



# Reading between the Yields: Extracting Market-Based Inflation Expectations from U.S. Breakeven Rates

Carola Giannini\*

## Abstract

This paper develops a regression-based framework to decompose the U.S. 10-year breakeven inflation rate, defined as the yield spread between nominal Treasury bonds and Treasury Inflation-Protected Securities, into three components: expected inflation, a liquidity premium, and an inflation risk premium. The analysis spans January 1999 to November 2025, covering several macro-financial regimes including the post-Global Financial Crisis period, the COVID-19 shock, the post-pandemic inflationary surge, and the subsequent tightening cycle.

The methodology proceeds in three steps. First, breakeven inflation is regressed on observable liquidity proxies, namely the off-the-run spread and the synthetic-cash breakeven spread, alongside survey-based inflation expectations and a cyclical control, to isolate the TIPS liquidity premium. Second, the predictability of excess holding-period returns on the liquidity-adjusted series is used to identify a time-varying inflation risk premium. Third, subtracting the estimated risk premium yields a market-implied measure of long-run inflation expectations, which is compared with the University of Michigan Survey of Consumers.

The results show that liquidity distortions are economically significant and regime-dependent: the 2008–2009 crisis generated a sharp surge in the liquidity premium, while the COVID-19 episode reflected a genuine downward revision in expectations. The comparison with survey data reveals a persistent gap, consistent with the upward bias in household expectations, and highlights that market-based and survey-based measures should be interpreted as complementary rather than substitutes. Overall, the framework provides a transparent and real-time alternative to structural affine term structure models, prioritising robustness and interpretability.

**Keywords:** Breakeven inflation; Inflation expectations; TIPS; Liquidity premium; Inflation risk premium.

**JEL Classification:** C22; E31; E43; E52; G12

---

*\*I would like to express my sincere gratitude to Professor Carlo Ambrogio Favero for his guidance and supervision throughout this thesis.*

# Contents

- 1 Introduction** **2**
- 1.1 The relevance of Breakeven Inflation . . . . . 3
- 1.2 The Anatomy of Breakeven Inflation . . . . . 8
- 1.3 Literature Review . . . . . 9
  
- 2 Data** **12**
- 2.1 Yield Data and Breakeven definition . . . . . 13
- 2.2 Liquidity Proxies . . . . . 14
- 2.3 Inflation Expectations Controls . . . . . 19
  
- 3 Model and regressions** **22**
- 3.1 Liquidity-Adjusted Breakeven Inflation . . . . . 27
- 3.2 Estimating the Inflation Risk Premium . . . . . 30
  
- 4 Comparison with Survey-Based Expectations** **43**
- 4.1 Other Non-Expectation Components of Breakeven Inflation . . . . . 49
  
- 5 Conclusions** **52**
  
- 6 Limitations** **54**
  
- References** **57**
  
- Appendix** **60**

# 1 Introduction

This thesis develops a simple regression-based framework to decompose the U.S. 10-year breakeven inflation rate, defined as the yield spread between nominal Treasury bonds and Treasury Inflation-Protected Securities (TIPS), into three components: expected inflation, a liquidity premium, and an inflation risk premium. The analysis spans the period January 1999 to November 2025, a sample that includes several distinct macro-financial regimes, including the post-GFC decade of unconventional monetary policy, the COVID-19 shock, the post-pandemic inflationary surge, and the subsequent tightening cycle. The analysis proceeds in three successive steps. First, ten-year breakeven inflation is regressed on observable liquidity proxies, alongside survey-based inflation expectations and a cyclical control, to isolate and remove the TIPS liquidity premium from observed breakevens. Second, the predictability of excess holding-period returns on this liquidity-adjusted series is exploited to identify and quantify a time-varying inflation risk premium, which is then aggregated to the 10-year horizon. Third, subtracting the estimated risk premium from the liquidity-adjusted breakeven yields a market-implied measure of long-run inflation expectations, which is compared against the University of Michigan Survey of Consumers to evaluate its informational content and plausibility as a policy indicator. The core motivation is that breakeven inflation is a noisy signal and during periods of financial stress, the non-expectational components may dominate movements in breakevens, precisely when policymakers rely most heavily on market-based indicators. Simply reading breakevens as "what markets expect for inflation" can therefore be misleading. This thesis contributes to the literature on market-implied inflation expectations in three ways. First, building on [Pflueger and Viceira \(2011\)](#), we extend their analysis of liquidity extraction and breakeven return predictability to a substantially longer sample, allowing us to reassess the robustness of their results across markedly different macro-financial environments and to evaluate newly relevant scenarios, which have not been explored extensively so far. Second, we use the fact that excess returns on breakevens are predictable, thereby implying the presence of a time-varying inflation risk premium, as a tool rather than an endpoint, and quantify the inflation risk premium to decompose liquidity-adjusted breakeven inflation

into a risk premium component and a residual that we interpret as expected inflation. Third, drawing on prior works, we compare our market-implied measure of inflation expectations with the Michigan Survey of Household Inflation Expectations, evaluating the informational content of our market-based measure relative to survey expectations and assessing its usefulness for policymakers.

## 1.1 The relevance of Breakeven Inflation

Inflation expectations play a central role in modern financial markets and monetary policy. The credibility of an inflation target, the transmission of policy to long-term interest rates, and the stability of macroeconomic outcomes all depend critically on how households, firms, and financial market participants form beliefs about future inflation. This matter has gained renewed centrality in recent periods, as central banks in several advanced economies and, most notably, the Federal Reserve, have faced heightened political pressure and, in some cases, explicit challenges to their operational independence. In such an environment, the anchoring of short, medium, and long-run inflation expectations becomes not only a key policy objective but also a visible indicator of institutional credibility. Strong institutional trust and well-anchored medium- to long-term inflation expectations are essential for ensuring that inflation shocks, if any, are perceived as temporary. When the public is confident in the central bank's ability to deliver on its inflation target and commitment, short-run fluctuations are less likely to trigger persistent revisions in expectations.

Inflation expectations do not merely reflect anticipated price developments, they also shape them. This self-reinforcing mechanism, often referred to as second-round effects, operates through the wage and price-setting process, whereby agents' anticipations of future inflation are incorporated into current nominal adjustments, thereby shaping realized inflation outcomes. When expectations remain well anchored, such self-reinforcing dynamics are mitigated, fostering macroeconomic stability and preserving the ability of firms and households to make intertemporal consumption and investment decisions. The empirical and theoretical examination of this relationship originates with the traditional Phillips Curve, which posited a short-run inverse relationship between inflation

and economic slack. Subsequent theoretical developments culminated in the formulation of the New Keynesian Phillips Curve (NKPC), which embeds rational, forward-looking expectations within a micro-founded price-setting environment. Building on the staggered price-adjustment mechanism proposed by [Calvo \(1983\)](#), the NKPC assumes monopolistic competition and nominal rigidities, leading to an expression where current inflation  $\pi_t$  depends on expected future inflation and by the output gap, according to

$$\pi_t = \beta E_t[\pi_{t+1}] + \kappa(y_t - y_t^*) + u_t.$$

In this framework, inflation expectations constitute a central propagation channel for monetary policy, determining how demand and supply shocks translate into realized inflation. The forward-looking term  $E_t[\pi_{t+1}]$  reflects the rational expectations hypothesis, implying that agents form beliefs consistent with the model's structural features. Hence, the credibility of monetary policy is fundamental to anchoring expectations and preventing persistent inflation deviations from the target. As a result, central banks closely monitor measures of inflation expectations at different horizons, precisely because shifts in expectations can either amplify or mitigate the inflationary pressures generated by cost-push shocks ( $u_t$ ) or fluctuations in the output gap, and thus serve as key indicators of the effectiveness and credibility of the inflation-targeting framework.

For these reasons, policymakers rely on a broad range of indicators intended to capture expectations across different horizons and across different populations. Given both the heterogeneity of agents and the practical difficulties in obtaining robust and coherent expectations from households, firms, and even professional forecasters over extended time periods, market-based measures derived from swaps, options, and inflation-linked securities assume a core role. These measures are available at high frequency, are inherently forward-looking, and reflect the beliefs of agents with direct financial exposure to inflation risk, making them particularly relevant inputs for monetary policy assessment. A natural and important question, therefore, is whether market-based measures provide a more accurate assessment of inflation expectations than those derived from surveys of the

broader population. A substantial body of literature has examined and compared the informational content, forecasting performance, and possible systematic biases of these two approaches.

In this framework, inflation-linked securities play a central role, with the U.S. market offering a deep, liquid, and well-structured environment that provides a robust foundation for our analysis. For the sake of our analysis, Treasury Inflation-Protected Securities (TIPS) play a central role.

TIPS are marketable Treasury securities whose principal amount is adjusted for inflation. Since their introduction in the US in 1997, they have become cornerstones in the US fixed-income market. They were first issued following the market's strong interest in the inflation-indexed asset class, driven by the need to protect long-term savers from the inflation volatility of the 1970s and 1980s. They have since that moment grown into a large, liquid segment of the Treasury market. By 2008, outstanding TIPS reached approximately \$500 billion, roughly 8% of the marketable Treasury debt stock. The Federal Reserve's large-scale asset purchase programs during and after the financial crisis included substantial TIPS purchases, which improved secondary market liquidity. As of the end of 2024, outstanding TIPS exceeded \$1.9 trillion. Despite this rapid growth, TIPS remain substantially less liquid than nominal Treasuries of comparable maturity, a feature that plays a central role in our empirical analysis and is examined in depth in the following sections. Unlike conventional nominal Treasuries, whose coupons and principal are fixed in nominal terms, TIPS cash flows are indexed to the non-seasonally adjusted urbanized Consumer Price Index (CPI-U), meaning that both the principal and the coupon payments are adjusted with realized inflation, allowing investors to lock in a real rate of return over the life of the bond, with maturities typically at 5, 10, and 30 years. The basic structure of a TIPS is straightforward. At issuance, the bond has an original principal  $P_0$  and a fixed coupon rate  $c$ . Over time, the reference CPI index evolves and if the CPI at time  $t$  is  $I_t$ , and  $I_0$  is the CPI at issuance, then the inflation-adjusted principal is given by

$$P_t = P_0 \cdot \frac{I_t}{I_0}.$$

The coupon payment at time  $t$  is then calculated as

$$c \cdot P_t.$$

When prices rise, the principal is scaled up, and coupon payments increase accordingly; when prices fall, the principal is scaled down, and coupon payments decline. At maturity, investors receive the higher of the original principal  $P_0$  and the inflation-adjusted principal, which guarantees that a prolonged deflation episode cannot reduce the redemption amount below par. Because of this structure, the yield on a TIPS bond is interpreted as a real yield, being the return an investor earns above realized inflation.

These characteristics have important implications for the role of TIPS in financial markets. First of all, it is now clear the relevance of those securities in hedging investors from inflation risk. But more importantly for the scope of this thesis and from the perspective of information extraction, the coexistence of nominal Treasuries and TIPS of similar maturities allows one to infer market-implied inflation compensation. Indeed, we can define the breakeven inflation rate as the difference between the yields of Nominal bonds ( $y_t^{\$,n}$ ) and TIPS ( $y_t^{TIPS,n}$ ) of the same maturity.

$$b_t^{(n)} = y_t^{\$,n} - y_t^{TIPS,n}$$

Ideally, in a frictionless and riskless market, this difference would be interpreted as the conditional expectation of average inflation over the horizon  $(t, t+n)$ ; however, in practice, this measure reflects compensation for inflation risk and the premium investors require to bear the liquidity risk of TIPS relative to nominal bonds. Formally, breakeven inflation at horizon  $n$  can be decomposed as

$$b_{n,t} = \mathbb{E}_t(\pi_{t,t+n}) + IRP_{n,t} - L_{n,t},$$

where  $\mathbb{E}_t(\pi_{t,t+n})$  denotes expected average inflation over the horizon,  $IRP_{n,t}$  the inflation risk premium, and  $L_{n,t}$  the relative liquidity premium embedded in TIPS yields. Because TIPS are less liquid than nominal Treasuries of comparable maturity, their yields

typically include a positive liquidity premium, mechanically biasing breakeven inflation downward. These non-expectational components, which are discussed in detail in Subsection 1.2, have been particularly salient in the history of the U.S. TIPS market. During the Great Financial Crisis, a flight to safety to the most liquid instruments coincided with a sharp drop in breakevens, at a time when underlying longer-run inflation expectations appeared much more stable (D’Amico et al. (2018)). Over the subsequent decade of large-scale asset purchases, abundant central bank liquidity compressed the relative illiquidity of TIPS and likely dampened their premia, partly because balance-sheet policies altered the composition of duration and inflation risk held by the private sector. More recently, the post-pandemic inflationary period and the fast tightening cycle starting in 2022 were accompanied by renewed strains in Treasury market functioning. In particular, as economic uncertainty surrounding COVID-19 escalated, investors initially engaged in a flight-to-quality, reallocating portfolios toward safe-haven sovereign bonds. However, by the second week of March 2020, this dynamic shifted into a global "dash for cash", with investors liquidating sovereign bonds, including U.S. Treasuries and TIPS, to raise liquidity redemptions, margin calls, and cash buffers, severely altering market functioning across jurisdictions. Barone et al. (2022) document that these disruptions were particularly acute in the U.S. Treasury market due to broader and more intense selling pressures compared to other advanced-economy sovereign bonds. This phenomenon suggests that fluctuations in breakeven inflation rates may not solely capture changes in underlying inflation expectations but may also reflect shifts in relative liquidity conditions across markets.

While we focus on breakeven inflation rates as a direct measure of market-based inflation compensation, they are not the only source of information on inflation expectations derived from the market. Financial markets provide several alternative indicators that can be used either as complements or, in our case, as cross-checks to breakevens. Among these, inflation swaps play a central role. These are derivative contracts in which one party agrees to swap fixed payments for floating payments tied to the inflation rate (the CPI in the US), for a given notional amount and period of time. In the most common zero-coupon structure, the net payment at maturity equals the difference between the

cumulative realized inflation over the life of the contract and the pre-agreed fixed swap rate. Then we have inflation caps and floors options that offer hedging against inflation being above or below a strike price, that is, a predetermined rate of inflation. Similar to TIPS, both these measures are used by investors to protect from inflation risk and to speculate on the course of inflation, and by market observers more broadly to infer inflation expectations.

## 1.2 The Anatomy of Breakeven Inflation

We now examine in greater detail the two non-expectational components embedded in breakeven inflation: the liquidity premium and the inflation risk premium.

The liquidity premium arises from structural differences in the market functioning of nominal Treasuries and TIPS. Nominal Treasuries constitute the deepest and most liquid segment of the global fixed-income market, widely used as benchmark securities and collateral instruments. TIPS, while sizeable and increasingly integrated into portfolios, remain comparatively less liquid, with lower turnover and a narrower investor base. Because TIPS are less liquid, investors require additional compensation to hold them, taking the form of a liquidity premium embedded in TIPS yields, which mechanically compresses observed breakeven inflation relative to expected inflation. Higher TIPS illiquidity therefore lowers breakevens, even if inflation expectations remain unchanged. Notably, liquidity conditions are not constant. They deteriorate during episodes of financial stress, when funding constraints tighten and investors reallocate toward the most liquid instruments. In such environments, liquidity premia can widen sharply, causing breakeven inflation to understate underlying inflation expectations precisely when policymakers are most attentive to market-based indicators.

The inflation risk premium represents compensation for bearing uncertainty regarding future inflation outcomes. Investors holding nominal bonds are exposed to real-return risk, as the purchasing power of fixed nominal cash flows depends on the realized inflation rate. When inflation uncertainty increases, investors typically demand greater compensation for this risk, especially for longer-maturity bonds. The sign and magnitude of the

IRP depend on the covariance between inflation and the marginal utility of consumption. When inflation is elevated during adverse economic conditions, nominal bonds underperform when consumption is low, leading to a positive inflation risk premium. Conversely, if nominal bonds provide a hedge against adverse states, such as during deflationary recessions, the IRP may be small or even negative. The sign of this correlation has not been constant over time, and the empirical evidence suggests that it shifted substantially around the early 2000s as the Federal Reserve established greater credibility for the 2% inflation target, moving from a positive to a near-zero or slightly negative correlation and compressing the IRP accordingly. Crucially, the IRP is distinct from expected inflation. A central bank may have long-run inflation expectations firmly anchored near the target, while investors simultaneously demand positive compensation for bearing inflation uncertainty, particularly if the distribution of inflation outcomes around that expectation is wide or skewed. This distinction has direct implications for interpreting breakeven dynamics: a rise in breakeven inflation can reflect either an upward revision in expected inflation, an increase in inflation uncertainty that raises the IRP, or both simultaneously. Without separating these components, movements in breakeven rates cannot be accurately interpreted. From a monetary policy perspective, the IRP provides information complementary to, but distinct from, expected inflation. A rising IRP may indicate that investors perceive an increase in inflation uncertainty without necessarily revising their central forecast upward, consistent with a widening of the distribution of future outcomes rather than a shift in its mean.

### 1.3 Literature Review

A large and sophisticated literature seeks to decompose breakeven inflation into expected inflation, risk premia, and liquidity effects using no-arbitrage term structure models. [D'Amico et al. \(2018\)](#) develop a no-arbitrage model that jointly prices nominal Treasuries, TIPS, and realized inflation, incorporating a TIPS-specific spread to capture illiquidity and technical factors. They demonstrate that models that ignore this spread produce implausible inflation risk premium estimates and expectations misaligned with surveys, whereas allowing for it yields reasonable decompositions that closely track SPF forecasts.

Christensen et al. (2010) estimate an arbitrage-free affine term structure model that jointly prices nominal and real yields, allowing for a decomposition of breakeven inflation into expected inflation and inflation risk premia. They document substantial time variation in inflation compensation that co-moves with macroeconomic conditions. Christensen and Gillan (2011) further highlight the role of TIPS-specific liquidity distortions in interpreting breakeven inflation, showing that divergences between inflation swaps and breakevens during stress episodes likely reflect liquidity premia rather than shifts in expectations. Gürkaynak et al. (2010) construct the full TIPS yield curve and analyze the dynamics of inflation compensation across maturities, providing a benchmark methodology widely used in subsequent works. Finally, Hördahl and Tristani (2014a), using a joint macro-finance model of nominal and real bonds, estimate time-varying inflation risk premia and emphasize the importance of integrating macroeconomic factors into breakeven decomposition. While these structural models differ in specification, they converge on a central insight: breakeven inflation cannot be interpreted as expected inflation without explicitly accounting for both liquidity premia and time-varying inflation risk compensation.

At a deeper level, this decomposition rests on the empirical rejection of the Expectations Hypothesis. Campbell and Shiller (1991) show that the slope of the nominal yield curve forecasts excess returns, rejecting the strict expectations hypothesis and implying that risk premia must vary over time. Cochrane and Piazzesi (2005) further demonstrate that a single common factor constructed from forward rates drives time variation in expected returns across the term structure, substantially strengthening the empirical rejection of the expectations hypothesis, which implies that excess returns should not be predictable. Building on this background, Pflueger and Viceira (2011) extend the analysis to a joint nominal–real setting, incorporating TIPS and breakevens. They show that both real-rate and inflation risk premia are important drivers of nominal bond return predictability and that TIPS liquidity premia are systematic and quantitatively large, especially in crisis episodes. We transform their evidence on return predictability into an operational procedure to extract inflation risk premium and then a market-based measure of inflation expectations.

Finally, extensive research compares market-based and survey-based inflation expectations, emphasizing their forecasting performance and behavioral differences. [Ang et al. \(2007\)](#) demonstrate that Survey of Professional Forecasters (SPF) expectations systematically outperform raw TIPS breakevens and inflation swaps, particularly at short horizons, due to markets' embedded risk and liquidity premia. [Bauer and McCarthy \(2015\)](#) confirm unadjusted breakevens underperform SPF forecasts, though liquidity adjustments substantially narrow this gap. Structural analyses reveal distinct dynamics. [Coibion and Gorodnichenko \(2015\)](#) show surveys anchor more firmly during macroeconomic volatility, while markets amplify premia-driven swings. [Dräger and Nghiem \(2021\)](#) document a persistent wedge reflecting markets' liquidity versus surveys' stability. [Grothe and Meyler \(2015\)](#) find that both market-based inflation swaps and the ECB Survey of Professional Forecasters provide information beyond naïve benchmarks, with swaps exhibiting somewhat superior short-horizon accuracy, whereas survey-based expectations dominate at longer forecast horizons.

Our primary methodological contribution is to develop a straightforward, regression-based framework for decomposing breakeven inflation that provides a practical complement to the affine term structure models that dominate the literature. While affine no-arbitrage models generate internally consistent yield-curve decompositions grounded in economic theory, they require strong assumptions about factor dynamics, pricing kernels, and the joint evolution of nominal and real yields. Our approach addresses these limitations by adopting a reduced-form methodology that exploits the predictability of breakeven returns to isolate time-varying inflation risk premia without imposing a full structural term structure specification. Instead of estimating latent state variables within a high-dimensional affine system, we rely on observable market data and standard predictive regressions to recover the inflation risk premium embedded in breakeven rates.

This framework offers several practical advantages. All inputs are directly observable in real time, estimation is based on standard econometric techniques, and the decomposition does not require latent factor extraction or distributional assumptions about the pricing kernel. As a result, the methodology is computationally simple, transparent, and

readily implementable, making it particularly suitable for real-time monitoring applications. To be clear, our approach is not intended as a substitute for structural models in all contexts. Affine term structure models remain the benchmark for fully specified yield-curve decompositions and counterfactual policy analysis. However, for the concrete task of monitoring inflation expectations in real time, especially during episodes of financial stress when structural assumptions may be violated, our regression-based framework provides a valuable complement that prioritizes transparency, and interpretability.

The remainder of the thesis is structured as follows. Section 2 describes the data, introduces the construction of the yield and breakeven series, and presents the liquidity proxies and expectations controls. Section 3 sets out the empirical framework, reports the estimation results for the liquidity premium and the implied inflation risk premium, and discusses their behavior across key historical episodes. Section 4 compares our measure of market-based inflation expectations with the Michigan survey, and Section 5 concludes by summarizing the main findings and outlining implications for the use of TIPS breakevens as indicators of inflation expectations and for future research.

## 2 Data

This thesis builds on and extends the empirical framework of [Pflueger and Viceira \(2011\)](#) by updating their analysis of breakeven inflation and liquidity premia using U.S. Treasury securities to incorporate recent market developments through the end of November 2025. The aim of this section is to build a dataset that allows for the empirical estimation of a liquidity-adjusted measure of breakeven inflation and that enables the decomposition of breakeven inflation into expected inflation, inflation risk premia, and liquidity premia. All series are sampled at a monthly frequency over the period January 1999 to November 2025

## 2.1 Yield Data and Breakeven definition

The analysis focuses exclusively on 10-year maturities, for which long time series are available, and market liquidity is most developed. The choice of the 10-year tenor follows standard practice in the literature, as these securities are the most actively traded and least affected by short-term market distortions. It is therefore characterized by relatively high liquidity and pricing efficiency. U.S. nominal and inflation-indexed zero-coupon yields are obtained from the off-the-run smoothed yield curves developed by [Gürkaynak et al. \(2007\)](#) for nominal Treasuries and by [Gürkaynak et al. \(2010\)](#) for Treasury Inflation-Protected Securities (TIPS).

Formally, let  $y_{n,t}^n$  denote the continuously compounded nominal zero-coupon yield with maturity  $n$  at time  $t$  and  $y_{n,t}^{TIPS}$  denote the corresponding inflation-indexed (real) yield. The breakeven inflation rate at maturity  $n$  and time  $t$  is defined as the difference between nominal and real yields:

$$b_{n,t} = y_{n,t}^{\$} - y_{n,t}^{TIPS}$$

To further exploit information contained in the term structure, yield spreads are constructed relative to the 1-year maturity. The nominal term spread is defined as

$$s_{n,t}^{\$} = y_{n,t}^{\$} - y_{1,t}^{\$}$$

and captures the slope of the nominal yield curve, while the real (TIPS) term spread is given by

$$s_{n,t}^{TIPS} = y_{n,t}^{TIPS} - y_{1,t}^{TIPS}$$

Analogously, the breakeven term spread is defined as

$$s_{n,t}^b = b_{n,t} - b_{1,t}$$

These spreads summarize expectations about future macroeconomic conditions and risk compensation across maturities. They will serve as state variables in our return pre-

dictability regressions, exploiting the well-established finding that when term spreads are wide, expected excess returns on long-maturity bonds are high, whereas when spreads are narrow or inverted, expected returns are low or negative. Beyond their predictive content for bond returns, term spreads also provide insight into the economic forces underlying breakeven inflation. Indeed, breakeven rates inherit information from both nominal and real yield curves and therefore reflect the joint dynamics of inflation expectations, risk compensation, and market liquidity conditions across maturities.

As we already stated, breakeven inflation can be expressed as the sum of three components:

$$b_{n,t} = \mathbb{E}_t(\pi_{t,t+n}) + \text{IRP}_{n,t} - L_{n,t},$$

where  $\mathbb{E}_t(\pi_{t,t+n})$  denotes expected inflation over the horizon  $(t, t+n)$ ,  $\text{IRP}_{n,t}$  is the inflation risk premium, and  $L_{n,t}$  is the liquidity premium.

The liquidity premium reflects the compensation investors require for holding relatively illiquid inflation-indexed securities relative to nominal Treasuries. The first goal of this thesis is to isolate and remove the liquidity premium from observed breakeven inflation. Once the estimated liquidity premium  $\widehat{L}_{n,t}$  is obtained, we will be able to compute liquidity-adjusted variables under the assumption that the relative liquidity premium is fully attributable to inflation-indexed bonds. This mirrors one of the main assumptions in [Pflueger and Viceira \(2011\)](#) and allows us to compute liquidity-adjusted TIPS yields and breakevens as

$$y_{n,t}^{TIPS,adj} = y_{n,t}^{TIPS} - \widehat{L}_{n,t},$$

and

$$b_{n,t}^{adj} = b_{n,t} + \widehat{L}_{n,t}.$$

## 2.2 Liquidity Proxies

To get to this result, several liquidity measures are constructed to proxy for relative liquidity conditions between nominal Treasuries and TIPS. The analysis focuses on four of the most widely used proxies in the literature: the off-the-run spread, the relative TIPS transaction volumes, the asset swap spread, and the synthetic-cash breakeven. The

off-the-run spread is computed as the difference between the 10-year off-the-run par yield from the [Gürkaynak et al. \(2007\)](#) nominal yield curve and the on-the-run 10-year nominal Treasury yield, proxied by the Bloomberg series `USGG10YR Index`. On-the-run Treasuries are the most recently issued securities of a given maturity and are typically the most liquid instruments in the Treasury market. Off-the-run securities, which are all but the newest issued ones, while nearly identical in terms of cash flows and credit risk, trade less frequently and therefore require a liquidity discount. That being said, this spread captures the specialness and liquidity advantage of on-the-run nominal Treasuries.

Relative transaction volume is defined as

$$\text{RTV}_t = \log \left( \frac{\text{Trans}_t^{\text{TIPS}}}{\text{Trans}_t^{\$}} \right),$$

where  $\text{Trans}_t^{\text{TIPS}}$  and  $\text{Trans}_t^{\$}$  denote, respectively, average monthly transaction volumes of inflation-indexed and long-term nominal Treasury zero-coupon bonds. Transaction volumes are constructed from primary dealers' data reported in the New York Fed's FR-2004 survey. For nominal Treasuries, we use transactions with remaining maturities of six years or longer. A reporting change in June 2001 shifted the maturity cutoff from five to six years; the post-break series is rescaled to ensure continuity at the break. For TIPS instead, given low trading activity at individual maturities, we aggregate volumes across all maturities to improve comparability in volume terms.

The series is smoothed using a three-month moving average. This measure captures the liquidity differential between TIPS and nominal bonds. Not surprisingly, the relative transaction volume is negative over the entire sample, reflecting the fact that TIPS persistently trade at lower volumes than nominal Treasury securities. For ease of interpretation, the series is normalized so that the highest observed value corresponds to zero. This normalization preserves the fact that more negative values indicate periods in which TIPS are particularly illiquid relative to nominal Treasuries, while values closer to zero correspond to episodes of relatively improved TIPS market liquidity.

The third liquidity measure is the asset-swap spread, which captures the difference between swap rates and sovereign yields and serves as a leading indicator of funding conditions, risk

aversion, and market liquidity. While acknowledging that a fully consistent measurement of asset-swap spread would require a detailed asset-pricing framework, this thesis follows the empirical literature in constructing a proxy based on readily available market data. In line with market convention and corporate research practice, the asset swap spread is defined as the difference between the yield on a sovereign bond and the fixed rate of an interest rate swap of the same maturity, indexed to a benchmark floating rate.

[Pflueger and Viceira \(2011\)](#) use the difference between TIPS and nominal Treasury asset-swap spreads to proxy for relative financing costs and liquidity conditions across the two markets. Accordingly, we define the nominal asset-swap spread proxy as

$$ASW_{10,t}^{\$} = y_{10,t}^{\$} - s_{10,t}^{\text{SOFR}}$$

where the  $s_{10,t}^{\text{SOFR}}$  is the fixed rate on a 10-year interest rate swap indexed to the Secured Overnight Financing Rate (SOFR), which has replaced LIBOR as the dominant floating-rate benchmark in U.S. dollar derivatives markets since 2017. An interest rate swap is a contractual agreement in which two counterparties exchange fixed-rate payments for floating-rate payments on a notional principal, with the fixed rate set such that the swap has zero value at inception.

The TIPS asset-swap spread proxy is defined as:

$$ASW_{10,t}^{\text{TIPS}} = y_{10,t}^{\text{TIPS}} - (s_{10,t}^{\text{SOFR}} - s_{10,t}^{\pi}).$$

where  $s_{10,t}^{\pi}$  is the 10-year zero-coupon inflation swap rate. An inflation swap is a derivative transaction in which one party agrees to pay a fixed rate, known as the inflation swap rate, on a notional principal, while the counterparty pays the realized cumulative inflation over the life of the contract, typically measured using a price index such as the CPI. Inflation swap rates and breakeven inflation rates should be equal in a frictionless market, but the former is often higher than the latter. Many studies compare the two series and attribute this difference to several reasons: a mispricing in the TIPS market compared to the nominal.

[Fleckenstein et al. \(2013\)](#) blame TIPS mispricing versus nominals, implying swaps better

capture pure expectations. [Christensen and Gillan \(2011\)](#) argue that differentials arise from liquidity premia in both TIPS (higher yields, lower breakevens) and swaps.

The asset-swap spread differential is then defined as

$$ASW_{10,t} = ASW_{10,t}^{TIPS} - ASW_{10,t}^{\$}.$$

These data are available from July 2004 onward. Before that date, the series is backfilled using the first available observation, approximately 38 basis points.

We acknowledge that the backfilling assumption is a simplification, and the early-sample estimates for specifications incorporating these variables should be interpreted with appropriate caution, as they introduce a potentially important limitation for the full-sample specifications. To assess the robustness of the baseline results to this choice, we verify that the sign, magnitude, and statistical significance of the coefficients in our baseline regression are qualitatively preserved when the estimation sample is restricted to the post-July 2004 period. The results, shown in the Appendix in Table 7, confirm that the core findings are not driven by the backfilled observations.

Finally, as an ulterior measure, the analysis considers the spread between synthetic and cash breakeven inflation <sup>1</sup>. The synthetic minus cash breakeven captures exactly the aforementioned differential between inflation compensation priced in derivative markets and that embedded in cash bonds. Although this spread could, in principle, be interpreted as reflecting mispricing or limited arbitrage between inflation-indexed bonds and inflation derivatives, the synthetic–cash breakeven spread has historically tracked funding differentials between nominal Treasuries and TIPS closely, suggesting that it primarily reflects relative financing costs and liquidity conditions rather than pure arbitrage distortions.

As for the ASW, this series is available from July 2004 onward, and it is backfilled with the first available observation, which is approximately 24 basis points. It is instructive to examine the time-series behavior of these liquidity measures. Figure 1 plots the main

---

<sup>1</sup>Synthetic breakeven is measured using the 10-year zero-coupon inflation swap rate from Bloomberg (USSWIT10 BGN Curncy), while Cash breakeven is obtained as the yield differential between 10-year nominal Treasuries and TIPS.

liquidity proxies over the sample period and provides a first descriptive assessment of their dynamics.

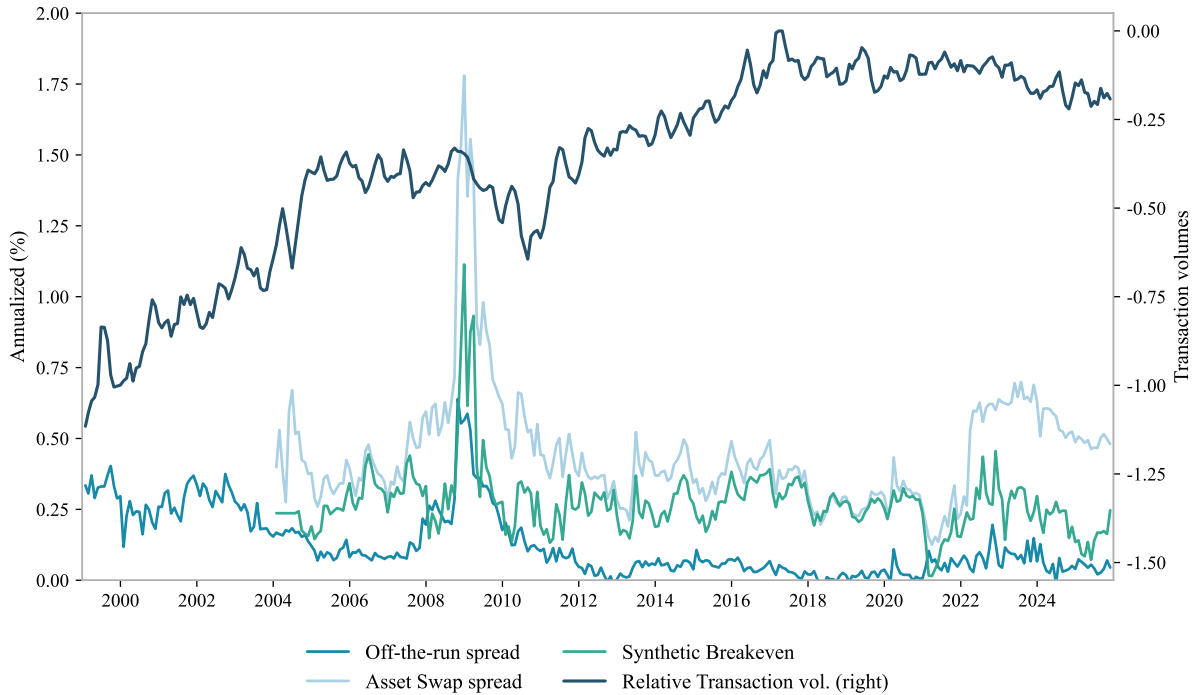


Figure 1: U.S. Liquidity proxies

By looking at Figure 1 it is possible to identify two main episodes of pronounced liquidity disruptions. The first and most significant movements occurred during the Great Financial Crisis. During this period, elevated risk aversion and a flight to liquidity were associated with a simultaneous spike across nearly all liquidity measures. The off-the-run spread increased sharply, reaching levels of approximately 50 basis points, reflecting a strong shift in investor demand toward on-the-run, newly issued Treasury securities. [Hu et al. \(2013\)](#) attribute this phenomenon to a scarcity of arbitrage capital, which created "noise" as dealers became unable to smooth price deviations between seasoned and newly issued securities.

At the same time, the asset-swap spread widened substantially, indicating that financing positions in TIPS became significantly more costly relative to nominal Treasuries. This widening reflected a pronounced deterioration in TIPS market liquidity and financing conditions, while nominal Treasuries retained their role as the preferred collateral in swap

and repo transactions.

Relative TIPS transaction volume displayed a distinct pattern, increasing steadily during the early part of the sample before stabilizing. However, following the financial crisis, it declined, reflecting a contraction in trading activity and reduced market participation. In the subsequent years, transaction volumes gradually recovered and normalized. In the decade following the GFC, liquidity measures declined from their crisis peaks and remained relatively compressed, reflecting a period of stabilized market functioning. This phase was characterized by sustained official sector support, among which the Federal Reserve’s large-scale asset purchase programs effectively provided a backstop to the market. By absorbing a significant share of Treasury and TIPS supply, these interventions reduced the duration risk held by the private sector and dampened the volatility of liquidity proxies (Engle et al., 2012). Consequently, both the off-the-run and asset-swap spreads fluctuated within narrow ranges, while relative transaction volumes stabilized, indicating a regime of abundant central bank liquidity that masked underlying structural fragilities.

A second episode of significant liquidity disruption occurred in the post-pandemic period, peaking in 2022. Several liquidity measures exhibited pronounced increases, consistent with a renewed “dash for cash” episode in which even high-quality assets were liquidated to meet margin calls and funding needs. This more recent volatility coincides with the transition to a high-inflation environment and a period of rapid monetary policy tightening by the Federal Reserve. The stress observed in 2022 is consistent with heightened interest rate uncertainty and the onset of Quantitative Tightening, which reduced the supply of central bank liquidity and increased the intermediation burden borne by private dealers.

### **2.3 Inflation Expectations Controls**

To control for long-run inflation expectations, the analysis includes the 10-year CPI inflation expectation from the Federal Reserve Bank of Philadelphia Survey of Professional Forecasters. Including this variable mitigates concerns that liquidity proxies might capture variation in inflation expectations rather than liquidity effects.

Finally, the dataset includes the Chicago Fed National Activity Index (CFNAI), to account for cyclical economic conditions and the possibility that short-term inflation expectations affect breakeven inflation, following [Stock and Watson \(1999\)](#). The CFNAI is a composite index constructed from a broad set of economic indicators, including measures of production, labor market conditions, personal consumption and housing activity, as well as indicators related to sales, orders, and inventories.

Table 1: Summary statistics (Monthly data)

Name	Symbol	Mean	Std	Min	Max
Nominal Yields	$y_{10,t}^{\$}$	3.54	1.41	0.55	6.70
Inflation-Indexed Yields	$y_{10,t}^{TIPS}$	1.39	1.33	-1.13	4.29
Breakeven	$b_{10,t}$	2.14	0.40	0.39	2.92
Survey Inflation	$\pi^E$	2.37	0.16	2.03	2.95
Chicago Fed Nat. Activity		-0.12	1.26	-18.22	6.32
Off-the-Run Spr.		0.12	0.12	-0.02	0.64
Relative Transaction Vol.		-0.37	0.26	-1.12	0.00
Asset-Swap Spr.		0.44	0.20	0.13	1.78
Synthetic - Cash		0.27	0.11	0.01	1.11

Table 1 reports summary statistics for our variables over the sample period spanning from 1999-1 to 2025-11. The results show that the average nominal Treasury yields are 3.54% per annum, while the average TIPS yields are 1.39%, implying an average breakeven inflation rate of 2.14%. The average survey-based inflation expectations amount to 2.37%, exceeding the breakeven rate by approximately 23 basis points. This systematic gap indicates that, not surprisingly, the breakeven inflation rate does not solely reflect expected inflation. In the absence of market frictions, breakeven inflation would be expected to align closely with survey-based expectations. The persistent difference observed in the data therefore, suggests the presence of non-expectational components embedded in breakeven rates, motivating the need to isolate and remove liquidity-related distortions using observable liquidity proxies.

Beyond means, the dispersion and extrema highlight the presence of distinct regimes. Breakeven inflation varies between 0.39% and 2.92%, which is large relative to its 0.40% standard deviation, implying that tail episodes as crises and inflation scares, are cent-

ral to identification. On the liquidity side, the Off-the-Run spread ranges from slightly negative values up to 64 bps, and the Synthetic–Cash breakeven ranges from near zero up to 111 bps. These upper-tail realizations matter because the liquidity premium is identified precisely from periods in which TIPS market functioning diverges sharply from the nominal Treasury market. A key implication is that a purely “average” interpretation of the sample is incomplete: the econometric contribution of liquidity proxies is concentrated in stress windows in which market segmentation, funding constraints, and dealer balance-sheet frictions become binding. Finally, the Relative Transaction Volume is negative throughout the sample by construction, reflecting persistently lower TIPS trading activity relative to nominal Treasuries. Because we normalize RTV so that the highest observed value corresponds to zero, movements toward more negative values can be read as episodes of particularly impaired TIPS liquidity. This normalization makes the direction of the regression coefficients economically transparent when we interpret the first-stage results below.

### 3 Model and regressions

In this section, we present the empirical estimation of the liquidity premium embedded in the breakeven inflation rate. By running a first-stage regression, we aim to isolate the portion of breakeven inflation driven by liquidity frictions rather than fundamental economic expectations.

The baseline regression model is specified as follows:

$$b_{n,t} = \alpha_1 + \alpha_2 X_t + \alpha_3 \pi_t^e + \alpha_4 CFNAI + \epsilon_t \quad (1)$$

where  $b_{n,t}$  represents the 10-year breakeven inflation (i.e., the difference between nominal and inflation-indexed bond yields),  $X_t$  is a vector of liquidity proxies,  $\pi_t^e$  represents inflation expectation proxies and CFNAI represents the cyclical economic condition.

Our empirical strategy follows a sequential specification that allows us to assess the marginal contribution of each additional regressor to the explanatory power of the model. Indeed, specification (1) of Table 2 includes only macroeconomic controls, namely survey inflation expectations and the CFNAI, both of which are highly statistically significant from the beginning. We then systematically add and expand the liquidity vector by incorporating the Off-the-Run spread, the ASW, and finally, including the Relative Log Transaction Volumes. In addition, Synthetic-Cash breakeven variable is included as an alternative proxy for the liquidity differential between TIPS and nominal bonds. This variable proves to be more informative than the raw ASW spread and is therefore preferred in the final specification. Finally, the last two columns replicate specification (5) using restricted samples to test for temporal stability and different time windows specifications. Notably, as we move from the macro-only specification to the full liquidity-adjusted model, the  $R^2$  increases substantially, rising from 0.28 in Specification (1) to 0.51 in Specification (5) and (5a). This improvement highlights the critical role played by liquidity variables in explaining breakeven inflation dynamics. This 0.23 gain demonstrates that liquidity conditions account for a substantial portion of breakeven monthly dynamics beyond business

cycle effects and survey expectations.

Table 2: First-Stage regression of 10-Year Breakeven Inflation on Liquidity and Macroeconomic Controls. Sequential Specifications. Newey–West HAC Standard Errors in Parentheses.

	$b_{n,t}$							
	(1)	(2)	(3)	(4)	(5)	(5a)	(6)	(7)
Off-the-Run Spr.		-1.16*** (0.39)	-0.93*** (0.35)	-1.69*** (0.49)	-0.94*** (0.33)	-0.77** (0.30)	-0.74* (0.40)	0.06 (0.88)
Asset-Swap Spr.			-0.28 (0.22)	-0.05 (0.23)				
Transaction Vol.				-0.36 (0.19)	-0.10 (0.14)		0.50*** (0.17)	0.19 (0.45)
Synthetic-Cash					-1.48*** (0.28)	-1.55*** (0.28)	-1.33*** (0.33)	-2.56*** (0.28)
Survey Inflation	0.99*** (0.22)	1.40*** (0.34)	1.42*** (0.32)	1.33*** (0.29)	1.35*** (0.28)	1.36*** (0.28)	0.95*** (0.33)	1.21*** (0.29)
CFNAI	0.19*** (0.07)	0.15*** (0.05)	0.13*** (0.04)	0.12*** (0.04)	0.09** (0.04)	0.10*** (0.04)	0.17** (0.08)	0.09*** (0.03)
$p$ -value	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
$R^2$	0.28	0.36	0.37	0.39	0.51	0.51	0.59	0.65
ADF of Residuals	-6.47***	-6.48***	-6.19***	-5.34***	-6.21***	-6.24***	-5.25***	-5.26***
Period	Full	Full	Full	Full	Full	Full	1999.03–2014.12	2015.01–2025.12

\*, \*\*, and \*\*\* denote statistical significance at the 10%, 5%, and 1% levels, respectively.

All regressions are estimated using HAC standard errors, specifically Newey–West standard errors with three lags, to account for potential heteroskedasticity and autocorrelation in the monthly bond liquidity data. To assess the validity of the long-run relationship, we conduct Augmented Dickey–Fuller tests on the residuals  $\epsilon_t$ . The strongly negative and statistically significant test statistics (e.g.,  $-6.24^{***}$  in column 5a) confirm that the variables are cointegrated, providing evidence against a spurious long-run relationship.

Results in Table 2 show that not all liquidity proxies are statistically significant across specifications. This outcome was to be expected given the high degree of collinearity among liquidity measures, but it does not undermine the empirical strategy. Since the objective of this regression is to construct a good proxy for the liquidity premium rather

than to identify the effects of individual variables on breakeven, we rely primarily on the model's overall explanatory power, as captured by the  $R^2$ .

In addition to the overall goodness of fit of the model, an important diagnostic for the first-stage regression is whether the estimated coefficients display signs that are consistent with economic intuition. This is the case across all the different specifications, with liquidity measures that capture relative illiquidity of TIPS vis-à-vis nominal Treasuries, such as the Off-the-Run spread, the Asset Swap spread, and the Synthetic-Cash breakeven, entering with negative coefficients. The negative sign of these coefficients indicates that an increase in these spreads signals worsening liquidity conditions for TIPS, prompting investors to demand higher required returns to hold them, thereby compressing the breakeven inflation rate. Notably, these are all price-based liquidity measures, which tend to be stationary throughout the sample and thus well-suited to full-sample estimation.

The Relative Transaction Volume measure warrants additional discussion. In specifications (4) and (5), where it is included alongside other liquidity proxies in the full sample, the coefficient is negative, which is at odds with the expected positive relationship between relative TIPS liquidity and breakeven inflation. However, when we split the sample into two sub-periods, 1999–2014 and 2015–2025, the coefficient turns positive in both cases and is statistically significant in the earlier sub-period. This sign reversal suggests that the negative full-sample coefficient is driven by the series' long-run structural trend rather than a genuine liquidity effect. As shown in Figure 1, the series exhibits a pronounced upward trend from 1999 to approximately 2018, reflecting the gradual deepening of the TIPS market relative to nominal Treasuries. In the full sample, OLS partially picks up this common trend rather than the true within-period liquidity relationship. Formal stationarity testing confirms this interpretation. The Augmented Dickey–Fuller test performed on this variable fails to reject the null of a unit root across all specifications, providing direct evidence of non-stationarity. To recover a stationary measure, we consider four alternative transformations: a first-differencing, a second-differencing, a linear detrending, and an HP-filtering, none of which produces a series satisfying the stationarity criterion. Moreover, a Chow test rejects parameter stability at the 1% level. We conclude, therefore, that the relative transaction volume measure is excluded from the preferred specification.

Specification (5a) retains the Off-the-Run spread and the Synthetic-Cash breakeven as the core liquidity proxies, both statistically significant, correctly signed, and stable across the full sample, with coefficients of  $-0.77^{**}$  and  $-1.55^{***}$ , respectively. The Asset-Swap spread is excluded, as its inclusion introduces multicollinearity without adding independent explanatory power. The model achieves an  $R^2$  of 0.51, the highest among full-sample specifications. At the same time, survey-based inflation expectations continue to enter with a positive and statistically significant coefficient ( $1.36^{***}$  in 5a), highlighting their dominant role in anchoring market-based inflation compensation.

The regression results in Table 2 confirm that the estimated coefficient on survey-based inflation expectations remains close to one in all specifications. Moreover, the null hypothesis  $\beta_{\pi^e} = 1$  cannot be rejected at conventional significance levels, as 1 always lies within the corresponding confidence intervals. From a theoretical standpoint, breakeven inflation decomposes into expected inflation and an inflation risk premium, plus a liquidity premium.

This structure implies a unit coefficient on expected inflation. Free estimation of this coefficient thus confirms the data's consistency with the decomposition. Since the unrestricted estimates uphold the restriction  $\beta_{\pi^e} = 1$ , we impose it and re-estimate the model using the spread between breakeven inflation and survey expectations ( $b_t - \pi_t^e$ ) as the dependent variable. This restriction yields two advantages. First, it isolates the non-expectational component, thereby giving the joint contribution of inflation risk and liquidity premia, preventing liquidity effects from being partially absorbed by small deviations in the estimated expectations coefficient. Second, imposing the unit restriction eliminates unnecessary variation associated with estimating the inflation expectations coefficient. Provided that the restriction holds in the data, this reduces the variance of the remaining parameter estimates and improves the precision of the liquidity coefficients.

The restricted first-stage results appear in Table 3. Coefficients on liquidity proxies remain economically and statistically similar to the unrestricted case, confirming that the restriction leaves core inferences intact. However,  $R^2$  declines modestly in some specifications, reflecting the tighter structural constraint imposed on the model. This trade-off, however,

delivers superior structural clarity and precision for the liquidity premium. Among the restricted full-sample specifications, we retain specification (5a) as our preferred model. This specification yields the highest  $R^2$  in the full-sample restricted regressions and incorporates the most informative set of liquidity proxies. The qualitative interpretation of the coefficients remains unchanged relative to the unrestricted case and the overall pattern of results is fully consistent with the previous discussion.

Table 3: First-Stage Regressions for the Non-expectational Breakeven Component ( $b_{n,t} - \pi_t^e$ ): Imposing  $\beta_{\pi^e} = 1$ . HAC (Newey–West) Standard Errors in Parentheses.

	$b_{n,t} - \pi_t^e$							
	(1)	(2)	(3)	(4)	(5)	(5a)	(6)	(7)
Off-the-Run Spr.		-0.88*** (0.30)	-0.66** (0.26)	-1.62*** (0.47)	-0.77** (0.32)	-0.51** (0.24)	-0.75** (0.37)	0.62 (0.67)
Asset-Swap Spr.			-0.25 (0.23)	0.02 (0.23)				
Transaction Vol.				-0.43** (0.19)	-0.14 (0.13)		0.50*** (0.17)	0.39 (0.46)
Synthetic-Cash					-1.47*** (0.28)	-1.56*** (0.28)	-1.33*** (0.33)	-2.53*** (0.28)
Survey Inflation	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
CFNAI	0.19*** (0.07)	0.16*** (0.05)	0.14*** (0.05)	0.13*** (0.04)	0.10** (0.04)	0.10*** (0.04)	0.17** (0.08)	0.09*** (0.03)
$p$ -value	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00
$R^2$	0.17	0.24	0.25	0.28	0.41	0.41	0.63	0.50
ADF of Residuals	-6.48***	-5.23***	-5.01***	-5.03***	-5.84***	-5.87***	-5.29***	-5.06***
Period	Full	Full	Full	Full	Full	Full	1999.03–2014.12	2015.01–2025.12

*\*, \*\*, and \*\*\* denote statistical significance at the 10%, 5%, and 1% levels, respectively. Survey Inflation coefficient restricted to unity*

Finally, a comparison of specifications (6) and (7), which partition the full sample at the end of 2014, reveals a notable shift in the structure of breakeven inflation dynamics. In the 1999–2014 period, the restricted model achieves an  $R^2$  of 0.63 with a Synthetic-Cash coefficient of  $-1.33^{***}$ . This suggests that, over this earlier regime, deviations of breakeven inflation from survey-based expectations were well captured by observable liquidity proxies and cyclical conditions. Liquidity distortions were economically meaningful, most

visibly during the 2008–2009 crisis, but largely episodic and mean-reverting, allowing the restricted specification to account for a substantial fraction of the variation. In contrast, during the 2015–2025 period, the Synthetic-Cash coefficient nearly doubles in magnitude to  $-2.53^{***}$ , indicating that breakeven inflation became significantly more sensitive to liquidity conditions. At the same time, however, the  $R^2$  declines to 0.50. The coexistence of a stronger liquidity loading and a lower overall fit suggests that additional forces, beyond survey expectations and measured liquidity proxies, played a larger role in driving breakeven fluctuations in the later regime. These forces could be interpreted with an increased importance of time-varying inflation risk premia and structural shifts in market functioning.

### 3.1 Liquidity-Adjusted Breakeven Inflation

We focus on the most informative restricted first-stage regression specification (5a) from Table 3 to construct the liquidity premium as:

$$\widehat{L}_{n,t} = -\widehat{\alpha}_2 X_t,$$

where  $X_t$  collects liquidity-related variables and  $\widehat{\alpha}_2$  denotes their estimated coefficients. Figure 2 plots the U.S. 10-year breakeven inflation alongside its liquidity-adjusted counterpart. This adjustment removes the liquidity differential in TIPS yields via

$$y_{n,t}^{\text{TIPS,adj}} = y_{n,t}^{\text{TIPS}} - \widehat{L}_{n,t},$$

$$b_{n,t}^{\text{adj}} = b_{n,t} + \widehat{L}_{n,t}.$$

yielding at

$$b_{n,t}^{\text{adj}} = \mathbb{E}_t(\pi_{n,t}) + IRP_{n,t}.$$

which captures expected inflation plus the inflation risk premium.

The comparison between raw and liquidity-adjusted breakeven highlights the importance of accounting for liquidity distortions, particularly during periods of market stress. Al-

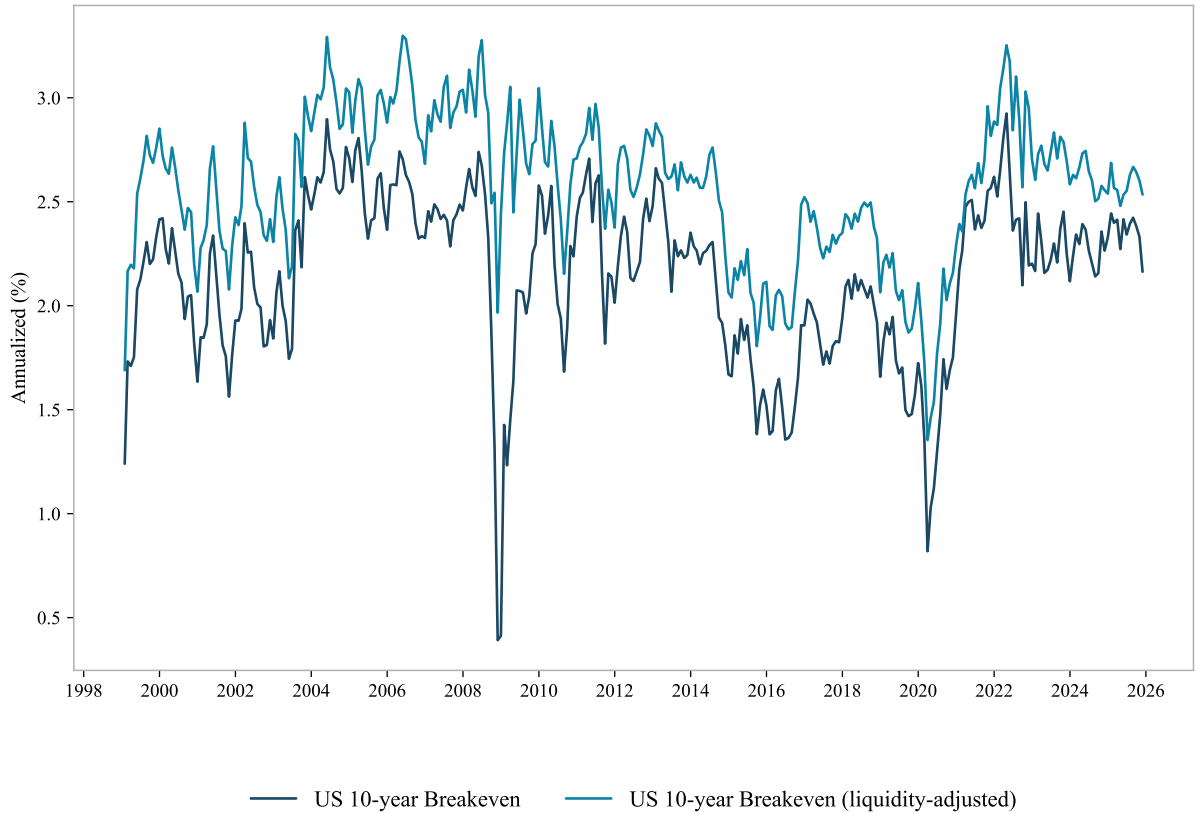


Figure 2: Liquidity-Adjusted U.S. 10 Year Breakeven

though the raw breakeven exhibits a sharp collapse during the 2008–2009 financial crisis, the liquidity-adjusted breakeven remains substantially more stable. This pattern suggests that a significant portion of the observed decline in breakeven inflation during the crisis is attributed to a surge in the TIPS liquidity premium, rather than a substantial downward revision of long-term inflation expectations. Consistent with this interpretation, the estimated liquidity premium reaches its maximum precisely during the crisis period and gradually reverts thereafter, as shown in Figure 3.

In contrast, the decline observed in early 2020, associated with the COVID-19 shock, displays a markedly different pattern. During this episode, both raw and liquidity-adjusted breakeven inflation fall sharply, indicating that the contraction in breakeven inflation cannot be primarily attributed to liquidity distortions. Indeed, the liquidity premium declines between 2020 and 2021 and only increases in subsequent years, suggesting that liquidity conditions in the TIPS market did not deteriorate as they did during the global

financial crisis. The simultaneous decline in both breakeven measures, therefore, points to a genuine downward revision in market-based inflation expectations, likely reflecting extreme macroeconomic uncertainty and heightened concerns about the inflation outlook at the onset of the pandemic. This collapse is consistent with fears of deflation or prolonged low inflation amid a severe global recession driven by an unprecedented economic environment with widespread lockdowns, disruptions to global supply chains, and a sudden contraction in aggregate demand. Immediately thereafter, breakeven inflation rebounded sharply, driven by a combination of policy interventions and shifts in inflation expectations. First, large-scale asset purchases by the Federal Reserve supported inflation-linked security prices and contributed to a compression of real yields. Second, as financial markets stabilized and unprecedented fiscal and monetary stimulus was deployed, inflation expectations began to recover from crisis lows, gradually incorporating concerns about a more inflationary environment.




Figure 3: Estimated U.S. Liquidity Premium

However, even after considering this adjustment,  $b_{n,t}^{adj}$  embeds a potentially time-varying

inflation risk premium ( $IRP_{n,t}$ ), reflecting investors' compensation for bearing inflation uncertainty. Assuming constant or negligible risk premia allows for interpreting liquidity-adjusted breakeven rates as a proxy for pure expectations, but time variation implies otherwise. To disentangle expected inflation from the inflation risk premium, we exploit the no arbitrage hypothesis: under rational expectations and constant premia, excess returns on nominal and inflation-indexed bonds, and, consequently, on breakeven rates, should not be predictable using information available at time  $t$ . Therefore, systematic predictability in breakeven excess return can be interpreted as evidence of time-varying risk premia, enabling the extraction of market-based expectations via

$$\mathbb{E}_t(\pi_{n,t}) = b_{n,t}^{\text{adj}} - IRP_{n,t}.$$

The next section implements this approach through return predictability regressions.

### 3.2 Estimating the Inflation Risk Premium

This section estimates the inflation risk premium embedded in the 10-year breakeven inflation rate using a return predictability framework grounded in the Expectations Hypothesis (EH) of the term structure. The identification strategy follows directly from no-arbitrage asset pricing principles. Under rational expectations and the absence of arbitrage, the yield on an  $n$ -period bond equals the average of expected future short-term rates plus a term premium. If this premium were constant over time, variation in long-term yields would reflect only revisions in expected future short rates, and excess holding-period returns would be unpredictable given information available at time  $t$ . Equivalently, under the EH, the slope of the yield curve carries no information about future excess returns. The same restriction applies to breakeven inflation: if the inflation risk premium were constant, the slope of the breakeven curve would reflect only variation in expected inflation and would have no predictive power over future breakeven excess returns. It follows that systematic predictability of liquidity-adjusted breakeven excess returns by the breakeven term spread constitutes direct evidence of a time-varying inflation risk premium.

Consistent with our analysis, we focus on bonds with a remaining maturity of ten years and define the one-period excess holding-period return as the difference between the realized return on a 10-year bond and the contemporaneous short-term risk-free rate. To assess the time variation in the inflation risk premium, we estimate predictive regressions of one-period liquidity-adjusted breakeven excess returns on their breakeven term spread. When excess returns are statistically predictable, the fitted values from these regressions are used as a proxy for the one-period inflation risk premium associated with holding a 10-year bond.

The 10-year inflation risk premium is then constructed by aggregating the estimated one-period premium over the remaining maturity of the bond. Finally, market-based inflation expectations for the 10-year horizon are recovered by subtracting the estimated inflation risk premium, if any, from liquidity-adjusted 10-year breakeven inflation. This procedure yields a measure of long-term expected inflation that is cleaned of both liquidity distortions and risk compensation, allowing for a direct comparison with survey-based inflation expectations.

The risk-free rate in the nominal bond excess return formula is the three-month Treasury bill rate, constructed by Campbell and Shiller and retrieved from the US Treasury database in the CRSP. Instead, to compute the excess returns of the 10-year TIPS relative to a real risk-free rate, we must address the non-observability of short-term real rates. That is because the US government does not issue inflation-indexed Treasury bills. Following the framework of [Campbell and Shiller \(1991\)](#) and [Pflueger and Viceira \(2011\)](#), we construct a synthetic ex-ante real short rate. To begin with, we define the ex-post real return on a three-month Treasury bill yield as the simple difference between the nominal short-term rate and the annualized overlapping three-month realized inflation rate realized through the CPI-U (Consumer Price Index for All Urban Consumers). Given that inflation is realized over the subsequent quarter, we are not able to anticipate inflation shocks. To isolate just the expected component, the realized real short rate is then modeled as a function of variables observable at time  $t$ , including the contemporaneous nominal short rate, the lagged realized real rate over the previous quarter, and year-over-year inflation observed at

time  $t$ . Since the dependent variable is constructed as a three-month overlapping return observed at monthly frequency, successive observations share common inflation innovations, generating serially correlated residuals. We therefore use Newey–West standard errors with three lags to correct the inference. Finally, to ensure an ex-ante interpretation, the real short rate used in the following analysis is estimated using an expanding-window out-of-sample forecasting procedure: at each date  $t$ , coefficients are estimated using only data available up to  $(t-1)$ . The fitted value at time  $t$  is computed using contemporaneous realization of the regressors.

$$r_t^s = \hat{\alpha}_{t-1} + \hat{\beta}_{1,t-1} y_{1,t}^s + \hat{\beta}_{2,t-1} r_{t-3,t}^{ex\ post} + \hat{\beta}_{3,t-1} \pi_t^{yoy} \quad (2)$$

The initial estimation window is set to 60 months (five years) of monthly observations, and the procedure produces OOS fitted values starting from January 2005, yielding a predictability sample of  $N = 246$  observations covering the period 2005–2025. This choice of warm-up window balances two competing objectives: a sufficiently long initial window to produce stable first-stage coefficient estimates, and a sufficiently early OOS start date to capture the pre-crisis 2005–2008 period, which is economically informative for breakeven dynamics. This fitted value represents the conditional expectation of the three-month real return given information available at time  $t$ , and is therefore interpreted as the synthetic ex-ante real short rate. This methodology isolates the real term premium by removing predictable movements in expected short rates and transitory inflation noise. The results of this regression are shown in Table 4.

Among the three predictors included in the real short rate forecasting equation above, the nominal short rate  $y_{1,t}^s$  emerges as the dominant driver, with a coefficient of 0.82\*\* that is precisely estimated and highly significant. The lagged ex-post real return enters with a positive but imprecisely estimated coefficient of 0.15, consistent with a modest degree of mean reversion in real rates. By contrast, year-over-year inflation  $\pi_t^{yoy}$  enters with a coefficient of  $-0.16$  that is not statistically distinguishable from zero, suggesting that, conditional on the nominal short rate and lagged real return, aggregate inflation adds limited incremental forecasting power for the expected real short rate. The in-sample

Table 4: Forecasted Real Short Rate

	$r_{1,t} = y_{1,t}^{\$} - \pi_{t+1}$
$y_{1,t}^{\$}$	0.82*** (0.20)
$y_{1,t-1}^{\$} - \pi_{t-1,t}$	0.15 (0.11)
$(\pi_{t-3} + \pi_{t-2} + \pi_{t-1} + \pi_t)/4$	-0.16 (0.28)
<i>Const.</i>	-0.02*** (0.004)
<i>p</i> -value	0.00
$R^2$	0.32
$N$	306
Sample	1999.01–2025.06

*Notes:* Standard errors (in parentheses) are HAC-robust with 3 Newey–West lags.

\*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

$R^2$  of 0.32 and the behaviour illustrated in Figure 4 confirm that the OOS fitted series behaves as a smooth, slowly-evolving component of the much more volatile ex-post real short rate, supporting its interpretation as the conditional expectation of the real return on short-term debt.

The ex-post series exhibits large spikes and troughs, down to about  $-10\%$  and up to around  $+14\%$ , reflecting sudden inflation and nominal rate shocks that are not forecastable with information available at time  $t$ . In contrast, the smoothed fitted series reacts gradually to regime shifts, supporting its interpretation as the conditional expectation of the 3-month real return given observable predictors. For this reason, we use this series as the synthetic ex-ante real short rate in the excess-return calculations.

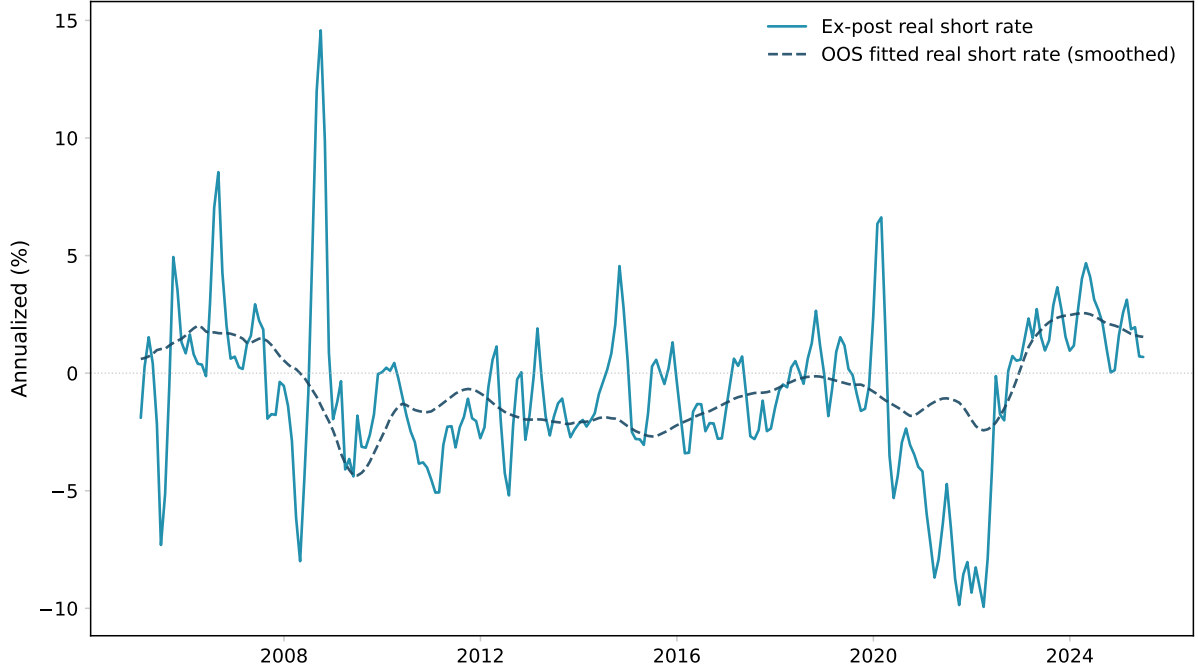


Figure 4: U.S. Ex-Post and Fitted Real Short-Term Rates

At this stage, we construct short-horizon excess returns on nominal Treasury bonds, TIPS, and breakeven inflation using monthly data and a three-month holding period. This choice comes from a theoretical approximation needed to construct short-horizon excess returns on the above using monthly data and a three-month holding period. Specifically, letting  $n=10$  denote maturity in years and  $h=3/12$  the holding period, the three-month excess return on a nominal bond is approximated as

$$xr_{t,t+h}^{\$} = n y_{n,t}^{\$} - (n - h) y_{n,t+h}^{\$} - h y_{1,t}^{\$},$$

where  $y_{n,t}^{\$}$  denotes the  $n$ -year nominal yield at time  $t$  and  $y_{1,t}^{\$}$  is the nominal short rate. Analogously, TIPS excess returns, net of the liquidity premium, are computed as follows:

$$xr_{t,t+h}^{TIPS-L} = n(y_{n,t}^{TIPS} - L_{n,t}) - (n - h)(y_{n,t+h}^{TIPS} - L_{n,t+h}) - h y_{1,t}^{TIPS},$$

where  $y_{1,t}^{TIPS}$  denotes the real short rate. Since yields with remaining maturity  $n - h$  are not directly observed, we proxy them using constant-maturity  $n$ -year yields at time  $t + h$ . Then, the breakeven excess return is the return on a portfolio long on nominal returns

and short on liquidity-adjusted TIPS returns with the same maturity  $m$ :

$$xr_{t,t+h}^b = xr_{t,t+h}^{\$} - xr_{t,t+h}^{TIPS-L},$$

For consistency, all returns are annualized and expressed as percentage points.

To identify time variation in risk compensation, we estimate predictability regressions relating one-period excess returns to slope measures of the yield curve, including the nominal, real, and breakeven term spreads. In models with time-varying risk premia, these slope variables summarize state information about the economy and therefore contain information about future excess returns. The corresponding term spreads are defined as

$$s_{n,t}^{\$} = y_{n,t}^{\$} - y_{1,t}^{\$},$$

$$s_{n,t}^{TIPS-L} = y_{n,t}^{TIPS-L} - y_{1,t}^{TIPS-L},$$

and the breakeven liquidity adjusted term spread as

$$s_{n,t}^b = b_{n,t} + L_{n,t} - b_{1,t}.$$

We estimate six specifications organized in Table 5, corresponding to three dependent variables (TIPS excess return, breakeven excess return, liquidity excess return) each regressed on a univariate slope specification and a multivariate specification including all three spread measures and the liquidity premium level  $L_{n,t}$  as an additional control, to account for any potential distortions arising from market frictions. Inference throughout relies primarily on moving block bootstrap p-values, with Newey–West standard errors reported alongside for comparability with the existing literature. This dual reporting strategy is motivated by well-documented limitations of standard asymptotic inference in this setting. Newey–West standard errors are consistent under heteroskedasticity and serial correlation, but their finite-sample performance deteriorates when regressors are highly persistent, as is the case for the breakeven term spread and the estimated liquidity premium, both of which display substantial autocorrelation over the sample. In such set-

tings, NW is known to under-correct for serial correlation, producing standard errors that are too small and correspondingly inflated t-statistics, a problem extensively documented by [Andrews \(1991\)](#) and [Hodrick \(1992\)](#). To address these finite-sample concerns, we compute for each coefficient, p-values using a moving block bootstrap under the null hypothesis of no predictability. The procedure resamples contiguous blocks of length  $b = 12$  months from the residuals of the restricted model and re-estimates the full regression on each bootstrap sample, producing an empirical null distribution for the HAC t-statistics. We use  $B = 2,000$  bootstrap replications and a fixed seed for reproducibility. P-values are computed as one-sided. Excess returns on breakeven rates display statistically significant predictability. In particular, the estimated coefficient on the liquidity-adjusted breakeven term spread is positive and statistically significant at the 5% level using bootstrap inference, indicating non-negligible time variation in inflation risk compensation. In contrast, excess returns on liquidity-adjusted TIPS exhibit weaker predictability. Overall, these results suggest that while return predictability in inflation-linked markets is attenuated in the post-2005 period, breakeven term spreads retain information about future excess returns. To extract the one-period inflation risk premium, we use the fitted values from specification (3), which regresses liquidity-adjusted breakeven excess returns on the liquidity-adjusted breakeven term spread alone. This choice is motivated by two considerations. First, specification (3) most directly isolates the component of breakeven excess returns attributable to the slope of the inflation-linked yield curve, which under the expectations hypothesis should be zero in the absence of time-varying risk premia. Second, excluding the liquidity premium from this regression avoids a mechanical channel through which our estimated liquidity measure could contaminate the IRP proxy. Since we assumed that predictable variation in breakeven excess returns reflects time variation in inflation risk compensation and we ruled out the role of liquidity in predictability, we can safely interpret the fitted values from the above regression, which we denote with  $\hat{\phi}_{t,t+1}$ , as a proxy for the one-period inflation risk premium associated with holding a bond of maturity  $T$  from  $t$  to  $t + 1$ .

To retrieve the long-horizon inflation risk premium from the estimated one-period measure, we assume the latter follows a stationary autoregressive process. Specifically, we fit

an AR(1) model to the estimated quarterly IRP series:

$$\hat{\phi}_t = \mu + \rho\hat{\phi}_{t-1} + \eta_t$$

estimated with Newey–West standard errors (2 lags).

The 10-year inflation risk premium is then constructed as the expected average of future one-period premia over the 40-quarter horizon of the bond:

$$\hat{\phi}_{t,t+T}^T = \frac{1}{T} \sum_{i=1}^T \mathbb{E}_t \hat{\phi}_{t+i-1,t+i}^T.$$

where the AR(1) projection formula

$$\mathbb{E}_t[\phi_{t+i}] = \frac{\mu}{1-\rho} + \rho^i \left( \hat{\phi}_t - \frac{\mu}{1-\rho} \right)$$

is applied iteratively for each quarter  $i = 1, \dots, 40$ . The AR(1) parameters from our sample are  $\hat{\mu} = 0.137$  and  $\hat{\rho} = 0.633$  ( $t(\rho) = 5.66$ ), confirming that the one-period IRP is persistent but mean-reverting, with an unconditional mean of approximately 0.37 percentage points (annualized).

Having obtained an estimate of the 10-year inflation risk premium  $\hat{\phi}_{t,t+T}^T$ , we recover market-based inflation expectations by subtracting this component from liquidity-adjusted breakeven inflation,

$$\mathbb{E}_t(\pi_{n,t}) = b_{n,t}^{\text{adj}} - \hat{\phi}_{t,t+T}^T.$$

This procedure yields a measure of long-term expected inflation purged of both liquidity distortions and compensation for inflation risk, allowing a clean decomposition of breakeven inflation into expectations and risk premia and facilitating direct comparison with survey-based measures of inflation expectations.

Given that we need the OOS real short rate to compute the excess returns, and that OOS fitted real rate is available in January 2004, the predictability sample then begins in January 2005. The full sample for the predictability regressions, therefore, covers 2005:01–2025:06. The estimated 10-year inflation risk premium over the 2005–2025 period

Table 5: Liquidity-Adjusted Bond Excess Return Predictability.

	(1)	(2)	(3)	(4)	(5)	(6)
	$xr_{n,t+1}^{TIPS-L}$	$xr_{n,t+1}^{TIPS-L}$	$xr_{n,t+1}^{b+L}$	$xr_{n,t+1}^{b+L}$	$r_{n,t+1}^L$	$r_{n,t+1}^L$
$(y_{n,t}^{TIPS} - L_{n,t}) - y_{1,t}^{TIPS}$	1.03	2.14		0.32		-0.64
Newey–West SE	(0.95)	(1.01)		(0.50)		(0.48)
Bootstrap p-value	18.6%	5.0%		25.4%		9.1%
$(b_{n,t} + L_{n,t}) - b_{1,t}$		3.79	1.86	1.92		-1.14
Newey–West SE		(1.77)	(0.77)	(0.98)		(0.77)
Bootstrap p-value		3.7%	1.2%	2.2%		6.9%
$L_{n,t}$		0.85		-2.46	15.80	15.95
Newey–West SE		(8.78)		(6.58)	(4.50)	(4.29)
Bootstrap p-value		47.0%		39.0%	0.4%	0.3%
Const.	-0.12	-2.76	-0.07	0.74	-7.05	-6.11
Newey–West SE	(1.56)	(4.55)	(0.80)	(3.08)	(1.86)	(2.11)
Bootstrap p-value	29.0%	22.6%	22.6%	52.9%	10.1%	0.2%
$p$ -value	0.28	0.10	0.02	0.08	0.00	0.00
$R^2$	0.01	0.06	0.04	0.04	0.21	0.23
$N$	246	246	246	246	246	246

is presented in Figure 6.

The series remains strictly positive for the whole sample, fluctuating around a mean of approximately 40 basis points. This suggests that investors generally demand positive compensation for exposure to long-horizon inflation uncertainty, even during periods in which inflation does not seem to cloud the economy. These findings are consistent with a large macro-finance literature documenting that long-horizon inflation risk premia are generally small but cyclical. Studies based on no-arbitrage term-structure models and inflation-linked bonds generally find that the 10-year US inflation risk premium fluctuates around a modest positive mean, on the order of a few tens of basis points: [Buraschi and Jiltsov \(2005\)](#) report 70 bps historically but lower recently; ([Hördahl et al. \(2010\)](#) and [Hördahl and Tristani \(2014b\)](#)) find similar modest means in low-inflation regimes; [Grishchenko and Huang \(2012\)](#) align with 10–50 bps averages adjusted for liquidity. Early contributions using pre-2000 data often report higher average premia, reflecting periods of less credible nominal anchors and more volatile inflation, whereas more recent estimates in the low-inflation regime emphasise compressed risk premia, thus meaning that variation in long-term breakeven rates is predominantly driven by revisions in expected inflation

rather than large swings in risk compensation, if not during episodes of severe systemic stress. The dynamics of the premium reveal economically significant shifts across different macroeconomic regimes, confirming its time-varying nature. The sharp decline in 2008 corresponds to the Global Financial Crisis, when the 10-year inflation risk premium dropped to an exceptionally low level due to a sudden contraction in global demand, and severe financial stress generated strong disinflationary pressures. The elevated values observed in 2010–2012 indicate a large inflation risk premium despite subdued realized inflation, consistent with heightened macroeconomic uncertainty, which produced pronounced volatility, and concerns related to the newly introduced unconventional monetary policies. It then declines steadily through the mid-2010s, reaching local lows around 2015–2016. This compression coincides with a shift towards a subdued inflation regime, well-anchored expectations, and strong monetary policy credibility. The onset of the COVID-19 pandemic induced a sharp, albeit transitory, contraction in the estimated premium. This decline reflects an acute flight-to-safety episode, where the sudden uncertainty led nominal Treasuries to serve as a hedge against idiosyncratic deflationary risks. Following this trough, the IRP experienced a marked expansion throughout 2021–2022. This surge coincided with the post-pandemic inflation spike, persistent supply-side bottlenecks, and a widening disagreement among market participants regarding the persistence of price pressures. This peak indicates that investors required substantial additional compensation as the credibility of the long-run inflation target was tested. The premium eventually moderated as aggressive monetary tightening by the Federal Reserve re-established investors' confidence in price stability. Broadly, the estimated series captures well-documented macroeconomic cycles and provides a plausible decomposition of liquidity-adjusted breakeven inflation. The modest absolute magnitude of the IRP, coupled with its countercyclical behavior, suggests that the variance in long-term breakeven rates is primarily driven by revisions in expected inflation rather than large-scale shifts in risk compensation, with the exception of periods marked by extreme systemic stress. Finally, the two charts highlight the distinction between the inflation risk premium and inflation expectations, which are not necessarily tightly or mechanically linked. In particular, there are periods in which long-term inflation expectations remain well anchored, while investors simultaneously de-

mand additional compensation for bearing inflation risk.




Figure 5: U.S. Liquidity-adjusted Breakeven decomposition




Figure 6: 10-year Inflation Risk Premium

As a robustness check, we re-estimate the original specification (5a) of Table 3 includ-

ing the inflation risk premium extracted in the previous steps ( $\text{IRP}_{10y}$ ) as an additional regressor. Under our theoretical framework, breakeven inflation decomposes as:

$$b_{10,t} = \mathbb{E}_t[\pi_{t,t+10}] + \text{IRP}_{10,t} - L_{10,t}$$

If our two-stage procedure successfully isolates the liquidity premium in the first stage and the inflation risk premium in the second stage, then regressing breakeven inflation on survey expectations (with the coefficient restricted to unity), liquidity proxies, and the extracted IRP should yield a coefficient on IRP equal to one. This provides a direct test of the internal consistency of our decomposition. Table 6 presents three specifications. Column (1) shows the baseline restricted regression estimated over the full 1999–2025 sample. Column (2) re-estimates the same specification over the 2005–2025 subsample to align with the period for which  $\text{IRP}_{10y}$  is available. Column (3) augments the 2005–2025 specification by including the extracted inflation risk premium as an additional regressor. The results strongly support our theoretical prediction. The coefficient on  $\text{IRP}_{10y}$  in column (3) is estimated at 0.90\*\*\* (HAC standard error = 0.34), statistically indistinguishable from unity. Formally testing the null hypothesis  $H_0: \beta_{\text{IRP}} = 1$  yields a t-statistic of  $-0.03$  and a p-value of 0.76, providing no evidence against the restriction. This near-exact unit coefficient confirms that our extracted inflation risk premium enters breakeven inflation one-for-one, precisely as the decomposition in the equation above predicts. Moreover, the inclusion of  $\text{IRP}_{10y}$  improves the model’s explanatory power. The  $R^2$  increases from 0.48 in column (2) to 0.51 in column (3), indicating that the extracted risk premium captures systematic movements in breakevens beyond those explained by liquidity and macroeconomic controls alone. This gain in explained variation, even if marginal, is economically meaningful and suggests that the inflation risk premium exhibits substantial time-variation over the 2005–2025 period. The off-the-run spread, which measures the liquidity advantage of recently issued nominal Treasuries and serves as a broad gauge of Treasury market functioning, enters with the expected negative sign and is statistically significant in the full sample ( $-0.51^{**}$ ,  $\text{SE} = 0.24$ ). It indicates that when nominal Treasury market liquidity deteriorates, breakeven inflation tends to compress, consistent

with correlated liquidity stress in the TIPS market. This relationship becomes statistically insignificant in the post-2005 subsample (0.48 and 0.46 in columns (2) and (3)), suggesting that once the sample is restricted, the off-the-run spread loses its explanatory power for breakeven variation, likely reflecting the Federal Reserve’s asset purchases compressing liquidity premia in ways that obscure this channel. The synthetic-cash measure dominates throughout the post-2005 period: it enters with a coefficient of  $-2.06^{***}$  (SE = 0.27) in column (2) and  $-1.89^{***}$  (SE = 0.28) in column (3), both significant at the 1% level, confirming that TIPS-specific illiquidity is the primary liquidity driver of breakeven compression. The  $IRP_{10y}$  enters positively and significantly ( $0.90^{***}$ , SE = 0.34), and the null hypothesis that its coefficient equals unity cannot be rejected, indicating that the inflation risk premium passes through into breakeven inflation nearly one-for-one, as theory would predict.

Table 6: Restricted Specification (5a): Robustness to the Inclusion of the Inflation Risk Premium.

	$b_{n,t} - \pi_t^e$		
	(1)	(2)	(3)
Off-the-Run Spr.	-0.51** (0.24)	0.48 (0.36)	0.46 (0.38)
Synthetic-Cash	-1.56*** (0.28)	-2.06*** (0.27)	-1.89*** (0.28)
$IRP_{10y}$			0.90*** (0.34)
Survey Inflation	1.00	1.00	1.00
CFNAI	0.10*** (0.04)	0.11*** (0.04)	0.10*** (0.04)
$p$ -value	0.00	0.00	0.00
Observations	320	246	246
$R^2$	0.41	0.48	0.51
Sample period	Full	2005.01–2025.06	2005.01–2025.06

Notes: HAC (Newey–West, 3 lags) standard errors in parentheses. The null hypothesis that the coefficient on IRP equals unity cannot be rejected ( $p = 0.76$ ).

\*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

## 4 Comparison with Survey-Based Expectations

As a final step in our analysis, we evaluate the economic plausibility of the market-based inflation expectations recovered from our decomposition by comparing them with standard survey-based measures of long-term inflation expectations. These surveys are typically interpreted as direct measures of agents' beliefs about future inflation and therefore provide a natural benchmark for our indicator. While the Survey of Professional Forecasters, which captures the median CPI inflation expectations over the next ten years constructed from forecasts by professional economists, is the most prominent and reliable of these measures, we refrain from using it here because it already enters our liquidity extraction model, and including it in this exercise could induce a mechanical correlation with our market-based expectations. Instead, we rely on an alternative, widely used survey indicator: the long-horizon inflation expectations from the University of Michigan Survey of Consumers. This series captures households' expectations for inflation over a 5–10 year horizon, elicited through the question “By about what percent per year do you expect prices to go up or down, on the average, during the next 5 to 10 years?”. Despite differences in respondents, horizon definitions, and information sets, both series are commonly used as reference measures of long-run inflation expectations and are of central importance to the Federal Reserve.

Survey-based measures of inflation expectations play a key role in the Federal Reserve's assessment of inflation dynamics and the credibility of monetary policy. Both indicators are closely watched as gauges of whether longer-run inflation expectations remain anchored near the Federal Reserve's inflation objective. Within the Federal Reserve System, they are regularly featured in policy briefings and public communications as cross-checks on longer-run anchoring, and are often contrasted with long-maturity breakeven rates and forward rates to infer the contribution of risk and liquidity premia to market-based measures. Persistent deviations of survey-based expectations from the target are viewed as signals of potential risks to inflation anchoring, while temporary fluctuations are typically given less weight in policy deliberations.

This prominent role of surveys in both policy monitoring and the academic literature

provides a natural backdrop for our comparison between household expectations and market-based measures. The theoretical case for expecting a persistent gap between household survey expectations and market-based measures is well established. Many papers perform the same comparison and we identified two streams of literature. On one side, [Ang et al. \(2007\)](#) provide the most comprehensive forecasting comparison, showing that survey-based measures, including the Michigan Survey and the Survey of Professional Forecasters, outperform term structure models, Phillips curve specifications, and ARIMA benchmarks in out-of-sample inflation forecasting. This result implies that surveys have information not fully priced into financial markets. At the same time, the Michigan Survey is known to embed a systematic upward bias relative to both realised inflation and the expectations of professional forecasters and financial market participants, leading us to the second stream of literature. [Coibion and Gorodnichenko \(2015\)](#) formalise this in the information rigidity framework: household respondents update their beliefs infrequently and with a lag, processing new macroeconomic information only partially each period. As a result, during extended periods of low inflation, household surveys chronically over-predict, as respondents fail to revise down toward the new steady state at the speed implied by full-information rational expectations.

The underlying mechanism has been identified by [D'Acunto et al. \(2021\)](#), who show that households form inflation expectations primarily from prices they encounter in their daily shopping experience, particularly food and energy, rather than from aggregate CPI developments. Because grocery prices exhibit higher average inflation and greater volatility than the overall price level, and because individuals systematically assign greater weight to price increases than to price decreases, household inflation expectations carry a structurally positive bias. Consistent with this, [Weber et al. \(2022\)](#) document that this upward bias is one of the most robust and replicated facts in the empirical inflation expectations literature, present across advanced economies characterised by low and stable inflation. In addition, the authors show that the bias also varies systematically across demographic groups: for instance, it is larger for women, lower-income households, and individuals with lower cognitive ability, all consistent with a mechanism driven by differential exposure to salient prices rather than by rational information processing.

In contrast, market-based measures reflect the aggregated beliefs of sophisticated financial market participants who are directly exposed to the financial consequences of forecast errors and who update continuously as new information arrives. This population of agents is more consistent with the professional forecasters tracked by the SPF, and far less consistent with the median Michigan Survey respondent. Consequently, a persistent level gap between the Michigan Survey and market-based expectations is not a sign of irrationality in either measure, but it just reflects the fact that they represent genuinely different populations with different information sets and updating abilities.

Figure 7 plots our market-based inflation expectations series alongside the Michigan Survey 5–10 year measure from January 2005 to November 2025. Three observations stand out. The first observation concerns the pre-crisis period of 2005–2008. Both series track each other relatively closely, with the Michigan Survey hovering near 2.8–3.2% and the market-based measure oscillating around 2.5–2.7%, consistent with a stable macroeconomic environment in which salient price experiences and financial market signals were broadly aligned. The spike in the Michigan Survey toward 3.4% around 2008 coincides with a period of uncertainty, while market-based measures were already beginning to price in the disinflationary consequences of the unfolding financial crisis. The most dramatic pre-2010 divergence occurs during the acute phase of the Global Financial Crisis. The market-based measure collapses sharply, bottoming near 1.6%, reflecting the combination of severe TIPS illiquidity, flight-to-safety flows into nominal Treasuries, and a pricing-in of deflationary tail risks. The Michigan Survey, by contrast, barely moves, remaining anchored near 2.8–3.0%, consistent with the information rigidity framework of [Coibion and Gorodnichenko \(2015\)](#): households, updating infrequently from salient everyday prices, did not register the deflationary dynamics being rapidly priced into financial markets. This episode is precisely the type our decomposition is designed to address, as raw breakeven inflation was severely distorted by liquidity premia during this period. For the remainder of the sample, the Michigan Survey lies approximately 50 to 80 basis points above the market-based measure. This wedge is entirely consistent with the aforementioned structural upward bias documented in the literature. In addition, the Michigan Survey hovers near 2.8–3.0% across the low-inflation decade of 2010–2019, while our market-based meas-

ure gravitates around 2.0–2.5%, more closely aligned with the Federal Reserve’s stated price stability objective, suggesting that their long-run component reflects a durable anchoring to an inflation norm shaped primarily by lived experience rather than continuous revision to macroeconomic fundamentals. But this pattern is not only a US peculiarity: an ECB Working Paper demonstrates that market-based and survey-based measures of inflation expectations are not directly comparable, owing to differences in construction and features, with household surveys systematically exceeding market-based and professional measures in low-inflation environments.

The second consideration is that, despite the persistent level gap, both series co-move meaningfully during periods of macroeconomic regime shifts, and the gap narrows substantially during the post-pandemic inflation surge of 2021–2022. This convergence leads to a natural interpretation. When inflation becomes sufficiently salient, even the slow-moving and inertia-prone household expectations process eventually catches up with the signals already embedded in financial markets. The information-rigidity framework discussed above predicts exactly this form of state dependence: the degree of rigidity falls during large and persistent macroeconomic shocks, as the salience of new information rises, and agents invest more resources in updating. The convergence in Figure 7 around 2022 is consistent with this mechanism and corresponds to a period in which the decades-long anchoring of household long-run expectations near the 2% norm appears to have come under meaningful strain. Notably, the market-based measure also briefly dips toward 1.0% in early 2020 at the onset of the COVID shock, again reflecting a combination of genuine deflationary fears and acute TIPS market illiquidity during the March 2020 liquidity crisis, an episode that, much like the GFC, illustrates the limits of raw breakeven inflation as a clean expectations signal.

Finally, the most striking feature of Figure 7 is the sharp terminal spike in the Michigan Survey toward approximately 4.4% in late 2025, while the market-based measure remains comparatively contained and shifted down after the 2022-pick, near 2.3–2.4%. This is the most pronounced divergence in the entire sample. The direction of this divergence is the opposite of what a simple information rigidity story would predict. If households were merely slow to update relative to financial markets, one would expect the survey to lag the

market-based measure rather than to overshoot it dramatically. The spike in the Michigan Survey is therefore unlikely to reflect a genuine revision in long-run inflation expectations of the kind that would be priced into bond markets. Several alternative explanations are more consistent with the data. The first explanation relates to the well-established grocery-price mechanism documented by [D'Acunto et al. \(2021\)](#). Tariff announcements in early 2025 generated immediate fear in price increases in consumer goods categories, such as electronics, automobiles, and household goods, precisely the type of salient price experiences that drive household inflation expectations upward, irrespective of their long-run macroeconomic significance. Households exposed to these signals would thus revise their long-run expectations upward in a way not replicated by bond market participants, who focus instead on aggregate demand, supply dynamics, and monetary policy credibility. By contrast, the stability of our market-based measure during this period is revealing. The decomposed inflation expectations series remains near 2.3%, broadly consistent with the Federal Reserve's inflation objective, indicating that financially sophisticated market participants did not interpret the 2025 tariff shock as a permanent regime shift in long-run inflation dynamics. This episode illustrates a broader point about these measures' distinct informational roles. While literature shows that the Professional Survey of Forecasters (SPF) outperforms term structure models in inflation forecasting, household surveys like the Michigan one, despite embedding the lived price experiences of a broad population, are sensitive to salient price shocks, sentiment, and political polarization during periods of uncertainty. When these forces dominate, as in late 2025, the Michigan Survey may reflect a mixture of genuine expectations and noise, while our market-based measure offers a cleaner signal of long-run inflation expectations relevant for monetary policy.

The second potential explanation relates to the well-documented sensitivity of the Michigan Survey to political sentiment. [Coibion and Gorodnichenko \(2025\)](#) document unprecedented partisan polarization in U.S. household inflation expectations over recent years, with notable gaps emerging between Republican- and Democrat-identifying respondents. In periods of elevated political salience, such as those surrounding trade policy, tariffs, and fiscal debates, the Michigan Survey median may partly reflect partisan sentiment alongside genuine probabilistic assessments of future inflation. The late 2025 spike coincides

with a time of considerable policy uncertainty tied to the new administration’s tariff proposals, during which available evidence suggests partisan gaps in inflation expectations widened further following the 2024 election, potentially contributing to upward pressure on the aggregate median without implying a broad reassessment of the macroeconomic outlook. Taken together, the comparison yields three main conclusions. First, the broad co-movement between our market-based measure and the Michigan Survey over most of the sample, particularly during 2013–2019 and the 2021–2023 convergence, provides substantive external validation that our decomposition successfully isolates genuine inflation expectations from non-expectational components. Second, the persistent level difference is informative about the distinct nature of these measures rather than a flaw in either. The Michigan Survey’s upward bias reflects the grocery-price mechanism documented by [D’Acunto et al. \(2021\)](#), not measurement error. By contrast, our market-based measure captures the beliefs of agents for whom forecast errors carry direct financial consequences, rendering it a cleaner proxy for the inflation expectations that drive asset prices and monetary transmission. Neither measure, therefore, serves as an absolute benchmark for the other, leading us to the third consideration: the two measures are best viewed as complementary rather than competitive. Our market-based expectations provide high-frequency updates, while the Michigan Survey, despite its structural bias, embeds the price experiences of a broad cross-section of households, offering insights into the distributional dynamics of inflation perceptions that influence consumption and wage-setting behavior.




Figure 7: Market-based versus survey-based long-run U.S. inflation expectations.

#### 4.1 Other Non-Expectation Components of Breakeven Inflation

Beyond the liquidity premium and the inflation risk premium, the observed breakeven inflation rate embeds several additional technical components that create a wedge between raw breakeven rates and true expected inflation. Although these components are typically small relative to the liquidity and risk premia estimated above, they are conceptually distinct and can become non-negligible in specific market environments. We do not explicitly estimate or remove these components from our preferred measure; instead, we document their expected direction and approximate magnitude to assess whether they are likely to affect our conclusions materially.

The first wedge arises from the convexity effect. Bond prices are convex functions of yields: by Jensen’s inequality, the expected return on a bond exceeds that implied by the expected yield path, as prices respond asymmetrically to yield changes, gaining more from declines than losing from increases. In breakeven inflation, this generates a wedge between nominal and real yield curves, unrelated to expectations or risk compensation. Nominal bonds exhibit greater convexity than comparable TIPS due to higher nominal yield volatility,

so nominals benefit more from this property, mechanically compressing breakevens below pure inflation expectations. Estimates vary across models: [Ang et al. \(2007\)](#) find effects below 1 basis point at long maturities, while a report by Lehman Brothers (2006) reports values in the 5–15 basis point range at the 10-year horizon under standard term-structure calibrations. Although model-dependent, the bias is small relative to the liquidity (40 bps) and inflation risk premia (10 to 60 bps) estimated in this thesis.

The second wedge we must account for is the indexation lag. Given that TIPS cash flows are adjusted for inflation and, specifically, they track the non-seasonally-adjusted consumer price index for all urban consumers, which is published with a delay, the Treasury cannot index TIPS to “current” inflation in real time, and it proceeds with a 3-month indexation lag for operational transparency. For this reason, since investors will effectively receive compensation for inflation realized prior to purchase while remaining exposed to inflation risk in the final 3 months before maturity, they will ask for additional compensation. As a result, breakeven inflation does not correspond exactly to the expected average inflation over the stated maturity, although the quantitative impact at the 10-year horizon is generally limited. [D’Amico et al. \(2014\)](#) estimate this wedge and find a small impact varying between  $-5$  and  $+3$  basis points at the 10-year maturity. Similar to the convexity wedge, the quantitative impact on our estimations is limited and therefore does not materially affect our empirical results.

The third and last wedge concerns the deflation floor. As described at the beginning of the thesis, U.S. TIPS guarantee repayment of at least par at maturity, even in the presence of cumulative deflation. This feature embeds an implicit put option on the price level, effectively providing insurance against outcomes below a floor, whose value is reflected in the lower yield investors accept on TIPS. When deflation risk is non-negligible, the option has positive value, reducing real yields and mechanically raising breakeven inflation relative to pure expectations. The deflation floor is highly state-dependent: it is negligible in normal inflation environments but becomes relevant during deflationary episodes, such as 2008–2009 and briefly in early 2020. In typical conditions, its impact amounts to only a few basis points. [D’Amico et al. \(2014\)](#) show that the deflation floor is small in normal times, ranging lower than 5 basis points in absolute value, but widened substantially to

about  $-20$  bps during the 2007–2009 financial crisis. In such episodes, the increase in the value of deflation protection may interact with liquidity dynamics in complex ways. Although it exerts upward pressure on breakevens, partially offsetting the convexity bias, the overall net effect varies with the inflation and interest-rate regime.

Taken together, these wedges are small relative to the liquidity and inflation risk premia and are unlikely to materially affect our results. We therefore abstract from these components in our baseline specification.

## 5 Conclusions

This thesis develops a transparent, regression-based framework for decomposing U.S. 10-year breakeven inflation into expected inflation, a liquidity premium, and an inflation risk premium. The analysis extends and updates the work of [Pflueger and Viceira \(2011\)](#) through November 2025, but moves beyond their original objectives in three successive steps. Where their contribution establishes a procedure for extracting the TIPS liquidity premium and documents the predictability of breakeven excess returns, we treat these results as a foundation rather than an endpoint. Drawing on the return predictability literature, we use the evidence of time-varying excess returns to quantify the inflation risk premium embedded in liquidity-adjusted breakevens, recover a market-implied measure of long-run expected inflation, and assess its informational content relative to the survey-based expectations of the Michigan Survey of Consumers. The primary objective is to evaluate the reliability of market-based breakeven rates in isolating underlying inflation expectations across distinct macro-financial regimes, including the post-GFC period of unconventional monetary policy, the COVID-19 shock, the post-pandemic inflationary surge, and the subsequent tightening cycle.

The analysis yields three principal findings. First, liquidity distortions in the TIPS market are economically significant, time-varying, and highly dependent on the macroeconomic regime. First-stage regressions show that observable liquidity proxies, especially the Synthetic-Cash breakeven spread and the Off-the-Run spread, explain much of the variation in breakevens beyond what survey expectations and cyclical conditions capture. The improvement in goodness of fit is striking: the preferred specification raises the  $R^2$  from 0.28 to 0.51, implying that liquidity conditions account for nearly a quarter of the total monthly breakeven variation, independently of macroeconomic fundamentals. Once adjusted, the decline in breakeven in 2008–2009 proves relatively stable. In contrast, during the COVID-19 episode, both raw and liquidity-adjusted breakevens declined together, indicating a greater contribution from downward revisions in long-term inflation expectations rather than impaired functioning of the TIPS market. This regime asymmetry shows that a large share of what is observed in raw breakeven inflation at any given point in time

may reflect liquidity conditions rather than underlying inflation beliefs. Second, liquidity-adjusted breakeven excess returns show strong and statistically significant predictability, providing direct evidence of a time-varying inflation risk premium. The estimated 10-year inflation risk premium fluctuates between 10 and 60 basis points, remains positive for the entire sample, and exhibits variation that is economically interpretable and consistent with the theoretical prior that the IRP reflects compensation for inflation uncertainty rather than for expected inflation itself.

Third, and most centrally to the empirical contribution of this thesis, the market-based inflation expectations series recovered from the decomposition is compared against the Michigan Survey of Consumer Expectations, an independent, widely monitored benchmark of long-run household inflation beliefs that plays a role in Federal Reserve policy deliberations. This comparison is the ultimate test of the decomposition: if the procedure successfully strips liquidity distortions and risk compensation from observed breakevens, what remains should behave like a genuine measure of inflation expectations, comoving with survey indicators in economically interpretable ways while preserving the high-frequency responsiveness that makes market-based measures valuable in the first place. The results of this comparison are striking across three distinct dimensions. The two series exhibit a stable gap of 50–80 basis points, reflecting the well-documented structural upward bias in household expectations. Households form inflation expectations primarily from the prices they encounter in daily consumption, particularly food and energy, which exhibit higher average inflation and greater volatility than the aggregate price level. Financially exposed market participants, by contrast, continuously update on the full information set relevant to monetary policy transmission. The persistent wedge between the two series is therefore a direct reflection of this population heterogeneity, not measurement error. The second dimension concerns the 2021–2022 convergence. As realised inflation surged and its persistence became undeniable, the gap between the two series narrowed sharply, with household expectations catching up toward levels already embedded in financial markets. The fact that the market-based measure anticipated this shift, rising earlier and more rapidly than the Michigan Survey, validates its usefulness as a leading indicator of expectation dynamics, even relative to a widely used survey designed

for exactly this purpose. The third and most striking dimension is the late-2025 divergence. The Michigan Survey spikes while the market-based expectation shows stability of measure during this same episode. This divergence is perhaps the clearest illustration in the entire sample of why the two measures are best understood as complements rather than substitutes: in normal times both provide useful and broadly consistent signals; in periods of political and economic turbulence, the market-based measure offers a cleaner read of the expectations that drive asset prices and monetary transmission, while the Michigan Survey captures the consumption-relevant perceptions that feed into wage-setting and household spending decisions, both of which matter for the full transmission of monetary policy, but in different ways and at different horizons.

Taken together, these findings provide a clear message regarding the use of breakeven inflation as a policy indicator. Raw breakevens are an imperfect and sometimes misleading signal of underlying inflation expectations, not because markets are uninformative, but because they include compensation for liquidity and inflation risk that varies with macroeconomic conditions. The regression-based decomposition developed here provides a computationally simple and real-time implementable way to remove these components, without requiring the latent factor extraction or pricing-kernel assumptions of structural affine models. While the latter remains the benchmark for fully specified yield-curve analysis, the present framework offers a practical complement that prioritises robustness and interpretability. These qualities are especially valuable during periods of financial stress, when structural assumptions are most likely to be strained, and policymakers most urgently need clear signals about the state of inflation expectations.

## 6 Limitations

We must recognize several limitations of the present analysis. First, the liquidity extraction in the first stage rests on the identifying assumption that survey-based inflation expectations, specifically the SPF 10-year CPI forecast, adequately proxy true underlying expected inflation. While this assumption is reasonable, given that this survey is widely

recognized as the most reliable measure of inflation expectations and is partially validated by the unit coefficient on survey expectations across all specifications, the surveys themselves exhibit well-documented biases. To the extent that the SPF systematically deviates from true expectations, the estimated liquidity premium will absorb part of this deviation, potentially combining misestimation of expectations with genuine liquidity distortions. Second, a limitation of this study is the counterintuitive result regarding relative transaction volumes. In our sample, this variable does not perform as expected, likely due to the absence of a proper non-stationary proxy. Future research should address this issue by developing and testing alternative specifications that more accurately capture the dynamic nature of transaction volumes. Finally, three technical components of the wedge between breakeven inflation and true expected inflation, the convexity effect, the indexation lag, and the deflation floor embedded in TIPS, are described and discussed in Section 4.1 but not explicitly estimated or removed from the preferred measure. We note that their combined impact is small relative to the liquidity and inflation risk premia under normal conditions, but it may become non-negligible in specific market environments, particularly during deflationary episodes when the deflation floor acquires option value. These limitations naturally point to several directions for future research. The most immediate extension would be to explicitly price the technical wedges identified above. Second, extending the framework to euro area inflation-linked bonds would enable a cross-country assessment of liquidity and risk premia in a market with distinct microstructural features and a different monetary policy regime, providing a natural robustness check on the generality of the results. Finally, a natural extension would consist of expanding the decomposition to multiple maturities, specifically the 5-year and 30-year TIPS points alongside the 10-year, would yield a term structure of liquidity and inflation risk premia, providing a substantially richer characterisation of how compensation for inflation uncertainty varies across horizons and regimes. The 5-year breakeven is more sensitive to near-term inflation dynamics and typically embeds a larger relative liquidity premium given the thinner TIPS market at that maturity, while the 30-year horizon is dominated by long-run anchoring dynamics and a more valuable deflation floor option. A joint decomposition across maturities would allow direct inference about the slope of the IRP

term structure, a quantity of direct relevance for monetary policy, as an inverted IRP term structure may signal that near-term inflation uncertainty dominates long-run anchoring concerns.

## References

- Andrews, D. W. K. (1991). Heteroskedasticity and autocorrelation consistent covariance matrix estimation. *Econometrica*, 59(3):817–858.
- Ang, A., Bekaert, G., and Wei, M. (2007). Do macro variables, asset markets, or surveys forecast inflation better? *Journal of Monetary Economics*, 54(4):1163–1212.
- Barone, J., Chaboud, A., Copeland, A., Kavoussi, C., Keane, F., and Searls, S. (2022). The global dash for cash: Why sovereign bond market functioning varied across jurisdictions in march 2020. Technical Report Staff Report 1010, Federal Reserve Bank of New York.
- Bauer, M. D. and McCarthy, E. (2015). Can we rely on market-based inflation forecasts? FRBSF Economic Letter 2015-26.
- Buraschi, A. and Jiltsov, A. (2005). Inflation risk premia and the expectations hypothesis. *Journal of Financial Economics*, 75(2):429–490.
- Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics*, 12(3):383–398.
- Campbell, J. Y. and Shiller, R. J. (1991). Yield spreads and interest rate movements: A bird’s eye view. *Review of Economic Studies*, 58(3):495–514.
- Christensen, J. H. E. and Gillan, J. M. (2011). Tips liquidity, breakeven inflation, and inflation expectations. Economic letter, Federal Reserve Bank of San Francisco.
- Christensen, J. H. E., Lopez, J. A., and Rudebusch, G. D. (2010). Inflation expectations and risk premiums in an arbitrage-free model of nominal and real bond yields. *Journal of Money, Credit and Banking*, 42(s1):143–178.
- Cochrane, J. H. and Piazzesi, M. (2005). Bond risk premia. *American Economic Review*, 95(1):138–160.
- Coibion, O. and Gorodnichenko, Y. (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review*, 105(8):2644–2678.
- Coibion, O. and Gorodnichenko, Y. (2025). Partisan bias in inflation expectations. *NBER Working Paper*. Working Paper.

- D'Acunto, F., Malmendier, U., Ospina, J., and Weber, M. (2021). Exposure to grocery prices and inflation expectations. *Journal of Political Economy*, 129(5):1615–1639.
- D'Amico, S., Kim, D. H., and Wei, M. (2014). Tips from tips: The informational content of treasury inflation-protected security prices. Finance and Economics Discussion Series (FEDS) 2014-024, Board of Governors of the Federal Reserve System, Washington, DC.
- D'Amico, S., Kim, D. H., and Wei, M. (2018). Tips from tips: The informational content of treasury inflation-protected security prices. *Journal of Financial and Quantitative Analysis*, 53(1):395–436.
- Dräger, L. and Nghiem, G. (2021). Are surveys and market-based measures of inflation expectations fully aligned? *International Journal of Central Banking*.
- Engle, R. F., Fleming, M. J., Mezrich, J. J., and Sarkar, A. (2012). Treasury yield volatility and the housing bubble. *Federal Reserve Bank of New York Staff Reports*.
- Fleckenstein, M., Longstaff, F. A., and Lustig, H. (2013). The TIPS-Treasury bond puzzle. *Journal of Finance*.
- Grishchenko, O. V. and Huang, J.-Z. (2012). Inflation risk premium: Evidence from the tips market and the cross-section of bonds. *Journal of Fixed Income*, 22(2):5–35.
- Grothe, M. and Meyler, A. (2015). Inflation forecasts: Are market-based and survey-based measures informative? Technical Report 1865, European Central Bank. Working Paper Series.
- Gürkaynak, R. S., Sack, B., and Wright, J. H. (2007). The U.S. Treasury yield curve: 1961 to the present. *Journal of Monetary Economics*, 54(8):2291–2304.
- Gürkaynak, R. S., Sack, B., and Wright, J. H. (2010). The TIPS yield curve and inflation expectations. *American Economic Journal: Macroeconomics*, 2(1):70–92.
- Hodrick, R. J. (1992). Dividends yields and expected stock returns: Alternative procedures for inference and measurement. *Review of Financial Studies*, 5(3):357–386.
- Hu, G. X., Pan, J., and Wang, J. (2013). Noise as information for illiquidity. *The Journal of Finance*, 68(6):2223–2272.
- Hördahl, P. and Tristani, O. (2014a). Inflation risk premia in the euro area and the us. *International Journal of Central Banking*.

- Hördahl, P. and Tristani, O. (2014b). Inflation risk premia in the euro area and the us. *International Journal of Central Banking*, 10(3):1–43.
- Hördahl, P., Tristani, O., and Vestin, D. (2010). Inflation risk premia in the term structure of interest rates. *Journal of Monetary Economics*, 57(6):685–702.
- Pflueger, C. E. and Viceira, L. M. (2011). Return predictability in the treasury market: Real rates, inflation, and liquidity. NBER Working Paper 16892, National Bureau of Economic Research.
- Stock, J. H. and Watson, M. W. (1999). Forecasting inflation. *Journal of Monetary Economics*, 44(2):293–335.
- Weber, M., D’Acunto, F., Gorodnichenko, Y., and Coibion, O. (2022). The subjective inflation expectations of households and firms. *Annual Review of Economics*, 14:37–61.

# Appendix

Table 7: First-stage regressions. HAC (Newey–West) standard errors in parentheses. Robustness check for backfilling hypotheses.

	$b_{n,t}$							
	(1)	(2)	(3)	(4)	(5)	(5a)	(6)	(7)
Off-the-Run Spr.		-1.16*** (0.39)	-0.93*** (0.35)	-1.69*** (0.49)	-0.94*** (0.33)	-0.77** (0.30)	-0.74* (0.40)	-0.48 (0.67)
Asset-Swap Spr.			-0.28 (0.22)	-0.05 (0.23)				-0.81*** (0.30)
Transaction Vol.				-0.36 (0.19)	-0.10 (0.14)	0.50*** (0.17)		-0.88*** (0.25)
Synthetic-Cash					-1.48** (0.28)	-1.55** (0.28)	-1.33** (0.33)	
Survey Inflation		0.99*** (0.22)	1.40*** (0.34)	1.42*** (0.32)	1.33*** (0.29)	1.35*** (0.28)	1.36*** (0.33)	1.40*** (0.23)
CFNAI		0.19*** (0.07)	0.15*** (0.05)	0.13*** (0.04)	0.12*** (0.04)	0.09** (0.04)	0.10** (0.08)	0.03 (0.02)
$p$ -value		0.00	0.00	0.00	0.00	0.00	0.00	0.00
$R^2$		0.28	0.36	0.37	0.39	0.51	0.51	0.48
ADF of Residuals		-6.47***	-6.48***	-6.19***	-5.34***	-6.21***	-6.24***	-5.25***
Period		Full	Full	Full	Full	Full	1999.03–2014.12	2004.07–2025.11

\*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$