + \pi_1\right)\tilde{\mu} + \frac{1}{2}\left(B'_{n-1,R} + \pi_1\right)\Sigma\left(B_{n-1,R} + \pi'_1\right) - \delta_{0,R} && (15)\\
B'_{n,R} &= \left(B'_{n-1,R} + \pi_1\right)\tilde{\Phi} - \delta_1 && (16)\\
A_{0,R} &= 0,\ B_{0,R} = 0_{K\times 1} && (17)
\end{aligned}$$

where we have defined $\delta_{0,R} = \delta_0 - \pi_0$.

2.3 Expected Inflation

The model can be used to compute expected inflation under both the risk-neutral and the physical measure at any horizon. The difference between the two measures of inflation expectations constitutes the inflation risk premium. Define the n-period average of expected log inflation under the physical measure as

$$\pi_t^{(n)} = -\frac{1}{n} \ln \mathbb{E}_t^{\mathbb{P}} \left[\frac{Q_t}{Q_{t+n}} \right]. \tag{18}$$

Assume that expected inflation is affine in the state variables according to

$$\ln \mathbb{E}_t^{\mathbb{P}} \left[\frac{Q_t}{Q_{t+n}} \right] = -A_{n,\pi} - B'_{n,\pi} X_t. \tag{19}$$

Then by expanding the left hand side, it can be shown that the coefficients follow the recursive relations

$$A_{n,\pi} = A_{n-1,\pi} + \left(\pi_1 + B'_{n-1,\pi}\right) \mu - \frac{1}{2} \left(\pi_1 + B'_{n-1,\pi}\right) \Sigma \left(\pi'_1 + B_{n-1,\pi}\right) + \pi_0 \tag{20}$$

$$B'_{n,\pi} = \left(\pi_1 + B'_{n-1,\pi}\right) \Phi \tag{21}$$

$$A_{0,\pi} = 0, \;\; B_{0,\pi} = 0. \tag{22}$$

Similarly, we can obtain the n-period average of expected log inflation under the risk-neutral measure as

$$\ln \mathbb{E}_t^{\mathbb{Q}} \left[\frac{Q_t}{Q_{t+n}} \right] = -A^{\star}_{n,\pi} - B^{\star\prime}_{n,\pi} X_t. \tag{23}$$

where

$$A^{\star}_{n,\pi} = A^{\star}_{n-1,\pi} + \left(\pi_1 + B^{\star\prime}_{n-1,\pi}\right) \tilde{\mu} - \frac{1}{2} \left(\pi_1 + B^{\star\prime}_{n-1,\pi}\right) \Sigma \left(\pi'_1 + B^{\star}_{n-1,\pi}\right) + \pi_0 \tag{24}$$

$$B^{\star\prime}_{n,\pi} = \left(\pi_1 + B^{\star\prime}_{n-1,\pi}\right) \tilde{\Phi} \tag{25}$$

Then, the inflation risk premium is simply

$$\varphi_t^{(n)} = A_{n,\pi} + B'_{n,\pi} X_t - A^{\star}_{n,\pi} - B^{\star\prime}_{n,\pi} X_t \tag{26}$$

2.4 TIPS Liquidity Effects

As documented by Pflueger and Viceira (2011), the liqudity of TIPS appears to be systematically priced. We denote this liquidity pricing as deviation by the term liquidity effects. Examples of the pricing of liquidity in TIPS have manifested themselves towards the beginning of the program prior to 2003, when the Treasury reaffirmed its commitment to the TIPS program (Sack and Elsasser, 2004), and in the aftermath of the Lehman bankruptcy in late 2008, when considerable TIPS inventory was sold into the market (Campbell, Shiller and Viceira, 2009). Liquidity effects are also present in nominal Treasuries, for example as captured by the on-the-run/off-the-run spread. However, measuring both absolute TIPS liquidity effects and absolute Treasury liquidity effects and taking the difference is equivalent to modeling the liquidity of these securities in purely relative terms. Our method of identification thus attributes the entire liquidity premium to TIPS yields.

Let L_t be a $K_l \times 1$ vector of liquidity factors which we assume to be observed. We can then model the liquidity effect $\ell_t^{(n)} = \ell_0 + \ell_1 L_t$ as a deviation from no arbitrage according to

$$P_{t,R}^{(n)} e^{-\ell_t^{(n)}} = \mathbb{E}_t^{\mathbb{Q}} \left[\exp(-r_t) P_{t+1,R}^{(n-1)} e^{-\ell_t^{(n-1)}} \right] \tag{27}$$

We can thus simply expand the state space to include L, so that

$$X_{T+1}^L = \left[X_{t+1}' \ L_{t+1} \right]'. \tag{28}$$

Augmenting the dimensions of $\Phi,\ \mu,\ \lambda,\ \delta$, and all other variables accordingly, we find that

$$\log \left(P_{t,R}^{(n)} e^{-\ell_t^{(n)}} \right) = A_{n,R} + B_{n,R}' X_t^L \tag{29}$$

where the same recursive restrictions on the TIPS yield parameters as in equations (15) and (16) above.[1] As mentioned earlier, our identification of liquidity effects assumes that Treasury yields do not have a liquidity component. We can impose this assumption in our model by treating the liquidity factors as 'unspanned' factors for Treasuries. In particular,

[1] With some abuse of notation we continue to use the parameters $A_{n,R}$ and $B_{n,R}$ despite the fact that we now use a larger state space.

this implies that $\tilde{\Phi}$ takes the block form

$$\begin{bmatrix} \Phi^{\star}_{XX} & 0 \\ \Phi^{\star}_{XL} & \Phi^{\star}_{LL} \end{bmatrix}$$

and that the elements in δ_1 corresponding to the liquidity factors are all zero. As discussed in Adrian, Crump and Moench (2012), this restriction may be imposed by a slight modification of the three-step estimator.

3 Estimation

Estimation of the model is in the spirit of Adrian, Crump and Moench (2012) using a three-step least-squares estimator for the parameters λ_0 and λ_1. This estimator may be viewed as a generalization of the Fama and MacBeth (1973) procedure allowing for time varying prices of risk. Contrary to traditional approaches to estimating Gaussian ATSMs, this approach uses excess holding period returns to estimate the model parameters.

3.1 Zero Coupon Excess Returns

Log excess one period holding returns on Treasuries are defined as

$$rx^{(n-1)}_{t+1} = \log P^{(n-1)}_{t+1} - \log P^{(n)}_{t} + \log P^{(1)}_{t}. \tag{30}$$

Following Adrian, Crump and Moench (2012), we expand this expression in terms of the state variables X_t and innovations ν_{t+1}

$$\begin{aligned} rx^{(n-1)}_{t+1} &= A_{n-1} + B'_{n-1}X_{t+1} - A_n - B'_n X_t - \delta_0 - \delta_1 X_t \\ &= a_{n-1} + c_{n-1}X_t + \beta'_{n-1}\nu_{t+1} \end{aligned} \tag{31}$$

where

$$a_{n-1} = B'_{n-1}\lambda_0 - \frac{1}{2}B'_{n-1}\Sigma B_{n-1} \tag{32}$$

$$c_{n-1} = B'_{n-1}\lambda_1 \tag{33}$$

$$\beta_{n-1} = B_{n-1}. \tag{34}$$

Log excess one period holding returns on inflation indexed securities are defined as

$$rx^{(n-1)}_{t+1,R} = \log \frac{Q_{t+1}}{Q_t} P^{(n-1)}_{t+1,R} - \log P^{(n)}_{t,R} + \log P^{(1)}_t. \tag{35}$$

A decomposition similar to the one above gives the expression

$$rx^{(n-1)}_{t+1,R} = a_{n-1,R} + c_{n-1,R}X_t + \beta'_{n-1,R}\nu_{t+1} \tag{36}$$

where

$$a_{n-1,R} = \left(B'_{n-1,R} + \pi_1\right)\lambda_0 - \frac{1}{2}\left(B'_{n-1,R} + \pi_1\right)\Sigma\left(B_{n-1,R} + \pi'_1\right) \tag{37}$$

$$c_{n-1,R} = \left(B'_{n-1,R} + \pi_1\right)\lambda_1 \tag{38}$$

$$\beta_{n-1,R} = B_{n-1,R} + \pi'_1. \tag{39}$$

3.2 Three Step Regression Approach

We now present our three-step procedure for estimating the model parameters using ordinary least squares exploiting the above representations for excess returns. This procedure parallels the estimation approach introduced by Adrian, Crump and Moench (2012) for Gaussian ATSMs without inflation-indexed bonds.

1. Given a set of pricing factors, the state equation (1) can be estimated using OLS to obtain $\hat{\mu}$, $\hat{\Phi}$, $\hat{\nu}$ and $\hat{\Sigma}$. We obtain the short-rate parameters $\hat{\delta}$ and $\hat{\pi}$ by regressing the nominal short rate and inflation on contemporaneous pricing factors.

2. Second, we estimate the following time series regressions for each of N selected Treasury

excess return maturities and N_R selected inflation-indexed excess return maturities:

$$rx_{t+1}^{(n)} = a_{n-1} + c_{n-1}X_t + \beta'_{n-1}\hat{\nu}_{t+1} + e_{t+1}^{(n-1)} \tag{40}$$

$$rx_{t+1,R}^{(n-1)} = a_{n-1,R} + c_{n-1,R}X_t + \beta'_{n-1,R}\hat{\nu}_{t+1} + e_{t+1,R}^{(n-1)}. \tag{41}$$

The estimated coefficients are stacked into the $N \times 1$ matrix $\hat{\mathbf{a}}$,the $N_R \times 1$ matrix $\hat{\mathbf{a}}_R$, the $N \times K$ matrices $\hat{\mathbf{c}}$ and $\hat{\beta}'$ and the $N_R \times K$ matrices $\hat{\mathbf{c}}_R$ and $\hat{\beta}'_R$. We also calculate the matrices $\hat{B}^\star = \left[\text{vec}\hat{\beta}_1\hat{\beta}'_1 \ldots \text{vec}\hat{\beta}_N\hat{\beta}'_N\right]'$ and the analogously-defined $\hat{B}^\star_R$.

3. Third, we estimate the price of risk parameters λ_0 and λ_1 via cross sectional regression. This is done by stacking the joint systems of excess return coefficients (32)-(33) and (37)-(38). This system of equations takes the form

$$\begin{bmatrix} \beta' \\ \beta'_R \end{bmatrix} [\lambda_0 \ \lambda_1] = \begin{bmatrix} \mathbf{a} + \frac{1}{2}\beta^\star \text{vec}(\Sigma) & \mathbf{c} \\ \mathbf{a}_R + \frac{1}{2}\beta^\star_R \text{vec}(\Sigma) & \mathbf{c}_R \end{bmatrix} \tag{42}$$

and can be more compactly written as

$$\mathbf{B}'\Lambda = [\mathbf{A} \ \mathbf{C}]. \tag{43}$$

Replacing these values with their sample estimates obtained in steps (1) and (2), we propose the estimator

$$\hat{\Lambda} = \left(\hat{\mathbf{B}}\hat{\mathbf{B}}'\right)^{-1} \hat{\mathbf{B}}' \left[\hat{\mathbf{A}} \ \hat{\mathbf{C}}\right]. \tag{44}$$

Adrian, Crump and Moench (2012) show that under standard distributional assumptions the estimator $\hat{\Lambda}$ is consistent and asymptotically normal. With the estimated parameters in hand we are now able to generate the bond price loadings via the recursions obtained in Section 2.2.

3.3 Iterative Estimation

Recall that real excess returns for equation (41) are calculated according to

$$rx_{t+1,R}^{(n-1)} = \log \frac{Q_{t+1}}{Q_t} P_{t+1,R}^{(n-1)} - \log P_{t,R}^{(n)} + \log P_t^{(1)}$$

Due to the CPI release lag, the price index Q_t for TIPS is the two-month-lagged value of the price index. However, when regressing lagged inflation on our contemporaneous pricing factors we find that the fit of that regression is poor.[2] We thus take as our measure of inflation a projection of π_t onto contemporaneous pricing factors according to equation (13). We refer to the fitted coefficients as $\hat{\pi}_1$. The long-term average of inflation is not well identified in our short sample covering less than 13 years of monthly data. We therefore fix π_0 in this regression to (2/12)%, thus implicitly assuming that the long run mean of annualized monthly inflation implied by our model is 2%. Note that while fixing π_0 does not impact the cross sectional fit of the model, it does influence the mean of the inflation risk premium implied by our model.

Excess returns on TIPS depend on π_1, which needs to be estimated. We employ an iterative procedure to do so. The procedure consists of an estimation algorithm which maximizes the model fit for TIPS yields. This is done in the following way. As we show in Appendix B the recursions for TIPS yields are quadratic functions of π_1. Given estimates of the remaining model parameters, we can thus perform a simple numerical optimization to find a new estimate $\widehat{\hat{\pi}}_1$ of the inflation factor loadings that minimizes the squared deviations between actual and model-implied TIPS yields. We update real excess returns using this new value for π and repeat the three-step estimation as described above. In practice, we iterate this procedure until convergence. This is generally achieved in fewer than 25 iterations, the whole process requiring approximately 6-8 seconds of computation on a standard personal computer.

[2] Kim (2009) also studies the relationship of inflation to the term structure in depth, finding that monthly changes to the CPI are unspanned by the term structure, whereas useful information about the trend component is embedded in yields. Our estimation supports this finding - the necessity of optimizing over the inflation parameters π_1 is unsurprising if we do not expect information about monthly CPI inflation to be spanned by the Treasury yield curve or the long-run TIPS yield curve.

4 Empirical Results

In this section, we provide results from the estimation of our joint term structure model. We first discuss the data sources as well as our choice of pricing factors. We next characterize the fit of the model for both the Treasury and the TIPS curve and present results on the estimated prices of risk. Finally, we decompose long-term breakeven inflation into its constituents: inflation expectations, the inflation risk premium, and a liquidity component.

4.1 Data and Factor Construction

We obtain our monthly zero coupon bond yields from the Gurkaynak, Sack and Wright (2007, 2010) datasets (GSW hereafter) which can be obtained from the Federal Reserve Board of Governors research data page.[3] These data are based on fitted Nelson-Siegel-Svensson curves, the parameters of which are published along with the estimated zero coupon curve. We use these parameters to back out the cross-section of nominal and real zero-coupon yields for maturities up to 10 years for TIPS and Treasuries, taking end-of-month observations from 1999:01-2012:08 for a total of $T = 163$ observations. We use the one-month Treasury yield from GSW as the nominal risk free rate. The price index Q_t used to calculate TIPS payouts is seasonally unadjusted CPI-U, which is available from the Bureau of Labor Statistics website.

In the regression (44) we select a cross section of $N = 11$ Treasury excess returns for maturities $n = 6, 12, 24, \ldots, 120$ months and $N_R = 8$ excess returns on TIPS with maturities $n = 36, \ldots 120$ months. As in Adrian, Crump and Moench (2012), we use as pricing factors the principal components extracted from yields. In addition, here we add liquidity factors to account for the deviations from no-arbitrage pricing observed in TIPS. In order to be able to disentangle the differential impact of yield curve factors and liquidity factors on TIPS yields, we partial out the latter before extracting principal components. Specifically, we first orthogonalize TIPS yields to the first three principal components of nominal Treasury yields. We then regress these orthogonalized TIPS yields onto the liquidity factors and finally extract principal components from the joint cross section of nominal yields and the residuals from this regression.

[3] See http://www.federalreserve.gov/econresdata/researchdata.htm

In our baseline specification we extract the first six principal components from the joint cross-section of Treasuries and regression residuals and use as our single liquidity factor the average absolute TIPS yield curve fitting error from the model of Gurkaynak, Sack and Wright (2010) and obtained from the Board of Governors.[4] Notably, this series shows a sharp spike in late 2008. Given this set of factors as well as the excess holding period Treasury returns and TIPS returns, we then estimate the parameters of our model according to the procedure described above.[5]

4.2 Model Fit and Parameter Estimates

We start by discussing the fit of the model for Treasury and TIPS yields and study the estimates of the price of risk parameters. We show that our model fits both yield curves extremely well and gives rise to substantial time variation in the prices of risk. Since traditional term structure models are estimated imposing nonlinear cross-equation restrictions, estimation of these models with a large number of factors is computationally demanding. In contrast, adding factors to our regression based approach comes at no computational cost.

Table 1 reports the time series properties of the yield pricing errors for the Treasury curve implied by the model. We see that the average yield pricing errors are very small, not exceeding 2.8 basis points in absolute value. They also exhibit little variability as the standard deviations of Treasury yield fitting errors are of the order of at most 7 basis points. While these Treasury yield fitting errors are slightly larger than those implied by the five factor model for Treasuries only of Adrian, Crump and Moench (2012), it is important to bear in mind that we are here jointly fitting the cross-section of Treasury and TIPS yields with only six yield factors. In line with the discussion of the relationship between yield and return pricing errors in Adrian, Crump and Moench (2012) we find that the yield pricing errors are strongly serially correlated whereas the return pricing errors show little autocorrelation. Table 2 reports analogous results for yield and return fitting errors of TIPS. Interestingly,

[4] We use monthly maturities from three months to ten years for nominal Treasuries and from three years to ten years for the TIPS regression residuals. We start at a maturity of three years because GSW do not publish TIPS maturities with very short maturities owing to the "carry" distortion caused by the CPI indexation lag.

[5] Note that we set $\mu = 0$ in the estimation since all factors have been demeaned.

our model fits TIPS yields better than Treasuries as the average yield pricing error for TIPS does not exceed 1.1 basis points across the maturity spectrum with similarly low standard deviations. As in the previous table, the yield pricing errors display a notable degree of serial correlation while the return pricing errors at least for shorter maturities are essentially serially uncorrelated.

Figures 2 and 4 provide a number of different visual diagnostics of the time series fit of the model for both curves. In particular, the upper two panels of these figures show the observed and model-implied time series of yields at different maturities for the Treasury and TIPS curve, respectively. Consistent with the results in Tables 1 and 2 discussed above, the two lines are almost visually indistinguishable in all of these charts. The lower two panels in the figures plot the observed and model-implied excess holding period returns. We also superimpose the model-implied expected excess returns to demonstrate that there is substantial time variation in the model-implied compensation for risk.

The upper two panels of Figures 3 and 5 provide plots of unconditional first and second moments of both yield curves as observed and fitted by the model. The charts reinforce that the model fits both moments very well for both yield curves. The lower left panel of Figures 3 and 5 provide plots of the estimated yield loadings B_n for Treasuries and $B_{n,R}$ for TIPS. These allow us to interpret the different factors according to their respective loadings on different sectors of both yield curves. In line with the previous literature, the first principal component clearly represents a level factor for both the Treasury and the TIPS term structure. Similarly, the third principal component represents a slope factor featuring positive loadings on long maturities and negative loadings on short maturities in both curves. The second and fourth principal components capture additional level and slope effects present in the two term structures that are not captured by the first and third components. The lower right panels of both figures show the corresponding excess return loadings $B_n\lambda_1$ for Treasuries and $B_{n,R}\lambda_1$ for TIPS. These show that the fifth and sixth principal components extracted from the joint cross-section are the main drivers of expected excess returns in our model. This is qualitatively consistent with the evidence presented in e.g. Cochrane and Piazzesi (2005) and Adrian, Crump and Moench (2012).

Table 3 provides the estimated market price of risk parameters for our model. These reinforce

the finding that the fifth and sixth principal components are important drivers of risk premia in our model. Interestingly, the liquidity factor does not significantly add to time variation in market prices of risk.

Given the pricing factors and the estimated model parameters, we can decompose breakeven inflation rates at any horizon into its three constituents: expected inflation, the inflation risk premium, as well as a liquidity adjustment. Figure 1 shows the time series of these components for both the 10 year breakeven inflation as well as the 5-10 year forward breakeven inflation rate. These charts highlight the two main conclusions from our model. First, while expected inflation explains some of the variation of average inflation over the next ten years, it is very close to constant at very long forward horizons. This implies that the bulk of the dynamics of long-term forward breakeven inflation rates is driven by inflation risk premia. Second, while liquidity effects have played only a minor role over most of the sample period, they have strongly contributed to the dynamics of breakeven rates in the recent financial crisis period. In particular, our model largely attributes the collapse of long-term forward breakeven inflation rates in the crisis to liquidity effects rather than changes in underlying inflation expectations.

5 Applications

In this section, we illustrate how the model can be used to extract quantities of interest to policy makers as well as fixed income investors. We start by studying the dynamics of our estimated inflation risk premium relative to relevant financial and economic time series. We then document that our breakeven inflation rates adjusted for the inflation risk premium provide a better indicator of the level of future inflation than observed breakevens. Finally, we derive the option value of the deflation floor embedded in TIPS.

5.1 Interpreting the Estimated Inflation Risk Premium

The inflation risk premium implied by our model is a linear combination of the six yield components as well as the liquidity factor. It is therefore difficult to directly interpret its

dynamics. Instead, we correlate our estimated inflation risk premium with a number of observable macroeconomic and financial variables the choices of which are motivated by economic theory. In particular, we consider 1) the cross-sectional standard deviation of individual inflation four quarters ahead from the Blue Chip Financial Forecasts survey; 2) the difference between the 85th and 15th percentile of one quarter ahead inflation forecasts from the same survey; 3) the three-month swaption implied Treasury volatility from Merrill Lynch; 4) the unemployment rate; 5) year-over-year core CPI inflation; and 6) consumer confidence as measured by the Conference Board survey.

Figure 4 plots the estimated two-year inflation risk premium along with each of these series. Perhaps not surprisingly, the inflation risk premium tends to comove with the two survey measures of forecaster disagreement about future inflation. This implies that market participants command higher inflation compensation at times when there is broad disagreement about the inflation outlook. We also find that swaption-implied Treasury volatility comoves with inflation risk premiums, suggesting that the expected volatility of Treasury securities to some extent reflects movements in the required compensation for bearing inflation risk demanded by investors. The estimated inflation risk premium further shows some comovement with the unemployment rate over our sample. This is consistent with the idea that risk premia are countercyclical. We see a somewhat smaller correlation between the inflation risk premium and core inflation which is somewhat surprising given that core inflation is generally perceived as a good measure of the underlying trend in inflation. Finally, there is a clear negative relationship between the inflation risk premium and the level of consumer confidence, which might loosely be interpreted as evidence for the welfare costs of price instability.

Table 4 provides estimation results for a number of regression specifications in which we relate the two-year inflation risk premium to the above variables. It should be emphasized that given the high degree of persistence in a number of these series, the reported standard errors need to be interpreted with caution. We therefore focus our discussion on assessing the degree of (partial) correlations between these variables. In order to do so we standardize the variables such that we can interpret the regression coefficients in relative terms. We first note that the estimated individual correlation coefficients are relatively similar in absolute

magnitude with consumer confidence having the strongest absolute correlation with the inflation risk premium. This result continues to hold true in a joint regression, where consumer confidence is most strongly negatively correlated with the inflation risk premium while the difference between the 85th and 15th percentile of one quarter ahead inflation forecasts from the Blue Chip survey is the most positively correlated.

5.2 Inflation Forecasting

Breakeven inflation rates are the primary market based measure of inflation expectations and are therefore of considerable interest to policy makers and market participants alike. However, breakeven inflation dynamics may also be influenced by changes in the degree of liquidity as well as inflation risk premia. Since our model allows us to separate these effects, it is natural to assess its performance in terms of predicting future inflation. We compute our model-implied average expected inflation over the next six, twelve, 24 and 36 months using the results from Section 2.3. We compare these predictions with those implied by actual observed zero-coupon breakeven inflation rates as well as those implied by a simple random walk. The results of this exercise are provided in Table 5 which shows the root mean squared error (RMSE) implied by the different forecasts for two sample periods: one that covers the entire estimation sample from $1999:01-2012:08$ and one that excludes crisis period from $2007:09-2009:05$. Considering the full sample, our model forecast outperforms both the unadjusted breakevens as well as the random walk forecast at all horizons, with the relative improvement over the unadjusted breakevens declining as the forecast horizon increases.[6] However, even at the three year horizon, the outperformance with respect to unadjusted breakevens is substantial with a reduction of RMSE of approximately 24%. When we exclude the crisis period, our model continues to forecast better than unadjusted breakevens and the random walk at short horizons. However, at forecast horizons beyond one year ahead the random walk now produces somewhat lower RMSEs. Since our model explicitly accounts for risk premia and liquidity effects we expect it to do well in crisis periods when

[6]Note that the actual breakeven inflation rates at the six and twelve month horizon are subject to substantial volatility resulting from issues related to the fitting of zero coupon TIPS yields by the Nelson-Siegel-Svensson methodology. Consequently, unadjusted breakeven forecasts at these horizons should be interpreted with caution.

these components play an important role. Figure 8 displays the inflation forecasts implied by the different models along with future average inflation over different forecast horizons. These charts show that our model forecasts for inflation are generally smoother than observed breakevens and visibly closer to actual inflation especially in the recent crisis period. Moreover, consistent with the above discussion, the random walk forecast actually rises through the crisis period when market-based measures of inflation were sharply declining.

5.3 TIPS Optionality

Although the principal of a TIPS is adjusted for monthly accrued inflation, the final payouts include an embedded option.[7] At maturity, the value of the bond will be the greater of the nominal principal ($100,000) or the principal adjusted for cumulative CPI-U inflation since issuance. Grishchenko, Vanden and Zhang (2011) exploit this feature to extract risk-neutral deflation probability forecasts from Treasury and TIPS prices and find that the time-varying value of this option has predictive content for several economic variables, most notably inflation. Following their approach, we can quantify the value of the embedded deflation floor within our model. Appendix A provides the derivation and gives an explicit formula for the value of the option in terms of the parameters of our model. Figure 10 plots the time series of the option value for the five and ten year TIPS under different assumptions about the remaining time to maturity. Of course, in periods when past inflation has been positive a shorter remaining time to maturity reduces the probability of seeing cumulative deflation over the lifetime of the bond. As the charts in Figure 10 show, except during the financial crisis the value of the embedded option is small, generally amounting to less than 1% of the principal. However, in the recent financial crisis when breakeven inflation expectations declined sharply the value of this option became much greater, approaching a maximum of 8% of principal among our observed maturities not taking into account the accrual of past inflation.

[7] It is important to note that TIPS coupon payments are not subject to this optionality.

6 Related Literature

The body of related literature extracting inflation expectations from the joint pricing of TIPS and Treasury yield curves is growing rapidly. Early papers on the topic include Chen, Liu and Cheng (2005), Grishchenko and Huang (2008), Hördahl and Tristani (2008), and Adrian and Wu (2009). Most of the papers - as we do - use the zero coupon yield curves of Gurkaynak, Sack and Wright (2007) and Gurkaynak, Sack and Wright (2010) to fit Gaussian ATSMs. Prices of zero-coupon bonds have cleaner theoretical behavior than coupon-bearing securities, and fitting a zero curve adjusts for the differences in duration due to the coupon structures of various issuances.

An important difference between our approach and alternative research is in the illiquidity adjustment of TIPS. Our illiquidity adjustment is most closely related to Pflueger and Viceira (2011), who use observable factors in order to adjust TIPS returns. While Pflueger and Viceira (2011) only conduct return forecasting using these illiquidity proxies, we are embedding the modeling of the illiquidity premium directly within the ATSM model. In fact, we allow our illiquidity factor to enter the pricing of risk of TIPS, which is equivalent to a return forecasting regression. However, we focus on only one illiquidity variable (the GSW absolute fitting error of TIPS relative to the absolute GSW fitting errors of the Treasury yield curve), as we find that the other variables of Pflueger and Viceira (2011) do not significantly improve our ability to fit. The absolute yield curve fitting errors have been used by practitioners and policy makers as a proxy of illiquidity for some time, first appearing in an academic paper by Hu, Pan and Wang (2010).

A number of papers have pointed out that TIPS have been less liquid than off the run Treasury securities until about 2003 (see Sack and Elsasser (2004) and Dudley, Roush and Ezer (2009)) which generated an illiquidity premium in breakeven inflation. Pflueger and Viceira (2011) attempt to capture this illiquidity premium by including the relative volume of TIPS versus Treasury securities as conditioning variable, while D'Amico, Kim and Wei (2010) model this type of illiquidity as a latent factor in an ATSM. The reported estimates by D'Amico, Kim and Wei (2010) end prior to the financial turbulence of 2008 so that it is not clear to what extent the illiquidity of the TIPS market in the Fall of 2008 would be

picked up by the model. In contrast, our illiquidity factor explicitly captures the market dislocations of 2008.

Christensen, Lopez and Rudebusch (2010) report estimates from an ATSM model with three nominal factors (level, slope and curvature) and one real factor (level). Their model is parsimonious, as the prices of risk are restricted so as to be consistent with a Nelson and Siegel (1987) yield curve (see Christensen, Diebold and Rudebusch (2011) for the relation between the ATSM models and the Nelson-Siegel curve). The parsimony of the approach of Christensen, Lopez and Rudebusch (2010) can be viewed as an alternative to our regression based approach to overcome the computational challenges of estimating ATSMs. However, the price of risk restrictions that are imposed in the setup of Christensen, Lopez and Rudebusch (2010) are likely rejected empirically. We find that our specification using six yield factors and a liquidity factor generates pricing errors that are considerably smaller than the ones reported by Christensen, Lopez and Rudebusch (2010). Our usage of a greater number of pricing factors is also justified from the literature on nominal yield curves, which points to the fact that the level, slope and curvature factors are not sufficient to explain the time series and cross section of nominal yields (see Cochrane and Piazzesi (2008) and Adrian, Crump and Moench (2012)). Furthermore, as Christensen, Lopez and Rudebusch (2010) do not adjust for TIPS illiquidity, their measure of expected inflation is hard to interpret during the financial crisis.

Haubrich, Pennacchi and Ritchken (2012) present a model that uses inflation swaps, actual inflation, and survey inflation in addition to the TIPS and Treasury yield curves. Similar to an earlier paper by Adrian and Wu (2009), Haubrich, Pennacchi and Ritchken (2012) allow for heteroskedasticity explicitly by estimating a GARCH model for the yield processes. Prices of risk are restricted to be functions of these estimated second moments. While the model of Haubrich, Pennacchi and Ritchken (2012) combines the different data sources elegantly, the resulting inflation risk premium differs sharply from our estimates. In fact, the inflation risk premium is close to constant over time, implying that far in the future breakeven forward rates reflect changes in inflation expectations. In contrast, in our model, the inflation risk premium varies substantially over time, while far in the future expected inflation is constant. We view our finding as a desirable feature, and it is indeed fully consistent with the intuitions

of Sack and Elsasser (2004) and Dudley, Roush and Ezer (2009) suggesting that variations in 5-10 year forward breakevens mainly reflects changes in inflation risk premia.

Chernov and Mueller (2012) present an ATSM model where a "hidden factor" is extracted from inflation surveys. They show that this hidden inflation survey factor is a significant price of risk factor. While we do not incorporate any survey inflation expectations, our framework would allow the introduction of Chernov and Mueller's hidden factor as an unspanned factor in a straightforward manner. Such unspanned factors would affect the pricing of risk, but not the cross sectional fit of the yield curve (i.e. it would change the $\mathbb{P}-$dynamics, but not the $\mathbb{Q}-$dynamics). We leave it to future research to include unspanned risk factors. D'Amico, Kim and Wei (2010), Haubrich, Pennacchi and Ritchken (2012), and Grishchenko and Huang (2012) also incorporate survey inflation expectations in their estimates of the inflation risk premium. However, those papers consider the inflation forecasts as true probability assessments, while Chernov and Mueller (2012) consider the forecasts of inflation to be subjective and possibly different from the ATSM implied inflation estimate.

A commonality among the alternative approaches by D'Amico, Kim and Wei (2010), Christensen, Lopez and Rudebusch (2010), Haubrich, Pennacchi and Ritchken (2012), and Chernov and Mueller (2012) is that they all use maximum likelihood estimation for the parameter estimates and a Kalman filter for the factor extraction. As discussed earlier, such estimates are computationally costly, and convergence to a global maximum is generally difficult to verify. In contrast, our approach relies only on linear regressions, which is numerically fast, computationally robust, and allows the straightforward extension of the model to include additional factors.

7 Conclusion

We present a joint Gaussian affine term structure model for the cross section of TIPS and Treasury securities that has a number of desirable features relative to the existing literature. We are adjusting for the relative illiquidity of TIPS during times of crisis by using the absolute pricing errors of TIPS relative to Treasuries from a Nelson-Siegel-Svenson curve in a model consistent way. Our estimation approach is regression based, allowing straightforward reestimation, and making the model easily adaptable to questions that require additional conditioning variables. Our methodology also allows us to fix the unconditional average of inflation, which is necessary given the short time series history for TIPS. Relative to other models in the literature, our pricing errors are extremely small, allowing the decomposition of breakeven inflation into an inflation risk premium, expected inflation, and an illiquidity premium with almost no error. Importantly, we find that the majority of variation in the 5-10 year forward breakeven inflation is due to variation in risk and liquidity premia, while the 10 year inflation rate is also considerably influenced by changes in inflation expectations.

A Price of Embedded Deflation Floors

We denote a the price of a deflation floor by

$$\begin{aligned} F_\pi(X_t, n) &= \mathbb{E}_t^{\mathbb{Q}}\left[\exp\left(-\sum_{i=1}^{n} r_{t+i-1}\right)\left(1-\frac{Q_{t+n}}{Q_t}\right)^+\right] \\ &= \mathbb{E}_t^{\mathbb{Q}}\left[\exp\left(-\sum_{i=1}^{n} r_{t+i-1}\right)\left(1-\frac{Q_{t+n}}{Q_t}\right)\chi_{\left\{\frac{Q_{t+n}}{Q_t}<1\right\}}\right] \end{aligned} \tag{45}$$

The joint distribution of the quantities $\frac{Q_{t+n}}{Q_t}$ and $-\sum_{i=1}^{n} r_{t+i-1}$ will suffice to calculate the price. We begin with the first term.

$$\begin{aligned} \frac{Q_{t+n}}{Q_t} &= \exp\left(\sum_{i=1}^{n}\pi_{t+i}\right) \\ &= \exp\left(n\cdot\pi_0+\pi_1\sum_{i=1}^{n}X_{t+i}\right) \\ &= \exp\left(n\cdot\pi_0+\pi_1\left[\sum_{i=1}^{n} i\cdot\tilde{\Phi}^{n-i}\tilde{\mu}+\sum_{i=1}^{n}\tilde{\Phi}^i X_t+\sum_{i=1}^{n}\sum_{j=1}^{i}\tilde{\Phi}^{j-1}\tilde{\nu}_{t+n-i+1}\right]\right). \end{aligned} \tag{46}$$

Under $\mathbb{Q}$ the quantity $\pi_1\sum_{i=1}^{n}\sum_{j=1}^{i}\tilde{\Phi}^{j-1}\tilde{\nu}_{t+n-i+1}$ is a normal random variable which we denote by V_n^π, and we write

$$\frac{Q_{t+n}}{Q_t} = \exp\left(n\cdot\pi_0+\pi_1\left[\sum_{i=1}^{n} i\cdot\tilde{\Phi}^{n-i}\tilde{\mu}+\tilde{\Phi}^i X_t\right]\right)\exp\left(V_n^\pi\right).$$

Similarly,

$$\begin{aligned} \exp\left(-\sum_{i=1}^{n} r_{t+i-1}\right) &= \exp\left(-n\cdot\delta_0-\delta_1\sum_{i=0}^{n-1}X_{t+i}\right) \\ &= \exp\left(-n\cdot\delta_0-\delta_1\left[\sum_{i=1}^{n-1} i\cdot\tilde{\Phi}^{n-i-1}\tilde{\mu}+\sum_{i=0}^{n-1}\tilde{\Phi}^i X_t+\sum_{i=2}^{n}\sum_{j=1}^{i-1}\tilde{\Phi}^{j-1}\tilde{\nu}_{t+n-i+1}\right]\right) \end{aligned} \tag{47}$$

and as before we denote the normal random variable $-\delta_1\sum_{i=2}^{n}\sum_{j=1}^{i-1}\tilde{\Phi}^{j-1}\tilde{\nu}_{t+n-i+1}$ by V_n^r. We characterize the joint distribution of V_n^r and V_n^π by their means, both zero, and their covariance matrix, denoted by

$$\begin{bmatrix} \xi_n^{r,r} & \xi_n^{r,\pi} \\ \xi_n^{\pi,r} & \xi_n^{\pi,\pi} \end{bmatrix}$$

which can be expressed in closed form. We can simplify the expectation (45)by noting that

$$\begin{aligned} \frac{Q_{t+n}}{Q_t} &< 1 \\ \exp\left(V_n^\pi\right) &< \exp\left(-n\cdot\pi_0-\pi_1\sum_{i=1}^{n}\left[i\cdot\tilde{\Phi}^{n-i}\tilde{\mu}+\tilde{\Phi}^i X_t\right]\right) \end{aligned}$$

Defining $d_{t,n} = -n \cdot \pi_0 - \pi_1 \sum_{i=1}^{n} \left[i \cdot \tilde{\Phi}^{n-i} \tilde{\mu} + \tilde{\Phi}^i X_t \right]$ we can write

$$\begin{aligned} F_\pi(X_t, n) &= \mathbb{E}_t^{\mathbb{Q}} \left[\exp\left(-\sum_{i=1}^{n} r_{t+i-1} \right) \chi_{\left\{ \frac{Q_{t+n}}{Q_t} < 1 \right\}} \right] - \mathbb{E}_t^{\mathbb{Q}} \left[\exp\left(\sum_{i=1}^{n} \pi_{t+i} - \sum_{i=1}^{n} r_{t+i-1} \right) \chi_{\left\{ \frac{Q_{t+n}}{Q_t} < 1 \right\}} \right] \\ &= \exp\left(S_{t,n}^1 \right) \cdot \mathbb{E}_t^{\mathbb{Q}} \left[\exp\left(V_n^r \right) \chi_{\{V_n^\pi < d_{t,n}\}} \right] - \exp\left(S_{t,n}^2 \right) \cdot \mathbb{E}_t^{\mathbb{Q}} \left[\exp\left(V_n^r + V_n^\pi \right) \chi_{\{V_n^\pi < d_{t,n}\}} \right]. \end{aligned} \tag{48}$$

where $S_{t,n}^1$ and $S_{t,n}^2$ combine the remaining terms from (46) and (47). The above expectations are of a standard form and can be computed explicitly. Denoting by $N(\cdot)$ the standard normal cumulative distribution function, by standard properties of the normal and lognormal distribution we have:

$$\mathbb{E}_t^{\mathbb{Q}} \left[\exp\left(Z_1 \right) \chi_{\{Z_2 < d\}} \right] = \exp\left(\mathbb{E}\left[Z_1 \right] + \frac{1}{2} \mathbb{V}\left[Z_1 \right] \right) N\left(\frac{d - \mathbb{E}\left[Z_2 \right] - \mathbb{C}\left[Z_1, Z_2 \right]}{\sqrt{\mathbb{V}\left[Z_2 \right]}} \right). \tag{49}$$

Applying equation (49) to the expectations in equation (48), we obtain

$$\mathbb{E}_t^{\mathbb{Q}} \left[\exp\left(V_n^r \right) \chi_{\{V_n^\pi < d_{t,n}\}} \right] = \exp\left(\frac{1}{2} \xi_n^{r,r} \right) N\left(\frac{d_{t,n} - \xi_n^{r,\pi}}{\sqrt{\xi^{\pi,\pi}}} \right)$$

and

$$\mathbb{E}_t^{\mathbb{Q}} \left[\exp\left(V_n^r + V_n^\pi \right) \chi_{\{V_n^\pi < d_{t,n}\}} \right] = \exp\left(\frac{1}{2} \left(\xi_n^{r,r} + 2\xi_n^{r,\pi} + \xi_n^{\pi,\pi} \right) \right) N\left(\frac{d_{t,n} - \xi_n^{r,\pi} - \xi_n^{\pi,\pi}}{\sqrt{\xi^{\pi,\pi}}} \right).$$

B Expansion for π Optimization

Given the structure of the model, one can estimate π_0 and π_1 by using information contained in the entire cross-section of TIPS yields rather than the real short rate alone.

B.1 Recursion Expansion

We flatten the no-arbitrage recursions as follows.

$$\begin{aligned} A_{n,R} &= A_{n-1,R} + B'_{n-1,R} \left(\tilde{\mu} + \Sigma \pi_1 \right) + \frac{1}{2} B'_{n-1,R} \Sigma B_{n-1,R} + \left(\frac{1}{2} \pi_1' \Sigma \pi_1 - \delta_{0,R} + \pi_1' \tilde{\mu} \right) \\ B'_{n,R} &= \left(B'_{n-1,R} + \pi_1' \right) \tilde{\Phi} - \delta_1' \\ A_{0,R} &= 0, \ B_{0,R} = 0_{k \times 1} \end{aligned}$$

Our updated formula for A is simply an expansion of the quadratic term.

$$\begin{aligned} A_{n,R} &= A_{n-1,R} + \left(B'_{n-1,R} + \pi_1' \right) \tilde{\mu} + \frac{1}{2} \left(B'_{n-1,R} + \pi_1' \right) \Sigma \left(B_{n-1,R} + \pi_1 \right) - \delta_{0,R} \\ &= A_{n-1,R} + B'_{n-1,R} \left(\tilde{\mu} + \Sigma \pi_1 \right) + \frac{1}{2} B'_{n-1,R} \Sigma B_{n-1,R} + \left(\frac{1}{2} \pi_1' \Sigma \pi_1 - \delta_{0,R} + \pi_1' \tilde{\mu} \right) \end{aligned}$$

We can express these as functions of the parameters $\tilde{\mu}, \tilde{\Phi}, \delta_0$, and δ_1.

$$\begin{aligned} B'_{1,R} &= \pi_1' \tilde{\Phi} - \delta_1' \\ B'_{2,R} &= \pi_1' \tilde{\Phi}^2 + \pi_1' \tilde{\Phi} - \delta_1' \tilde{\Phi} - \delta_1' = \left(\pi_1' \tilde{\Phi} - \delta_1' \right) \left(\tilde{\Phi} + I \right) \\ &\vdots \\ B'_{n,R} &= \left(\pi_1' \tilde{\Phi} - \delta_1' \right) \zeta_n' \end{aligned}$$

Where we have defined

$$\left(\sum_{i=0}^{n-1}\tilde{\Phi}^i\right)' = \zeta_n.$$

We will make use of the following series:

$$\begin{aligned}
G_n &= \sum_{i=1}^{n-1} B_{i,R}'\Sigma B_{i,R} \\
G_1 &= 0 \\
G_n &= \sum_{i=1}^{n-1} B_{i,R}'\Sigma B_{i,R} \\
&= \sum_{i=1}^{n-1}\left(\pi_1'\tilde{\Phi}-\delta_1'\right)\zeta_n'\Sigma\zeta_n\left(\tilde{\Phi}'\pi_1-\delta_1\right) \\
&\vdots \\
&= \pi_1'\tilde{\Phi}\Xi_n^1\tilde{\Phi}'\pi_1 - 2\delta_1'\Xi_n^1\tilde{\Phi}'\pi_1 + \delta_1'\Xi_n^1\delta_1
\end{aligned}$$

Define the quantity

$$\begin{aligned}
\Xi_n^2 &= \sum_{i=0}^{n-1}\zeta_i \\
&= \left(\sum_{i=0}^{n-2}(n-1-i)\,\tilde{\Phi}^i\right)
\end{aligned}$$

Now:

$$\begin{aligned}
A_{1,R} &= \pi_1'\tilde{\mu} - \delta_{0,R} + \frac{1}{2}\pi_1'\Sigma\pi_1 \\
A_{2,R} &= 2\left(\pi_1'\tilde{\mu} - \delta_{0,R} + \frac{1}{2}\pi_1'\Sigma\pi_1\right) + \left(\pi_1'\tilde{\Phi}-\delta_1\right)(\tilde{\mu}+\Sigma\pi_1) + \frac{1}{2}G_2 \\
&\vdots \\
A_{n,R} &= n\left(\pi_1'\tilde{\mu} - \delta_{0,R} + \frac{1}{2}\pi_1'\Sigma\pi_1\right) + \left(\pi_1'\tilde{\Phi}-\delta_1'\right)\left(\Xi_n^2\tilde{\mu}+\Xi_n^2\Sigma\pi_1\right) + \frac{1}{2}G_n
\end{aligned}$$

Hence,

$$\begin{aligned}
p_{t,R}^{(n)} &= A_{n,R} + B_{n,R}'X_t \\
&= n\left(\pi_1'\tilde{\mu} - \delta_{0,R} + \frac{1}{2}\pi_1'\Sigma\pi_1\right) + \left(\pi_1'\tilde{\Phi}-\delta_1'\right)\left(\Xi_n^2\tilde{\mu}+\Xi_n^2\Sigma\pi_1\right) + \frac{1}{2}G_n + \left(\pi_1'\tilde{\Phi}-\delta_1'\right)(\zeta_n)'X_t
\end{aligned}$$

or

$$
\begin{aligned}
y_{t,R}^{(n)} &= -\frac{1}{n}p_t^{(n)} = -\frac{1}{n}\left(A_n + B_n' X_t\right) \\
&= -\frac{1}{n}\left(n\left(\pi_1'\tilde{\mu} - \delta_{0,R} + \frac{1}{2}\pi_1'\Sigma\pi_1\right) + \left(\pi_1'\tilde{\Phi} - \delta_1'\right)\left(\Xi_n^2\tilde{\mu} + \Xi_n^2\Sigma\pi_1\right) + \frac{1}{2}G_n + \left(\pi_1'\tilde{\Phi} - \delta_1'\right)\left(\zeta_n\right)' X_t\right) \\
&= -\pi_0 + \left(-\tilde{\mu}'\left(I + \frac{1}{n}\Xi_n^{2\prime}\tilde{\Phi}'\right) + \frac{1}{n}\delta_1'\left(\Xi_n^2\Sigma + \Xi_n^1\tilde{\Phi}'\right) - \frac{1}{n}X_t'\zeta_n\tilde{\Phi}'\right)\pi_1 \\
&\quad -\pi_1'\left(\left(\frac{1}{2}I + \frac{1}{n}\tilde{\Phi}\Xi_n^2\right)\Sigma + \frac{1}{2n}\tilde{\Phi}\Xi_n^1\tilde{\Phi}'\right)\pi_1 + \frac{1}{n}\left(\tilde{\mu}'\Xi_n^{2\prime} + X_t'\zeta_n\right)\delta_1 - \frac{1}{2n}\delta_1'\Xi_n^1\delta_1 + \delta_0.
\end{aligned}
$$

This formula for $y_{t,R}^{(n)}$ shows that real yields are quadratic in the inflation loadings π_1. It is straightforward to setup a quadratic least squares optimization problem, minimizing the difference between observed yields and the fitted yields as a function of π_1.

References

Adrian, Tobias and Hao Wu. 2009. The Term Structure of Inflation Expectations. Staff Reports 362 Federal Reserve Bank of New York.

Adrian, Tobias, Richard Crump and Emanuel Moench. 2012. "Pricing the Term Structure with Linear Regressions." *Journal of Financial Economics* Forthcoming.

Campbell, John Y., Robert J. Shiller and Luis M. Viceira. 2009. "Understanding Inflation-Indexed Bond Markets." *Brookings Papers on Economic Activity* 40(1):79–138.

Chen, Ren-Raw, Bo Liu and Xiaolin Cheng. 2005. Inflation, Fisher Equation, and the Term Structure of Inflation Risk Premia: Theory and Evidence from TIPS. Working paper.

Chernov, Mikhail and Philippe Mueller. 2012. "The Term Structure of Inflation Expectations." *Journal of Financial Economics* 106(2):367–394.

Christensen, Jens H. E., Jose A. Lopez and Glenn D. Rudebusch. 2010. "Inflation Expectations and Risk Premiums in an Arbitrage-Free Model of Nominal and Real Bond Yields." *Journal of Money, Credit and Banking* 42(s1):143–178.

Christensen, Jens H.E., Francis X. Diebold and Glenn D. Rudebusch. 2011. "The Affine Arbitrage-Free Class of Nelson-Siegel Term Structure Models." *Journal of Econometrics* 164(1):4–20.

Cochrane, John and Monika Piazzesi. 2005. "Bond Risk Premia." *American Economic Review* 95:138–160.

Cochrane, John and Monika Piazzesi. 2008. "Decomposing the Yield Curve." *working paper, University of Chicago* .

D'Amico, Stefania, Don H Kim and Min Wei. 2010. Tips from TIPS: the Informational Content of Treasury Inflation-Protected Security Prices. Finance and Economics Discussion Series 19 Board of Governors of the Federal Reserve.

Dudley, William C., Jennifer Roush and Michelle S. Ezer. 2009. "The Case for Tips: an Examination of the Costs and Benefits." *Economic Policy Review* 15(1):1–17.

Duffee, Greg. 2002. "Term Premia and Interest Rate Forecasts in Affine Models." *Journal of Finance* 57(1):405–443.

Fama, Eugene F and James D MacBeth. 1973. "Risk, Return, and Equilibrium: Empirical Tests." *Journal of Political Economy* 81(3):607–36.

Grishchenko, Olesya V. and Jing-Zhi Huang. 2008. Inflation Risk Premium: Evidence from the Tips Market. Working paper Penn State University.

Grishchenko, Olesya V. and Jing-Zhi Huang. 2012. Inflation Risk Premium: Evidence from the TIPS market. Feds working paper Federal Reserve Board.

Grishchenko, Olesya V., Joel M. Vanden and Jianing Zhang. 2011. The Information Content of the Embedded Deflation Option in TIPS. Finance and Economics Discussion Series 2011-58 Board of Governors of the Federal Reserve System (U.S.).

Gurkaynak, Refet S., Brian Sack and Jonathan H. Wright. 2007. "The U.S. Treasury Yield Curve: 1961 to the Present." *Journal of Monetary Economics* 54(8):2291–2304.

Gurkaynak, Refet S., Brian Sack and Jonathan H. Wright. 2010. "The TIPS Yield Curve and Inflation Compensation." *American Economic Journal: Macroeconomics* 2(1):70–92.

Haubrich, Joseph, George Pennacchi and Peter Ritchken. 2012. "Inflation Expectations, Real Rates, and Risk Premia: Evidence from Inflation Swaps." *Review of Financial Studies* 25(5):1588–1629.

Hördahl, Peter and Oreste Tristani. 2008. Inflation Risk Premia in the U.S. and the Euro Area. Working paper Bank for International Settlements.

Hu, Xing, Jun Pan and Jiang Wang. 2010. Noise as Information for Illiquidity. Working Paper 16468 National Bureau of Economic Research.

Kim, Don H. 2009. "Challenges in Macro-Finance Modeling." *Federal Reserve Bank of St. Louis Review* (Sep):519–544.

Nelson, Charles R and Andrew F Siegel. 1987. "Parsimonious Modeling of Yield Curves." *Journal of Business* 60(4):473–89.

Pflueger, Carolin E. and Luis M. Viceira. 2011. An Empirical Decomposition of Risk and Liquidity in Nominal and Inflation-Indexed Government Bonds. NBER Working Papers 16892 National Bureau of Economic Research, Inc.

Piazzesi, Monika. 2003. Affine Term Structure Models. In *Handbook of Financial Econometrics*. Vol. 1 Elsevier chapter 12, pp. 691–758.

Sack, Brian and Robert Elsasser. 2004. "Treasury Inflation-Indexed Debt: a Review of the U.S. Experience." *Economic Policy Review* 10(1):47–63.

Singleton, Kenneth J. 2006. *Empirical Dynamic Asset Pricing*. Princeton University Press.

C Tables and Figures

Table 1: **Treasuries: Fit Diagnostics**

This table summarizes the time series properties of the pricing errors implied by our benchmark model. The sample period is 1999:01-2012:08. "Mean", "std", "skew", and "kurt" refer to the sample mean, standard deviation, skewness, and kurtosis of the errors; $\rho(1), \rho(6)$ denote their autocorrelation coefficients of order one and six. Panel 1 reports properties of the yield pricing errors and Panel 2 reports properties of the excess return pricing errors.

	n = 12	n = 24	n = 36	n = 60	n = 84	n = 120
Panel 1: Yield Pricing Errors						
mean	-0.019	-0.021	-0.028	-0.023	-0.004	-0.002
std	0.069	0.047	0.043	0.034	0.027	0.034
skew	-0.982	0.407	0.207	-0.192	-0.158	0.126
kurt	4.458	2.925	3.479	3.169	2.479	2.815
$\rho(1)$	0.890	0.813	0.868	0.848	0.897	0.801
$\rho(6)$	0.654	0.477	0.577	0.682	0.778	0.503
Panel 2: Return Pricing Errors						
mean	-0.000	-0.000	-0.000	-0.000	-0.000	-0.000
std	0.000	0.001	0.001	0.001	0.001	0.002
skew	0.270	-0.255	0.074	0.022	0.179	-0.186
kurt	6.054	4.051	4.095	3.479	3.361	3.533
$\rho(1)$	-0.131	-0.169	-0.107	-0.202	-0.271	-0.224
$\rho(6)$	0.143	0.147	0.029	0.147	0.067	0.049

Table 2: **TIPS: Fit Diagnostics**

This table summarizes the time series properties of the pricing errors implied by our benchmark model. The sample period is 1999:01-2012:08. "Mean", "std", "skew", and "kurt" refer to the sample mean, standard deviation, skewness, and kurtosis of the errors; $\rho(1), \rho(6)$ denote their autocorrelation coefficients of order one and six. Panel 1 reports properties of the yield pricing errors and Panel 2 reports properties of the excess return pricing errors.

	n = 36	n = 60	n = 84	n = 120
		Panel 1: Yield Pricing Errors		
mean	-0.004	0.011	0.004	-0.006
std	0.041	0.013	0.013	0.024
skew	1.970	1.190	-1.488	-1.835
kurt	15.168	5.036	7.225	8.446
$\rho(1)$	0.718	0.770	0.803	0.857
$\rho(6)$	0.279	0.470	0.351	0.358
		Panel 2: Return Pricing Errors		
mean	0.001	0.001	0.001	0.001
std	0.007	0.004	0.002	0.002
skew	-0.317	-0.987	-0.548	-0.725
kurt	17.295	10.659	3.173	3.708
$\rho(1)$	0.143	-0.004	0.765	0.815
$\rho(6)$	0.130	0.136	0.585	0.623

Table 3: **Market Prices of Risk**

This table summarizes the estimates of the market price of risk parameters λ_0 and λ_1 for the benchmark specification. t-statistics are reported in parentheses. The standard errors have been approximated using the DAPM formulas found in Adrian, Crump and Moench (2012). Wald statistics for tests of the rows of Λ and of λ_1 being different from zero are reported along each row, with the corresponding p-values in parentheses below.

	λ_0	$\lambda_{1.1}$	$\lambda_{1.2}$	$\lambda_{1.3}$	$\lambda_{1.4}$	$\lambda_{1.5}$	$\lambda_{1.6}$	$\lambda_{1.7}$	W_Λ	W_{λ_1}
X_1	0.189	0.001	0.039	0.124	-0.252	-0.503	0.484	-0.003	10.797	9.910
t-stat	(0.691)	(0.078)	(1.635)	(1.399)	(-1.663)	(-1.115)	(0.913)	(-0.071)	(0.213)	(0.194)
X_2	0.197	0.030	**-0.163**	0.024	**0.796**	-0.262	-0.097	0.074	**22.849**	**21.159**
t-stat	(0.357)	(0.826)	(-3.432)	(0.135)	(2.611)	(-0.288)	(-0.091)	(0.746)	(0.004)	(0.004)
X_3	-0.039	**-0.010**	**0.017**	-0.005	-0.016	0.055	0.036	-0.010	**16.097**	13.709
t-stat	(-0.522)	(-2.071)	(2.654)	(-0.206)	(-0.387)	(0.448)	(0.252)	(-0.752)	(0.041)	(0.057)
X_4	-0.132	-0.002	0.019	**-0.105**	**-0.208**	0.248	-0.407	0.013	**28.903**	**27.777**
t-stat	(-1.154)	(-0.311)	(1.892)	(-2.825)	(-3.296)	(1.318)	(-1.840)	(0.612)	(0.000)	(0.000)
X_5	-0.010	0.000	-0.004	-0.006	-0.023	**-0.260**	0.004	-0.001	**21.680**	**21.411**
t-stat	(-0.279)	(0.165)	(-1.201)	(-0.527)	(-1.130)	(-4.259)	(0.056)	(-0.155)	(0.006)	(0.003)
X_6	-0.001	0.002	0.002	-0.014	-0.003	-0.074	**-0.160**	-0.002	**20.333**	**20.026**
t-stat	(-0.029)	(1.008)	(0.918)	(-1.571)	(-0.219)	(-1.672)	(-3.067)	(-0.462)	(0.009)	(0.006)

Table 4: **Inflation Risk Premium Regressions**

This table displays results from linear regressions of the 2-year inflation risk premium implied by our model on various observable macroeconomic and financial indicators. These are 1) the cross-sectional standard deviation of individual inflation four quarters ahead from the Blue Chip Financial Forecasts survey; 2) the difference between the 85th and 15th percentile of one quarter ahead inflation forecasts from the same survey; 3) the three-month swaption implied Treasury volatility from Merrill Lynch; 4) the unemployment rate; 5) year-over-year core CPI inflation; and 6) consumer confidence as measured by the Conference Board survey. The sample period is January 1999-August 2012.

	2yIRP	2yIRP	2yIRP	2yIRP	2yIRP	2yIRP	2yIRP	2yIRP
4-qtr ahead BC st. dev.	0.48						0.07	-0.05
	(0.10)						(0.10)	(0.08)
1-qtr ahead BC disagreement		0.56					0.39	0.27
		(0.10)					(0.11)	(0.10)
3-month swaption MOVE			0.49				0.26	0.17
			(0.08)				(0.08)	(0.08)
Unemplyoment rate				0.42				-0.28
				(0.07)				(0.12)
Core CPI annual inflation					-0.33			-0.20
					(0.06)			(0.07)
Consumer confidence						-0.62		-0.56
						(0.07)		(0.12)
Constant	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	(0.07)	(0.07)	(0.07)	(0.07)	(0.07)	(0.06)	(0.06)	(0.06)
R^2	0.23	0.32	0.24	0.18	0.11	0.39	0.38	0.51
N	163	163	163	163	163	163	163	163

Note: Robust heteroskedasticity adjusted standard errors in parenthesis. Variables are standardized so that regression coefficients are partial correlations.

Table 5: **Inflation Forecasting**

This table compares the root mean squared error of three methods for predicting future inflation. The first method uses the model-implied inflation expectations derived in Section 2.3. The second method takes TIPS breakevens as a predictor of continously-compounded inflation. The third method, a 'random walk' forecast, takes average realized inflation over the prior n months as a prediction of average inflation over the next n months. Forecasts are performed over horizons from 6 to 36 months, and forecasting errors are computed using overlapping observations. The first panel reports forecasting RMSE for the full sample from January 1999-August 2012 whereas the second panel reports the forecasting RMSE excluding the crisis period, taken to be September 2007-May 2009.

	6m	12m	24m	36m
Panel A: Full Sample				
Model Forecast	2.641	1.605	1.055	0.884
Breakevens	3.435	1.798	1.240	1.155
Random Walk Forecast	4.268	2.106	1.336	1.130
Panel B: Excluding Crisis Period				
Model Forecast	2.097	1.418	1.101	1.018
Breakevens	2.428	1.458	1.086	1.061
Random Walk Forecast	3.169	1.473	0.981	0.839

Figure 2: **Treasuries: Observed and Model-Implied Time Series**

This figure provides time series plots of observed and model-implied Treasury yields and excess returns. The upper panels plot zero coupon Treasury yields at 2-year and 10-year maturities and the bottom panels plot excess holding period returns at 2-year and 10-year maturities. The observed yields and returns are plotted by solid lines, whereas dashed green lines correspond to model-implied yields and returns. Dashed red lines in the lower panels track model-implied expected excess holding period returns.

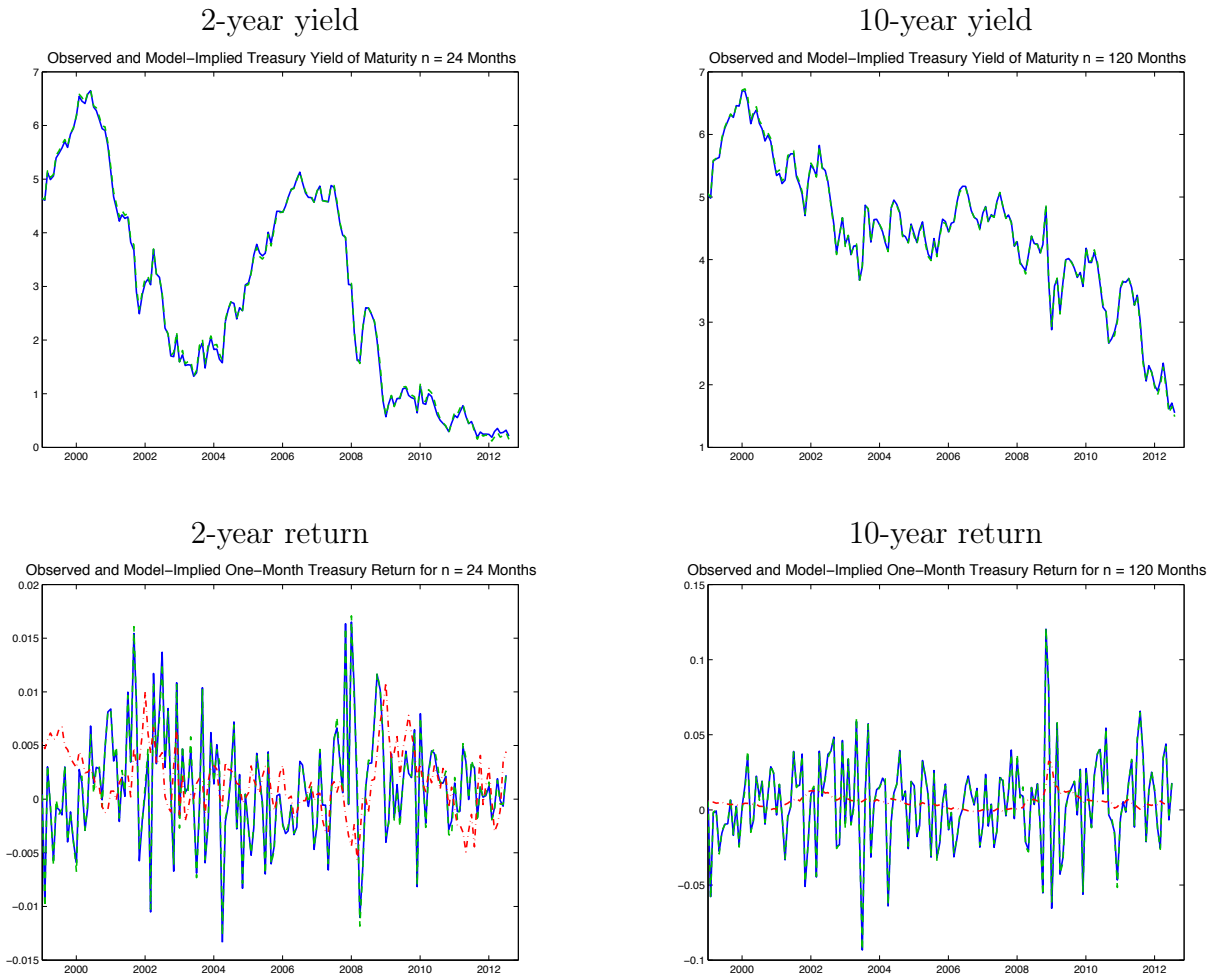

Figure 3: **Treasuries: Cross Sectional Diagnostics**

This figure provides graphs exhibiting the cross-sectional fit and interpretation of the factors as drivers of Treasury yields. The upper two panels plot the unconditional means and standard deviations of observed yields against those implied by the model. The lower left panel plots the implied yield loadings $-\frac{1}{n}B_n$. These coefficients can be interpreted as the response of the n-month yield to a contemporary shock to the respective factor. The lower right panel plots the expected return loadings $B_n'\lambda_1$. These coefficients can be interpreted as the response of the expected one-month excess holding return on an n-month bond to a contemporary shock to the respective factor.

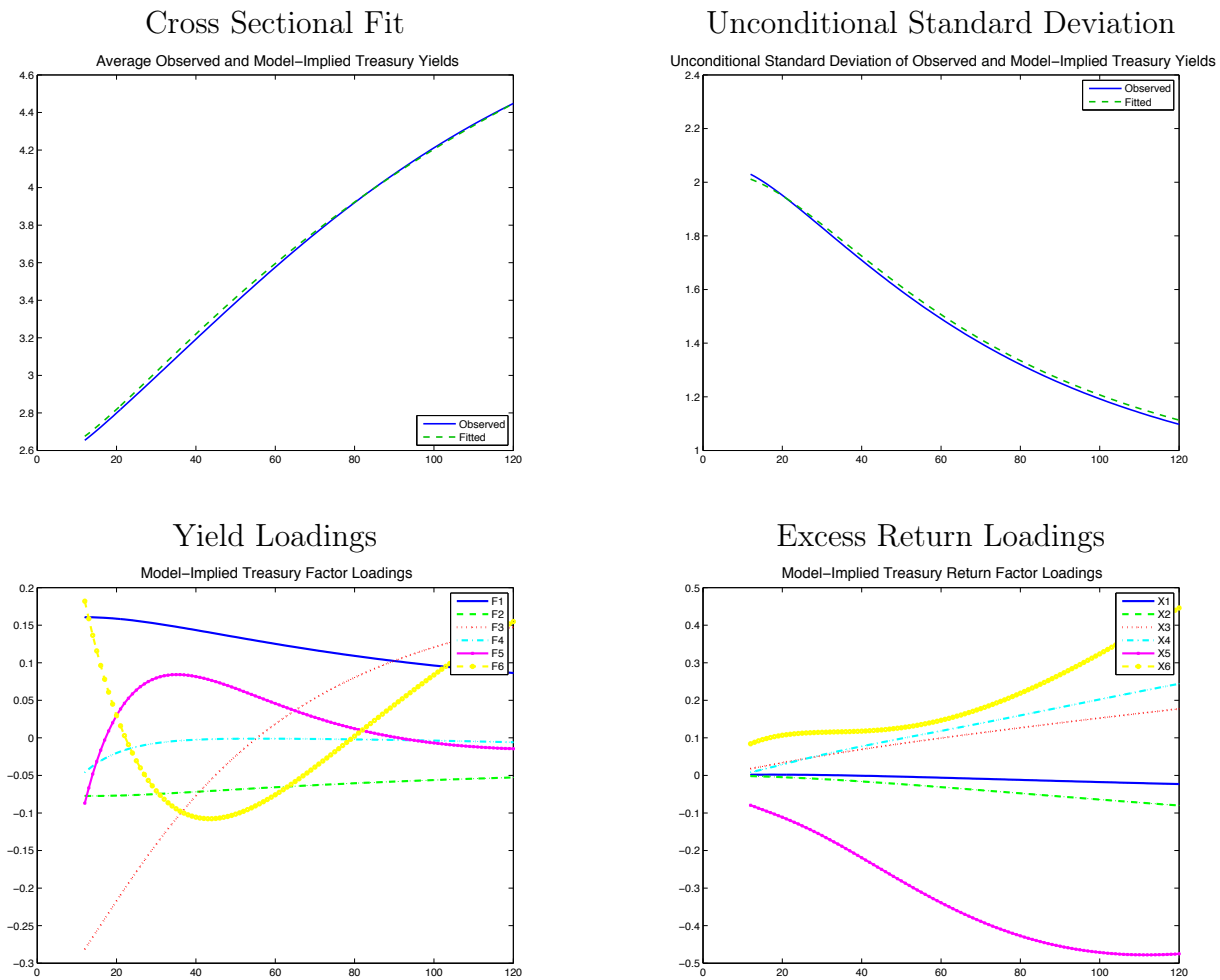

Figure 4: **TIPS: Observed and Model-Implied Time Series**

This figure provides time series plots of observed and model-implied TIPS yields and excess returns. The upper panels plot zero coupon yields at 2-year and 10-year maturities and the bottom panels plot excess holding period returns at 2-year and 10-year maturities. The observed yields and returns are plotted by solid blue lines, whereas dashed green lines correspond to model-implied yields and returns. Dashed red lines in the lower panels track model-implied expected excess holding period returns.

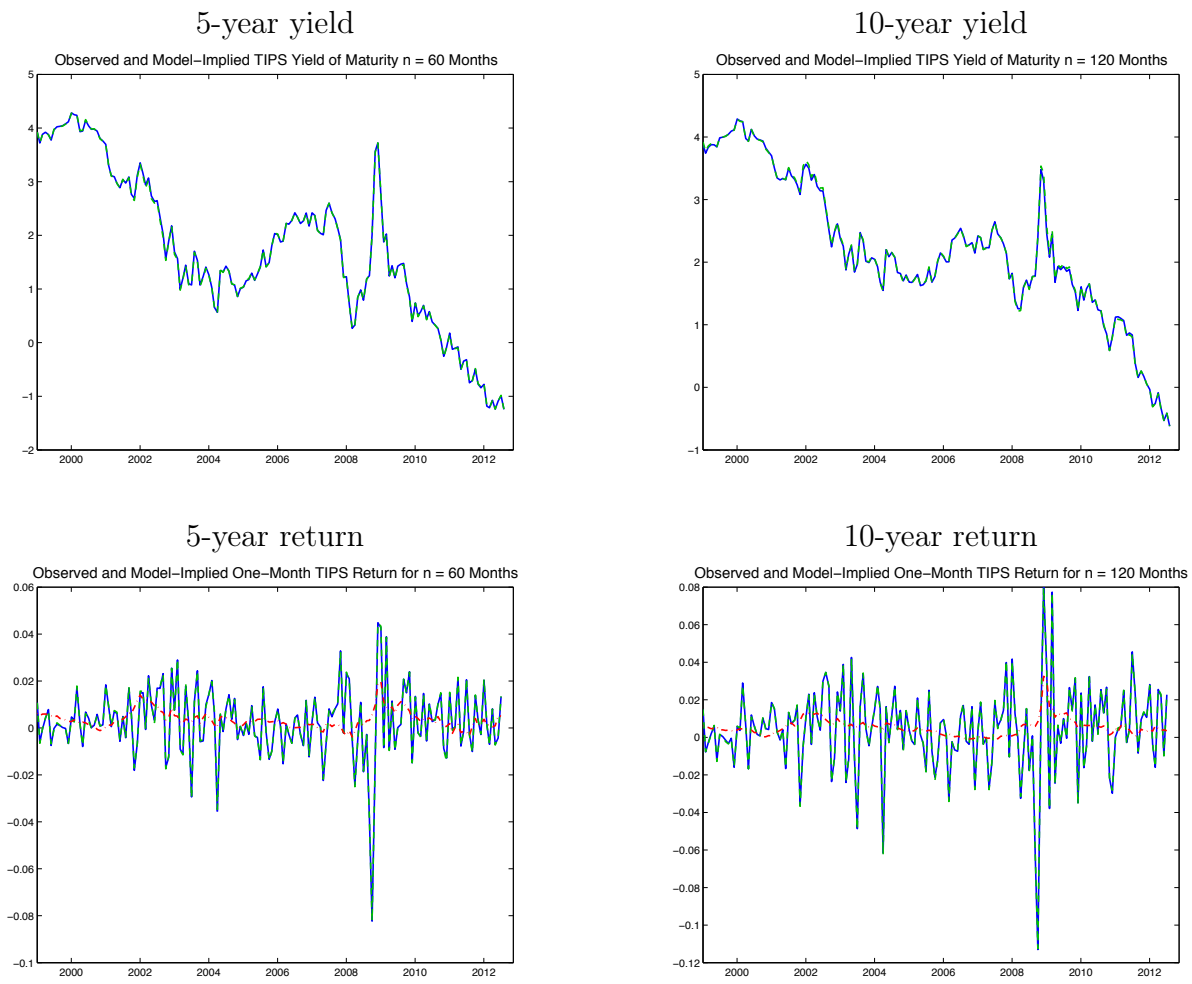

Figure 5: **TIPS: Cross Sectional Diagnostics**

This figure provides graphs exhibiting the cross-sectional fit and interpretation of the factors as drivers of TIPS yields. The upper two panels plot unconditional means and standard deviations of yields against those implied by the model. The lower left panel plots the implied yield loadings $-\frac{1}{n}B_{n,R}$. These coefficients can be interpreted as the response of the n-month yield to a contemporary shock to the respective factor. The lower right panel plots the expected return loadings $B'_{n,R}\lambda_1$. These coefficients can be interpreted as the response of the expected one-month excess holding return on an n-month bond to a contemporary shock to the respective factor.

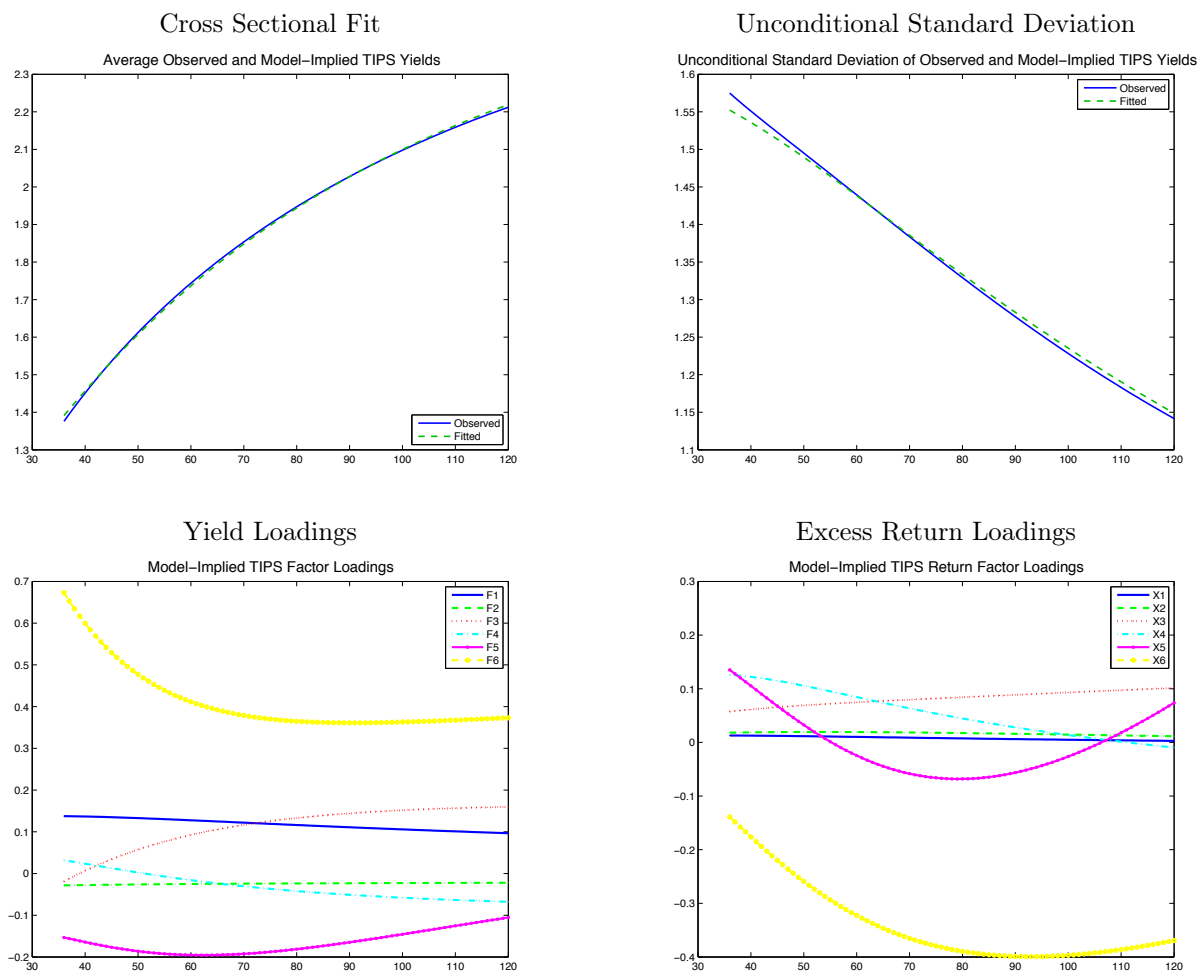

Figure 6: **Treasury and TIPS Term Premia**

This figure provides plots of the Treasury term premium and TIPS term premium across maturities. The upper panels plot decompositions of ten year Treasury and TIPS yields into the expected future short rate and the respective nominal and real term premium. A black dotted line represents the TIPS liquidity adjustment. The lower left panel plots the ten year Treasury and TIPS term premia together. The difference between the measures is the inflation risk premium. The lower right panel plots the decomposition of Treasury-TIPS breakevens into expected inflation in green and the inflation risk premium in red.

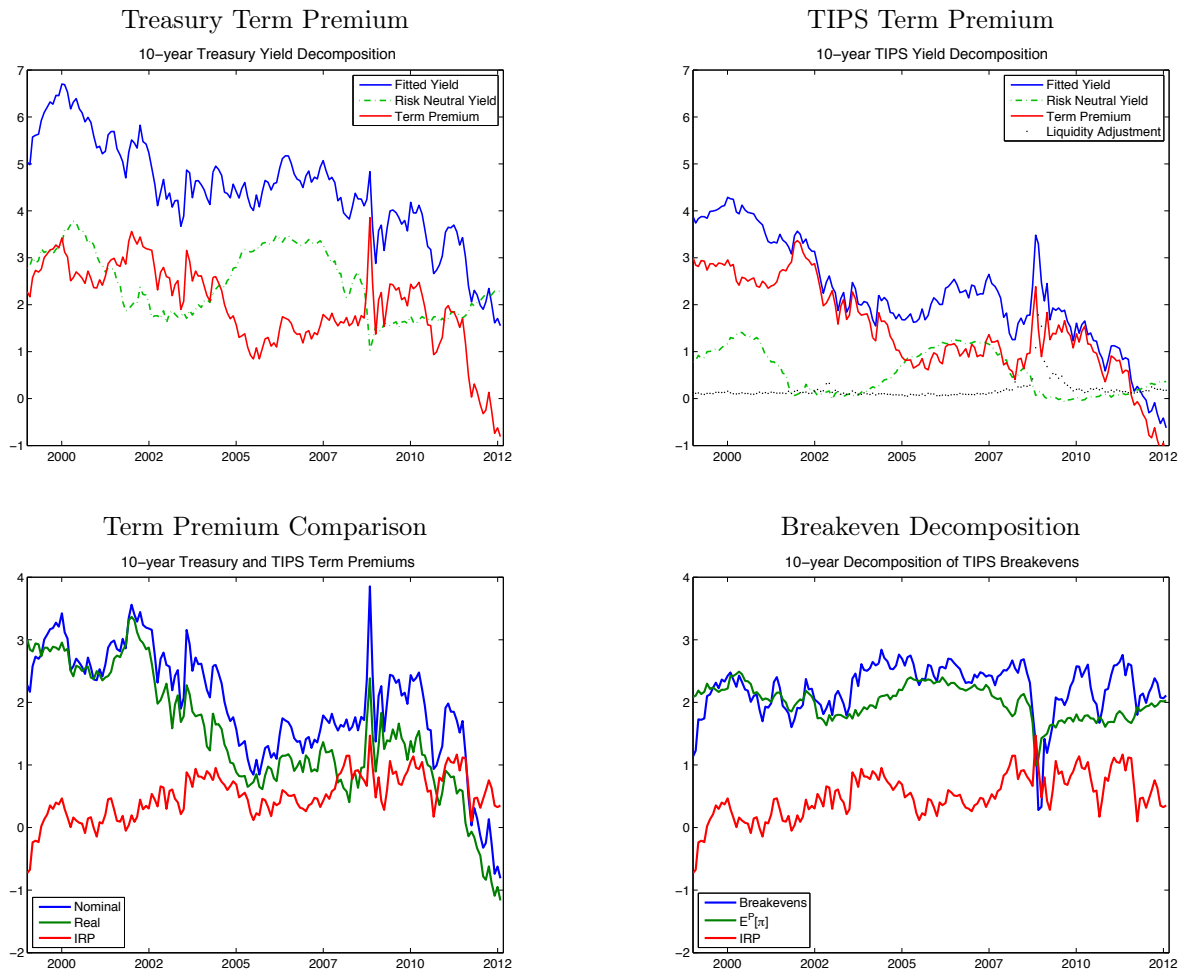

Figure 7: **Pricing Factors: Observed Time Series**

This figure plots the time series of each of the six principal components extracted from the cross section of Treasury and TIPS yields.

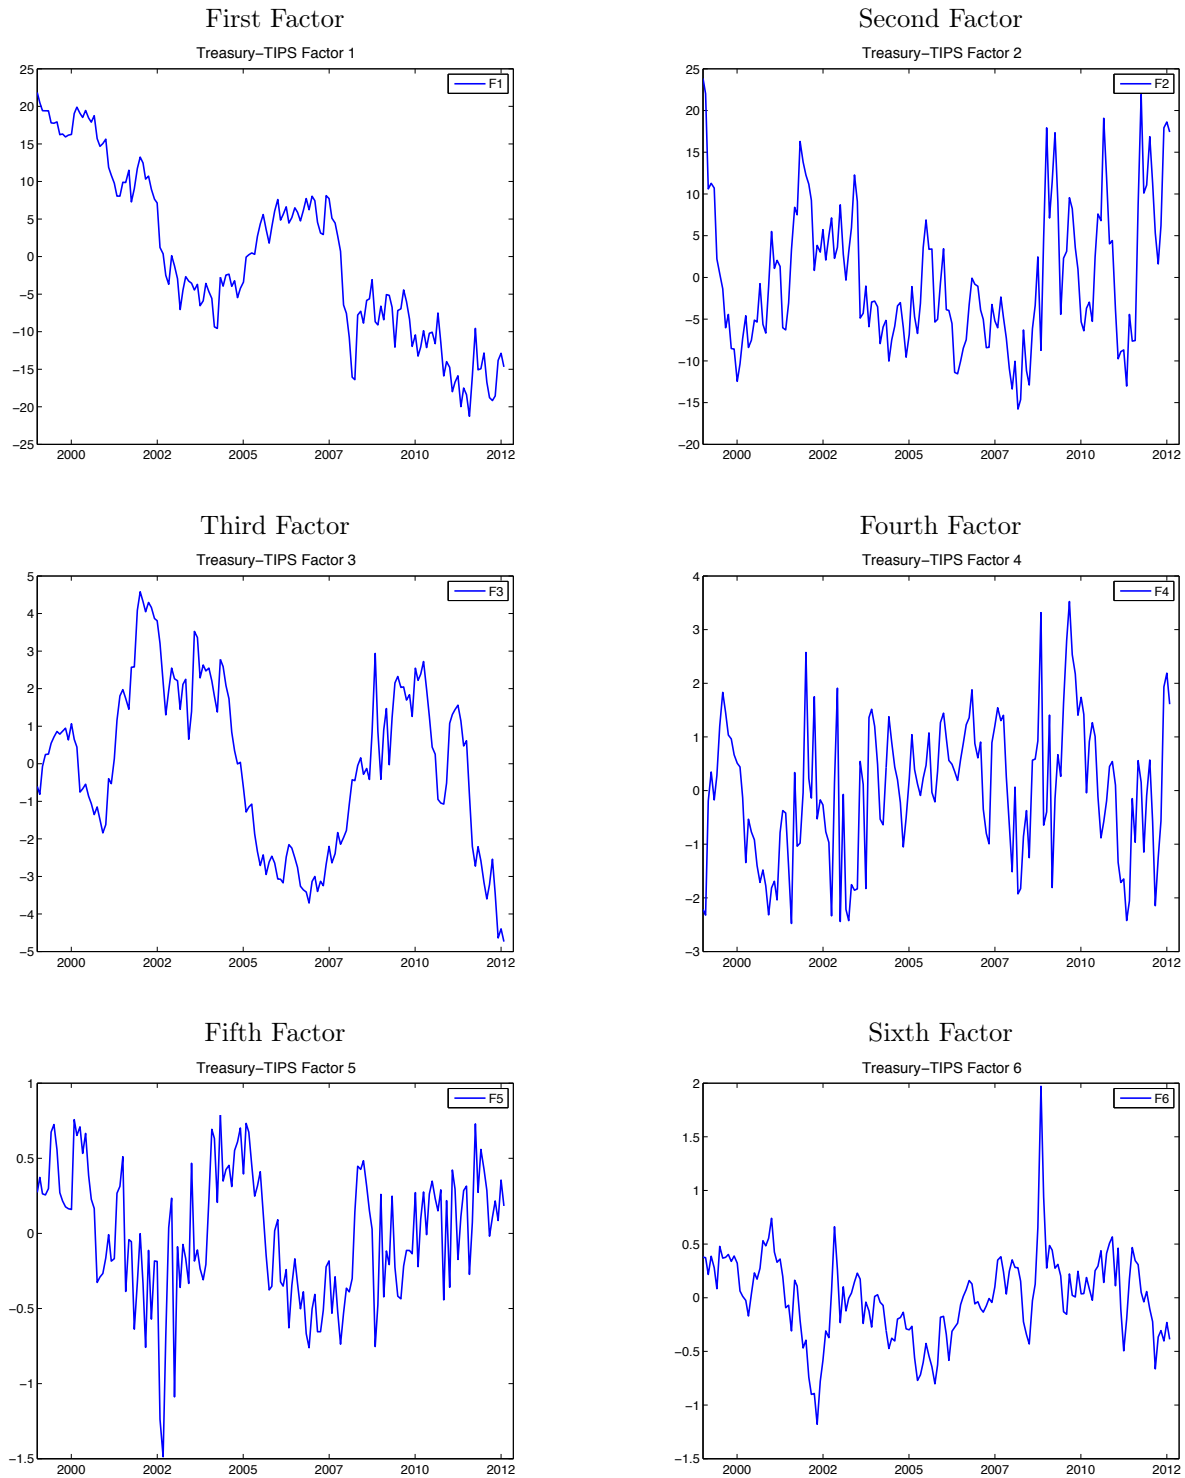

Figure 8: **Inflation Forecasting**

This figure plots the results of the inflation forecasting exercises; each panel displays the result at a different forecasting horizon. The first method, plotted in blue, is the forecast generated by our model and given by equations (20)-(21). The second method, plotted in green, take zero-coupon Treasury-TIPS breakevens as estimates of future inflation. The final forecasting methodology, plotted in red, is a random walk forecast. This method takes realized inflation over the previous n months as a forecast of average n-month future inflation. Realized CPI inflation is plotted in black. The crisis period of September 2007 - May 2009 is shaded in gray.

24 Months

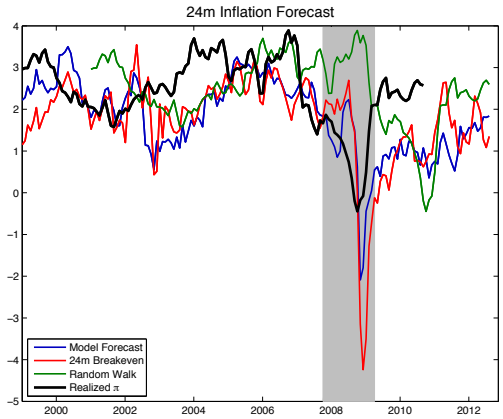

36 Months

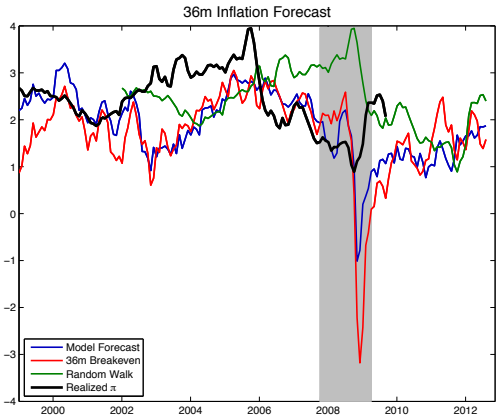

Figure 9: **Interpretation of Inflation Risk Premium**

This figure plots the two-year inflation risk premium generated by our benchmark specification beside several macroeconomic series. The upper two panels use the disagreement in inflation forecasts; the left plots the difference between the 85th percentile and the 15th percentile of the 1-quarter ahead Blue Chip forecasts, the right plots the standard deviation of the 4-quarter ahead Blue Chip forecasts. The middle two panels plot the three-month Merrill Lynch Swaption Volatility Expectations and the unemployment rate. The lower two panels plot 12-month Core CPI inflation and the Consumer Confidence Index.

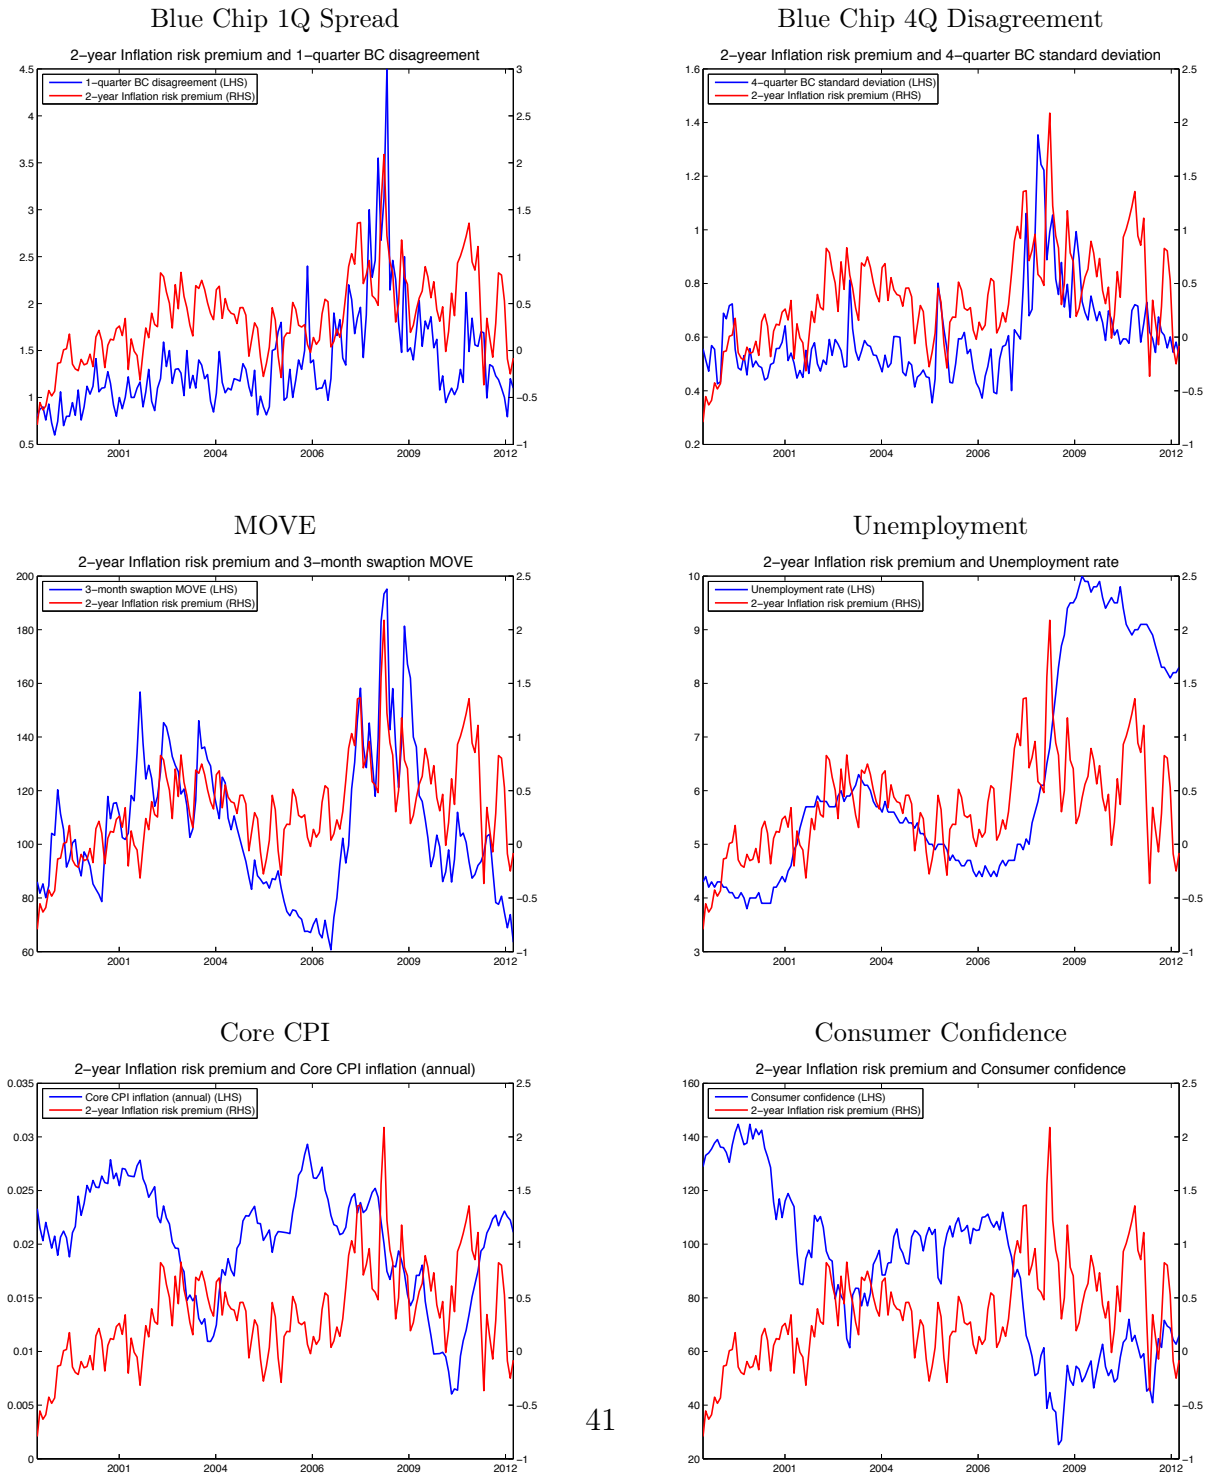

Figure 10: **TIPS Optionality**
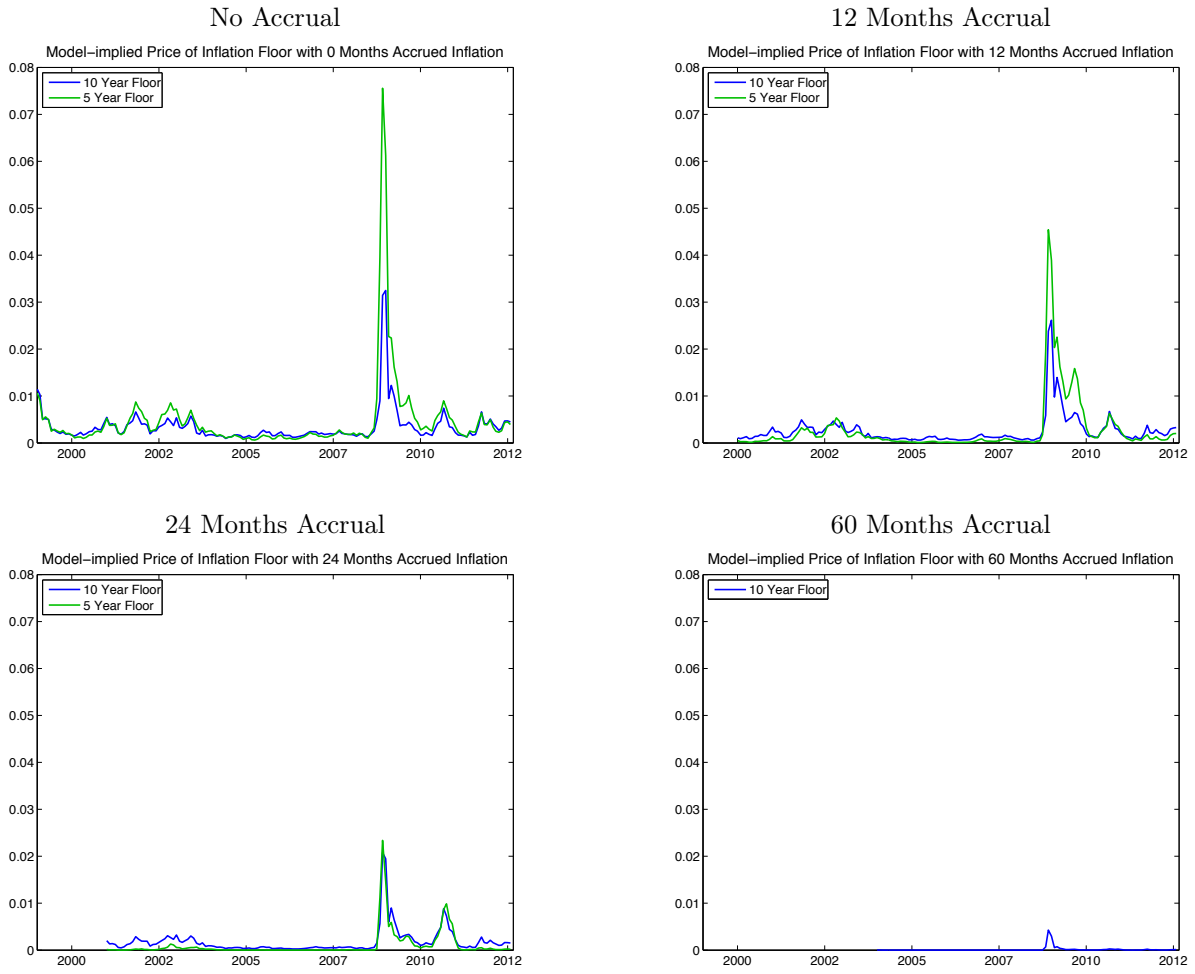
This figure plots the value of the inflation floors embedded in TIPS as a fraction of the principal. The ten-year inflation floor is plotted in blue and the five-year inflation floor is plotted in green. The upper left panel plots the option value in the case that the bond is a new issuance. The upper right panel plots the option value in the case that the bond was issued 12 months ago, and the inflation realized over that period has been incorporated into the option. Similarly, the lower panels plot the price of floors with 24 months and 36 months of accrued inflation.