FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

Could the U.S. Treasury Benefit from Issuing More TIPS?

Jens H.E. Christensen
Federal Reserve Bank of San Francisco

James M. Gillan
Federal Reserve Bank of San Francisco

June 2012

Working Paper 2011-16
http://www.frbsf.org/publications/economics/papers/2011/wp11-16bk.pdf

Could the U.S. Treasury Benefit from Issuing More TIPS?

Jens H. E. Christensen†

and

James M. Gillan

Federal Reserve Bank of San Francisco
101 Market Street, Mailstop 1130
San Francisco, CA 94105

Abstract

Yes. We analyze the economic benefit of Treasury Inflation Protected Securities (TIPS) issuance by estimating the inflation risk premium that penalizes nominal Treasuries vis-à-vis TIPS and the cost derived from TIPS liquidity disadvantage. To account for the latter, we introduce a novel model-independent range for the liquidity premium in TIPS exploiting additional information from inflation swaps. We also adjust our model estimates for finite-sample bias. The resulting measure provides a lower bound to the benefit of TIPS, which is positive on average. Thus, our analysis suggests that the Treasury could save billions of dollars by significantly expanding its TIPS program.

JEL Classification: E43, G12, H63.
Keywords: Term structure modeling, bias correction, market liquidity, government debt management, inflation swaps.

This paper has circulated in previous versions under the title "A Model-Independent Maximum Range for the Liquidity Correction of TIPS Yields." We thank Michael Bauer, Ib Hansen, John Krainer, Simon Kwan, Jose Lopez, Glenn Rudebusch, Xiaopeng Zhang, and seminar participants at the Federal Reserve Bank of San Francisco, the Board of Governors, the Danish National Bank, the 21st Derivatives Securities and Risk Management Conference, and the IBEFA-WEAI Summer Meeting in San Diego 2011 for helpful comments on earlier drafts of the paper. A special thank goes to Óscar Jorda for many helpful comments and suggestions. The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of San Francisco or the Federal Reserve System.
†Corresponding author: jens.christensen@sf.frb.org
This version: June 11, 2012.

1 Introduction

In light of the large deficits in recent years and the budgetary outlook for the years ahead, the question of how the U.S. government should finance its debt has become a pressing issue. Such a challenging fiscal environment highlights the importance of minimizing the cost of debt in order to reduce future tax burdens and the related economic inefficiencies. Historically, the U.S. government financed its deficits by issuing regular Treasury securities that pay fixed coupons and a fixed nominal amount at maturity. However, in 1997 the Treasury began issuing an alternative set of securities known as Treasury Inflation Protected Securities (TIPS) that deliver a real payoff because their principal and coupon payments are adjusted for inflation.[1] Since its inception, the market for TIPS has grown steadily and, as of the end of 2010, it represented 7.0 percent of the total market for tradeable U.S. Treasury securities.[2]

The policy question we address in this paper is straightforward: does the U.S. Treasury benefit from issuing TIPS instead of regular nominal Treasury securities and, if so, to what extent? The most direct benefit of issuing TIPS is that the government can avoid paying the premium on inflation uncertainty that investors demand in return for the risk of inflation exceeding its expected path. This inflation risk premium tends to push nominal Treasury yields up relative to TIPS yields, but estimating its level can be complicated due to the fact that the market for TIPS is smaller and less liquid than the market for Treasuries. As a result, bond investors require a liquidity premium to compensate them for the greater frictions to trade they might encounter from trading a less liquid asset during times of market uncertainty. In order to assess the economic benefit of issuing TIPS, we need to quantify this tradeoff between the advantage of the inflation risk premium and the disadvantage of the liquidity premium.

We determine the inflation risk premium using the dynamic term structure model of nominal and real Treasury yields developed in Christensen, Lopez, and Rudebusch (2012, henceforth referred to as the CLR model). This model allows us to decompose the difference between the two yield curves—also known as breakeven inflation (BEI)—into two components: bond investors' expected inflation and the associated inflation risk premium. The model we use is a slightly restricted version of the canonical affine term structure models classified by Dai and Singleton (2000) and is a well established approach within the empirical term structure literature to extracting risk premiums from yield curve data.[3]

[1] The Treasury uses the change in the headline Consumer Price Index (CPI) from the time when a series was issued to account for inflation compensation.

[2] According to the Bureau of the Public Debt, the total outstanding notional of TIPS was $616 billion as of the end of December 2010, which should be compared to $8,863 billion in total marketable Treasury securities. The data is available at: www.TreasuryDirect.gov

[3] We provide the basic model details in the paper, but those interested in a more thorough description should refer to CLR and the related paper by Christensen, Lopez, and Rudebusch (2010).

In order to account for the TIPS liquidity premium, we use a novel approach that exploits additional information in the market for inflation swaps. Instead of attempting to estimate the liquidity premium directly, we apply a series of no-arbitrage assumptions about the pricing of TIPS and inflation swaps relative to nominal Treasury bonds. The result is a model-independent range that allows for any assumption about the size of the TIPS liquidity premium that does not violate the joint pricing information in regular Treasuries, TIPS, and inflation swaps. The fundamental observation underlying our range is that, in a world without frictions to trade, BEI should equal the inflation swap rate. However, in reality, the observed BEI and inflation swap rates are not the same. We attribute the difference between the two to non-negative liquidity premiums in both the TIPS and inflation swap markets that reflect the distance these markets are from the ideal frictionless outcome.[4] In support of our approach, we find that our measure of the sum of TIPS and inflation swap liquidity premiums is highly correlated not only with existing estimates of TIPS liquidity premiums, but also with measures of liquidity premiums in other fixed-income markets.

This model-independent range allows us to correct TIPS yields for any admissible assumption about the level of TIPS liquidity premiums prior to model estimation. We use the CLR model to obtain estimates of the inflation risk premium at both extremes of the range. At the lower extreme TIPS yields contain no liquidity premiums and at the upper they contain the maximum. This produces a range of inflation risk premium estimates that maps one-to-one to any admissible assumption about the level of TIPS liquidity premiums. We also account for the finite-sample bias that occurs in the estimates of the model factor dynamics that drive the BEI decomposition.[5]

We combine our inflation risk and liquidity premium analysis to construct a measure that provides a quantitative lower bound to the economic benefit of TIPS issuance. In order to construct our measure, we first take the minimum of the estimated range of inflation risk premiums for each observation date and at each maturity. We then deduct the corresponding maximally admissible TIPS liquidity premium.[6] We refer to the resulting measure as the *minimum liquidity-adjusted inflation risk premium*. The average of our measure is positive over our sample with means of 5.5 and 10.5 basis points at the five- and ten-year maturity,

[4]Note that, due to collateral posting, the credit risk in inflation swap contracts is negligible and can be neglected for pricing purposes. Also, we assume the default risk of the U.S. government to be negligible, which is warranted for our sample that ends in 2010 well before the downgrade of U.S. Treasury debt in August 2011. However, even for this later period, any significant credit risk premium is not likely to bias an analysis like ours as it would affect regular Treasury and TIPS yields in the same way, leaving BEI effectively unchanged.

[5]Bauer, Rudebusch, and Wu (2012) provide a complete discussion of the finite-sample bias in empirical affine Gaussian term structure models. Our bias correction represents an adaptation of their approach to non-Gaussian models.

[6]We also adjust the result for the fact that the U.S. Treasury issues regular Treasury securities at low, so-called on-the-run Treasury yields, which differ from the seasoned, or so-called off-the-run, Treasury yields we use in model estimation.

respectively.[7] This suggests that the inflation risk premium has more than outweighed the liquidity disadvantage of TIPS on net and that the U.S. Treasury has benefitted from its TIPS program during our sample period.

To put the cost difference into perspective, we calculate the savings to the Treasury of dedicating a larger portion of its debt portfolio to TIPS instead of nominal Treasuries using our ten-year *minimum liquidity-adjusted inflation risk premium*. We find that over the next ten years at current debt levels the Treasury would save over $24.6 billion if the proportion of TIPS in its portfolio were increased from the current seven percent to one third. More realistically, if the debt continues to grow at the rate it has for the past ten years, restructuring the Treasury's portfolio to contain one third in TIPS would save $50.7 billion over ten years. Given the conservative nature of our measure and the fact that increasing the volume of TIPS would likely reduce their liquidity disadvantage and further increase the net premium investors are willing to forgo for inflation protection, the actual savings could prove to be substantively higher.

The rest of the paper proceeds as follows. Section 2 briefly describes the *minimum liquidity-adjusted inflation risk premium* and presents our results of assessing the benefit of TIPS issuance to the U.S. government. Section 3 describes the theory construction of the admissible range for the liquidity premium in TIPS yields, while Section 4 introduces the CLR model and our estimation of the inflation risk premium penalty of regular Treasuries. Section 5 concludes the paper. Appendices contain additional technical details.

2 The Benefit of TIPS Issuance to the U.S. Treasury

In this section, we summarize the result of our analysis of the benefit of TIPS issuance to the U.S. Treasury, while subsequent sections describe the various components that went into generating the result.

To assess the benefit to the U.S. Treasury of issuing TIPS, we calculate the most conservative estimate possible of a liquidity-adjusted inflation risk premium based on our bias-corrected model estimation results. To construct this measure for a given maturity, we first take the minimum for each observation date of the range of inflation risk premium estimates produced by the CLR model while allowing for any admissible TIPS yield liquidity correction. Second, we deduct the top of our range of admissible TIPS liquidity premiums. A positive difference at this point implies that even the smallest inflation risk premium estimate produced by the CLR model is larger than the largest feasible TIPS liquidity penalty, independent of any assumption about the specific level of the TIPS liquidity premium. We also make a final adjustment for the liquidity differential between the seasoned Treasury yields we use in the

[7] Most TIPS issuance takes place at the five- and ten-year maturities.

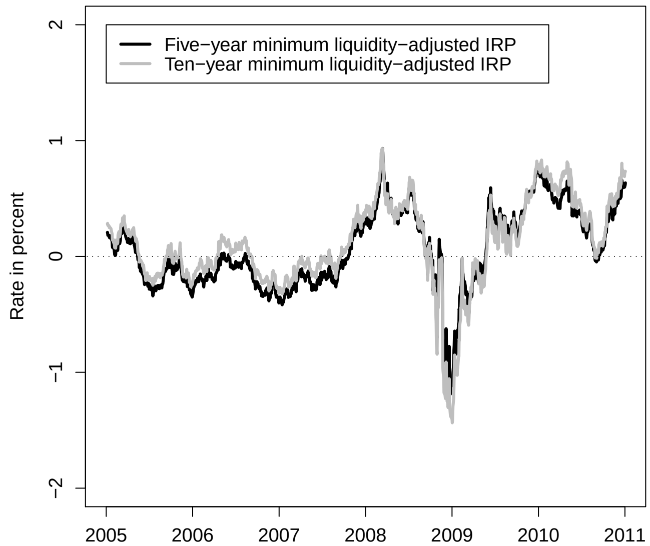

Figure 1: **Minimum Liquidity-Adjusted Inflation Risk Premiums.**
Illustration of the five- and ten-year minimum liquidity-adjusted inflation risk premiums based on the CLR model.

model estimation and the lower yields of newly issued Treasuries. In order to correct for this effect, we deduct the spread between off-the-run and on-the-run Treasury yields of the same maturity from the difference above to produce our final measure that we refer to as the *minimum liquidity-adjusted inflation risk premium* shown in Figure 1.

While this liquidity-adjusted measure of the smallest possible inflation risk premium is the most conservative estimate of the cost advantage of TIPS we can generate using the CLR model, it is still positive for 49.2 and 58.9 percent of the sample at the five- and ten-year maturities, respectively. The respective mean values are 5.5 and 10.5 basis points, while the respective maxima are 93.3 and 93.1 basis points.[8] Taken with the fact that, as we will argue later, the TIPS liquidity premium is likely in the lower half of the admissible range, far below the value assumed in the calculation of our measure, the actual benefit of TIPS is most likely higher than the estimate we provide. We believe this is a strong indication that the U.S. Treasury *has* benefitted from issuing TIPS for extended periods during this six-year sample.

As an alternative way of illustrating the benefits of TIPS, we attempt to put the cost savings of TIPS into dollar terms. This requires assumptions about both the amount of Treasury debt and its composition. Here, we limit the exercise to two simple calculations to

[8]For robustness, we repeated the analysis using the Gaussian model of nominal and real yields introduced in Christensen et al. (2010) and obtained similar results. These are available upon request.

exemplify the potential savings. The first scenario we consider is that the total amount of tradeable Treasury securities remains constant over the next ten years at their level as of the end of 2010 ($8,863 billion) and that the proportion of TIPS is increased from the current seven percent to one third.[9] Using the average of our ten-year *minimum liquidity-adjusted inflation risk premium* as the cost saving rate, we find that the Treasury would save $24.6 billion over the next ten years. Another scenario is the somewhat more realistic assumption that the Treasury debt will continue to grow at the rate it has for the past ten years. If we assume the Treasury restructures its portfolio to contain one third in TIPS, it would lead to a cost reduction of the order of $50.7 billion dollars over ten years.[10] It is also important to note that these are likely not to be exaggerated numbers. We first point to the conservative nature of our measure that imposes the maximally admissible liquidity penalty on TIPS. Second, the development of the TIPS market and its liquidity since its inception in 1997 suggests that increasing the TIPS issuance volume as implied by the envisioned policy change above would likely reduce their liquidity costs. As a result, TIPS would become even more attractive to investors and the cost reductions would be even larger.

However, we emphasize that this is a narrow cost-benefit analysis of the tradeoff on any given day between issuing a nominal or a real bond based on bond investors' outlook for inflation and their associated premium for being exposed to that risk. Other considerations such as the overall composition of the Treasury's liabilities or differences between the Treasury's inflation outlook and that of market participants might lead to a different conclusion,[11] but such aspects are outside the scope of this paper.

The remainder of the paper is devoted to detailing the components that went into achieving the result above. In Section 3, we describe our model-independent measure of the liquidity disadvantage of TIPS, while Section 4 describes our usage of the CLR model to generate estimates of the inflation risk premium penalty of nominal Treasuries.

3 Measuring the Liquidity Premium in TIPS Yields

The existence of TIPS liquidity premiums is well documented. Fleming and Krishnan (2012) report that the market characteristics of TIPS trading indicate smaller trading volume, longer turn around time, and wider bid-ask spreads than is normally observed in the nominal Treasury bond market (see also Campbell, Shiller, and Viceira 2009, Dudley, Roush, Steinberg Ezer 2009, Gürkaynak, Sack, and Wright 2010, Sack and Elsasser 2004). However, the degree

[9]We choose this ratio because it is close to the ratio of real to total government debt that the U.K currently issues

[10]This assumes the amount of marketable Treasuries continues to grow for ten years at the average rate of 13.1 percent, the yearly rate it has averaged from 2001 through 2011. Additionally, the cost comparison assumes a constant ratio of TIPS of one third versus seven percent.

[11]We thank Ib Hansen for pointing out this distinction.

to which they bias TIPS yields remains a topic of debate because attempts to estimate the TIPS liquidity premium directly have resulted in varying results. D'Amico, Kim, and Wei (2010, henceforth DKW) estimate the liquidity premium using a latent-factor affine term structure model of nominal and real Treasury yields. They find that the premium was 100 to 200 basis points during the early years of TIPS issuance, but that it has since declined to average about 50 basis points at the ten-year maturity.[12] Pflueger and Viceira (2011, henceforth PV) estimate the TIPS liquidity premium by regressing BEI on several measures of liquidity in bond markets. Their estimate follows a similar overall pattern to the DKW estimates, but averages 70 basis points. Both estimates show sharp spikes in the liquidity premium during the financial crisis that exceed 250 basis points.

This liquidity disadvantage of TIPS is the major point of criticism for those who believe that the TIPS program is ultimately costly for the Treasury. In a recent paper, Fleckenstein, Longstaff, and Lustig (2010, henceforth FLL) go so far as to claim that the poor liquidity in the TIPS market has resulted in severe instances of arbitrage that could have otherwise been avoided if the Treasury had issued nominal debt instead. However, compared to this earlier research, our assessment is arguably preferable in that we do not have to rely on model-based estimates of the TIPS liquidity premiums, which have until now been inconsistent across studies. Our approach is novel in that we use the additional information in inflation swap rates to create a bounded range for the TIPS liquidity premium. Further analysis may provide insight into the level of the premium, but for the purpose of our cost-benefit analysis, the range we provide is ideal since it is model-independent. The following sections introduce the theory, data, and characteristics of the range we use to account for the TIPS liquidity disadvantage in our assessment.

3.1 The Maximum Range for the Liquidity Correction of TIPS Yields

To begin, we first detail our theoretical thinking and the information set that underlie it. We observe a set of nominal and real Treasury zero-coupon bond yields, which we denote $\widehat{y}_t^N(\tau)$ and $\widehat{y}_t^R(\tau)$, respectively, where τ is the number of years to maturity. Also, we observe a corresponding set of zero-coupon rates on inflation swap contracts denoted $\widehat{IS}_t(\tau)$. Due to microstructure frictions, such as bid-ask spreads and discrete-time trading in discrete denominations, these rates differ from the unobserved values that would prevail in a frictionless world without any obstacles to continuous trading, which we denote $y_t^N(\tau)$, $y_t^R(\tau)$, and $IS_t(\tau)$, respectively. Assuming investors are forward looking and utility-maximizing, there should be no arbitrages in such a world and the resulting asset prices should only reflect the underlying

[12] This figure excludes the abnormal periods of market illiquidity that occurred early in the series and during the financial crisis. The average is calculated with data beginning in January 2004 after the decline had flattened out and excludes the crisis period from August 2007 to May 2009.

economic fundamentals.

Implicit in the usage of the word "premium" (or penalty) is the notion that a clean, unobserved price would prevail if only some, not necessarily well-identified, market microstructure frictions did not bias the prices actually observed. We define the absolute liquidity premium as the price difference between the observed and the unobservable "frictionless" market outcome of a given asset. As is common in the literature, we only seek to identify the relative liquidity premium between two comparable assets since it is extremely difficult to identify the unobservable frictionless price of an asset directly. In this sense the liquidity premiums we derive represent the total cost of all frictions to trade (wider bid-ask spreads, lower trading volume, etc.) of the less liquid asset beyond those of the more liquid asset against which it is being compared.

3.1.1 Three Fundamental Assumptions

In order to derive the bound for the admissible range of liquidity premiums in TIPS yields, we introduce three fundamental assumptions:

(i). The nominal Treasury yields we observe are very close to the unobservable nominal yields that would prevail in a frictionless world. Thus, we assume $\widehat{y}_t^N(\tau) = y_t^N(\tau)$ for all t and all relevant τ. Even if not exactly true (say, for example, during the financial crisis of 2008 and 2009), this is not critical to the analysis as the story is ultimately about the relative liquidity between securities that pay nominal and real yields.

(ii). TIPS are no more liquid than nominal Treasury bonds. As a consequence, the TIPS yields we observe contain a time-varying liquidity premium denoted $\delta_t^R(\tau)$, which generates a wedge between the observed TIPS yields and its frictionless counterpart given by $\widehat{y}_t^R(\tau) = y_t^R(\tau) + \delta_t^R(\tau)$ with $\delta_t^R(\tau) \geq 0$ for all t and all relevant τ.

(iii). Inflation swaps are no more liquid than nominal Treasury bonds. As a consequence, observed inflation swap rates can also be different from their frictionless counterpart with the difference given by $\widehat{IS}_t(\tau) = IS_t(\tau) + \delta_t^{IS}(\tau)$ and $\delta_t^{IS}(\tau) \geq 0$ for all t and all relevant τ.

Provided the above assumptions are valid, there exists an upper bound to the unobservable liquidity premium in TIPS yields tied to the difference between observed inflation swap and BEI rates, which we define as

$$\Delta_t(\tau) \equiv \widehat{IS}_t(\tau) - \widehat{BEI}_t(\tau) = \widehat{IS}_t(\tau) - [\widehat{y}_t^N(\tau) - \widehat{y}_t^R(\tau)]. \tag{1}$$

In a world without any frictions to trade, arbitrageurs would chisel away any difference between the inflation swap rate and breakeven inflation. Thus, the frictionless inflation swap rate must equal the frictionless breakeven inflation rate such that $IS_t(\tau) = y_t^N(\tau) - y_t^R(\tau)$ for all t and all τ.[13] By implication, the difference $\Delta_t(\tau)$ is a measure of how far the observed market rates are from the frictionless outcome. Using our three assumptions the difference can be re-written as

$$\Delta_t(\tau) = IS_t(\tau) + \delta_t^{IS}(\tau) - [y_t^N(\tau) - (y_t^R(\tau) + \delta_t^R(\tau))] \quad (2)$$

$$= \delta_t^R(\tau) + \delta_t^{IS}(\tau) \geq 0. \quad (3)$$

This shows that, under our assumptions, the difference between the observed inflation swap and breakeven inflation rates is non-negative and equal to the *sum* of the liquidity premium in TIPS and the liquidity premium in inflation swaps. Since we assume the liquidity premiums to be non-negative, the difference provides a bound on the TIPS liquidity premium:

$$\min \delta_t^R(\tau) = 0 \quad \text{and} \quad \max \delta_t^R(\tau) = \Delta_t(\tau).$$

At the lower bound, observed TIPS prices contain no liquidity premium because the deviation from the frictionless market outcome is entirely attributable to liquidity premiums in the observed inflation swap rate. In terms of the notation introduced above, $\delta_t^R(\tau) = 0$ and $\delta_t^{IS}(\tau) = \Delta_t(\tau)$, implying that the observed TIPS yields are equal to the frictionless real yields, $y_t^R(\tau) = \widehat{y}_t^R(\tau)$. This extreme represents the assumption implicit in papers like Adrian and Wu (2010) and CLR where TIPS yields are taken at face value without any correction for liquidity.

Conversely, at the upper bound, TIPS are priced with the maximum liquidity premium while inflation swap rates carry none. In this case, $\delta_t^R(\tau) = \Delta_t(\tau)$ and $\delta_t^{IS}(\tau) = 0$, meaning the observed inflation swap rates reflect the frictionless breakeven rates, $\widehat{IS}_t(\tau) = IS_t(\tau)$, while $\widehat{y}_t^R(\tau) = y_t^R(\tau) + \Delta_t(\tau)$ or, equivalently, $y_t^R = \widehat{y}_t^N(\tau) - \widehat{IS}_t(\tau)$. Thus, the corrected frictionless TIPS yields are the observed nominal yields less the observed inflation swap rates. This assumption is implicit in the analysis performed by FLL as all opportunities for arbitrage documented in that paper are assigned to the market for TIPS rather than the market for inflation swaps. Finally, it should be noted that the continuum in between the two extremes gives all admissible combinations of $\delta_t^R(\tau)$ and $\delta_t^{IS}(\tau)$.

3.1.2 The Liquidity Premium in Inflation Swap Rates

We believe our first two assumptions are uncontroversial because the market for nominal U.S. Treasury bonds is one of the most liquid fixed-income markets. In comparison, TIPS are

[13] In a frictionless world, to buy one nominal zero-coupon bond at t that yields $y_t^N(\tau)$ produces the same cash flow as buying one real zero-coupon bond that yields $y_t^R(\tau)$ and going short, at zero cost, a τ-year inflation swap contract with a fixed rate of $IS_t(\tau)$.

Receipts	$\widehat{IS}_t(\tau)\tau+\widehat{y}_t^R(\tau)\tau+\frac{CPI(t+\tau)}{CPI(t)}$-1 +[LIBOR+$\beta_t^N(\tau)$]$\tau$
Payments	$\widehat{y}_t^N(\tau)\tau$ + $\frac{CPI(t+\tau)}{CPI(t)}$-1 +[LIBOR+$\beta_t^R(\tau)$]$\tau$
Net receipts	$[\widehat{IS}_t(\tau) - (\widehat{y}_t^N(\tau) - \widehat{y}_t^R(\tau)) + \beta_t^N(\tau) - \beta_t^R(\tau)]\tau$

Table 1: **Cash Flow of Investment Strategy that Hedges a Short Position in a Zero-Coupon Inflation Swap Contract.**
Illustration of the cash flows involved in the investment strategy that hedges a short position in an inflation swap. It involves: 1) The inflation swap position itself. 2) A long asset swap position in the τ-year zero-coupon TIPS. 3) A short asset swap position in the τ-year zero-coupon Treasury bond.

widely considered to be less liquid. However, the assumption that the observed inflation swap rates are above their ideal frictionless rate is less obvious and merits elaboration.

The mechanics of hedging activity in the inflation swap market, as described in Campbell et al. (2009), suggest that the observed inflation swap rate should be marked up from the unobserved frictionless rate due to the financing and transaction costs of replicating cash flows in related asset swap markets. In practice, there are two strategies for generating the CPI-linked floating cash flows in the inflation swap contract. The first is to buy the TIPS with the desired maturity. This requires funding and implies receiving cash flows on all coupon dates of that security, which investors may not find to be desirable. The alternative is to enter into a zero-coupon inflation swap of the desired maturity. There are no funding costs in a zero-coupon inflation swap as its value is zero at inception. As a consequence, investors should be willing to pay an extra premium to avoid such funding costs, which explains why the inflation swap rate can be above BEI in equilibrium. However, the size of the inflation swap rate markup is primarily determined from the supply side. The counter party to the inflation swap (typically a hedge fund or investment bank) generates CPI-linked cash flows by going long in TIPS and short in nominal Treasury bonds through the asset swap market. Thus, the markup represents the compensation the counter party requires for assuming the liquidity risk of multiple transactions on the backside of the contract.

This hedging activity creates a connection between our range for the TIPS liquidity premiums and asset swap rates. In an asset swap, the party long the contract pays LIBOR plus a spread while receiving the cash flow of a specific bond without exchange of the principal. In the inflation swap, the party short the contract typically generates CPI-linked cash flows by making the following set of transactions at time t to hedge the assumed risk:

- A short position in the τ-year zero-coupon inflation swap struck at $\widehat{IS}_t(\tau)$, that is, the investor will receive $\widehat{IS}_t(\tau)\tau$ at maturity in return for delivering the net change in the price level $\frac{CPI(t+\tau)}{CPI(t)} - 1$.

- A long asset swap position for the τ-year zero-coupon TIPS, that is, agree to paying [LIBOR+$\beta_t^R(\tau)$]τ in order to receive the fixed accrued coupon $\widehat{y}_t^R(\tau)\tau$ and the accrued inflation compensation $\frac{CPI(t+\tau)}{CPI(t)} - 1$.[14]

- A short asset swap position in the τ-year zero-coupon Treasury bond, that is, agree to paying the nominal Treasury yield $\widehat{y}_t^N(\tau)\tau$ in order to receive [LIBOR+$\beta_t^N(\tau)$]τ.

Here, $\beta_t^N(\tau)$ and $\beta_t^R(\tau)$ denote the asset swap spreads for the nominal Treasuries and TIPS, respectively. As all transactions involve swaps on zero-coupon assets, there is no outlay upon inception because they all have zero net value and payments are only exchanged at maturity τ years later. Table 1 summarizes the outlays and receipts from this set of transactions at maturity. The net receipt to the party short the inflation swap is

$$[\widehat{IS}_t(\tau) - (\widehat{y}_t^N(\tau) - \widehat{y}_t^R(\tau)) + \beta_t^N(\tau) - \beta_t^R(\tau)]\tau \geq 0. \tag{4}$$

Note that this strategy is really a hedge as the value on the left-hand side of Equation 4 is deterministic and set at the inception of the contract. For the leveraged investor to be willing to participate in the inflation swap market this value must be non-negative as indicated. Since we defined the maximum TIPS liquidity premium by $\Delta_t(\tau) = \widehat{IS}_t(\tau) - (\widehat{y}_t^N(\tau) - \widehat{y}_t^R(\tau))$, the inequality in Equation (4) can also be written as

$$\Delta_t(\tau) + \beta_t^N(\tau) - \beta_t^R(\tau) \geq 0. \tag{5}$$

Campbell et al. (2009) note that, normally, the asset swap spreads, $\beta_t^N(\tau)$ and $\beta_t^R(\tau)$, are negative and more so for the nominal Treasuries, that is,

$$\beta_t^N(\tau) < \beta_t^R(\tau) \leq 0.^{15} \tag{6}$$

Under competitive circumstances (zero cost of entry to the inflation swap market etc.), we expect Equation (5) to hold with equality. Using the inequality in Equation (6), we can then re-write Equation (5) as

$$\Delta_t(\tau) = \beta_t^R(\tau) - \beta_t^N(\tau) > 0. \tag{7}$$

Thus, our maximum range for the TIPS liquidity premium equals the difference between the TIPS and nominal Treasury asset swap spreads and is strictly positive. Of course, this

[14] Here, we are neglecting the value of the deflation protection in the TIPS in that the actual payment on the TIPS asset swap is $\max[\frac{CPI(t+\tau)}{CPI(t)} - 1, 0]$. We thank Xiaopeng Zhang for pointing this out. Thus, the calculations are accurate provided the value of the deflation protection for the particular TIPS in the asset swap is negligible. If not, they can be corrected by calculating its value in a way similar to the one described in CLR.

[15] During the financial crisis of 2008 and 2009, $\beta_t^R(\tau)$ turned positive, but the relative relationship between $\beta_t^N(\tau)$ and $\beta_t^R(\tau)$ remained as indicated by the first inequality.

idealized calculation is based on zero-coupon bonds, but the difference between asset swap spreads on TIPS and regular Treasuries should still provide a good approximation to our maximum range.

We re-emphasize that, without additional information, any combination of non-negative $\delta_t^R(\tau)$ and $\delta_t^{IS}(\tau)$ that satisfies the condition $\Delta_t(\tau) = \delta_t^R(\tau) + \delta_t^{IS}(\tau)$ is admissible and cannot be legitimately excluded ex ante. Also, we underscore that our construction is valid for any sample of nominal Treasury and real TIPS yields as long as our three key assumptions are satisfied by the data. This observation implies that the size and shape of the range depend on the underlying pool of bonds, the method used in the yield curve construction etc., but that the range still represents the admissible band for the TIPS liquidity premium for the specific sample under consideration.

3.1.3 Construction and Characteristics of the Range

We use daily estimates of zero-coupon nominal and real Treasury bond yields as constructed by Gürkaynak, Sack, and Wright (2007, 2010, henceforth GSW) for our observed bond yields. For our inflation swap rates, we use daily quotes from Bloomberg. The rates we use are for zero-coupon inflation swap contracts, meaning they have no exchange of payment upon issuance and a single cash flow exchanged at maturity. The quoted rates represent the payment of the fixed leg at an annual rate, which we convert into continuously compounded rates using the formula $\widehat{IS}_t^c(\tau) = \ln(1 + \widehat{IS}_t(\tau))$ to make them comparable to the other interest rates. Bloomberg begins reporting quotes on inflation swap rates in early 2004, but the data is not consistently populated until the end of the year. As a result, we begin the sample period on January 4, 2005 and eliminate the few days during the sample period where quotes are not available for all maturities, which leaves us with a sample of 1,482 observations.

To construct the admissible range for the TIPS liquidity premium, we calculate zero-coupon BEI rates at the five-, six-, seven-, eight-, nine- and ten-year maturities and deduct them from the inflation swap rates of the corresponding maturities to obtain $\Delta_t(\tau)$. We begin the sample period on January 4, 2005 in order to align the TIPS sample with the inflation swap data resulting in 1,482 observations.[16] Figure 2 illustrates the upper bound to the TIPS liquidity premium at the five-, seven-, and ten-year maturities, while Table 2 reports the summary statistics for all six series. Consistent with our theoretical assumptions, the values of $\Delta_t(\tau)$ are strictly positive for all six maturities. Furthermore, the term structure of the maximum TIPS liquidity premiums tends to be downward sloping with maturity.

We note that the size and variation of the maximum range do not simply reflect bid-ask spreads of the underlying bonds and inflation swap contracts. Figure 3 shows the bid-ask

[16] For the purpose of consistency and to avoid filtering, we eliminate the nominal and real yields on days where the inflation swap rates are not available.

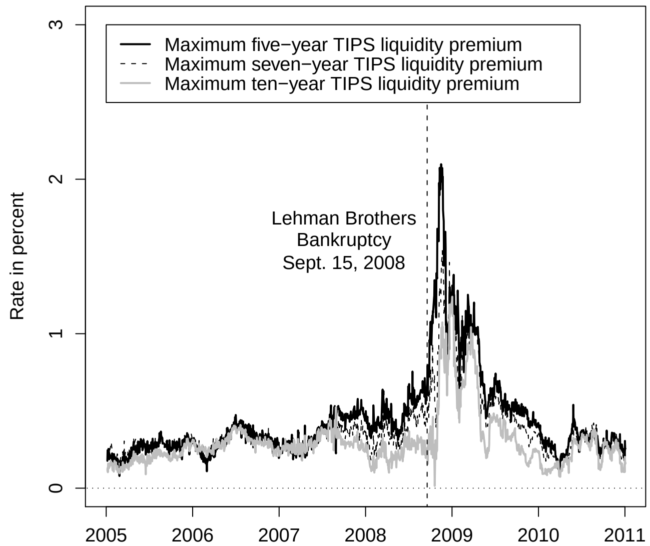

Figure 2: **The Maximum TIPS Liquidity Premiums.**
Illustration of the difference between zero-coupon inflation swap rates and comparable zero-coupon BEI rates that constitute the upper bound of the admissible range for TIPS liquidity premiums.

Maturity in months	Mean in bps	Std. dev. in bps	Min. in bps	Max. in bps	Skewness	Kurtosis
60	44.16	29.11	7.88	209.73	2.62	11.54
72	41.58	24.67	13.49	171.87	2.49	9.93
84	38.56	21.69	10.82	153.81	2.50	9.62
96	35.38	19.17	9.87	138.24	2.56	9.93
108	32.29	18.02	8.55	129.03	2.60	10.58
120	29.90	17.43	1.44	123.97	2.54	10.44

Table 2: **Summary Statistics for the Maximum TIPS Liquidity Premiums.**
Summary statistics for the sample of differences between zero-coupon inflation swap rates and comparable zero-coupon breakeven inflation rates covering the period from January 4, 2005 to December 31, 2010, a total of 1,482 daily observations.

spread as reported by Bloomberg for the most recently issued, or so-called on-the-run, five- and ten-year TIPS. The average over our sample period is 1.4 basis point for the five-year TIPS and 0.7 basis point for the ten-year TIPS, which should be compared to averages of 9.1 and 7.7 basis points for the five- and ten-year inflation swap, respectively. This is of the order of transaction costs in the inflation swap market reported by FLL based on their conversations with traders. For comparison, Fleming and Mizrach (2009) report the inside bid-ask spreads for regular Treasuries average one basis point of par at the ten-year maturity.

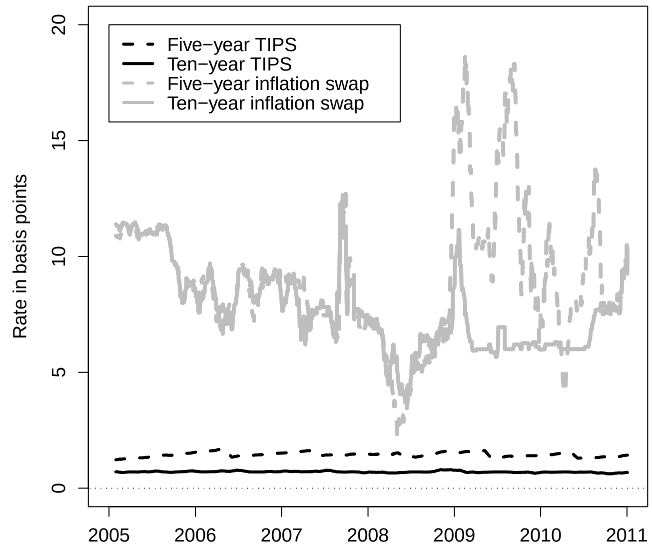

Figure 3: **Bid-Ask Spreads in the TIPS and Inflation Swap Markets.**
Illustration of the bid-ask spread as reported by Bloomberg for the most recently issued, or so-called on-the-run, five- and ten-year TIPS. Shown are also the bid-ask spreads from the inflation swap market for the five- and ten-year zero-coupon inflation swap contracts. All series are smoothed four-week moving averages and measured in basis points.

This implies even smaller trading costs in terms of yield to maturity, the return comparable to the inflation swap rate. Thus, bid-ask spreads only account for a small fraction of the sum of liquidity premiums in TIPS and inflation swaps that represents our maximum range for the TIPS liquidity premium.

3.2 Comparison to Other Liquidity Premium Estimates

Our measure of combined liquidity premiums in TIPS yields and inflation swap rates is highly correlated with several other commonly used measures of bond liquidity. Notably, the ten-year off-the-run over on-the-run nominal Treasury bond yield difference and the credit spread of AAA-Rated U.S. industrial bonds have simple correlation coefficients of 0.67 and 0.73, respectively.[17] The former represents the premium investors place on seasoned nominal treasuries over the highly liquid newly issued series of the same maturity. The latter represents

[17]The spread between the ten-year off-the-run Treasury par-coupon bond yield from the GSW (2007) database and the ten-year Treasury par-coupon bond yield from the H.15 series at the Board of Governors is compared with the maximum ten-year TIPS liquidity premium. The spread between the one-year AAA-rated U.S. industrial corporate bond yield and the comparable Treasury yield is compared with the five-year TIPS liquidity premium.

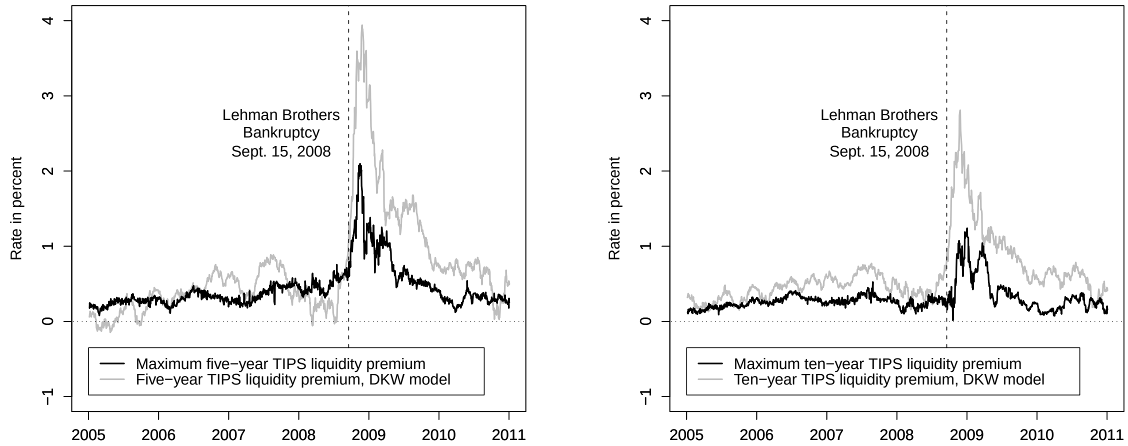

(a) Five-year maturity.

(b) Ten-year maturity.

Figure 4: **Comparison to the TIPS Liquidity Premiums from the DKW Model.** Comparison of the model-independent maximum for the admissible liquidity premium in TIPS yields to the estimated TIPS liquidity premiums from the DKW model.

the liquidity advantage of government bonds over corporate bonds that contain negligible credit risk. These comparisons are meant to illustrate that more general variation in bond market liquidity is apparent in our maximum range.

We also compare our measure to the TIPS liquidity premiums estimated in DKW and PV. DKW estimate the liquidity premium in TIPS yields by constructing a joint model of nominal and real Treasury yields and CPI inflation and including a TIPS-specific factor in order to capture the liquidity disadvantage of TIPS yields explicitly. Figure 4 illustrates the five- and ten-year DKW TIPS liquidity premium series and compares them to our maximum range for the admissible five- and ten-year TIPS liquidity premiums. The DKW estimate and the maximum range are highly correlated, with correlation coefficients of 0.86 and 0.77 at the five- and ten-year maturities, respectively. While this suggests that both measures are capturing similar variation in the market yields, the magnitude of the DKW estimate is not consistent with the maximum of the admissible range.

At the five-year maturity, the DKW TIPS liquidity premium is above our maximum for about half of the sample period, and at the ten-year maturity their estimate is systematically outside the range. This might result from the fact that DKW attributes all variation in the TIPS-specific factor they identify to liquidity effects, leading to excessively large estimates of the TIPS liquidity premium. The impact this overstatement of the liquidity premium has on the decomposition of BEI rates is not entirely clear for the inflation risk premium. However, results from our CLR model exercise to be discussed later suggest that the DKW estimation

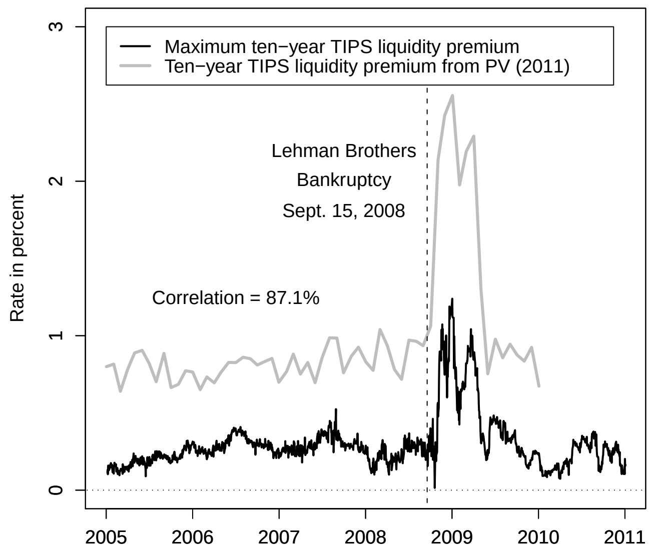

Figure 5: **Comparison to the TIPS Liquidity Premium from PV.**
Comparison of the model-independent maximum for the admissible liquidity premium in the ten-year TIPS yield to the estimated ten-year TIPS liquidity premium from PV.

of expected inflation may be too high because the DKW model overcorrects TIPS yields for liquidity differentials. Finally, it is worth noting that DKW are not likely to disagree with the fundamental assumptions that produce our model-independent maximum range because their approach relies on the same intuition, namely that TIPS are less liquid than nominal Treasuries, and they use the same GSW nominal and real Treasury yields.

PV take a different approach to estimating TIPS liquidity premiums in that they regress BEI on several measures of bond liquidity (including the off- over the on-the-run yield spread discussed previously) to generate estimates of the liquidity premium differential between TIPS and nominal yields. Figure 5 compares our maximum range for the ten-year TIPS liquidity premium to their estimated ten-year TIPS liquidity premium for the overlapping sample period.[18] The level of their liquidity premium is well above the maximum range, but they share a very high correlation. This suggests that, while regression-based estimates of liquidity premiums using economic fundamentals can be off in terms of the level of the premium, they may capture the time variation reasonably well.[19] As was the case for DKW, we emphasize that the analysis in PV is based on BEI rates from the same GSW databases used in the

[18]We are very thankful to Luis Viceira for sharing their data.

[19]Technically, this requires the additional assumption that the actual, unobserved TIPS liquidity premium located inside the range be highly correlated with the maximum of the range used in the comparison here.

construction of our upper bound for the TIPS liquidity premium. Hence, their estimated liquidity premium is comparable to our range.

To summarize, the information in nominal Treasury and real TIPS yields can be used in junction with inflation swap rates to derive a model-independent range for the admissible size of TIPS liquidity premiums. The constructed range relies on two assumptions, namely that TIPS and inflation swap contracts are each less liquid than nominal Treasuries, which we think are widely accepted and well documented in the literature. Furthermore, the range is highly correlated with existing estimates of TIPS liquidity premiums. Thus, with an upper bound for the liquidity disadvantage of TIPS in place, we can turn our attention to the benefit of TIPS, i.e., the saved inflation risk premiums. We address this subject in the following section.

4 Measuring the Inflation Risk Premium in Nominal Yields

While we use a model-independent measure of the liquidity effects in TIPS yields because researchers have yet to arrive to a consensus on their size, the practice of empirically extracting inflation risk premiums from breakeven inflation using term structure models is much more established within the bond pricing literature. To that end, we employ the model introduced in CLR, but use it in a novel fashion to exploit the information in the model-independent maximum range for the TIPS liquidity premium described in the previous section.

To begin the analysis, we first provide a brief theoretical discussion of how an arbitrage-free term structure model of nominal and real yields can decompose the difference between the two into inflation expectations and inflation risk premiums, where the latter is the excess yield a nominal bond has to pay for not providing inflation protection. We then describe the model introduced in CLR before proceeding to the empirical results.

4.1 Decomposing Breakeven Inflation

To begin, let the nominal and real stochastic discount factors be denoted by M_t^N and M_t^R, respectively. The requirement of no arbitrage enforces a consistency of pricing for any security over time. Specifically, the price of a nominal bond that pays one dollar at time $t+\tau$ and the price of a real bond that pays one unit of the consumption basket at time $t+\tau$ must satisfy

$$P_t^N(\tau) = E_t^P\left[\frac{M_{t+\tau}^N}{M_t^N}\right] \quad \text{and} \quad P_t^R(\tau) = E_t^P\left[\frac{M_{t+\tau}^R}{M_t^R}\right]. \tag{8}$$

Given their payment structure, the no-arbitrage condition also requires a consistency between the prices of nominal and real bonds such that the price of the consumption basket, denoted as the overall price level Π_t, is the ratio of the nominal and real stochastic discount factors:

$$\Pi_t = \frac{M_t^R}{M_t^N}. \tag{9}$$

We assume that the nominal and real stochastic discount factors have the standard dynamics $dM_t^N/M_t^N = -r_t^N dt - \Gamma_t' dW_t^P$ and $dM_t^R/M_t^R = -r_t^R dt - \Gamma_t' dW_t^P$, where r_t^N and r_t^R are the instantaneous, risk-free nominal and real rates of return, respectively, and Γ_t is a vector of premiums on the risks represented by W_t^P. Then, by Ito's lemma, $d\Pi_t = (r_t^N - r_t^R)\Pi_t dt$. Thus, with the absence of arbitrage, the instantaneous growth rate of the price level is equal to the difference between the instantaneous nominal and real risk-free rates. Now, it easily follows that we can express the price level at time t+τ as

$$\Pi_{t+\tau} = \Pi_t e^{\int_t^{t+\tau}(r_s^N - r_s^R)ds}. \tag{10}$$

The relationship between the yields and inflation expectations can be obtained by decomposing the price of the nominal bond as follows:

$$\begin{aligned} P_t^N(\tau) &= E_t^P\left[\frac{M_{t+\tau}^N}{M_t^N}\right] = E_t^P\left[\frac{M_{t+\tau}^R/\Pi_{t+\tau}}{M_t^R/\Pi_t}\right] = E_t^P\left[\frac{M_{t+\tau}^R}{M_t^R}\frac{\Pi_t}{\Pi_{t+\tau}}\right] \\ &= E_t^P\left[\frac{M_{t+\tau}^R}{M_t^R}\right] \times E_t^P\left[\frac{\Pi_t}{\Pi_{t+\tau}}\right] + cov_t^P\left[\frac{M_{t+\tau}^R}{M_t^R}, \frac{\Pi_t}{\Pi_{t+\tau}}\right] \\ &= P_t^R(\tau) \times E_t^P\left[\frac{\Pi_t}{\Pi_{t+\tau}}\right] \times \left(1 + \frac{cov_t^P\left[\frac{M_{t+\tau}^R}{M_t^R}, \frac{\Pi_t}{\Pi_{t+\tau}}\right]}{E_t^P\left[\frac{M_{t+\tau}^R}{M_t^R}\right] \times E_t^P\left[\frac{\Pi_t}{\Pi_{t+\tau}}\right]}\right). \end{aligned}$$

Converting this price into a yield-to-maturity using

$$y_t^N(\tau) = -\frac{1}{\tau}\ln P_t^N(\tau) \quad \text{and} \quad y_t^R(\tau) = -\frac{1}{\tau}\ln P_t^R(\tau),$$

the connection between nominal and real zero-coupon yields and expected inflation can be readily expressed as

$$y_t^N(\tau) = y_t^R(\tau) + \pi_t^e(\tau) + \phi_t(\tau), \tag{11}$$

where the market-implied rate of inflation expected at time t for the period from t to $t+\tau$ is[20]

$$\pi_t^e(\tau) = -\frac{1}{\tau}\ln E_t^P\left[\frac{\Pi_t}{\Pi_{t+\tau}}\right] = -\frac{1}{\tau}\ln E_t^P\left[e^{-\int_t^{t+\tau}(r_s^N - r_s^R)ds}\right]. \tag{12}$$

[20] Appendix A explains how the expected inflation is calculated within the CLR model.

The corresponding inflation risk premium is

$$\phi_t(\tau) = -\frac{1}{\tau}\ln\left(1 + \frac{cov_t^P\left[\frac{M_{t+\tau}^R}{M_t^R}, \frac{\Pi_t}{\Pi_{t+\tau}}\right]}{E_t^P\left[\frac{M_{t+\tau}^R}{M_t^R}\right] \times E_t^P\left[\frac{\Pi_t}{\Pi_{t+\tau}}\right]}\right). \quad (13)$$

This last equation highlights that the inflation risk premium can be positive or negative. It is positive if and only if

$$cov_t^P\left[\frac{M_{t+\tau}^R}{M_t^R}, \frac{\Pi_t}{\Pi_{t+\tau}}\right] < 0. \quad (14)$$

That is, the riskiness of nominal bonds relative to real bonds depends on the covariance between the real stochastic discount factor and inflation and is ultimately determined by investor preferences, but for our analysis we will not need to specify those.

4.2 The CLR Model of Nominal and Real Yields

As is clear from the theoretical section above, the market-implied expected inflation and the associated risk premium is measured by modeling the instantaneous nominal and real rates, r_t^N and r_t^R. The model we use for that purpose was first introduced in CLR and is centered around the arbitrage-free Nelson-Siegel (AFNS) model framework derived in Christensen, Diebold, and Rudebusch (2011), but modified to allow for stochastic yield volatility. The four factors that determine the rates are explained by the state vector $X_t = (L_t^N, S_t, C_t, L_t^R)$. L_t^N and L_t^R represent level effects in the nominal and real yield curves, respectively, while S_t and C_t represent slope and curvature effects common to both yield curves. The instantaneous nominal and real risk-free rates are defined by:

$$r_t^N = L_t^N + S_t, \quad (15)$$

$$r_t^R = L_t^R + \alpha^R S_t. \quad (16)$$

The differential scaling of real rates to the common slope (and indirectly the common curvature) factor is captured by the parameter α^R. To preserve the Nelson-Siegel factor loading structure in the nominal and real yield functions, the dynamics of the state variables under the pricing, or risk-neutral, probability measure (traditionally referred to as the Q-

measure) are, in general, given by the stochastic differential equations:[21]

$$\begin{pmatrix} dL_t^N \\ dS_t \\ dC_t \\ dL_t^R \end{pmatrix} = \begin{pmatrix} \kappa_{L^N}^Q & 0 & 0 & 0 \\ 0 & \lambda & -\lambda & 0 \\ 0 & 0 & \lambda & 0 \\ 0 & 0 & 0 & \kappa_{L^R}^Q \end{pmatrix} \left[\begin{pmatrix} \theta_{L^N}^Q \\ 0 \\ 0 \\ \theta_{L^R}^Q \end{pmatrix} - \begin{pmatrix} L_t^N \\ S_t \\ C_t \\ L_t^R \end{pmatrix} \right] dt$$
$$+ \begin{pmatrix} \sigma_{11} & 0 & 0 & 0 \\ 0 & \sigma_{22} & 0 & 0 \\ 0 & 0 & \sigma_{33} & 0 \\ 0 & 0 & 0 & \sigma_{44} \end{pmatrix} \begin{pmatrix} \sqrt{L_t^N} & 0 & 0 & 0 \\ 0 & \sqrt{1} & 0 & 0 \\ 0 & 0 & \sqrt{1} & 0 \\ 0 & 0 & 0 & \sqrt{L_t^R} \end{pmatrix} \begin{pmatrix} dW_t^{L^N,Q} \\ dW_t^{S,Q} \\ dW_t^{C,Q} \\ dW_t^{L^R,Q} \end{pmatrix}.$$

This structure implies that nominal yields are[22,23]

$$y_t^N(\tau) = g^N\left(\kappa_{L^N}^Q\right) L_t^N + \left(\frac{1-e^{-\lambda\tau}}{\lambda\tau}\right) S_t + \left(\frac{1-e^{-\lambda\tau}}{\lambda\tau} - e^{-\lambda\tau}\right) C_t - \frac{A^N(\tau;\kappa_{L^N}^Q)}{\tau}, \quad (17)$$

while real yields are given by

$$y_t^R(\tau) = g^R\left(\kappa_{L^R}^Q\right) L_t^R + \alpha^R \left(\frac{1-e^{-\lambda\tau}}{\lambda\tau}\right) S_t + \alpha^R \left(\frac{1-e^{-\lambda\tau}}{\lambda\tau} - e^{-\lambda\tau}\right) C_t - \frac{A^R(\tau;\kappa_{L^R}^Q)}{\tau}. \quad (18)$$

To link the risk-neutral and real-world dynamics of the state variables, we use the extended affine risk premium specification introduced by Cheridito, Filipović, and Kimmel (2007). The maximally flexible affine specification of the P-dynamics is thus

$$\begin{pmatrix} dL_t^N \\ dS_t \\ dC_t \\ dL_t^R \end{pmatrix} = \begin{pmatrix} \kappa_{11}^P & 0 & 0 & \kappa_{14}^P \\ \kappa_{21}^P & \kappa_{22}^P & \kappa_{23}^P & \kappa_{24}^P \\ \kappa_{31}^P & \kappa_{32}^P & \kappa_{33}^P & \kappa_{34}^P \\ \kappa_{41}^P & 0 & 0 & \kappa_{44}^P \end{pmatrix} \left[\begin{pmatrix} \theta_1^P \\ \theta_2^P \\ \theta_3^P \\ \theta_4^P \end{pmatrix} - \begin{pmatrix} L_t^N \\ S_t \\ C_t \\ L_t^R \end{pmatrix} \right] dt$$
$$+ \begin{pmatrix} \sigma_{11} & 0 & 0 & 0 \\ 0 & \sigma_{22} & 0 & 0 \\ 0 & 0 & \sigma_{33} & 0 \\ 0 & 0 & 0 & \sigma_{44} \end{pmatrix} \begin{pmatrix} \sqrt{L_t^N} & 0 & 0 & 0 \\ 0 & \sqrt{1} & 0 & 0 \\ 0 & 0 & \sqrt{1} & 0 \\ 0 & 0 & 0 & \sqrt{L_t^R} \end{pmatrix} \begin{pmatrix} dW_t^{L^N,P} \\ dW_t^{S,P} \\ dW_t^{C,P} \\ dW_t^{L^R,P} \end{pmatrix}.$$

To keep the model arbitrage-free, the two level factors must be prevented from hitting the lower zero-boundary. This positivity requirement is ensured by imposing the Feller condition

[21] We follow CLR and limit our focus to specifications with a diagonal Σ volatility matrix.

[22] Analytical formulas for $g^N\left(\kappa_{L^N}^Q\right)$, $g^R\left(\kappa_{L^R}^Q\right)$, $A^N(\tau;\kappa_{L^N}^Q)$, and $A^R(\tau;\kappa_{L^R}^Q)$ are provided in CLR.

[23] In our implementation, we follow CLR and fix $\kappa_{L^N}^Q = \kappa_{L^R}^Q = 10^{-7}$ to get a close approximation to the uniform level factor loading in the original Gaussian AFNS model.

on their factor dynamics under both probability measures, i.e.,

$$\kappa_{11}^P\theta_1^P+\kappa_{14}^P\theta_4^P > \frac{1}{2}\sigma_{11}^2, \quad 10^{-7}\cdot\theta_{L^N}^Q > \frac{1}{2}\sigma_{11}^2, \quad \kappa_{41}^P\theta_1^P+\kappa_{44}^P\theta_4^P > \frac{1}{2}\sigma_{44}^2, \quad \text{and} \quad 10^{-7}\cdot\theta_{L^R}^Q > \frac{1}{2}\sigma_{44}^2.$$

Furthermore, to have well-defined processes for L_t^N and L_t^R, the sign of the effect that these two factors have on each other must be positive, which requires the restrictions that

$$\kappa_{14}^P \leq 0 \quad \text{and} \quad \kappa_{41}^P \leq 0.$$

These conditions ensure that the two square-root processes will be non-negatively correlated.

Finally, the model is estimated with the standard Kalman filter, see CLR for details.

4.3 Model Specification and Parameter Sensitivity

Our model application exercise follows CLR and uses nominal yields with three-month, six-month, one-, two-, three-, five-, seven-, and ten-year maturities and the six real yields with five- to ten-year maturities. For the nominal yields, we use a sample period from January 3, 1995 to December 31, 2010 with a total of 3,972 observations. For the real yields, the sample begins on January 4, 2005 so that it aligns the TIPS sample with the inflation swap data, resulting in 1,482 observations.

We use the admissible range for the TIPS liquidity premium derived in Section 3 to correct for the liquidity effects in TIPS yields prior to model estimation. This allows us to study the sensitivity of the CLR model output to the assumed level of the TIPS liquidity premium. The bound provided by the admissible range of TIPS liquidity premiums converts into a range of admissible estimates for the output produced by the model. Thus, this exercise is telling of the relationship between TIPS liquidity premiums and the BEI decomposition as captured through the lens of the model. Since the specification of the P-dynamics is an important element in determining a model's decomposition of BEI rates into inflation expectations and associated risk premiums, we conduct a careful evaluation of various model specifications.

We start the model selection from the unrestricted specification of the mean-reversion matrix K^P, and pare down this matrix using a general-to-specific strategy that restricts the least significant parameter (as measured by the ratio of the parameter value to its standard error) to zero and then re-estimate the model. This strategy of eliminating the least significant coefficients continues to the most parsimonious specification, which has a diagonal K^P matrix. In order to select our preferred specification, we use the Bayes information criterion, which is commonly used for model selection and is defined as BIC $= -2\log L + k\log T$, where T is the number of data observations (see e.g., Harvey 1989).[24]

[24]We have 3,972 nominal yield and 1,482 real yield daily observations. We interpret T as referring to the longest data series and fix it at 3,972.

The first case we consider assumes that there are no liquidity premiums priced into TIPS yields. In this case we use the GSW yields without any corrections, which is implicitly assumed in papers like Adrian and Wu (2010) and CLR. We then evaluate the model specification assuming that there are no liquidity risk premiums in the inflation swap rates. That is, the model estimations are performed on TIPS yields that have been maximally corrected for liquidity differentials as described in Section 3. The preferred specification according to the BIC has a mean-reversion matrix K^P under both assumptions given by

$$K^P_{BIC} = \begin{pmatrix} \kappa^P_{11} & 0 & 0 & 0 \\ 0 & \kappa^P_{22} & \kappa^P_{23} & 0 \\ 0 & 0 & \kappa^P_{33} & 0 \\ 0 & 0 & 0 & \kappa^P_{44} \end{pmatrix}. \tag{19}$$

This matrix is identical to the specification preferred in CLR.[25]

Most parameter estimates and corresponding standard deviations for the P-dynamics are practically indistinguishable and not significantly different across the two assumptions about the size of TIPS liquidity premiums, except for a markedly higher estimated value of κ^P_{44} when TIPS are assumed subject to the maximum liquidity premium. This suggests a very limited impact on the estimated P-dynamics of correcting for TIPS liquidity premiums. The most notable difference relates to the estimated Q-dynamics where the estimated value of α^R is raised from 0.3860 to 0.5890 in response to a lower estimated path for L^R_t as TIPS yields are lower when maximally corrected for liquidity premiums. As a consequence, the variation of the common slope and curvature factors plays a bigger role for real TIPS yields in this case, which ultimately affects the model's decomposition of BEI into expected inflation and inflation risk premiums as we will see later. Finally, Table 3 reports the summary statistics of the fitted errors across maturities based on the two estimations and provides further documentation of the accuracy of the model under either assumption.[26]

4.4 Finite-sample Bias Correction

The main limitation of the model results discussed so far is one that generally plagues the estimation of any dynamic term structure model. Because interest rates are highly persistent, empirical autoregressive models, including dynamic term structure models like ours, suffer from potentially substantial finite-sample estimation bias. Specifically, model estimates will generally be biased toward a dynamic system that displays much less persistence than the

[25] The full details of the specification evaluations and the estimated preferred specification under both the zero and maximum TIPS liquidity premium assumptions are provided in Appendix B.

[26] As expected, the fit of the nominal yields is practically identical and the dispersion in the fit of the real yields is of modest magnitude.

Maturity in months	No TIPS liquidity premium			Maximum TIPS liquidity premium		
Nom. yields	Mean	RMSE	$\widehat{\sigma}_\varepsilon(\tau_i)$	Mean	RMSE	$\widehat{\sigma}_\varepsilon(\tau_i)$
3	0.81	19.58	19.57	0.59	19.59	19.58
6	-0.12	8.25	8.25	-0.24	8.26	8.25
12	0.00	0.00	0.00	0.00	0.00	0.00
24	0.44	1.54	1.54	0.48	1.55	1.55
36	0.00	0.00	0.00	0.00	0.00	0.00
60	-0.27	1.28	1.34	-0.30	1.31	1.38
84	0.21	0.69	1.05	0.25	0.81	1.15
120	-1.26	4.51	4.65	-1.20	4.52	4.66
TIPS yields	Mean	RMSE	$\widehat{\sigma}_\varepsilon(\tau_i)$	Mean	RMSE	$\widehat{\sigma}_\varepsilon(\tau_i)$
60	-0.88	17.93	17.98	-0.17	14.81	14.95
72	-0.05	10.70	10.77	-0.14	8.61	8.85
84	0.30	4.90	5.04	0.08	3.82	4.36
96	0.09	0.60	1.31	0.16	2.59	3.27
108	-0.67	4.26	4.40	-0.18	5.77	5.98
120	-1.96	8.01	8.09	-1.54	9.63	9.67
Max log L	253,073.70			252,343.70		

Table 3: **Summary Statistics of the Fitted Errors.**
The mean fitted errors, the root mean squared fitted errors (RMSE), and the estimated fitted error standard deviations from the preferred specification of the CLR model of nominal and real yields are shown with the minimum and the maximum liquidity correction of TIPS yields prior to model estimation. All numbers are measured in basis points. The nominal yields cover the period from January 3, 1995, to December 31, 2010, while the real TIPS yields cover the period from January 4, 2005, to December 31, 2010.

true process (so estimates of the real-world mean-reversion matrix, K^P, are upward biased). In our context of decomposing BEI into expected inflation and the associated inflation risk premium, such bias tends to imply that the expected inflation is too stable and the inflation risk premium too volatile. For our sample period, which is dominated by the financial crisis of 2008-2009 and the ensuing low inflation outcomes from sharp declines in energy prices, the upward bias manifests itself as excessively high estimates of the expected inflation component of BEI and corresponding excessively low estimates of the inflation risk premium. This could lead to serious bias in the assessment of the benefit of issuing TIPS versus regular Treasury securities. For these reasons a correction for the finite-sample bias is critical to our analysis.

In a recent paper, Bauer, Rudebusch, and Wu (2012) provide a complete discussion of the finite-sample bias in empirical affine Gaussian term structure models. However, as our model is non-Gaussian, we cannot follow their exact procedure as it relies on the ability to rapidly estimate the model, which is feasible for affine Gaussian term structure models due to the work of Joslin, Singleton, and Zhu (2011) and Hamilton and Wu (2012), but not for non-Gaussian

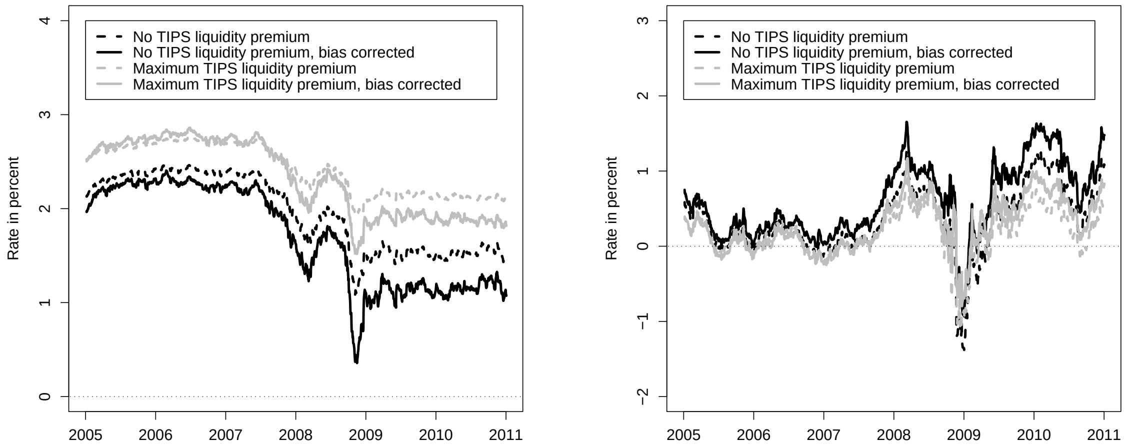

(a) Inflation expectations.

(b) Inflation risk premiums.

Figure 6: **Range of Ten-Year Inflation Expectations and Risk Premiums due to TIPS Yield Liquidity Corrections.**
Illustration of the range of viable ten-year inflation expectations and associated inflation risk premiums due to TIPS yield liquidity corrections based on the CLR model with and without correction for the parameter estimation finite-sample bias. In each chart, the black lines represent the case with no TIPS liquidity premiums, while the grey lines represent the reverse case where TIPS yields contain the maximum liquidity premium.

models like ours as the underlying statistical limit results do not apply. Instead, we make an adaptation of the approach of Bauer et al. (2012) to non-Gaussian models that preserves their median-based bias correction, but relies on a parameter-by-parameter adjustment.[27]

4.5 Estimates of Inflation Expectations and Risk Premiums

In this section we analyze the impact the assumed level of the liquidity premium in TIPS yields has on the CLR model decomposition of BEI rates into expected inflation and the associated risk premium. We also illustrate the effect on the decomposition from the bias correction.

Figure 6(a) shows the wedge of the expected inflation obtained from the two extreme cases in the assumed level of the TIPS liquidity premium, both with and without bias correction.[28] We note the significant negative effect on the estimates of the expected inflation component

[27] The full details of the bias correction are provided in Appendix C along with the complete table of the bias-corrected parameters under both assumptions about the size of TIPS liquidity premiums.

[28] In order to correct each estimate, we first use the bias-corrected parameter set along with the observed yield data as input into the Kalman filter to obtain the corresponding filtered factor paths. In a second step, the expected inflation is calculated as described in Appendix A and deducted from the fitted BEI (difference between fitted nominal and real yields) from the Kalman filter to produce the bias-corrected estimates of the inflation risk premium.

from the bias correction. Since the inflation risk premium is determined as a residual by deducting the expected inflation from BEI, there is a similar effect of the bias correction on the estimated inflation risk premiums, but with the opposite sign as shown in Figure 6(b). This underscores the importance of addressing the finite-sample bias in the parameter estimation when it comes to assessing the benefit of TIPS issuance over issuing regular Treasury securities. Second, we note the large difference in the estimates of the expected inflation when we move from the case with no liquidity correction of TIPS yields before estimation to the case with a maximum correction for TIPS liquidity premiums. The latter case produces estimates that are systematically higher as BEI is larger under that assumption.

For the range of estimated ten-year inflation risk premiums shown in Figure 6(b), we note that the differences are typically much smaller and with no systematic sign, indicating that the assumed level of the TIPS liquidity premium has a much smaller effect on estimates of the inflation risk premium. Also worth noting are the occasional simultaneously negative values in both series, suggesting that inflation risk premiums can be negative even when TIPS yields are maximally corrected for liquidity premiums.[29] Theoretically, a number of circumstances could lead to this result. One possibility is that the CPI figure used in the indexation of TIPS tends to overstate the true inflation because substitution effects are not adequately accounted for. Another explanation is that the marginal investor who determines the tradeoff between nominal and real yields might have a personal price index different from headline CPI. If so, TIPS only provide a partial hedge for inflation risk and, as a result, the investor demands a premium for being exposed to an imperfect hedge. In both cases, TIPS yields can be higher than they otherwise might be absent such risk effects. Assuming inflation expectations do not change, which is reasonable as such risk effects should not alter investors' outlook for inflation, this would result in negative estimates for the inflation risk premiums.

To recapitulate the construction of the *minimum liquidity-adjusted inflation risk premium* shown in Figure 1, we first take the minimum of the estimated range of inflation risk premiums (shown in Figure 6(b) for the ten-year maturity) for each observation date and at each maturity. We then deduct the corresponding maximally admissible TIPS liquidity premium introduced in Section 3. Lastly, we adjust the result for the fact that the U.S. Treasury issues regular Treasury securities at low, so-called on-the-run Treasury yields, which differ from the seasoned, or so-called off-the-run, Treasury yields we use in model estimation. The final series represent the most conservative assessment of the economic benefit of TIPS over nominal Treasuries that our approach allows us to generate. Despite its conservative nature, the measure is still positive on average over our sample. This suggests that the U.S. Treasury could reduce its debt servicing costs by increasing its TIPS program. Our own calculation

[29] Chernov and Mueller (2011) also report occasional negative inflation risk premium estimates in the most recent part of their sample that covers the 1971-2008 period.

assuming a constant debt level and an increase in the ratio of TIPS to one third showed that savings of the order of \$24 billion over the next ten years could be obtained by such a shift in debt management policy.

5 Conclusion

In this paper, we provide a measure we refer to as the minimum liquidity-adjusted inflation risk premium in order to quantitatively assess the economic benefit to the U.S. Treasury of issuing TIPS. It is constructed by taking the smallest inflation risk premiums produced by the CLR model and subtracting the maximum admissible TIPS liquidity premium in addition to the on-the-run premium for nominal Treasury bonds. The measure itself is conservative in nature, yet it still indicates that the U.S. Treasury has benefitted from issuing TIPS since the beginning of our sample period in 2005. Provided other considerations are secondary, our cost-benefit analysis suggests that the U.S. Treasury could save billions of tax payer dollars over the next ten years by expanding its issuance of TIPS relative to nominal Treasuries.

To assess the costs of issuing TIPS, which is their liquidity disadvantage relative to regular Treasury securities, we introduce a model-independent maximum range for the admissible liquidity premium in TIPS using the joint pricing information in the markets for nominal Treasury securities, TIPS, and inflation swaps. The theoretical assumptions we make in constructing the range are simple, realistic, and consistent with findings from studies on the microstructure of these three markets. Furthermore, the range is highly correlated with other widely accepted proxies of bond liquidity.

As a consensus on the precise level of the TIPS liquidity premium has yet to develop, the range also provides a means by which current and future models may be assessed. To exemplify, we compare the range to the TIPS liquidity premiums estimated in DKW and PV and find the levels of their estimates to be largely outside the admissible range. The strategies employed by DKW and PV seem to be problematic. By attributing variation in TIPS yields largely to liquidity, they overestimate its role in the pricing of TIPS yields. Thus for the purpose of our cost-benefit assessment, our maximum range is ideal since it does not depend on model-generated results.

To assess the benefits of TIPS, which is the inflation risk premium that nominal Treasuries are penalized with for *not* providing inflation protection, we use the CLR model of nominal and real yields. We follow the most recent literature and adjust the estimated model parameters for finite-sample bias. We also demonstrate the sensitivity of the CLR model decomposition of BEI to variations in the TIPS liquidity correction by estimating the model at both extremes of the maximum range for the TIPS liquidity premium. This exercise can be done with any term structure model to assess its robustness under different assumptions about liquidity effects in TIPS pricing.

Our analysis provides insight into the economic costs and benefits of the U.S. Treasury's TIPS issuance with the most concrete result being that it is very likely that the U.S. Treasury gets a cost reduction in the form of saved inflation risk premiums from TIPS that surpasses the liquidity penalty investors price into their rates. We roughly calculate that these savings could be augmented by tens of billions of dollars if the Treasury were to allocate a larger portion of its portfolio to TIPS over the next ten years and, given the conservative nature of the assumptions in the construction of our metric, these figures likely understate the benefit of issuing TIPS.

Appendices

A). **Calculation of the Expected Inflation and the Inflation Risk Premium**

From the theoretical discussion in Section 4.1, it follows that the model-implied expected inflation is given by

$$\pi_t^e(\tau) = -\frac{1}{\tau} \ln E_t^P \Big[e^{-\int_t^{t+\tau} (r_s^N - r_s^R) ds} \Big]. \tag{20}$$

For the affine framework, we are working within, it holds that

$$E_t^P \Big[e^{-\int_t^{t+\tau} (r_s^N - r_s^R) ds} \Big] = \exp(B^\pi(\tau)' X_t + A^\pi(\tau)), \tag{21}$$

where $B_1^\pi(t,T)$, $B_2^\pi(t,T)$, $B_3^\pi(t,T)$, and $B_4^\pi(t,T)$ are the unique solutions to the following system of ODEs

$$\begin{pmatrix} \frac{dB_1^\pi(t,T)}{dt} \\ \frac{dB_2^\pi(t,T)}{dt} \\ \frac{dB_3^\pi(t,T)}{dt} \\ \frac{dB_4^\pi(t,T)}{dt} \end{pmatrix} = \begin{pmatrix} 1 \\ 1-\alpha^R \\ 0 \\ -1 \end{pmatrix} + \begin{pmatrix} \kappa_{11}^P & 0 & 0 & 0 \\ 0 & \kappa_{22}^P & 0 & 0 \\ 0 & \kappa_{23}^P & \kappa_{33}^P & 0 \\ 0 & 0 & 0 & \kappa_{44}^P \end{pmatrix} \begin{pmatrix} B_1^\pi(t,T) \\ B_2^\pi(t,T) \\ B_3^\pi(t,T) \\ B_4^\pi(t,T) \end{pmatrix}$$

$$-\frac{1}{2}\sum_{j=1}^{4} \left[\begin{pmatrix} \sigma_{11} & 0 & 0 & 0 \\ 0 & \sigma_{22} & 0 & 0 \\ 0 & 0 & \sigma_{33} & 0 \\ 0 & 0 & 0 & \sigma_{44} \end{pmatrix} \begin{pmatrix} (B_1^\pi)^2 & B_1^\pi B_2^\pi & B_1^\pi B_3^\pi & B_1^\pi B_4^\pi \\ B_1^\pi B_2^\pi & (B_2^\pi)^2 & B_2^\pi B_3^\pi & B_2^\pi B_4^\pi \\ B_1^\pi B_3^\pi & B_2^\pi B_3^\pi & (B_3^\pi)^2 & B_3^\pi B_4^\pi \\ B_1^\pi B_4^\pi & B_2^\pi B_4^\pi & B_3^\pi B_4^\pi & (B_4^\pi)^2 \end{pmatrix} \begin{pmatrix} \sigma_{11} & 0 & 0 & 0 \\ 0 & \sigma_{22} & 0 & 0 \\ 0 & 0 & \sigma_{33} & 0 \\ 0 & 0 & 0 & \sigma_{44} \end{pmatrix} \right]_{j,j} (\delta^j)',$$

where δ is given by[30]

$$\delta = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

and the boundary condition is $B^\pi(T,T) = 0$, while $A^\pi(t,T)$ is the unique solutions to the following ODE

$$\frac{dA^\pi(t,T)}{dt} = -B^\pi(t,T)' K^P \theta^P - \frac{1}{2}\sigma_{22}^2 B_2^\pi(t,T)^2 - \frac{1}{2}\sigma_{33}^2 B_3^\pi(t,T)^2, \quad A^\pi(T,T) = 0.$$

Here, K^P, θ^P, and Σ reflect the P-dynamics of the state variables. Now, the ODEs above are solved with a standard fourth order Runge-Kutta method.

By Equation (11), we can easily calculate the inflation risk premiums, once we have the corresponding expected inflation, $\pi_t^e(\tau)$, and the fitted nominal and real yields from Equations (17) and (18), respectively.

B). **Model Specification Evaluation**

We start the model selection from the unrestricted specification of the mean-reversion matrix K^P, which provides the maximally admissible flexibility in fitting the data given our restrictions on the state variable volatility structure. We then pare down this matrix using a general-to-specific strategy that restricts the least significant parameter (as measured by the ratio of the parameter value to its standard error) to zero and then re-estimate the model. This strategy of eliminating the least significant coefficients continues to the most parsimonious specification, which has a diagonal K^P matrix.

Each estimated specification is listed with its maximum log likelihood, its number of estimated parameters (k), and the p-value from a likelihood ratio test of the hypothesis that it differs from the specification with

[30] In the equation above, δ^j denotes the jth row of δ.

Alternative specifications	Goodness-of-fit statistics			
	Max $\log L$	k	p-value	BIC
(1) Unrestricted K^P	253,078.4	38	n.a.	-505,841.9
(2) $\kappa^P_{14} = 0$	253,078.4	37	1.0000	-505,850.2
(3) $\kappa^P_{14} = \kappa^P_{41} = 0$	253,078.4	36	1.0000	-505,858.5
(4) $\kappa^P_{14} = \kappa^P_{41} = \kappa^P_{24} = 0$	253,077.9	35	0.3173	-505,865.8
(5) $\kappa^P_{14} = \ldots = \kappa^P_{32} = 0$	253,076.5	34	0.0943	-505,871.2
(6) $\kappa^P_{14} = \ldots = \kappa^P_{31} = 0$	253,076.3	33	0.5271	-505,879.1
(7) $\kappa^P_{14} = \ldots = \kappa^P_{21} = 0$	253,074.4	32	0.0513	-505,883.6
(8) $\kappa^P_{14} = \ldots = \kappa^P_{34} = 0$	253,073.7	31	0.2367	**-505,890.5**
(9) $\kappa^P_{14} = \ldots = \kappa^P_{23} = 0$	253,067.3	30	0.0003	-505,886.0

Table 4: **Evaluation of Alternative Specifications of the CLR Model With No TIPS Liquidity Premiums.**
Nine alternative estimated specifications of the CLR model of nominal and real Treasury bond yields are evaluated assuming no liquidity premiums in TIPS yields. Each specification is listed with its maximum log likelihood (Max $\log L$), number of parameters (k), the p-value from a likelihood ratio test of the hypothesis that the specification differs from the one directly above that has one more free parameter. The Bayes information criterion (BIC) is also reported, and its minimum value is given in boldface.

one more parameter—that is, comparing specification (s) with specification ($s-1$). We also report the Bayes information criterion commonly used for model selection, which is defined as BIC $= -2\log L + k \log T$, where T is the number of data observations (see e.g., Harvey 1989).[31]

We start with the case where it is assumed that there are no liquidity premiums priced into TIPS yields. In that case we use the GSW yields without any corrections in the model estimation. The result of the model selection in this case is shown in Table 4. The preferred specification according to BIC has a mean-reversion matrix K^P given by

$$K^P_{BIC} = \begin{pmatrix} \kappa^P_{11} & 0 & 0 & 0 \\ 0 & \kappa^P_{22} & \kappa^P_{23} & 0 \\ 0 & 0 & \kappa^P_{33} & 0 \\ 0 & 0 & 0 & \kappa^P_{44} \end{pmatrix}.$$

This is identical to the specification preferred in CLR. Table 5 contains the estimated parameters in this case.

In Table 6 we evaluate various model specifications assuming that there are no liquidity risk premiums in the inflation swap rates, that is, the model estimations are performed on TIPS yields that have been maximally corrected for liquidity differentials as described in Section 3. Under this assumption, the preferred specification of the mean-reversion matrix is again given by

$$K^P_{BIC} = \begin{pmatrix} \kappa^P_{11} & 0 & 0 & 0 \\ 0 & \kappa^P_{22} & \kappa^P_{23} & 0 \\ 0 & 0 & \kappa^P_{33} & 0 \\ 0 & 0 & 0 & \kappa^P_{44} \end{pmatrix}.$$

Table 7 contains the estimated parameters in this case.

[31]We have 3,972 nominal yield and 1,482 real yield daily observations. We interpret T as referring to the longest data series and fix it at 3,972.

K^P	$K^P_{\cdot,1}$	$K^P_{\cdot,2}$	$K^P_{\cdot,3}$	$K^P_{\cdot,4}$	θ^P		Σ
$K^P_{1,\cdot}$	1.0188	0	0	0	0.0421	σ_{11}	0.0634
	(0.2481)				(0.0035)		(0.0003)
$K^P_{2,\cdot}$	0	0.6073	-0.5678	0	-0.0110	σ_{22}	0.0126
		(0.1610)	(0.1285)		(0.0131)		(0.0001)
$K^P_{3,\cdot}$	0	0	0.7530	0	-0.0065	σ_{33}	0.0315
			(0.2168)		(0.0114)		(0.0004)
$K^P_{4,\cdot}$	0	0	0	1.4246	0.0148	σ_{44}	0.0554
				(0.2281)	(0.0017)		(0.0004)

Table 5: **Parameter Estimates for the Preferred Specification of the CLR Model With No TIPS Liquidity Premiums.**
The estimated parameters of the K^P matrix, θ^P vector, and diagonal Σ matrix are shown for the preferred specification of the CLR model when there are no liquidity premiums in TIPS yields. The Q-related parameters are estimated at: $\lambda = 0.5980$ (0.0011), $\alpha^R = 0.3860$ (0.0032), $\theta^Q_{L^N} = 32{,}600$ (18.09), and $\theta^Q_{L^R} = 15{,}368$ (32.50). The numbers in parentheses are the estimated parameter standard deviations. The maximum log likelihood value is 253,073.70.

Alternative specifications	Goodness-of-fit statistics			
	Max log L	k	p-value	BIC
(1) Unrestricted K^P	252,354.8	38	n.a.	-504,394.7
(2) $\kappa^P_{24} = 0$	252,354.7	37	0.6547	-504,402.8
(3) $\kappa^P_{24} = \kappa^P_{32} = 0$	252,351.5	36	0.0114	-504,404.7
(4) $\kappa^P_{24} = \kappa^P_{32} = \kappa^P_{41} = 0$	252,351.3	35	0.5271	-504,412.6
(5) $\kappa^P_{24} = \ldots = \kappa^P_{21} = 0$	252,349.6	34	0.0652	-504,417.4
(6) $\kappa^P_{24} = \ldots = \kappa^P_{14} = 0$	252,348.8	33	0.2059	-504,424.1
(7) $\kappa^P_{24} = \ldots = \kappa^P_{31} = 0$	252,346.9	32	0.0513	-504,428.6
(8) $\kappa^P_{24} = \ldots = \kappa^P_{34} = 0$	252,343.7	31	0.0114	**-504,430.5**
(9) $\kappa^P_{24} = \ldots = \kappa^P_{32} = 0$	252,337.1	30	0.0003	-504,425.6

Table 6: **Evaluation of Alternative Specifications of the CLR Model With the Maximum TIPS Liquidity Premiums.**
Nine alternative estimated specifications of the CLR model of nominal and real Treasury bond yields are evaluated for the case when TIPS yields are assumed to contain the maximum liquidity premium. Each specification is listed with its maximum log likelihood (Max log L), number of parameters (k), the p-value from a likelihood ratio test of the hypothesis that the specification differs from the one directly above that has one more free parameter. The Bayes information criterion (BIC) is also reported, and its minimum value is given in boldface.

C). The Simulations-Based Finite-Sample Bias Correction

In this appendix, we detail the classic parametric bootstrap simulation exercise we use to correct for the finite-sample bias of the parameter estimates in the CLR model. As in the model estimation, we consider both extremes: (1) Assuming no TIPS liquidity premiums and (2) assuming the maximally admissible TIPS liquidity premium.

To begin the exercise, we consider the model parameters estimated as of December 31, 2010 and reported in Tables 5 and 7, respectively, as the "true" parameters in each case. In the next step, we use the model to

K^P	$K^P_{\cdot,1}$	$K^P_{\cdot,2}$	$K^P_{\cdot,3}$	$K^P_{\cdot,4}$	θ^P		Σ
$K^P_{1,\cdot}$	0.8707	0	0	0	0.0436	σ_{11}	0.0573
	(0.2257)				(0.0036)		(0.0003)
$K^P_{2,\cdot}$	0	0.5858	-0.5610	0	-0.0124	σ_{22}	0.0122
		(0.1857)	(0.1351)		(0.0136)		(0.0001)
$K^P_{3,\cdot}$	0	0	0.7531	0	-0.0074	σ_{33}	0.0320
			(0.1681)		(0.0117)		(0.0003)
$K^P_{4,\cdot}$	0	0	0	2.1022	0.0129	σ_{44}	0.0609
				(0.2299)	(0.0011)		(0.0003)

Table 7: **Parameter Estimates for the Preferred Specification of the CLR Model With the Maximum TIPS Liquidity Premiums.**
The estimated parameters of the K^P matrix, θ^P vector, and diagonal Σ matrix are shown for the preferred specification of the CLR model when TIPS yields are assumed to contain the maximum liquidity premium. The Q-related parameters are estimated at: $\lambda = 0.5937$ (0.0010), $\alpha^R = 0.5890$ (0.0030), $\theta^Q_{L^N} = 30{,}120$ (18.48), and $\theta^Q_{L^R} = 18{,}559$ (32.37). The numbers in parentheses are the estimated parameter standard deviations. The maximum log likelihood value is 252,343.70.

simulate $N = 100$ artificial data sets identical to the actual sample of nominal and real U.S. Treasury yields. Specifically, (i) the number of observation dates and the time in between observations are identical to the original sample, (ii) the yield maturities on each observation date are identical to those in the original sample, and (iii) i.i.d. errors are added with a maturity-specific standard deviation given by the estimated values of $\sigma_\varepsilon(\tau)$ from the original sample (reported in the third and sixth column of Table 3, respectively). Third, we use the $N = 100$ artificial data samples as input into the original Kalman filter estimation, whereby we obtain $N = 100$ alternative optimal parameter sets. Importantly, if the model is true, these alternative parameter sets are each statistically just as likely as the "true" parameter set we got from the original data set. Thus, for each parameter, we can rank the estimated values and take out the median. It is worth noting that this approach implicitly handles any correlation there might be between various parameter estimates, in particular joint outliers naturally occur.

The details of the simulation of the factor paths are provided in the following. The continuous-time P-dynamics are, in general, given by

$$dX_t = K^P(\theta^P - X_t)dt + \Sigma D(X_t)dW_t^P.$$

For both restricted square-root processes and unconstrained processes we approximate the continuous-time process using the Euler approximation.[32] To exemplify, for a restricted square-root process,

$$dX_t^i = \kappa_{ii}^P(\theta_i^P - X_t^i)dt + \kappa_{ij}^P(\theta_j^P - X_t^j)dt + \sigma_{ii}\sqrt{X_t^i}dW_t^{P,i},$$

the algorithm is

$$X_t^i = X_{t-1}^i + \kappa_{ii}^P(\theta_i^P - X_{t-1}^i)\Delta t + \kappa_{ij}^P(\theta_j^P - X_{t-1}^j)\Delta t + \sigma_{ii}\sqrt{X_{t-1}^i}\sqrt{\Delta t}z_t^i, \quad z_t^i \sim N(0,1).$$

We fix Δt at a uniform value of 0.0001, which is equivalent to approximately 27 shocks to the X_t^i-process via the Brownian motion per day. As Feller conditions and other non-negativity requirements are imposed in the estimations performed with the observed bond yields, the parameter sets used in the simulations naturally sat-

[32] Thompson (2008) is a recent example.

Parameter	1^{st} iteration		2^{nd} iteration		3^{rd} iteration		4^{th} iteration	
	Estimate	Median	Value	Median	Value	Median	Value	Median
κ_{11}^P	**1.0188**	1.2712	0.7664	0.9441	0.8412	1.0865	0.7735	1.0748
κ_{22}^P	**0.6073**	0.7498	0.4648	0.5745	0.4976	0.6033	0.5016	0.6076
κ_{23}^P	**-0.5678**	-0.6007	-0.5348	-0.5871	-0.5155	-0.5377	-0.5456	-0.5919
κ_{33}^P	**0.7530**	1.0178	0.4883	0.6139	0.6273	0.8132	0.5672	0.7958
κ_{44}^P	**1.4246**	1.9298	0.9193	1.4826	0.8613	1.5144	0.7716	1.3592
σ_{11}	**0.0634**	0.0637	0.0631	0.0634	0.0632	0.0635	0.0631	0.0633
σ_{22}	**0.0126**	0.0126	0.0126	0.0126	0.0126	0.0126	0.0126	0.0126
σ_{33}	**0.0315**	0.0314	0.0315	0.0316	0.0315	0.0314	0.0315	0.0314
σ_{44}	**0.0554**	0.0553	0.0556	0.0557	0.0554	0.0555	0.0554	0.0554
θ_1^P	**0.0421**	0.0428	0.0413	0.0411	0.0423	0.0423	0.0422	0.0420
θ_2^P	**-0.0110**	-0.0126	-0.0095	-0.0103	-0.0102	-0.0081	-0.0131	-0.0132
θ_3^P	**-0.0065**	-0.0062	-0.0067	-0.0047	-0.0085	-0.0062	-0.0087	-0.0098
θ_4^P	**0.0148**	0.0145	0.0151	0.0152	0.0147	0.0143	0.0152	0.0145
λ	**0.5980**	0.5979	0.5981	0.5981	0.5979	0.5980	0.5979	0.5982
α^R	**0.3860**	0.3874	0.3847	0.3846	0.3861	0.3828	0.3893	0.3860
θ_1^Q	**32,600**	32,683	32,517	32,634	32,484	32,623	32,461	32,585
θ_4^Q	**15,368**	15,341	15,396	15,511	15,253	15,398	15,223	15,367
σ_ε^N(3m)	**0.0020**	0.0020	0.0020	0.0020	0.0020	0.0020	0.0020	0.0020
σ_ε^N(6m)	**0.0008**	0.0008	0.0008	0.0008	0.0008	0.0008	0.0008	0.0008
σ_ε^N(1yr)	**0.0000**	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
σ_ε^N(2yr)	**0.0002**	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002
σ_ε^N(3yr)	**0.0000**	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
σ_ε^N(5yr)	**0.0001**	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001
σ_ε^N(7yr)	**0.0001**	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001
σ_ε^N(10yr)	**0.0005**	0.0005	0.0005	0.0005	0.0005	0.0005	0.0005	0.0005
σ_ε^R(5yr)	**0.0018**	0.0018	0.0018	0.0018	0.0018	0.0018	0.0018	0.0018
σ_ε^R(6yr)	**0.0011**	0.0011	0.0011	0.0011	0.0011	0.0011	0.0011	0.0011
σ_ε^R(7yr)	**0.0005**	0.0005	0.0005	0.0005	0.0005	0.0005	0.0005	0.0005
σ_ε^R(8yr)	**0.0001**	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001
σ_ε^R(9yr)	**0.0004**	0.0004	0.0004	0.0004	0.0004	0.0004	0.0004	0.0004
σ_ε^R(10yr)	**0.0008**	0.0008	0.0008	0.0008	0.0008	0.0008	0.0008	0.0008

Table 8: **Summary Statistics of the Finite-Sample Bias Correction Assuming No Liquidity Correction of TIPS Yields.**
The table reports the summary statistics of the first four steps of the finite-sample bias correction of the CLR model with no liquidity correction of TIPS yields. The results of each step of the algorithm is based on $N = 100$ simulated data sets identical to the original data sample.

isfy all non-negativity requirements, so the "true" underlying continuous-time process never becomes negative P-a.s. However, for the discretely observed process above there is always a positive, but usually very small, probability that the approximation will become negative. Whenever this happens, we truncate the simulated square-root processes at 0 similar to what we do in the model estimations.

As for the starting point, X_0, of the simulation algorithm, we ideally want to draw it from the unconditional joint distribution of the state variables. However, due to the non-Gaussian property of the model, the unconditional distribution of the state variables is unknown. We overcome this problem by taking the estimated value of the state variables at the end of the observed bond yield sample and simulate the state variables according to the algorithm above for 100 years and repeat this $N = 100$ times. This effectively gives us random draws from the joint unconditional distribution of X_t.

The linear bias correction itself proceeds as follows. Denote the "true" parameter set obtained from the

Parameter	1^{st} iteration		2^{nd} iteration		3^{rd} iteration		4^{th} iteration	
	Estimate	Median	Value	Median	Value	Median	Value	Median
κ^P_{11}	**0.8707**	0.9421	0.7992	0.9961	0.6738	0.8599	0.6846	0.8291
κ^P_{22}	**0.5858**	0.6994	0.4723	0.5586	0.4995	0.6584	0.4270	0.5440
κ^P_{23}	**-0.5610**	-0.5633	-0.5586	-0.6232	-0.4964	-0.5223	-0.5351	-0.5626
κ^P_{33}	**0.7531**	0.9153	0.5909	0.7877	0.5563	0.7573	0.5520	0.7764
κ^P_{44}	**2.1022**	2.6977	1.5067	2.0320	1.5769	2.0197	1.6594	1.9662
σ_{11}	**0.0574**	0.0575	0.0572	0.0575	0.0571	0.0574	0.0571	0.0572
σ_{22}	**0.0122**	0.0122	0.0122	0.0122	0.0122	0.0122	0.0122	0.0122
σ_{33}	**0.0320**	0.0319	0.0320	0.0320	0.0320	0.0319	0.0321	0.0319
σ_{44}	**0.0609**	0.0609	0.0610	0.0612	0.0608	0.0609	0.0608	0.0606
θ^P_1	**0.0436**	0.0434	0.0437	0.0432	0.0440	0.0437	0.0439	0.0425
θ^P_2	**-0.0124**	-0.0112	-0.0136	-0.0110	-0.0149	-0.0138	-0.0135	-0.0154
θ^P_3	**-0.0074**	-0.0059	-0.0089	-0.0065	-0.0097	-0.0085	-0.0086	-0.0102
θ^P_4	**0.0129**	0.0130	0.0128	0.0129	0.0128	0.0126	0.0131	0.0127
λ	**0.5937**	0.5935	0.5940	0.5942	0.5935	0.5932	0.5940	0.5940
α^R	**0.5890**	0.5905	0.5875	0.5839	0.5925	0.5890	0.5925	0.5883
θ^Q_1	**30,120**	30,146	30,094	30,226	29,988	30,114	29,995	30,033
θ^Q_4	**18,559**	18,545	18,574	18,758	18,375	18,589	18,345	18,412
σ^N_ε(3m)	**0.0020**	0.0020	0.0020	0.0019	0.0020	0.0020	0.0020	0.0019
σ^N_ε(6m)	**0.0008**	0.0008	0.0008	0.0008	0.0008	0.0008	0.0008	0.0008
σ^N_ε(1yr)	**0.0000**	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
σ^N_ε(2yr)	**0.0002**	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002
σ^N_ε(3yr)	**0.0000**	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
σ^N_ε(5yr)	**0.0001**	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001
σ^N_ε(7yr)	**0.0001**	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001
σ^N_ε(10yr)	**0.0005**	0.0005	0.0005	0.0005	0.0005	0.0005	0.0005	0.0005
σ^R_ε(5yr)	**0.0015**	0.0015	0.0015	0.0015	0.0015	0.0015	0.0015	0.0015
σ^R_ε(6yr)	**0.0009**	0.0009	0.0009	0.0009	0.0009	0.0009	0.0009	0.0009
σ^R_ε(7yr)	**0.0004**	0.0004	0.0004	0.0004	0.0004	0.0004	0.0004	0.0004
σ^R_ε(8yr)	**0.0003**	0.0003	0.0003	0.0003	0.0003	0.0003	0.0003	0.0003
σ^R_ε(9yr)	**0.0006**	0.0006	0.0006	0.0006	0.0006	0.0006	0.0006	0.0006
σ^R_ε(10yr)	**0.0010**	0.0010	0.0010	0.0010	0.0010	0.0010	0.0010	0.0010

Table 9: **Summary Statistics of the Finite-Sample Bias Correction Assuming the Maximum Liquidity Correction of TIPS Yields.**
The table reports the summary statistics of the first four steps of the finite-sample bias correction of the CLR model with maximum liquidity correction of TIPS yields. The results of each step of the algorithm is based on $N = 100$ simulated data sets identical to the original data sample.

estimation based on the observed data by ψ_1. In the first step of the bias correction algorithm, we use ψ_1 to simulate the model as described above. The median of the $N = 100$ estimation results from the simulated data is denoted $med(\psi_1)$. As a result, the initial bias is given by

$$bias(\psi_1) = med(\psi_1) - \psi_1.$$

To determine the parameter set to be used in the second step of the algorithm, we assume that the shape of the bias function for each parameter is approximately linear, monotonically increasing, and independent of all other model parameters. As a consequence, the best guess of the bias-corrected parameter set given the information available from the first step is

$$\psi_2 = \psi_1 - bias(\psi_1).$$

Parameter	No TIPS liquidity premium			Maximum TIPS liquidity premium		
	Estimate	Bias corrected	Median	Estimate	Bias corrected	Median
κ_{11}^P	**1.0188**	0.7735	1.0748	**0.8707**	0.6846	0.8291
κ_{22}^P	**0.6073**	0.5016	0.6076	**0.5858**	0.4270	0.5440
κ_{23}^P	**-0.5678**	-0.5456	-0.5919	**-0.5610**	-0.5351	-0.5626
κ_{33}^P	**0.7530**	0.5672	0.7958	**0.7531**	0.5520	0.7764
κ_{44}^P	**1.4246**	0.7716	1.3592	**2.1022**	1.6594	1.9662
σ_{11}	**0.0634**	0.0631	0.0633	**0.0574**	0.0571	0.0572
σ_{22}	**0.0126**	0.0126	0.0126	**0.0122**	0.0122	0.0122
σ_{33}	**0.0315**	0.0315	0.0314	**0.0320**	0.0321	0.0319
σ_{44}	**0.0554**	0.0554	0.0554	**0.0609**	0.0608	0.0606
θ_1^P	**0.0421**	0.0422	0.0420	**0.0436**	0.0439	0.0425
θ_2^P	**-0.0110**	-0.0131	-0.0132	**-0.0124**	-0.0135	-0.0154
θ_3^P	**-0.0065**	-0.0087	-0.0098	**-0.0074**	-0.0086	-0.0102
θ_4^P	**0.0148**	0.0152	0.0145	**0.0129**	0.0131	0.0127
λ	**0.5980**	0.5979	0.5982	**0.5937**	0.5940	0.5940
α^R	**0.3860**	0.3893	0.3860	**0.5890**	0.5925	0.5883
θ_1^Q	**32,600**	32,461	32,585	**30,120**	29,995	30,033
θ_4^Q	**15,368**	15,223	15,367	**18,559**	18,345	18,412
$\sigma_\varepsilon^N(3\text{m})$	**0.0020**	0.0020	0.0020	**0.0020**	0.0020	0.0019
$\sigma_\varepsilon^N(6\text{m})$	**0.0008**	0.0008	0.0008	**0.0008**	0.0008	0.0008
$\sigma_\varepsilon^N(1\text{yr})$	**0.0000**	0.0000	0.0000	**0.0000**	0.0000	0.0000
$\sigma_\varepsilon^N(2\text{yr})$	**0.0002**	0.0002	0.0002	**0.0002**	0.0002	0.0002
$\sigma_\varepsilon^N(3\text{yr})$	**0.0000**	0.0000	0.0000	**0.0000**	0.0000	0.0000
$\sigma_\varepsilon^N(5\text{yr})$	**0.0001**	0.0001	0.0001	**0.0001**	0.0001	0.0001
$\sigma_\varepsilon^N(7\text{yr})$	**0.0001**	0.0001	0.0001	**0.0001**	0.0001	0.0001
$\sigma_\varepsilon^N(10\text{yr})$	**0.0005**	0.0005	0.0005	**0.0005**	0.0005	0.0005
$\sigma_\varepsilon^R(5\text{yr})$	**0.0018**	0.0018	0.0018	**0.0015**	0.0015	0.0015
$\sigma_\varepsilon^R(6\text{yr})$	**0.0011**	0.0011	0.0011	**0.0009**	0.0009	0.0009
$\sigma_\varepsilon^R(7\text{yr})$	**0.0005**	0.0005	0.0005	**0.0004**	0.0004	0.0004
$\sigma_\varepsilon^R(8\text{yr})$	**0.0001**	0.0001	0.0001	**0.0003**	0.0003	0.0003
$\sigma_\varepsilon^R(9\text{yr})$	**0.0004**	0.0004	0.0004	**0.0006**	0.0006	0.0006
$\sigma_\varepsilon^R(10\text{yr})$	**0.0008**	0.0008	0.0008	**0.0010**	0.0010	0.0010

Table 10: **Finite-Sample Bias-Corrected Parameter Estimates.**
The table reports the summary statistics of the fourth step of the finite-sample bias correction of the CLR model, first assuming no liquidity correction of TIPS yields, second assuming the maximum liquidity correction of TIPS yields. In each case, the results are based on $N = 100$ simulated data sets identical to the original data sample in each step of the bias correction algorithm.

Now, ψ_2 is used in the simulation of the model and everything is repeated. In general, after step i, the remaining bias is

$$bias(\psi_i) = med(\psi_i) - \psi_1,$$

and the parameter set to be used in step $i + 1$ is

$$\psi_{i+1} = \psi_i - bias(\psi_i).$$

This algorithm is continued until a satisfactory level of accuracy is obtained. In terms of determining the stopping point, it should be kept in mind that there is an unavoidable base level of uncertainty inherent in the bias correction tied to the structure of the observed data as well as the number of simulations. Here, we choose to stop the bias-correction algorithm after the fourth step as the remaining bias is within the

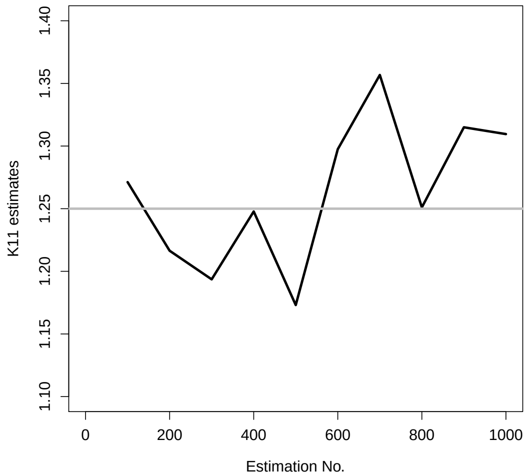

(a) κ_{11}^{P} estimates.

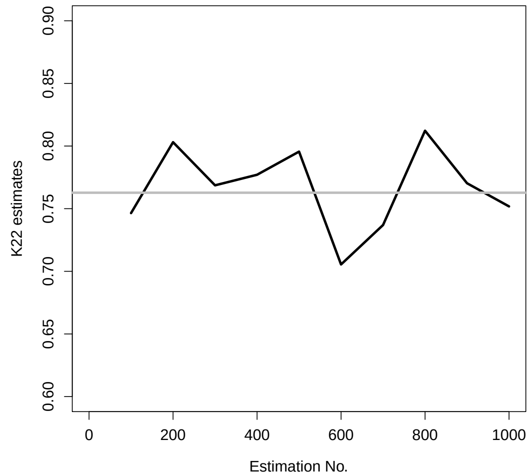

(b) κ_{22}^{P} estimates.

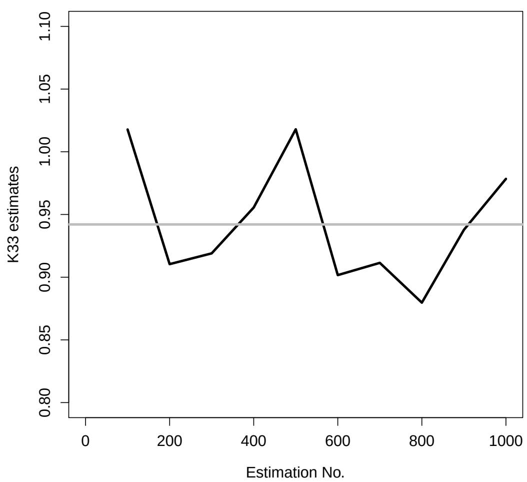

(c) κ_{33}^{P} estimates.

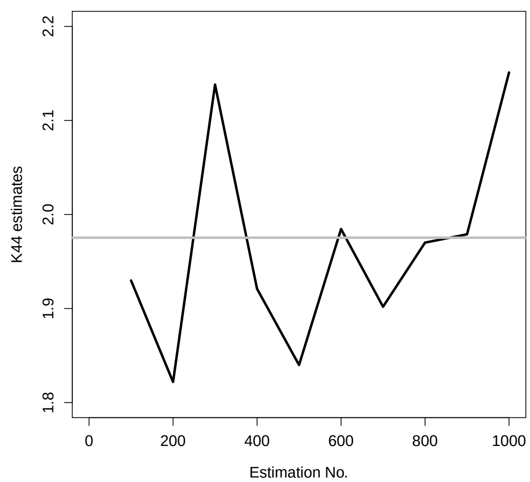

(d) κ_{44}^{P} estimates.

Figure 7: **Variation in Medians of Estimated K^P Parameters.**
Illustration of the medians of estimated parameters in the K^P matrix in the CLR model from subsamples containing 100 parameter estimates each (1-100, 101-200 etc.) out of a total of 1,000 parameter estimates from simulated data as described in the text. In each chart, the solid grey line indicates the median of all 1,000 estimates.

base level uncertainty. The results of the first four steps of the bias correction for the case where no TIPS liquidity premiums are assumed are reported in Table 8, while the corresponding results for the case where the maximum TIPS liquidity premiums are assumed are reported in Table 9. Finally, the summary result of the bias correction with the bias-corrected parameter sets are reported in Table 10.

To illustrate the inherent base uncertainty in the median-based bias correction we employ in this paper, we simulate $N = 1{,}000$ artificial samples based on the “true” parameters from the observed data for the case where no TIPS liquidity premiums are assumed. Using the simulated data in model estimations, gives us a total of $N = 1{,}000$ alternative parameter sets that we split into 10 subgroups each containing 100 parameter sets (1-100, 101-200 etc.). Now, we calculate the median estimates of each parameter from the ten subgroups and compare them to the median of all $N = 1{,}000$ parameter estimates. The result is shown in Figure 7 for the four parameters in the diagonal of K^P that are most severely affected by the finite-sample bias. The important thing to note is that there is a tradeoff between accuracy and computation time to be made when choosing the number of simulations to use in the bias correction, and 100 simulations in each iteration appear sufficient as the median is not that sensitive to the number of simulations.[33] Also, in light of the statistical uncertainty of the parameter estimates we obtain from the observed data, we think that $N = 100$ strikes a reasonable balance in terms of the tradeoff. Importantly, we do not want to leave the impression that the bias-corrected parameters are estimated with greater accuracy than the "true" parameters we get from the data.

[33]This differs from situations where the tails of each parameter distribution are of interest, say, for the generation of confidence bands. In those cases, a thousand or more simulations could be required to reach an adequate level of accuracy.

References

Adrian, Tobias, and Hao Wu, 2010, The Term Structure of Inflation Expectations, Federal Reserve Bank of New York Staff Reports No. 362.

Bauer, Michael D., Glenn D. Rudebusch, and Jing (Cynthia) Wu, 2012, Correcting Estimation Bias in Dynamic Term Structure Models, forthcoming in *Journal of Business and Economic Statistics*.

Campbell, John Y., Robert J. Shiller, and Luis M. Viceira, 2009, Understanding Inflation-Indexed Bond Markets, *Brookings Papers on Economic Activity*, Spring, 79-120.

Chernov, Mikhail, and Philippe Mueller, 2011, "The Term Structure of Inflation Expectations," forthcoming *Journal of Financial Economics*.

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*, Vol. 164, 4-20.

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*, Supplement to Vol. 42, No. 6, 143-178.

Christensen, Jens H. E., Jose A. Lopez, and Glenn D. Rudebusch, 2012, Pricing Deflation Risk with U.S. Treasury Yields, Working Paper #2012-07, Federal Reserve Bank of San Francisco.

Dai, Qiang and Kenneth J. Singleton, 2000, Specification Analysis of Affine Term Structure Models, *Journal of Finance*, Vol. 55, 1943-1978.

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 No. 19, Federal Reserve Board.

Dudley, William C., Jennifer Roush, and Michelle Steinberg Ezer, 2009, The Case for TIPS: An Examination of the Costs and Benefits, *Federal Reserve Bank of New York Economic Policy Review*, Vol. 15, No. 1, 1-17.

Duffee, Gregory R., 2002, Term Premia and Interest Rate Forecasts in Affine Models, *Journal of Finance*, Vol. 57, 405-443.

Fleckenstein, Mathias, Francis A. Longstaff, and Hanno Lustig, 2010, Why Does the Treasury Issue TIPS? The TIPS-Treasury Bond Puzzle, Working Paper #16358, NBER Working Paper Series.

Fleming, Michael, 2003, Measuring Treasury Market Liquidity, *Federal Reserve Bank of New York Economic Policy Review*, Vol. 9, No. 3, 83-108.

Fleming, Michael, and N. Krishnan, 2012, The Microstructure of the TIPS Market, *Federal Reserve Bank of New York Economic Policy Review*, Vol. 18, No. 1, 27-45.

Fleming, Michael, and Bruce Mizrach, 2009, The Microstructure of a U.S. Treasury ECN: The BrokerTec Platform, Federal Reserve Bank of New York Staff Reports No. 381.

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright, 2007, The U.S. Treasury Yield Curve: 1961 to the Present, *Journal of Monetary Economics*, Vol. 54, 2291-2304.

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright, 2010, The TIPS Yield Curve and Inflation Compensation, *American Economic Journal: Macroeconomics*, Vol. 2, No. 1, 70-92.

Hamilton, James D., and Jing Cynthia Wu, 2012, Identification and Estimation of Gaussian Affine Term Structure Models, *Journal of Econometrics*, Vol. 168, 315-331.

Harvey, A.C., 1989, *Forecasting, structural time series models and the Kalman filter* (Cambridge University Press, Cambridge).

Hurd, Matthew, and Jon Relleen, 2006, New information from inflation swaps and index-linked bonds, *Bank of England Quarterly Bulletin*, Spring, 24-34.

Joslin, Scott, Kenneth Singleton, and Haoxiang Zhu, 2011, A New Perspective on Gaussian DTSMs, *Review of Financial Studies*, Vol. 24, 926-970.

Nelson, Charles R., and Andrew F. Siegel, 1987, Parsimonious Modeling of Yield Curves, *Journal of Business*, Vol. 60, 473-489.

Pflueger, Carolin E., and Luis M. Viceira, 2011, An Empirical Decomposition of Risk and Liquidity in Nominal and Inflation-Indexed Government Bonds, Working Paper #16892, NBER Working Paper Series.

Sack, Brian, and Robert Elsasser, 2004, Treasury Inflation-Indexed Debt: A Review of the U.S. Experience, *Federal Reserve Bank of New York Economic Policy Review*, Vol. 10, No. 1, 47-63.

Thompson, Samuel, 2008, Identifying Term Structure Volatility from the LIBOR-Swap Curve, *Review of Financial Studies*, Vol. 21, No. 2, 819-854.