Risk-Based Interest Rate Expectations

Robert Rogers*

First version: October 2025

This version: January 2026

Click here for the latest version

*London School of Economics. Email: r.rogers@lse.ac.uk. I am grateful to Dimitri Vayanos and Martin Oehmke for their continued support as advisors and to Ian Martin and Christopher Polk for their extensive comments and advice. I also thank Péter Kondor, Manav Chaudhary, Christian Julliard, Per Strömberg, Oliver Ashtari Tafti and seminar participants at the London School of Economics for their valuable input.

Risk-Based Interest Rate Expectations

Abstract

I develop a method to extract interest-rate expectations from options markets by connecting interest-rate risk premia to interest-rate variance risk premia, with no assumptions on interest-rate stationarity. I document that historical excess bond returns were driven by risk premia, not forecast errors. Second, risk-based forecasts outperform surveys and term-structure models out of sample. Third, risk-neutral variance drops sharply around FOMC announcements, suggesting a risk-premium explanation for the FOMC-announcement decline in long-term rates. Fourth, the recent positive stock-yield correlation is driven by expectations and not risk premia.

Keywords: term premium, term structure, bond risk premium, variance risk premium, return predictability, interest-rate risk, duration

JEL Classification: G12, G21

In January 2010 the yield curve sloped sharply upward. A 1-year loan could be made at an interest rate of just 0.6%, but the same loan agreed to start in 1 year would pay 2%.[1] Were investors expecting inflation and interest rates to spike? Or was this just compensation for duration risk? Distinguishing between these possibilities and, more broadly, measuring interest-rate expectations is a critical task for central banks and investors.

The long secular decline in interest rates makes this task particularly challenging. The most obvious statistical approach is to fit a model that predicts future interest rates based on the current yield curve. But such an approach assumes the relationship between predictors and expectations stayed constant even as long-term rates dropped from 15% in 1981 to 1% in 2020.

This paper proposes instead to measure interest-rate expectations and risk premia from options markets. I derive a no-arbitrage relationship that expresses interest-rate expectations as a function of interest-rate variance risk premia from options markets, without assuming a stable interest-rate process. I measure variance and variance risk premia with simple statistical tools and use the relationship to construct a "risk-based" measure of expected interest rates. This measure predicts out-of-sample changes in interest rates with a high degree of economic and statistical significance, outperforming traditional term-structure models. The measured expectations show that high historical bond returns have been driven by risk premia, not surprises, that bond risk premia are negatively correlated with stock returns, and that risk premium may account for much of the decline in long-term rates around monetary policy announcements.

The key identity relates an investor's risk to their expected return. If investors expect high returns on bonds, they will take on more interest-rate risk. A simple no-arbitrage identity tells us that the interest-rate risk premium RP_t perceived by an investor equals the investor's interest-rate risk "exposure" λ_t, multiplied by the risk-neutral variance σ_t^{*2}:

$$RP_t = \lambda_t \times \sigma_t^{*2}$$

[1]The 1-year LIBOR swap rate on January 1st was 0.6%, and the 1y-in-1y swap forward rate was 2%.

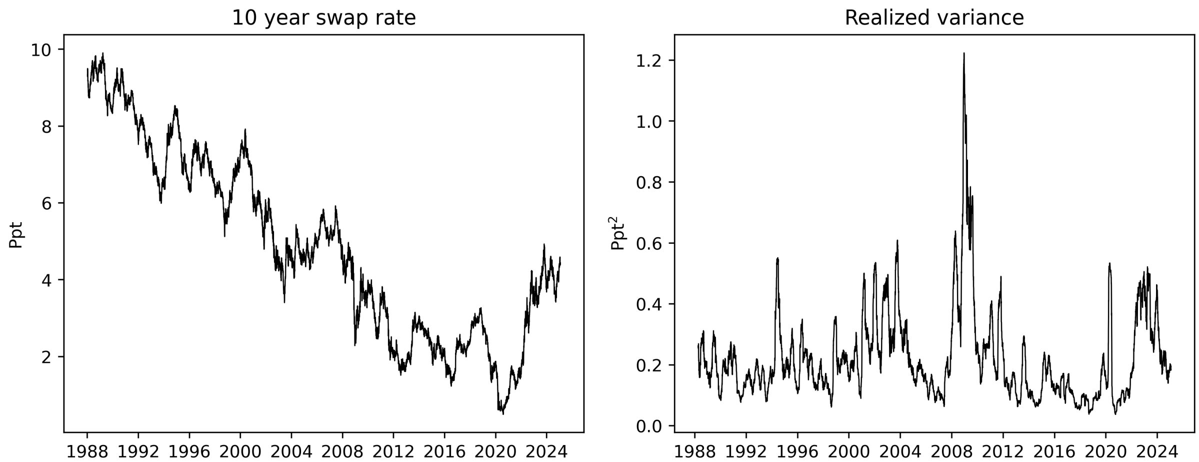

Figure 1: Stationarity of interest rate levels vs variance. The left-hand side plots the 10-year swap rate and the right-hand side plots the quarterly realized variance of the same rate, 1988-2025.

The exposure, λ_t, represents the investor's sensitivity to interest rate increases. For a bond investor with constant relative risk aversion, this is approximately equal to portfolio duration times risk aversion (and exactly equal to duration in the log-utility special case).

To identify exposure, I use a parallel identity for squared interest-rate changes. After one simplifying assumption, discussed shortly, the risk premium on contracts that pay off squared changes in interest rates, $RPSq_t$, is equal to the level of interest-rate exposure multiplied by the risk-neutral third moment:

$$RPSq_t \approx \lambda_t \times \text{ThirdMoment}_t^*$$

Intuitively, when the level of interest-rate risk exposure, λ_t, is high, the equilibrium price of "insuring against" large changes in interest rates, $RPSq$, should also be high. We can use the price of this insurance to learn about the level of risk exposure. This relationship will be stronger and the insurance more valuable when the third moment is high, because large increases in rates will be more common than large declines.

Expected squared changes depend on variance, and so this identity allows us to back out the exposure and risk premium based on a variance forecast rather than a level forecast. Forecasting variance offers two advantages: first, the variance, unlike the level, is clearly stationary (see Figure 1). Second, variance can be learned quickly from daily data, and hence variance-forecasting models have a strong track record of empirical success (e.g. Bollerslev, 1986; Corsi, 2009).

The exact identity includes an additional term capturing risks that are uncorrelated with interest rates but covary with squared interest-rate changes. In my benchmark results for 1-quarter-ahead risk premium, I assume this term is zero. I validate this assumption by demonstrating that non-interest-rate risks, such as those from equities or credit spreads, have too little correlation with squared rate changes over short horizons to explain the results. I also show that allowing for nonlinearities stemming from interest-rate risks — for example, from higher risk aversion or multi-factor models — may scale the estimates up or down modestly but leaves the time-series dynamics intact. In all specifications, the estimated risk premium remains >95% correlated with the main results.

The forecasts that this approach generates are correlated with surveys and traditional risk premium measures, capturing well-known effects such as the association of slope with risk premium. However, they significantly outperform existing methods. At one-quarter horizons, out-of-sample MSE is 28% lower than the Survey of Professional Forecasters and 10% lower than the popular Adrian, Crump, and Moench (2013) term-structure model (see Table 4). A trading strategy based on these forecasts generates a Sharpe ratio of 0.7, vs -0.2 for a traditional term-structure model.

Beyond forecasting performance, the risk-based approach provides new perspectives on three puzzles in fixed income markets. First, do the high returns on long-term bonds over the past 25 years represent an anticipated risk premium, or a series of surprises? Forward rates have tended to point upward, while rates stayed low, leading to large returns for bond investors. The causes of these returns are of central interest to finance research: van Binsbergen (2020) suggests they can potentially account for the entire equity risk premium in this period.

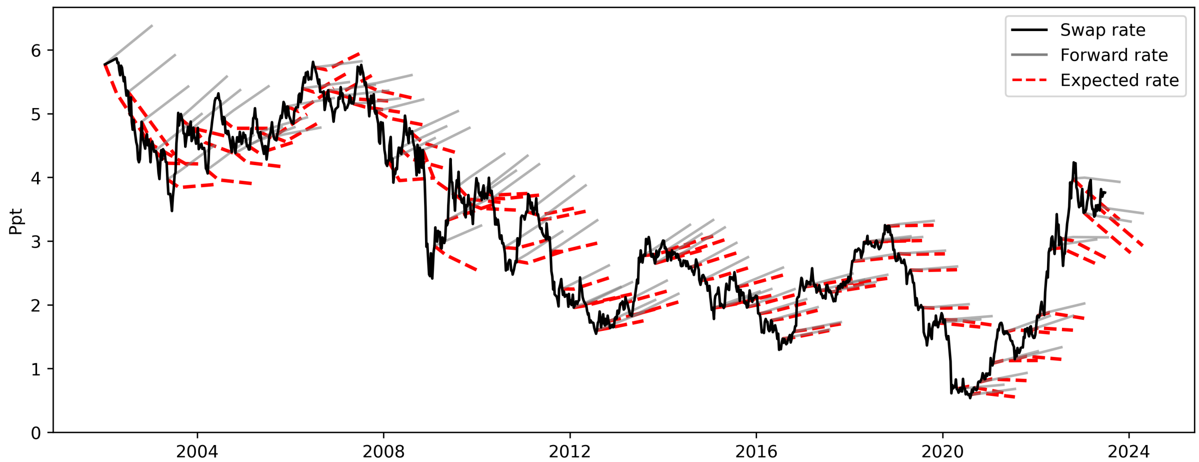

Figure 2: Forward rates vs risk-based forecasts of the 10-year swap rate. The black line plots the 10-year swap rate from 2002 to 2023. The light gray lines each quarter plot the 1-quarter and 1-year forward rate. The red dashed lines plot the 1-quarter and 1-year expected interest rates calculated using this paper's methodology.

Standard approaches struggle to answer this question because of the secular decline in rates. Models with a stationary interest-rate process assume rates will return to the historical average, leading to low or negative risk premia in recent years. If one instead allows the long-term "end point" to shift, then a wide range of different risk premia are compatible with a reasonable set of priors (Farmer, Nakamura, and Steinsson, 2024).

My method gathers information instead from the variance of interest rates, which, in contrast, exhibits mean reversion with no obvious long-term trend, as shown in Figure 1. I find a highly significant variance risk premium that implies an interest-rate risk premium of 43 bp per year. This risk premium is similar to the realized premium of 34 bp, and not statistically significantly different from the 28 bp average risk premium under a "random walk model" where interest rates are always expected to stay at their current levels. Figure 2 plots interest rate forwards versus my measure of expectations at each quarter from 2002–2023. When forward rates pointed upward, as in 2002 or 2010, risk premia were large enough that expected rates stayed flat or declined. The duration risk premium appears to have been expected risk compensation rather than forecasting errors.

Second, do investors learn their long-term interest-rate expectations from the Fed? Hillenbrand (2025) documents that the entire secular decline in long-term rates occurred during three-day FOMC windows, interpreting this as markets learning about long-run fundamentals from the Federal Reserve. Standard term-structure models are poorly suited to decompose these announcement effects into expectations versus risk premium because they assume a constant long-run level, ruling out any learning.

My risk-based decomposition suggests an alternative explanation. Risk-neutral variance falls around FOMC meetings by amounts far exceeding mechanical effects from uncertainty resolution, and then reverses between meetings. This pattern is consistent with temporary risk premium compression. The implied declines in risk premium could potentially account for the entire FOMC-window rate puzzle. While I cannot rule out shifts in physical variance expectations, the transitory nature and magnitude of these moves point toward monetary policy systematically compressing risk premia.

Third, why are stocks now positively correlated with yields? The stock-yield correlation flipped from negative to positive in the late 1990s. Is this driven by correlated risk premia, with flight-to-safety shocks driving up stock risk premia and down bond risk premia (e.g. Antolin-Diaz, 2025) or correlated expectations about growth (e.g. Campbell, Sunderam, and Viceira, 2017)? Standard models show positive correlation of interest-rate risk premia with equity returns, but this may result from the stationarity assumptions. Using the risk-based approach, I find that the interest-rate risk premium correlates negatively with equity returns. When stocks decline investors seem to demand higher returns on long term bonds. The positive yield-stock correlation therefore appears to reflect correlation of stocks with expected future rates more than flight-to-safety effects.

Related Literature

This paper extends techniques from the literature on option-based expected returns (Martin, 2017; Kremens and Martin, 2019; Chabi-Yo and Loudis, 2020; Tetlock, McCoy, and Shah, 2024). The existing literature has assumed that the SDF is a function of stock market returns. We cannot take a similar approach for interest rates because there is no liquid instrument that reveals the risk-neutral covariance of interest rates with equity returns. Nor is there an obvious theoretical benchmark for what exposure to interest rates a representative investor should have.

I therefore develop a different approach that instead takes the assumption that the investor's exposure to interest rates are not "too-nonlinear," and then uses the variance risk premium to find exposures and expected returns. My methodology is related to that used in Tetlock et al. (2024), who link expected returns on the S&P 500 to the variance risk premium and the risk-neutral third and higher moments extracted from S&P 500 options.

This paper also contributes to the extensive literature on bond risk premia and term structure by offering a new methodology to forecast rates. The large positive term premium estimates can be interpreted as supporting evidence for non-stationary shifting end-point models e.g., Kozicki and Tinsley (2001); van Dijk, Koopman, van der Wel, and Wright (2014); Bauer and Rudebusch (2020). The use of risk-neutral moments complements existing empirical evidence from Bauer and Chernov (2024) that links skewness to risk premium, as well as evidence on unspanned stochastic volatility, e.g., Collin-Dufresne and Goldstein (2002); Collin-Dufresne, Goldstein, and Jones (2009). My finding of a positive interest rate variance risk premium is in line with Trolle and Schwartz (2014) and Choi, Mueller, and Vedolin (2017).

Organization of the paper

The remainder of the paper details the derivation (Section 1), data and estimation (Section 2), results and forecasting performance (Section 3), applications to FOMC announcements and stock-bond correlations (Section 4), and validation of the key identifying assumption (Section 5). Section 6 concludes and an appendix contains extensions to other maturities, robustness tests, and technical details.

1 Measuring expectations from risk

If investors are highly exposed to interest-rate risk, the equilibrium price to insure against variance in rates will be high. The price of that insurance can therefore inform us about their interest-rate exposure and the expected interest rate. This section formalizes this intuition by deriving an identity relating the interest-rate risk premium to the variance risk premium and one unobservable quantity. I then discuss the key assumption needed to operationalize this relationship.

1.1 Defining interest rates and interest-rate risk premium

This paper aims to measure expectations of next period's interest rate:

$$E_t(y_{t+1})$$

where y_{t+1} could represent any rate or yield. In practice, I will focus on the 10 year swap rate one quarter or one year in the future. I will use payoffs on linear interest rate forwards as my basic unit of theoretical analysis. At time $t+1$ the buyer of a linear interest rate forward receives the level of the interest rate minus the pre-agreed forward price F_t. His payoff is equal to:

$$\Delta y_{t+1} = y_{t+1} - F_t$$

We can construct this interest rate forward price F_t from the observable yield curve.[2] Framing the analysis around interest rate forwards rather than zero coupon bond returns allows us to work with swap rates or par yields which have observable option prices.

Using standard notation we can also describe the forward rate as the "risk-neutral expectation" of interest rates: $E_t^*(y_{t+1})$. If markets were priced by a risk-neutral agent, this would be equal to the physical expectation. We call this hypothesis that $E_t^*(y_{t+1}) = E_t(y_{t+1})$ the "expectations hypothesis." In practice, we know that it is not true, and there are large gaps

[2]We do not exactly observe linear forward prices for most interest rates, but their value can be constructed with a minuscule convexity adjustment described in sections 2 and Appendix I.2.

between the forward rates and interest-rate expectations (Fama and Bliss, 1987; Campbell and Shiller, 1991).

I will define the "interest-rate risk premium," RP_t, as the difference between the forward rate and the physical expected rate. This is equal to the expected payoff from selling the linear interest-rate forward:

$$RP_t = F_t - E_t(y_{t+1}) = E_t^*(y_{t+1}) - E_t(y_{t+1}) = -E_t(\Delta y_{t+1})$$

This risk premium is closely related to the expected excess return on bonds. If we assume that y_{t+1} is the yield on a bond or swap with maturity T, then the expected excess return on a bond with maturity T+1 is approximately its duration times the interest-rate risk premium:[3]

$$E_t(R_{t+1}^{(T+1)} - R_{f,t}) \approx D \times RP_t$$

where $R_{f,t}$ denotes the gross risk free rate known at time t from time t to time $t+1$.

1.2 Measuring the interest-rate risk premium without assumptions on the rate process

I will assume that there is no arbitrage and the fundamental theorem of asset pricing holds throughout this paper, and denote the stochastic discount factor (SDF) by M_{t+1}. In that case, the following result allows us to learn about expectations without specifying an interest-rate process:

Proposition 1. *If no arbitrage holds, the expectation of any payoff X is given by:*

$$E_t^*(X) - E_t(X) = -\frac{1}{R_{f,t}} cov_t^* \left(\frac{1}{M_{t+1}}, X \right)$$

Where E^* represents risk-neutral expectations. From Martin and Wagner (2019); Chabi-Yo

[3] To derive this, simply take a first-order approximation of the bond price $P_{t+1}(y_{t+1})$ around $y_{t+1} = F_t$. Assume the bond forward approximately equal to the linear forward, so that $\frac{P_{t+1}(F_t)}{P_0} \approx R_{f,t}$

and Loudis (2020)[4].

Applying this identity to the interest rate y_{t+1} yields:

$$RP_t = -\frac{1}{R_{f,t}} cov_t^* \left(\frac{1}{M_{t+1}}, y_{t+1} \right)$$

The risk premium is revealed by the risk-neutral covariance of interest rates with the SDF. This result resembles the more familiar asset pricing identity that relates expected returns with the *physical* covariance with the SDF. However, it has the advantage of working with directly observable quantities. Risk-neutral covariances are potentially observable from asset prices.

Hence if we can make some assumptions about the nature of the SDF (i.e., what constitutes a good or bad state) and find the price of an asset whose payoffs are linked to interest rates but also this SDF, then we can calculate interest-rate expectations.

1.3 Projecting the SDF onto yields

The existing literature that uses proposition 1 to measure expectations has taken the perspective of an equity investor for whom $\frac{1}{M_{t+1}}$ is a function of stock market returns (Martin, 2017; Martin and Wagner, 2019; Kremens and Martin, 2019; Chabi-Yo and Loudis, 2020; Tetlock et al., 2024). In this case, expected returns are given by the risk-neutral covariance with stock market returns.

This approach will not work for interest rates. We do not observe any liquid assets that reveal the risk-neutral covariance of interest rates with stocks (Martin, 2025). Additionally, even if such an asset were available, we might question whether equity returns are a reasonable proxy for fixed-income investor wealth. The efforts to find joint risk factors across stocks and bonds have not always been successful.

I therefore propose a new approach that does not rely on equity returns. Instead, I

[4]The proposition can be proven simply by calculating the price of an asset with payoff $Z = \left(\frac{1}{M_{t+1}} - R_{f,t} \right) X$ using both SDF and risk-neutral pricing notation

will consider the SDF as a function of interest rates and residual terms. Without loss of generality, we can consider the projection of any SDF onto interest rates (Δy_{t+1}) under the risk-neutral measure:

Definition 1. *Projection of the inverse SDF*

$$\frac{1}{M_{t+1}} = R_{f,t}\left(1 - \lambda_t \Delta y_{t+1} + \varepsilon_{t+1}\right)$$

where $E_t^*(\varepsilon_{t+1}) = cov_t^*(\varepsilon_{t+1}, \Delta y_{t+1}) = 0$

This is simply a linear projection under the risk-neutral measure. The SDF can be written in this fashion for any model with no arbitrage. The intercept must be $R_{f,t}$ because $E_t^*\left(\frac{1}{M_{t+1}}\right) = R_{f,t}$ and $E_t^*(\Delta y_{t+1}) = 0$. In the case of a log investor who holds exactly linear exposures to the interest rate the inverse SDF is linear and $\varepsilon_{t+1} = 0$.

In general, I will refer to λ_t as the investor's "exposure" to interest rates, because it captures how much the investor suffers from a rise in interest rates. For a CRRA investor with linear interest-rate exposure and risk aversion coefficient γ:

$$\lambda_t \approx \gamma \times \text{duration}_t$$

For example, with a CRRA coefficient of 2 and a linear exposure of duration 20, $\lambda \approx 40$. In the special case of log utility where $\gamma = 1$, this approximation is exact and λ_t is the duration. More generally, there will be a small approximation error. Appendix E tests this approximation up to $\gamma = 4$ and shows it is accurate within 4% for the average risk-neutral distribution, and 8% on the most extreme day.

1.4 Risk premium as a function of exposure

Applying Proposition 1 to measure interest-rate expectations and plugging in the projection for $\frac{1}{M_{t+1}}$ yields an expression for interest-rate risk premium.

Proposition 2. *If no arbitrage holds, the interest-rate risk premium is given by:*

$$RP_t = \lambda_t \times var_t^*(\Delta y_{t+1}) \equiv \lambda_t \sigma_t^{*2}$$

In words, Proposition 2 tells us the risk premium is equal to exposure times risk-neutral variance. Loosely speaking, this is because in equilibrium, the amount of risk agents take should tell us both the rewards to taking interest-rate risk and the price to insure against large changes in interest rates (the risk-neutral variance).

Now we have an expression for the interest-rate risk premium in terms of an observable variable. The risk-neutral variance represents the price of a contract that pays out squared interest-rate changes and can be easily measured from option prices, using the methodology described in Section 2. However, we still need to know our exposure term, λ_t before we know interest-rate expectations.

1.5 Exposure as a function of variance

We want to use the value of λ_t to measure expected payoffs on interest-rate forwards. But the same logic can be applied in reverse: if we can estimate any expected payoff, we can use it to learn about λ_t. I will measure the expected payoff on a contract that pays out Δy_{t+1}^2 and use this expected payoff to learn about λ_t. This approach resembles the strategy that Tetlock et al. (2024) apply to equity markets.

Variance has three important characteristics that make it an attractive payoff to forecast. First, interest-rate variance does not display the same long-term non-stationary trend as the level (see Figure 1). Second, variance can be measured easily with high frequency data. For any function involving higher moments, this becomes more challenging due to volatility clustering.[5] Third, there is a long track record of empirical success with simple variance

[5]Intuitively, the daily autocorrelation of changes in rates should be low. Otherwise there would be large potential trading profits. But the daily autocorrelation of squared changes in rates, or the last period's change in rates with next period's squared change may be very high. Variance is persistent. Hence we cannot find a good estimate for third or higher moments by simply dividing up the sample into small pieces and calculating the moments of the short observations.

forecasting models (e.g., Bollerslev, 1986; Corsi, 2009).

Applying Proposition 1 to the payoff Δy^2_{t+1} yields an expression for the risk premium on squared interest rate contracts and variance risk premium.

Proposition 3. *The risk premium on squared interest rate contracts is given by:*

$$RPSq_t \equiv \sigma^{*2}_t - E_t(\Delta y^2_{t+1}) = \lambda_t E^*_t(\Delta y^3_{t+1}) - cov^*_t(\Delta y^2_{t+1}, \varepsilon_{t+1}) \tag{1}$$

Or, rearranged in in terms of the "variance risk premium:"

$$VRP_t \equiv \sigma^{*2}_t - \sigma^2_t = \lambda_t E^*_t(\Delta y^3_{t+1}) + \sigma^{*4}_t \lambda^2_t - cov^*_t(\Delta y^2_{t+1}, \varepsilon_{t+1}) \tag{2}$$

To move from Equation (1) to (2), I have used the fact that the sum of squared changes is equal to the variance plus the expected change (i.e., risk premium) squared:

$$E_t(\Delta y^2_{t+1}) = \sigma^2_t + RP^2_t = \sigma^2_t + \sigma^{*4}_t \lambda^2_t$$

The terminology of "variance risk premium" for the left-hand side of Equation (2) — the difference between the risk-neutral and physical variance — is consistent with the existing literature, e.g., Bollerslev, Tauchen, and Zhou (2009); Drechsler and Yaron (2011).

Proposition 3 tells us the expected profit from selling variance contracts is equal to exposure scaled by the risk-neutral third moment minus an unobservable term that I will call the "residual coskew." The risk-neutral third moment can be measured from option data in exactly the same way as the risk-neutral variance. So if we can make an assumption about the size of the residual coskew term, then a forecast of physical variance will yield a value for the exposure and interest-rate risk premium.

The exposure will be better-identified by the physical variance in periods when the risk-neutral third moment, $E^*_t(\Delta y^3_{t+1})$, is high. If the distribution is highly right-skewed (i.e., the third moment is high), then the variance contract will mostly pay off in high-rate states of the world where the investor is poor and hence will be highly valuable. If, on the other hand, the risk-neutral third moment is very low, it will take a large amount of exposure to deliver much variance risk premium.

1.6 The residual coskew assumption

The residual term ε_{t+1} represents the portion of the inverse SDF that is uncorrelated with interest rates under the risk-neutral measure. The inverse SDF, loosely speaking, captures how "good" states are for the investor (inverse of marginal utility). So the residual coskew term, $cov_t^*(\Delta y_{t+1}^2, \varepsilon_{t+1})$ represents the extent to which the investor tends to be better or worse off in states of the world with large interest-rate changes, regardless of their sign.

This quantity is unobservable. We do not observe ε, and even if we did, we might not observe the option prices that would reveal the relevant risk-neutral moment. Producing an estimate of the interest-rate risk premium therefore requires us to make an assumption on the size of this covariance. The simplest assumption is simply that the value is zero. This assumption will approximately hold over short horizons, for example, in standard dynamic term-structure models which assume the log SDF is linear in interest-rate factors, and thus the inverse SDF is nearly linear for short intervals.[6]

Setting the residual coskew to zero amounts to an assumption that the risk premium on variance contracts results from directional interest-rate risk by the investor (i.e., λ_t) rather than a nonlinear exposure to rates. In other words, a representative investor would choose not to sell variance contracts despite their positive payoff because of her level of interest-rate risk. This does not impose that other asset classes (e.g. equities) are not important to the investor, but only that the directional exposure of those assets to interest-rate risk is more important than exposure to squared changes.

I will take this as my main benchmark in developing a measure of 1-quarter-ahead interest-rate expectations. It turns out that despite the simplicity of this assumption, it

[6] It is possible that other interest rate factors would display residual coskewness with the single interest rate in the univariate model described in this section. I address this possibility in Section 5.

produces forecasts that perform very well out of sample and line up with the variables we expect to be associated with interest-rate risk premia (e.g. the slope of the yield curve), as shown in Section 3.

In Section 5 I will measure the closest observable data and argue the assumption is empirically reasonable over the one-quarter horizon. I calculate the physical covariance of ε with squared changes in rates under a few common specifications for the SDF and find they are too small to explain the size of the variance risk premium. I also produce estimates of exposure and risk premium after loosening the assumption to allow for residual coskewness from higher CRRA risk aversion, from exposure to multiple interest rates, and from exposure to equity and corporate bond risk factors. In all cases, the simple zero-residual-coskew estimates remain >95% correlated with the true values, although they could overestimate the true risk premium by up to 30% or underestimate the true risk premium by up to 50%.

This assumption becomes more dubious as the forecast horizon becomes longer because the residual coskewness is likely to involve higher moments that grow at a faster than linear rate. For example, if the true inverse SDF is a quadratic function of rates, and rates follow a Brownian motion, then residual coskew will be proportional to the forecast horizon squared.[7] For this reason, I focus my main results on the one-quarter horizon. Where it is necessary to produce one-year-ahead forecasts, as in Section 3, I relax the zero residual coskew assumption by allowing the inverse SDF to have a constant loading on Δy_{t+1}^2 using the methodology described in Section 2.3.5.

2 Data and estimation

Section 1 shows how exposure, and hence interest-rate risk premium, can be identified by the physical variance, the risk-neutral variance, and the risk-neutral third moment of interest rates (Proposition 3). This section describes how I operationalize this relationship to calculate interest-rate expectations. I first present the data sources that the paper

[7] $cov_t(y_{t+h}^2, y_{t+h}^2) \propto h^2$ if y_{t+h} is normally distributed

employs. I then describe how I estimate the exposure and variance in a simple two-equation nonlinear regression framework.

2.1 Data

I measure realized variance from daily interest-rate swap rates obtained from Refinitiv, and risk-neutral moments from swaptions quotes obtained from Bloomberg. The sample covers USD LIBOR swaps from December 2001 to June 2023, when LIBOR ceased publication. I track six tenors (1, 2, 5, 10, 20, and 30 years) at three forecast horizons (3 months, 1 year, and 5 years). I construct swap forwards from swap rates and treasury yields using the procedure described in Appendix H.

The results of this paper should be interpreted as risk premium on swap rates, rather than treasury rates. Arguably, this is the more relevant rate for many market participants who are more likely to borrow at rates linked to swap rates or hedge using swaps. The differences are unlikely to be large. The realized quarterly risk premium on 10-year swaps vs treasuries during my sample period differs by just 0.7 basis points. The spread between 10-year swaps and the equivalent (off the run) treasury yield averages just 12 basis points during the sample period and it explains only 3% of the variation of the 10-year yield (unconditional variance ratio). Since 2008, this spread has often been negative and is thought to be related to leverage constraints and capital requirements rather than bank credit risk (Boyarchenko, Gupta, Steele, and Yen, 2018).

I use swaptions to calculate risk-neutral interest-rate moments. A swaption is a contract that gives its holder the right to enter an interest-rate swap at a predetermined rate. Swaption markets represent the most liquid interest option market, with extensive buy-side and interdealer trading (ISDA, 2014; Barnes, 2024). While swaptions trade over-the-counter rather than on exchanges, collateralization through initial and variation margin is standard practice and limits counterparty credit risk. Treasury options data from CME offers an alternative source of interest-rate option data. While this data source has a longer history than swaptions, it covers fewer tenors, has lower trading volume, and trades as American options at a single, short and time-varying horizon of 1-3 months

Out-of-the-money swaption quotes are not available from Bloomberg before July 2011. However, we can learn risk-neutral variance and skewness from three sources: first, Trolle and Schwartz (2014) calculate and display the risk-neutral variance and skewness of 10y-in-1y swaption contracts from 2001 to 2010 using data from an interdealer broker. Second, Bloomberg provides quotes for at-the-money swaption variance, which has an $> 99.5\%$ correlation with risk-neutral variance. Third, Bauer and Chernov (2024) calculate the risk-neutral variance and skewness of closely-related 10y treasury options. I therefore calculate pre-2011 risk-neutral variance using at-the-money quotes supplemented with Trolle and Schwartz (2014) data (digitized from their Figure 1), and use risk-neutral skewness directly from Trolle and Schwartz (2014) and Bauer and Chernov (2024). I describe this procedure in Appendix I.1 and demonstrate that following a similar process for the periods where we have all data yields accurate predictions of the 2nd and 3rd risk-neutral moments. Data before 2007 is only available on a weekly basis.

.

2.2 Constructing risk-neutral moments

I follow Carr and Madan (1998) and Martin (2017) to extract model-free measures of risk-neutral moments from option prices. The idea behind this approach is that the risk-neutral variance is the forward price of an instrument that pays off Δy_{t+1}^2, by definition. By combining a bundle of many different options on y_{t+1} at different strike prices, we can construct a portfolio whose payoff is Δy_{t+1}^2 (Breeden and Litzenberger, 1978). Since we observe all of the different option prices, we can calculate the forward price of this portfolio, and hence the risk-neutral variance. The third moment can be calculated in the same way.

This result holds exactly when we can trade options on the underlying rate itself. However, swaptions are options on swap values rather than swap rates. The value of a swap depends on the price of an annuity, not just the swap rate. In Appendix I.2, I show that we can nevertheless derive a close approximation of the risk-neutral moments of the swap rate itself. If we assume that changes in the 10-year annuity yield are the same as changes in the 10-year swap rate, then we can calculate the risk-neutral moments of the swap rate itself. Historical correlation of quarterly or annual changes in the swap rate with the annuity rate are $> 98\%$, so this is unlikely to introduce large approximation errors. I can also use this strategy to calculate the risk-neutral expectation of the interest rate. This is the linear forward price discussed in Section 1, which can theoretically differ from the ordinary swap or bond forward rate. In practice I find the differences are only 2 bp on average, with a standard deviation of 1 bp, and hence immaterial for the purpose of this paper.

Table 1 presents summary statistics for the second and third risk-neutral moments, and Figure 9 in Appendix D plots the time series. Variance is high during and after the financial crisis, during the 2022 inflation, and, more surprisingly, in 2002-4. The high option-implied interest-rate volatility of the early 2000s was noted at the time as puzzling by central banks (Fornari, 2005; ECB, 2005).

2.3 Estimation

2.3.1 Procedure

Section 1 describes how exposure (λ_t) can be measured from physical variance (σ_t^2) under the zero residual coskew assumption. To solve for these quantities I make two further assumptions. First, I assume that realized variance is an unbiased measure of conditional physical variance:

$$RV_{t+1} = \sigma_t^2 + \varepsilon_{t+1} \tag{3}$$

where RV_{t+1} is the sum of squared daily changes in interest rates over the subsequent quarter or year. This is a standard assumption used in the variance forecasting literature. It will hold as the observation window (here, 1 day) approaches zero either if there is no autocorrelation in changes in rates, or if there are no jumps in rates.

Second, I assume that exposure and physical variance are a constant linear function of state variables X_t and Z_t:

$$\lambda_t = \lambda' X_t \tag{4}$$

$$\sigma_t^2 = \beta' Z_t \tag{5}$$

These parameterization choices are described in the following two subsections and based on standard formulations of variance forecasting models and term structure models. The methodology is robust to relaxing this assumption or to reasonable changes in the choice of state variable.

Combining Proposition 3 with the above two assumptions (Equations 3, 4, and 5) yields a two-equation nonlinear regression setup to estimate exposure and conditional variance:

$$\sigma_t^{*2} - \beta' Z_t = (\lambda' X_t) \times E_t^*(\Delta y_{t+1}^3) + (\lambda' X_t)^2 \times \sigma_t^{*4} + \eta_t \tag{6}$$

$$RV_{t+1} = \beta' Z_t + \varepsilon_{t+1} \tag{7}$$

The first equation is the variance-risk-premium pricing equation from Proposition 3. The second equation is a physical variance prediction equation. Since σ_t^{*2} can be directly observed from option data, this model can be estimated by finding the values of λ and β that minimize the sum of squared errors $\sum \eta_t^2$ and $\sum \varepsilon_t^2$. I estimate $\hat{\beta}$ by ordinary least squares on equation 7 and then $\hat{\lambda}$ by a standard nonlinear least squares procedure on equation 6.[8] For out-of-sample figures I produce estimates using expanding windows.

To address heteroskedasticity and reduce the influence of a few high-variance states, I weight the sum of the variance prediction errors $\sum \eta_t^2$ by the inverse of risk-neutral variance (similar to Tetlock et al., 2024). This is equivalent to dividing equation (6) through by σ_t^* and improves the stability of estimates, reducing the excess kurtosis of the left-hand variable from 4.2 to 0.1.

[8] I choose to solve for $\hat{\beta}$ and $\hat{\lambda}$ separately so that the variance estimator resembles standard variance forecasting tools more closely. Jointly solving for the values that minimize the sum of all standard errors yields similar results.

2.3.2 Parameterizing variance

To produce conditional estimates of σ_t^2, I assume that conditional variance is a linear function of past realized variance.

$$
\begin{aligned}
RV_{t+1} &= \beta' Z_t + \varepsilon_{t+1} \\
Z_t &= [1,\ RV_{t-5d\to t},\ RV_{t-21d\to t},\ RV_{t-63d\to t}]'
\end{aligned}
$$

where $RV_{t\to t+Hd}$ is the realized variance calculated from daily data from period t to $t+H$ business days. I assume 21 days in a month and 63 in a quarter.

Under this assumption, conditional variance can be estimated by a simple linear regression of realized variance on lagged realized variance. This approach is similar to the heterogeneous autoregressive realized variance (HAR-RV) model of Corsi (2009), widely used in the variance risk premium literature (e.g., Bollerslev et al., 2009; Drechsler and Yaron, 2011). I use weekly, monthly, and quarterly lags rather than the standard daily, weekly, and monthly from Corsi (2009) to better suit the longer forecast horizons in my application, and because I do not have intra-day swap rate data.

The time series of quarterly estimated variance risk premium is plotted in Figure 3. Results are robust to including interest-rate levels or allowing variance to depend on rate levels through a constant elasticity of variance specification.

The model performs well out of sample, delivering 20% lower quarterly mean squared error than simply using the risk-neutral variance to forecast realized variance, as shown in Panel (A) of Table 2. This predictability confirms that it is easier to learn about variance than means, one of the key advantages of the risk-based forecasting approach.

To test that I successfully measure the variance risk premium, I regress my measured variance risk premium, $\sigma_t^{*2} - \widehat{\sigma}_t^2$, out-of-sample on the realized variance risk premium $\sigma_t^{*2} - RV_{t+1}$. Panel (B) of Table 2 shows that the predictor is highly significant, with a t-statistic of over 4 and R^2 of over 23% at a quarterly horizon. The coefficient of 0.63 is somewhat lower than the value of 1 that we would expect if the prediction was perfect. This difference could imply that the regression modestly overestimates the true variance risk premium, or that the realized variance risk premium in the sample period was somewhat lower than the ex-ante variance risk premium due to unexpected shocks. Because of this discrepancy I use the whole-sample estimates when measuring the average size of exposure and risk premium in Section 3.1.

2.3.3 Parameterizing exposure

The variance risk premium identity in Proposition 3 and the conditional variance forecasts from Section 2.3.2 allow us to solve for the exposure, λ_t, in each period. However, in periods where the risk-neutral third moment is low, the exposure will be poorly identified and highly sensitive to noise in the variance and skew estimates, as discussed in Section 1.5. To estimate exposure for the whole time series, I therefore assume that λ_t is a smooth function of state variables. This allows us to use the periods where the third moment is high to identify the plausible exposure when it is low. In Appendix B, I consider the alternative approach that estimates λ_t period-by-period without functional form assumptions. The results resemble the main results, indicating that the choice of functional form is not critical.

I parameterize λ_t as a linear function of the first three principal components of the yield curve and the two risk-neutral moments used to construct my estimates (variance and skewness):

$$\lambda_t = \lambda' X_t \tag{8}$$

$$X_t = [1,\ f_{1,t},\ f_{2,t},\ f_{3,t},\ \sigma_t^*,\ skew_t^*]' \tag{9}$$

where f_1, f_2, and f_3 are the first three principal components of yields, which capture the level, slope, and curvature of the yield curve. This can be interpreted as an assumption that investors choose their level of exposure depending on the shape of the yield curve and risk-neutral distribution of future yields, or that the shape of the yield curve and the risk-neutral distribution responds to investor choices and exposure.

The first three principal components explain $> 99\%$ of variation in the shape of the yield curve and are empirically well known to be related to interest-rate risk premium. It has been standard practice in term-structure modeling since Duffee (2002) to assume that the price of risk is a function of yield curve factors.

I include the risk-neutral volatility and skewness for the sake of completeness because they are used in construction of the estimator, because empirical work has suggested they contain additional information on interest-rate risk premium that is not included in the yield curve (Joslin and Konchitchki, 2018; Bauer and Chernov, 2024), and because they tend to appear in the price of risk in term structure models with stochastic volatility (e.g Creal and Wu, 2015). Ultimately I find these are not significant explainers of the variance risk premium (shown in Table 6 in Section 3), but I allow my estimator to determine this rather than assuming this relationship a priori.

This paper's methodology is flexible, and more state variables could easily be added in the future to improve prediction or test theoretical predictions.

2.3.4 Standard errors

I calculate GMM standard errors that account for the variance introduced by both ε_t and η_t, and allow for autocorrelation of errors up to one year using the standard GMM procedure with Newey-West kernel. While my main point estimates use weekly data, I estimate standard errors using monthly data for more conservative figures that avoid very long Newey-West lags.

The variance risk premium pricing error term η_t captures at least three sources of error: measurement error in risk-neutral moments, time variation in λ_t and σ_t orthogonal to the state variables, and the residual coskew term. Some caution is warranted in interpreting these standard errors, as the autocorrelation structure of η_t depends on the relative importance of these three components. The main validation of the results of this method come from its out-of-sample performance rather than in-sample standard errors. Section 5 considers estimates that allow the second order coefficient to be time-varying.

2.3.5 Annual forecast horizons

The main results of this paper focus on one-quarter-ahead forecasts. However, for the sake of comparison with other models I also produce 1-year-ahead forecasts.

Residual coskew is likely to be larger over one-year horizons, because it involves higher moments that grow at a faster than linear rate in forecast horizon. As a simple way of allowing for residual coskewness in these annual horizon forecasts, I consider a projection of the inverse SDF onto the first two moments of interest rate changes, and assume the loading on the second moment to be constant:

$$\frac{1}{M_{t+1}} = R_{f,t}\left(1 - \lambda_t^{(1)}\Delta y_{t+1} - \lambda^{(2)}(\Delta y_{t+1} - \sigma_t^{*2}) + \varepsilon_{t+1}\right)$$

If we assume that ε_{t+1} has zero risk-neutral covariance with Δy_{t+1}^3, and use the same parameterization assumptions for $\lambda^{(1)}$ as described in Section 2.3.3, the variance risk premium estimation equation (6) becomes:

$$\sigma_t^{*2} - \beta' Z_t = \lambda' X_t E_t^*(\Delta y_{t+1}^3) + \lambda^{(2)}(E_t^*(\Delta y_{t+1}^4) - \sigma_t^{*4}) + \left(\lambda' X_t \sigma_t^{*2} + \lambda^{(2)} E_t^*(\Delta y_{t+1}^3)\right)^2 + \eta_t$$

The λ and $\lambda^{(2)}$ parameters can be estimated using the same nonlinear least squares approach described in Section 2.3.1. This approach is closely related to the approach of Tetlock et al. (2024) in equity markets.

This second-order approximation approximately doubles the estimated annual risk premium versus a first order approximation, bringing it in line with the size of the quarterly risk premium (i.e., approximately 4× larger). In contrast, applying the same constant second-order term to the quarterly forecast would only increase the average estimated risk premium by 20%, which is within the confidence intervals of the zero-residual-coskew estimates. Section 5 considers estimates that alow for a time-varying second order coefficient.

3 Risk-based interest-rate expectation results

This section reports the results and forecasting performance of the risk-based interest-rate expectation measure. I first show that the average interest-rate risk premium was large. I then show that conditional forecasts outperform existing measures out of sample and capture well-known relationships between the yield curve and risk premium. I close by showing evidence that equity betas and aggregate bond-market duration may drive some of the variance in investors' interest rate exposure.

This paper's results focus on the 10-year swap rate, the most liquid tenor. The methodology can be readily applied to other tenors of swap rate with similar results, although a few additional complications are introduced by the zero lower bound for short tenors. Appendix A shows the results from forecasts of the 1-, 2-, 5-, 20-, and 30-year interest-rate risk premium, assuming λ_t is constant across tenors.

3.1 Average risk premium

In my baseline estimation, I find the average 10y-in-1q risk premium is 11 bp with a standard error of 3 bp (or 43 bp and 10 bp annualized). This amounts to a very strong rejection of the expectations hypothesis with a t-statistic of 4.3. The statistical significance is high enough that a simple difference in priors is unlikely to be able to account for the apparent risk premium, as suggested by Farmer et al. (2024) for the level-based forecasts.

The level of risk premium is similar to the realized quarterly risk premium over this period of 9 bp, or 34 bp annualized. Markets seem to have been demanding compensation for bearing interest-rate risk, rather than making systematic mistakes. The lower end of the confidence interval also covers the average level of quarterly risk premium implied by a model in which agents always expect interest rates to stay at their current level (i.e., a random walk model): 7 bp, or 27 bp annualized.

These estimates are substantially larger than those implied by the Survey of Professional Forecasters or stationary term structure models such as Adrian et al. (2013). On average

2002-2023, both of these sources find large negative interest rate risk premium. There is some suggestive evidence that these negative risk premia are not representative of the beliefs of large institutional investors. For example, Couts, Gonçalves, and Loudis (2023) document that, on average, investment consultants and managers expect a 1.1% long-term excess return on government bonds 1987–2022 (see their Table 3). This is roughly equivalent to an average long-term expected interest rate risk premium of 0.28, assuming an average bond market duration of 4.

The average estimated exposure (λ_t) is 35, with a standard error of 11, corresponding to a duration of 35 years for a log-utility fixed-income investor, or approximately $35/\gamma$ for a CRRA investor. This high sensitivity to interest-rate changes could be consistent with models where leveraged intermediaries are the marginal investors in fixed-income markets. For example, Kekre, Lenel, and Mainardi (2024) suggest durations of 10 to 30 years for fixed-income arbitrageurs. Section 5 repeats this estimation allowing for the residual coskewness introduced by non-log CRRA preferences.

3.2 Forecasting performance

Moving beyond the average level of risk premium, in this subsection I demonstrate that the conditional forecasts successfully predict changes in interest rates out of sample better than alternative measures.

3.2.1 Prediction of interest rates

I first show that the measured risk premium predicts changes in interest rates Δy_{t+1}. Table 3 reports regressions of out-of-sample realized rate changes onto the conditional risk premium estimates at various horizons:

$$\Delta y_{t,t+h} = \alpha - \beta \times \widehat{\mathrm{RP}}_{t,h} + \varepsilon_{t+h}$$

where $\widehat{\mathrm{RP}}_{t,h}$ denotes the predicted risk premium from time t to $t+h$. For monthly horizons, I use the quarterly risk premium estimates divided by three. For annual forecasts I use the second-order approximation methodology described in section 2.3.5. If the forecasts are perfect, we should expect that $\alpha = 0$ and $\beta = 1$.

The estimated β is highly significant at monthly (t-statistic 3), quarterly (t-statistic 3.0), and annual (t-statistic 3.3) horizons. I also cannot reject $\beta = 1$ or $\alpha = 0$ at any horizon. This suggests the measure captures risk premia accurately in levels, not just direction. The R^2 values are high even at short horizons (6% monthly, and 7% quarterly), where term-structure models have traditionally struggled. The short-term performance allows for high trading profits as I will demonstrate in Section 3.2.3.

This predictive power is robust to controlling for other common bond return predictors, splitting up the subsample, or variance weighting. Panel (B) shows that the risk premium measure subsumes the predictive power of the term spread and the Cochrane & Piazzesi factor. I use the 10-year minus 3-month treasury yield term spread measure from FRED, and calculate the out-of-sample Cochrane & Piazzesi factor from rolling regressions of forward rates onto zero coupon bonds as described in Cochrane and Piazzesi (2005). After including my interest-rate risk premium measure, both factors are insignificant. Panel (B) also shows that the effects are robust to only focusing on the first or second half of the sample period and to using WLS, weighting the observations by the inverse of the physical variance estimated in section 2.3.2. The coefficient size is larger in the second half of the sample than the first, but in neither case can we reject $\beta = 1$, so the difference may simply be a matter of small sample size.

3.2.2 Performance vs other predictors

This predictive power is greater than that of professional forecasters, traditional term-structure models, or regressions based on other common bond return predictors. Table 4 reports out-of-sample relative R^2, or reduction in mean squared error, of the risk-based model vs other models. Relative R^2 is calculated as $1 - \sum \varepsilon_t^2 / \sum \nu_t^2$, where ε_t is the error when the risk-based model is used to forecast Δy_{t+1} and ν_t is the error when the alternative model is used. This represents the improvement in forecast accuracy from using the risk-based measure versus alternatives. A positive number implies that the risk-based measure delivers a more accurate forecast than the alternative.

I compare my forecasts of the 10-year yield to the expectations hypothesis, the survey of professional forecasters (SPF), out-of-sample forecasts from the two most commonly used dynamic term-structure models (DTSMs) (Adrian et al., 2013; Kim and Wright, 2005), out-of-sample forecasts generated from the two most common bond return prediction factors (the term spread and the Cochrane & Piazzesi factor), the Bauer and Rudebusch (2020) non-stationary "observed shifting endpoint" model, and a simple forecast of no change in rates ("random walk"). The SPF and Bauer and Rudebusch (2020) data are produced quarterly, and Bauer and Rudebusch (2020) forecasts are only produced up to 2018. I use the pre-crisis Kim & Wright calibration provided by the Federal Reserve Board to ensure most forecasts are out of sample.[9] Where forecasts of the 10 year par rate 1 quarter or 1 year away are not available, I use the closest available horizon of forecast. Details on the construction of the alternative benchmarks are included in Appendix J.

The risk-based interest-rate expectation measure delivers more accurate forecasts than all alternative benchmarks at all horizons. Improvements are highly statistically significant vs the expectations hypothesis, the stationary DTSMs, and the Cochrane & Piazzesi factor at the 1-month and 1-year horizons. The improvements versus the Bauer and Rudebusch (2020) model are weaker (7% improvement in quarterly R^2), and with just 18 years of data, the difference is not statistically significant.

The power of the risk-based forecasts comes from the use of theory to derive the level of interest-rate risk premium. If instead we were to take our predictor and regress it on changes in interest rates in each period, we would get a substantially worse out-of-sample performance, as the coefficient loadings swing around based on past performance. The last row of Table 4 compares the out-of-sample R^2 from the risk-based measure to a regression-based forecast constructed in this fashion and finds a substantial and significant

[9]The Federal Reserve board has noted that this older calibration performed poorly during the zero rate period, predicting large negative term premium

improvement in R^2.

The random walk forecast (i.e., the current rate) and the forecast based solely on the term spread remain surprisingly competitive, beating all benchmarks except for my risk-based forecast. However, ex-ante, there was no obvious theoretical justification to favor a forecast based only on the slope of the yield curve.

3.2.3 Economic significance

To demonstrate that the R^2 from the predictive regressions in Section 3.2 are economically meaningful, Figure 6 shows cumulative returns from trading strategies based on different interest-rate risk premium estimates. The strategy buys 10-year swap rate forwards to achieve a duration in each period equal to the estimated risk premium divided by double the risk-neutral variance:

$$D_t = \frac{\widehat{RP}_t}{2\sigma_t^{*2}}$$

Positions are scaled up based on measured risk premium, and down based on risk-neutral variance. Proposition 2 tells us this is approximately the level of interest rate exposure a CRRA investor with risk aversion coefficient 2 would choose if she believed in this risk premium estimate. In the case of my risk-based estimates, $D_t = \lambda_t/2$.

The risk-based strategy generates a cumulative return of approximately $30\times$ over the sample period, with a Sharpe ratio of 0.7, compared to less than 1 using SPF forecasts or Adrian et al. (2013) (ACM). Much of the outperformance comes from 2013–2021, when stationary models consistently predicted rising rates that did not arrive. Sharpe ratios are the same using different levels of risk aversion coefficient, although the amount of leverage and cumulative returns will differ.

These results should be interpreted cautiously. The strategies exhibit high volatility, and a substantial portion of the returns are driven by a few sharp movements, such as the sudden drop in interest rates in 2009 which almost doubles wealth. Nevertheless, the magnitude of outperformance shows that the statistical improvements in forecasting are economically meaningful.

3.3 Comparison with surveys and term-structure models

The time-series of interest rate forecasts correlates with other surveys and models, despite the different underlying methods and sources of information. Table 5 shows the correlation of the forecasts to alternative benchmarks including the Survey of Professional Forecasters (SPF) and the widely used Adrian et al. (2013) (ACM) affine term-structure model. The risk-based estimate is $> 50\%$ correlated with the SPF, ACM, and Term Spread forecasts. All of these forecasts load heavily on the slope of the yield curve.

The key difference between the risk-based and stationary DTSM forecasts lies in the behavior at low rates. The stationary assumption built into standard DTSMs mean that they always forecast rates to return to their long-term averages eventually. So when current and forward rates were low, as in 2016–2020, the ACM model forecast a large negative risk premium. Figure 4 plots the ACM risk premium time series versus my estimates. At their lowest, ACM forecasts implied investors were expecting a loss of 1.5% on 10-year bonds in a single quarter.[10] In contrast, my estimates make no assumption about stationarity of interest-rate levels, and end up forecasting almost no interest-rate risk premium in this period. These forecasts are more in line with the suggestive evidence from surveys of institutional investors who mostly do not seem to report negative expected excess returns on bonds (Dahlquist and Ibert, 2024; Couts et al., 2023).

3.4 Investor exposure to rates

Besides a risk premium estimate, the variance-based approach delivers an estimate of λ_t: the exposure to interest-rate increases of the log investor who chooses a linear interest-rate exposure. Figure 5 shows the evolution of the estimated (in-sample) λ_t over time. The series shows a gradual decline from approximately 70 before 2008 to -10 in 2020, followed by an increase to 35 in 2023. To provide suggestive evidence on the economic determinants of investor exposure, I repeat the conditional risk premium estimation exercise (i.e., equation

[10] Assuming duration of 8, and 10y-in-1q interest-rate risk premium of -0.2

6) allowing λ_t to be a function of economically-motivated variables instead of yield curve principal components. Table 6 reports coefficients from these alternative specifications. I find two significant relationships in Column (6).

First, λ_t decreases with the sensitivity of equity markets to interest-rate changes: the coefficient on the estimated sensitivity of stock returns to yield changes ("Equity β") is -2.5. This is consistent with a model in which a representative investor holds the aggregate wealth portfolio. In such a model exposure should approximately represent how much an increase in risk aversion lowers wealth, adjusted for risk aversion.[11] However, the large intercept in λ_t suggests that most interest-rate exposure comes from non-equity sources.

Second, λ_t appears to rise following increases in aggregate bond supply. A 1% increase in aggregate bond duration (scaled by GDP) raises λ by 0.7, or approximately 1.5%. This relationship aligns with models where slow-moving capital forces specialized arbitrageurs to absorb supply shocks, increasing their exposure and the compensation they demand. The long-term decline in exposure could be related to the long-term growth of arbitrageur capital relative to the aggregate bond duration.

In general, economically-motivated variables do a poor job explaining exposure relative to the simple principal components. The above two variables are not statistically significant on their own, and only add 5 ppt of R^2 over the constant-only model. The underlying causes of the reduction in sensitivity of investors to interest rates remain an open question.

4 Applications

This section considers two applications that traditional interest-rate forecasting methods are poorly suited to answer. I consider the average decline in long term rates around FOMC meetings and the positive stock-yield correlation of 2000–2021 and ask whether they are driven by interest-rate expectations or interest-rate risk premium. Both of these questions

[11] It is also consistent with an arbitrageur whose equity value is correlated with the overall market (e.g. a bank). For example, the rolling sensitivity to interest-rate changes of stock returns of the primary dealers from the He, Kelly, and Manela (2017) closely resembles that of the overall market.

require measuring short-term changes in risk premium. I therefore start by developing a simple measure of daily or weekly changes in risk premium that loosens the exposure-parameterization assumptions from Section 2.

4.1 Measuring short-term changes in interest-rate risk premium

To measure short-term changes in conditional interest-rate risk premium, we could simply take the daily changes of the risk premium estimates from Section 3. However, we might be concerned that the functional form we have imposed for λ_t does not accurately capture small, high-frequency changes in the relationship between interest-rate risk premium and the yield curve, particularly around shocks such as FOMC announcements.

To find an alternative measure with looser functional form assumptions, we can combine Propositions 2 and 3 to state interest-rate risk premium in terms of variance risk premium. Differentiating the relationship and then solving for dRP yields:

$$dRP_t \approx \frac{d\left(\sigma_t^{*2} - \sigma_t^2\right)}{\frac{E_t^*\left(\Delta y_{t+1}^3\right)}{\sigma_t^{*2}} + 2RP_t}$$

Change in interest-rate risk premium is revealed by change in variance risk premium, scaled by skewness and risk premium.[12]

This equation naturally suggests an estimator for short-term changes in risk premium. Simply plug in the fitted values for interest-rate risk premium, $\widehat{RP}_t$, and physical variance, $\Delta\hat{\sigma}_t^2$, from Section 3:

$$\widehat{\Delta RP}_t \;\approx\; \frac{\Delta\sigma_t^{*2} - \Delta\hat{\sigma}_t^2}{\frac{E_t^*\left(\Delta y_{t+1}^3\right)}{\sigma_t^{*2}} + 2\widehat{RP}_t}$$

This estimator no longer depends on short-term changes in exposure exactly aligning with changes in the interest-rate principal components used to parameterize λ_t. The functional form assumptions are still embedded in the risk premium estimates, $\widehat{RP}_t$, under the assumption that these are still a reasonable characterization of the level of exposure, even if

[12] I omit a term involving changes in skewness, $-E_t(\Delta y_{t+1}) d\frac{E_t^*(\Delta y_{t+1}^3)}{\sigma_t^{*2}}$, which is empirically negligible.

they don't capture short-term changes well.

It is not possible in general to measure short-term changes in physical variance, $\Delta\hat{\sigma}_t^2$, around shocks with any accuracy. I therefore only consider long term averages and covariances of this change in risk premium measure, rather than considering specific days. In the case of FOMC announcement window results I will treat these results as indicative. I also provide further evidence in Appendix G that there is not a large average decline in physical variance during these windows

4.2 Interest-rate risk premium during FOMC announcement windows

Hillenbrand (2025) documents that the entire secular decline in long-term rates from the 1990s to 2020 occurred during three-day windows around FOMC meeting and interprets this fact as evidence that markets primarily learn about long-run interest-rate levels from the Federal Reserve. Standard dynamic term-structure models are not well-suited to decompose announcement effects. They either inherently assume constant long-term expectations (e.g. Adrian et al., 2013) or update long-run expectations at lower frequencies (e.g. Bauer and Rudebusch, 2020).

In this section I provide suggestive evidence that this FOMC-window decline may reflect falling risk premia rather than expectations. I first document that there is a large decline in risk-neutral variance around FOMC meetings that appears to be consistent with a change in risk premium. I then show that the size of the declines is large enough to be potentially consistent with full FOMC-window effect being driven by interest-rate risk premium rather than expectations.

More generally, these findings address the debate over how much of the response of the long end of the yield curve to monetary policy announcements is due to the market learning information from the Fed (e.g. Nakamura and Steinsson, 2018) or due to term premium and demand (e.g. Hanson and Stein, 2015). My findings are consistent with the literature that argues that negative monetary shocks to short-term interest rates increase demand for long term bonds and push down term premia e.g., due to reaching for yield (Hanson and Stein, 2015), improvements in intermediary capital (Kekre et al., 2024), or MBS or bank deposit duration hedging effects (Hanson, 2014; Rogers, 2024). The demand effects coming from the negative shocks to short-term rates during the sample period could explain the FOMC-window decline in term premium and long term rates.

4.2.1 The FOMC-window variance decline

Table 7 shows that risk-neutral variance of interest rates declines across tenors and maturities during FOMC windows. The proportional declines are similar whether we look at short or long tenors, 3-month or 1-year horizons. Even five-year variance shows substantial FOMC-window declines.

This decline is much larger than the net change in physical or risk-neutral variance during the sample period. For example, over the full sample, five-year risk-neutral variance for the 10-year rate actually increased from 3.0 to 4.2 percentage points squared. Yet during FOMC windows alone, it declined by a cumulative 12.2 percentage points squared. For variance to end up higher despite falling dramatically during FOMC meetings requires an offsetting increase of 13.4 percentage points squared between meetings.

The evidence for the idea that the market primarily learns about the long term from the Federal reserve was the fact that the size FOMC-window and aggregate declines in long-term forward rates coincide. This is not the case for the declines in long-term risk-neutral variance. If there is information in the FOMC announcement windows about variance, it must be systematically countered by other information received outside the FOMC windows. It seems likely that instead, these FOMC-window declines in risk-neutral variance are a result of declines in risk premia that reverse outside the windows.

In Appendix G I show two further pieces of evidence to support the idea that these declines come from variance risk premium rather than physical variance. First, they are too large to be explained by simple "mechanical effects" from the high-variance FOMC window dropping out of the forecast period. Second, using the information from these changes does not improve on the regression-based forecast of interest-rate variance.

4.2.2 Measuring the implied risk premium decline

To understand how much interest-rate risk premium decline could be implied by the declines in variance risk premium, I will find the result if we assume there is no change in the physical variance during the estimation windows, i.e., $\Delta\sigma_t^2 = 0$. This is roughly what is implied by the regression-based variance stimates from Section 3.

From 2007–2018 (the end of the sample period from Hillenbrand, 2025), the average FOMC-window decline in the 10y-in-1y forward rate is 2.8 bp per meeting. The estimation procedure above yields a decline in 10y-in-1y risk premium of 3.5 bp, with a standard error of 1.2 bp, using daily data with standard errors clustered by meeting.

In other words, the entire FOMC-window decline in long-term rates could be attributable to declines in risk premia that are expected to be realized over the course of a year, with no change in long-term expectations at all. To the extent that physical variance declined during these windows as well, the decline in risk premium could be smaller. It is likely that the true decline involved some combination of risk premium and expectations.

4.3 Stock yield correlation

The correlation between stocks and long-term yields famously switched from negative to positive in the late 90s. A substantial literature seeks to explain the causes of the modern positive correlation.

By some explanations, the source is risk premium: when stocks go up, bond risk premium and yields go up. For example, Antolin-Diaz (2025) proposes flight to safety shocks move institutional investors out of stocks into long-term bonds.[13] By others, the covariance is expectations driven: news that leads to low stock returns lowers expected rates. For example, Campbell et al. (2017) propose inflation shocks now correlate with stock returns, or

[13] The convenience yield covariance from Acharya and Laarits (2023) and the "hedging premium" shocks from Cieslak and Pang (2021) should also involve positive correlation of stock returns with bond risk premia.

the framework from Gormsen and Lazarus (2025) implies positive long-term growth news will push stocks and yields in the same direction.

This paper lets us test between these explanations. Does the 1-year interest-rate risk premium correlate positively or negatively with stock returns? I show that the correlation is negative. Risk premium shocks cannot explain the positive stock-yield correlation 2002-2022. Table 8 shows the correlation of estimated changes in 10y-in-1y risk premium with stocks. I show results for daily, weekly, and monthly observation horizons, using the short-term change in risk premium measure developed in Section 4.1. All show a negative correlation, although the weekly figure is not significant. Risk-neutral variance tends to increase on days with bad stock returns more than the regression-estimated physical variance, suggesting risk premium rises.

In contrast, a traditional term-structure model finds the opposite. The second shows the correlation with changes in the Kim and Wright (2005) term-structure model, which are all positive. The model loads heavily on the level of interest rates due to the stationary structure, and hence interprets a portion of all increases in the level of interest rates to be due to risk premium, yielding a positive term-premium-yield correlation.

5 Validating the residual coskew assumption

The main theoretical result of this paper is that the interest-rate risk premium is a function of the variance risk premium and an unobservable term, the "residual coskew," that captures the correlation of squared rate changes with the inverse SDF. For my main empirical results, I assume that the residual coskew is zero. In this section I demonstrate that this is an empirically reasonable assumption, and that after adding plausible sources of residual coskewness, the main results of the paper may modestly over- or underestimate true risk premium, but retain a $> 95\%$ correlation.

I first consider a general second-order approximation of the inverse SDF with respect to interest rates. This approach allows for nonlinearities introduced by a wide range of different sources, including, for example, intermediary-specific exposures from selling options, equity market exposures, or high risk aversion.

I then consider four specific sources of residual coskewness. First, the inverse SDF could be nonlinear in interest rates because risk aversion is higher than log. Second, nonlinearity could be introduced by convexity of fixed income investments. Third, the inverse SDF could be a function of multiple different interest rates or interest rate factors as in typical term structure models. Fourth, the inverse SDF could include non-interest-rate exposure, such as equities, that are uncorrelated with rates but correlated with squared rate changes. I can measure the first three sources using risk-neutral swaption moments, and estimate the fourth using the empirical co-moments of equities and interest rates.

I re-estimate the interest-rate risk premium after allowing for each of these sources of residual coskewness. Table 9 shows the results, plotted in Figure 7. In all cases the risk premium remains $> 95\%$ correlated with the main estimates of this paper, although the average level of risk premium could be underestimated by up to 50% on average for the case of high risk aversion (CRRA $\gamma = 4$) or overestimated by 30% in the case of multiple interest rate factors. The remainder of this section describes how I quantify these sources of residual coskewness.

5.1 Second-order approximation

As an alternative to the zero residual coskewness model, we can allow the inverse SDF to have a quadratic exposure to rates. This would allow for nonlinear exposure of the representative investor to interest rates in general terms. For example, we might expect concave exposures if the SDF is driven primarily by options-sellers, as suggested by Dew-Becker and Giglio (2025) for equities markets before 2011.

If we write the inverse SDF as a projection onto a second-order function of changes in rates, then we can write the inverse SDF and variance risk premium as:

$$\frac{1}{M_{t+1}} = R_{f,t}\left(1 - \lambda_t^{(1)}\Delta y_{t+1} - \lambda_t^{(2)}(\Delta y_{t+1}^2 - \sigma_t^{*2}) + \varepsilon_{t+1}\right)$$
$$VRP_t = \lambda_t^{(1)} E_t^*(\Delta y_{t+1}^3) + \lambda_t^{(2)} E_t^*((\Delta y_{t+1}^4) - \sigma_t^{*4}) + RP_t^2$$

where I have assumed that the risk-neutral covariance of the residual ε_{t+1} with cubed changes in rates is zero.

This relationship suggests that the loading of the inverse SDF on squared changes in rates ($\lambda_t^{(2)}$) can be identified by the time-series covariance of the variance risk premium with the fourth risk-neutral moment. This is the approach taken by Tetlock et al. (2024) for equity markets.

I therefore repeat the same estimation process described in section 2, with this additional exposure $\lambda_t^{(2)}$. I parameterize $\lambda_t^{(2)}$ based on state variables that are likely to correlate with concavity of dealer exposure to interest rates, as well as variables based on the term structure literature. Specifically, I assume:

$$\lambda_t^{(2)} = \lambda^{(2)\prime} Z_t^{(2)}$$
$$Z_t^{(2)} = [1,\ pre2011,\ CAgg,\ \sigma_t^*,\ Skew_t^*]'$$

The definition and motivation behind each of these variables is as follows. $pre2011$ is a dummy variable for observations during and before 2010. Dew-Becker and Giglio (2025) suggest that dealers had very concave portfolios in equity markets (i.e., were "short gamma") up to 2010, but not thereafter. $CAgg$ is the aggregate convexity of the bond market (using the Bloomberg Agg index), and is meant to proxy for investor demand for swaptions. Swaptions are typically used to hedge nonlinearities in portfolios, for example from mortgage backed security or callable bond prepayment options. The demand should therefore be greatest when the "negative convexity" of these options is at its greatest. I include the risk-neutral volatility and skewness because these typically appear in the formula for the price of risk associated with volatility shocks in term structure models with stochastic volatility (e.g. Creal and Wu, 2015).

Table 9 shows that this second-order approximation does not change the results dramatically. The resulting 10y-in-1q risk premium estimates are $> 97\%$ correlated with the main estimates, and approximately 4 bp larger. This difference is well within the confidence interval of the estimated mean.

5.2 Risk aversion

If an investor is more risk averse than the log-utility case, then the inverse SDF will be a convex function of interest rates, and the residual coskew will therefore be positive. This will lead the main estimates of this paper to underestimate the true interest-rate risk premium associated with a given level of variance risk premium.

Intuitively, two investors can generate the same interest-rate risk premium through different combinations of risk aversion and leverage. A low-risk-aversion investor needs higher leverage to reach the same risk premium as a high-risk-aversion investor. But the low-risk-aversion, high-leverage investor fears far-away states more than the high-risk-aversion investor, because a large move brings her close to insolvency. Hence the variance risk premium associated with any given level of interest-rate risk premium should be lower in a high-risk-aversion world than a low-risk-aversion world. Appendix F offers a proof of this general statement for CRRA utility with risk aversion greater than 1.

We can measure the exact size of this under-estimation for a given level of risk aversion using the higher risk-neutral moments of Δy_{t+1}. For example, for a CRRA investor with $\gamma = 2$ and linear exposure to interest rates, the residual coskew will depend on the fourth moment. To quantify the impact of the residual coskew, I re-estimate the conditional risk-premium figures from Section 3, for $\gamma = 2$ using the fourth moment data.

The results are shown in the first row of Table 9 and plotted in Figure 7. If the $\gamma = 2$ model is true, then the main estimates of this paper would underestimate risk premium by 34%, but would retain a 99% correlation with the true values. The methodology and derivations are described in Section F, and shows that higher levels of risk aversion are likely to only modestly increase the magnitude of this underestimation.

5.3 Interest rate convexity

Another possible reason for nonlinearity of the SDF is if the payoffs are not linear in rates. For example, bond prices are convex in rates and mortgage-backed security prices can be concave in rates. In this case, residual covariance will also depend on higher risk-neutral moments of Δy_{t+1}. However, it turns out that the degree of convexity present in diversified fixed income portfolios is too small to affect the λ estimates. In Appendix F I allow the investor to have payoff convexity relative to duration similar to that of the aggregate bond index, as would be the case if their exposure to interest rates came from a levered investment in the aggregate bond portfolio. The resulting estimates for quarterly interest-rate risk premium only differ by 0.6 bp on average and are $> 99\%$ correlated (see the second row of Table 9).

5.4 Multiple interest rate factors

So far I have only considered a single interest rate Δy_{t+1}. But an investor could have different exposures to long-term and short-term rates. I therefore conduct the same estimation approach allowing for multiple interest rate factors. In the case of linear exposures to N different interest rates, the inverse SDF would then become:

$$\frac{1}{M_{t+1}} = R_{f,t}\left(1 - \sum_{i=1}^{N} \lambda_t^{(i)} \Delta y_{t+1}^{(i)}\right)$$

where $\lambda_t^{(i)}$ represents the investor's exposure to interest rate $y_{t+1}^{(i)}$.

To address this concern, I generalize my Propositions 2 and 3 to apply to a setting with multiple interest rates in Appendix F. I then re-estimate the model in a multiple-interest rate setting, assuming that all swap rates can be exactly decomposed into three principal components with constant loadings. This is a relatively innocuous approximation: the first three principal components explain 99.9% of quarterly variance in swap rates.

Estimating this multi-factor model requires measuring the 3×3 risk-neutral factor covariance matrix and the $3 \times 3 \times 3$ third moment "cube." Fortunately this is possible, because swaption markets allow us to measure risk-neutral variance for many different tenors at the same time. I calculate the risk-neutral variance and third moment for six different tenors of swap rate (1, 2, 5, 10, 20, and 30 years). Because a 3×3 covariance matrix has six independent elements, the entire risk-neutral factor covariance matrix can be identified by solving for the values that match each tenor's variance. The third moment cube can be estimated similarly, although it is not fully identified because it has ten independent elements. I therefore select the value that matches the observed third moments and minimizes the sum of squared co-third-moments. An alternative approach using a Kalman filter for identification yields similar results.

The resulting estimated 10y-in-1q risk premia are 96% correlated with this paper's main estimates, but are 3 bp lower on average, as shown in Table 9 and plotted in Figure 7. The estimated average risk premium of 0.09 ppt per quarter (0.36 annualized) would still imply that the realized bond returns are almost entirely attributable to risk premium, as found in section C. I use the univariate model instead of the multi-factor model for the main results of this paper for simplicity of exposition and to avoid the potential noise and estimation error introduced by the estimation of the risk-neutral skewness cube and the six different physical variances for each interest rate.

5.5 Non-interest-rate factors

An unconstrained investor's payoffs are likely to also include non-interest-rate risk, for example from equities or corporate bonds or other risk factors. If these factors are negatively correlated with Δy_{t+1}^2 after removing any directional interest rate exposure, then this could explain the variance risk premium without any interest-rate risk.

I cannot measure the risk-neutral covariance of interest rates with equities or other interest risk factors directly because there are no derivatives that reveal these co-moments. As a second-best approach, I instead measure the physical covariance of the SDF with Δy_{t+1}^2 under various linear factor models, after removing directional interest-rate risk from these models. This measures the variance risk premium that would be expected under these models that cannot be explained by directional interest-rate risk. I consider CAPM, the Fama & French three factor model, the Fama & French three factor model plus a corporate bond return factor,[14] and a seven factor model consisting of the Fama & French five factor model plus the momentum factor plus corporate bond excess returns. Equity factors are taken from the Ken French data library.

For each model, I calculate the physical residual coskew as follows. I first calculate the price of risk associated with each factor using the standard formulas applied to daily data 2002-2023 (i.e., using mean factor returns and the sample factor covariance matrix). On a rolling quarterly basis, I then calculate the residuals of this SDF after regressing on daily changes in interest rates. I then calculate the quarterly coskewness of the sum of the residual SDF realizations ε with the squared interest rate changes. If the interest rate changes and the residual SDF are approximately martingales on a daily horizon, then this quarterly coskew can be measured from daily data as:

$$cov\left(\sum_{i=0}^{62}\varepsilon_{t-i},\left(\sum_{t=0}^{62}\Delta y_{t-i}\right)^2\right) = 63E(\varepsilon_t\Delta y_t^2)+\sum_{i=1}^{62}(63-i)\left(cov(\varepsilon_{t-i},\Delta y_t^2)+2cov(\Delta y_{t-i},\varepsilon_t\Delta y_t)\right)$$

as shown in Neuberger and Payne (2020). Using the log instead of simple SDF yields similar results, as does the equivalent calculation using monthly or quarterly observations on a rolling annual basis, albeit with lower power.

The resulting residual coskew estimates are too small to explain the observed variance risk premium. Table 10 expresses the measured residual coskews as a percentage of the physical quarterly variance of rates. This represents the relative variance risk premium that

[14] I use the return on the IBOXX IG corporate index minus the IBOXX treasury index as quoted on LSEG workspace

would be explained by the orthogonal-to-interest-rates components of each factor model. While the observed variance risk premium is 17%, the explained portion is only 1% – 3%. Unsurprisingly, re-running the full term premium estimation allowing for the equity-contribution to residual coskew does not change the results meaningfully. The second row of Table 9 shows that adjusting the observed variance risk premia for the estimated CAPM residual coskew yields interest-rate risk premia that are > 99% correlated with the main estimates and differ by just 0.6 bp.

It may seem surprising that residual equity coskew is so small. Over the long run, equities do seem to covary with squared interest rate changes: stocks dropped both when interest rates fell sharply in 2009 and when they rose sharply in 2022. However, over shorter horizons, within each quarter, most of this effect is absorbed by the directional interest-rate risk: stocks were correlated negatively with rates during the 2022 rate rise and positively during the 2009 rate fall.

6 Conclusion

This paper introduces a new approach to measuring interest-rate expectations and risk premia. By shifting focus from the non-stationary level of interest rates to their stationary variance, we can extract expectations from options data and test hypotheses about interest-rate dynamics in a new way.

The core theoretical result links the interest-rate risk premium perceived by an investor to the variance risk premium perceived by the same investor and an unobserved "residual coskew" quantity related to the nonlinearity of the his marginal utility with respect to interest rates. If we assume as a benchmark that this residual coskew term is zero, an assumption that does not appear unreasonable based on the data, then we can derive the exact perceived interest-rate risk premium from the variance risk premium.

The empirical findings are fourfold. First, the historical interest-rate risk premium is large. The strong average bond returns over the past 25 years appear to be mostly due to risk premium rather than forecasting errors. Second, I construct a conditional measure of the interest-rate risk premium that outperforms standard dynamic term-structure models and survey-based measures in forecasting excess returns at short horizons, with a trading strategy based on the measure generating large economic profits. Third, I find that the secular decline in long-term rates concentrated in FOMC announcement windows may be better explained by a compression of risk premia than by learning about the long-run level of rates. Risk-neutral variance falls by a large amount during these windows, a pattern more consistent with temporary changes in risk appetite or exposure than with information updates. Finally, I show that the positive stock-yield correlation of the 2000s onward is not driven by bond risk premium. This suggests explanations based on inflation or growth expectations may have more promise that than explanations driven by investor flight to quality.

References

Acharya, Viral V., and Toomas Laarits, 2023, When do treasuries earn the convenience yield? — a hedging perspective, *NBER Working Paper* 31863.

Adrian, Tobias, Richard K Crump, and Emanuel Moench, 2013, Pricing the term structure with linear regressions, *Journal of Financial Economics* 110, 110–138.

Antolin-Diaz, Juan, 2025, How did government bonds become safe?, *MIT Sloan Working Paper* 7314.

Barnes, Chris, 2024, Swaption volumes by strike, *Clarus FT* Q2.

Bauer, Michael D., and Mikhail Chernov, 2024, Interest rate skewness and biased beliefs, *Journal of Finance* 79, 173–217.

Bauer, Michael D., and Glenn D. Rudebusch, 2020, Interest rates under falling stars, *American Economic Review* 110, 1316â54.

Bollerslev, Tim, 1986, Generalized autoregressive conditional heteroskedasticity, *Journal of Econometrics* 31, 307–327.

Bollerslev, Tim, George Tauchen, and Hao Zhou, 2009, Expected stock returns and variance risk premia, *Review of Financial Studies* 22, 4463–4492.

Boyarchenko, Nina, Pooja Gupta, Nick Steele, and Jacqueline Yen, 2018, Negative swap spreads, *Federal Reserve Bank of New York Economic Policy Review* 24.

Breeden, Douglas T., and Robert H. Litzenberger, 1978, Prices of state-contingent claims implicit in option prices, *The Journal of Business* 51, 621–651.

Campbell, John Y., and Robert J. Shiller, 1991, Yield spreads and interest rate movements: A bird's eye view, *The Review of Economic Studies* 58, 495–514.

Campbell, John Y., Adi Sunderam, and Luis M. Viceira, 2017, Inflation bets or deflation hedges? the changing risks of nominal bonds, *Critical Finance Review* 6, 263–301.

Carr, Peter, and Dilip Madan, 1998, Towards a theory of volatility trading, in Robert A. Jarrow, ed., *Volatility: New Estimation Techniques for Pricing Derivatives*, chapter 29, 417–427 (Risk Books, London).

Chabi-Yo, Fousseni, and Johnathan Loudis, 2020, The conditional expected market return, *Journal of Financial Economics* 137, 752–786.

Choi, Hoyong, Philippe Mueller, and Andrea Vedolin, 2017, Bond variance risk premiums, *Review of Finance* 21, 987–1022.

Cieslak, Anna, and Hao Pang, 2021, Common shocks in stocks and bonds, *Journal of Financial Economics* 142, 880–904.

Cochrane, John H., and Monika Piazzesi, 2005, Bond risk premia, *American Economic Review* 95, 138â160.

Collin-Dufresne, Pierre, and Robert S. Goldstein, 2002, Do bonds span the fixed income markets? theory and evidence for unspanned stochastic volatility, *The Journal of Finance* 57, 1685–1730.

Collin-Dufresne, Pierre, Robert S. Goldstein, and Christopher S. Jones, 2009, Can interest rate volatility be extracted from the cross section of bond yields?, *Journal of Financial Economics* 94, 47–66.

Corsi, Fulvio, 2009, A simple approximate long-memory model of realized volatility, *Journal of Financial Econometrics* 7, 174–196.

Couts, Spencer J., Andrei S. Gonçalves, and Johnathan Loudis, 2023, The subjective risk and return expectations of institutional investors Working paper.

Creal, Drew D., and Jing Cynthia Wu, 2015, Estimation of affine term structure models with spanned or unspanned stochastic volatility, *Journal of Econometrics* 185, 60–81.

Dahlquist, Magnus, and Markus Ibert, 2024, Equity return expectations and portfolios: Evidence from large asset managers, *The Review of Financial Studies* 37, 1887–1928.

Dew-Becker, Ian, and Stefano Giglio, 2025, The decline of the variance risk premium: Evidence from traded and synthetic options, *FRB of Chicago Working Paper* 2025.

Drechsler, Itamar, and Amir Yaron, 2011, What's vol got to do with it, *Review of Financial Studies* 24, 1–45.

Duffee, Gregory R, 2002, Term premia and interest rate forecasts in affine models, *Journal of Finance* 57, 405–443.

ECB, 2005, Box 2: The term structure of implied volatilities, *ECB Monthly Bulletin* December.

Fama, Eugene F., and Robert R. Bliss, 1987, The information in long-maturity forward rates, *The American Economic Review* 77, 680–692.

Farmer, Leland E., Emi Nakamura, and Jón Steinsson, 2024, Learning about the long run, *Journal of Political Economy* 132, 3334–3377.

Fornari, Fabio, 2005, The rise and fall of us dollar interest rate volatility: evidence from swaps, *BIS Quarterly Review* September.

Gormsen, Niels Joachim, and Eben Lazarus, 2025, Interest rates and equity valuaitons, *Working paper* .

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright, 2007, The u.s. treasury yield curve: 1961 to the present, *Journal of Monetary Economics* 54, 2291–2304.

Hanson, Samuel G., 2014, Mortgage convexity, *Journal of Financial Economics* 113, 270–299.

Hanson, Samuel G., and Jeremy C. Stein, 2015, Monetary policy and long-term real rates, *Journal of Financial Economics* 115, 429–448.

He, Zhiguo, Bryan Kelly, and Asaf Manela, 2017, Intermediary asset pricing: New evidence from many asset classes, *Journal of Financial Economics* 126, 1–35.

Hillenbrand, Sebastian, 2025, The fed and the secular decline in interest rates, *Review of Financial Studies* 38, 981–1013.

ISDA, 2014, Size and use of the non-cleared derivatives market, *An ISDA study* April.

Joslin, Scott, and Yaniv Konchitchki, 2018, Interest rate volatility, the yield curve, and the macroeconomy, *Journal of Financial Economics* 128, 344–362.

Kekre, Rohan, Moritz Lenel, and Federico Mainardi, 2024, Monetary policy, segmentation, and the term structure NBER Working Paper No. 32324.

Kim, Don H, and Jonathan H Wright, 2005, An arbitrage-free three-factor term structure model and the recent behavior of long-term yields and distant-horizon forward rates, Technical Report 2005-33, Board of Governors of the Federal Reserve System.

Kozicki, Sharon, and P A Tinsley, 2001, Shifting endpoints in the term structure of interest rates, *Journal of Monetary Economics* 47, 613–652.

Kremens, Lukas, and Ian Martin, 2019, The quanto theory of exchange rates, *American Economic Review* 109, 810â43.

Martin, Ian, 2017, What is the expected return on the market?, *The Quarterly Journal of Economics* 132, 367–433.

Martin, Ian, and Christian Wagner, 2019, What is the expected return on a stock?, *Journal of Finance* 74, 1887–1929.

Martin, Ian W. R., 2025, Information in derivatives markets: Forecasting prices with prices, *Annual Review of Financial Economics* .

Nakamura, Emi, and Jón Steinsson, 2018, High-frequency identification of monetary nonneutrality: The information effect*, *The Quarterly Journal of Economics* 133, 1283–1330.

Neuberger, Anthony, and Richard Payne, 2020, The skewness of the stock market over long horizons, *The Review of Financial Studies* 34, 1572–1616.

Rogers, Robert, 2024, Bye bye beta: deposit duration with fixed rates, *Working paper* .

Tetlock, Paul C., Jack McCoy, and Neel Shah, 2024, The implied equity premium Working paper.

Trolle, Anders B., and Eduardo S. Schwartz, 2014, The swaption cube, *Review of Financial Studies* 27, 2307–2353.

van Binsbergen, Jules, 2020, Duration-based stock valuation, CEPR Discussion Papers 14904, C.E.P.R. Discussion Papers.

van Dijk, Dick JC, Siem Jan Koopman, Michel van der Wel, and Jonathan H Wright, 2014, Forecasting interest rates with shifting endpoints, *Journal of Applied Econometrics* 29, 693–712.

Tables

Table 1: **Summary statistics**

The table reports summary means and standard deviations for the key model inputs and outputs. All statistics are quoted for January 2002 – June 2023 based on weekly data. Interest rates are reported in *ppt*, variances in ppt^2, and the third moment in ppt^3. Outputs are shown as produced by expanding-windows for out-of-sample analysis.

		Quarterly		Annual	
		Mean	Std Dev	Mean	Std Dev
Data	10 year rate (y_t)	3.18	1.36	3.18	1.36
	Change in 10y rate (Δy_{t+1})	-0.09	0.48	-0.33	0.94
	R.n variance (σ_t^{*2})	0.27	0.18	1.03	0.52
	R.n skewness ($skew_t^*$)	0.29	0.23	0.30	0.23
	R.n third moment ($E_t(\Delta y_{t+1}^3)$)	0.05	0.08	0.39	0.48
	Realized variance (RV_t)	0.23	0.18	0.92	0.59
Outputs	Conditional variance ($\hat{\sigma}_t^2$)	0.22	0.10	0.89	0.29
(OOS forecasts)	Variance risk prem. ($\widehat{VRP}_t$)	0.04	0.11	0.14	0.35
	Interest-rate risk prem. ($\widehat{RP}_t$)	0.12	0.13	0.32	0.35

Table 2: **Interest-rate variance forecasting performance**

Panel (A) reports out-of-sample relative R^2 from forecasts of 10 year swap rate realized variance compared to two benchmarks: risk-neutral variance, and lagged realized variance (random walk). Relative R^2 is calculated as $1 - MSE(A)/MSE(B)$, where $MSE(A)$ is the mean-squared error of the regression forecast and $MSE(B)$ of the alternative. Panel (B) reports coefficients from a regression of realized variance risk premium (i.e., risk-neutral variance minus realized variance) on out-of-sample forecast variance risk premium (i.e., risk-neutral variance minus forecast variance). Forecasts are calculated using an expanding window starting in 1988. Sample period is 2002–2023.

Panel (A): R^2 of regression-forecast vs alternative benchmarks

	Quarterly	Annual
Risk neutral variance	0.200	0.239
Random walk	0.136	0.327
Observations	1064	244
Frequency	Weekly	Monthly

Panel (B): Variance risk premium (VRP) forecasting regressions

	Dependent variable: Realized VRP	
	Quarterly	Annual
Predicted VRP	0.625***	0.653***
	(0.142)	(0.134)
Intercept	0.009	0.010
	(0.010)	(0.014)
R2	0.235	0.271
Observations	1064	244
Frequency	Weekly	Monthly
Newey-West lags	13	12
Note:	*p<0.1; **p<0.05; ***p<0.01	

Table 3: **Interest-rate level forecasting regressions**

Results from regressions of realized changes in 10-year swap rates on predicted risk premium: $\Delta y_{t,t+h} = \alpha - \beta \times \widehat{RP}_{t,h} + \varepsilon_{t+h}$. Panel (A) reports results by different forecast period lengths. Panel (B) reports results robustness tests, including controls for the 3m-10y term spread and the Cochrane and Piazzesi (2005) factor (column 1), first and second half subsamples (columns 2 & 3), and WLS weighting by inverse estimated variance (column 3). Sample period is March 2002– June 2023. Newey-West standard errors in parentheses.

Panel (A): Regressions by forecast horizon

	1m	1q	1y
Expected change in rates	1.472***	0.980***	0.848***
	(0.49)	(0.33)	(0.25)
Intercept	0.010	0.024	-0.078
	(0.02)	(0.05)	(0.20)
R^2	0.057	0.074	0.115
Observations	1077	1064	244
Frequency	Weekly	Weekly	Monthly
NW lags	5	13	12

Panel (B): Quarterly robustness tests

	With controls	Sample: 1st half	Sample: 2nd half	WLS
	(1)	(2)	(3)	(4)
Interest-rate risk premium	0.812**	0.851*	1.759**	1.004***
	(0.356)	(0.446)	(0.830)	(0.331)
Term Spread	-0.017			
	(0.056)			
Cochrane-Piazzesi	-0.852			
	(0.710)			
Intercept	0.020	-0.010	0.057	0.011
	(0.098)	(0.110)	(0.072)	(0.055)
R^2	0.078	0.049	0.027	0.065
Observations	253	510	555	1064
Frequency	Monthly	Weekly	Weekly	Weekly
NW lags	3	13	13	13

Note: *p<0.1; **p<0.05; ***p<0.01

Table 4: **Interest-rate level forecasting improvement vs alternatives**

This table reports the relative R^2 of the risk-based interest-rate risk premium measure versus alternative benchmarks, 2002–2023. Relative R^2 is calculated as $1 - MSE(A)/MSE(B)$, where $MSE(A)$ is the mean squared error of the risk-based forecast of changes in interest rates and $MSE(B)$ of the alternative. Positive numbers denote the risk-based measure has lower MSE. P-values from one-sided Diebold-Mariano tests for an improvement in forecast accuracy are included in square brackets. Weekly or monthly frequencies are not available for the SPF or Bauer and Rudebusch (2020). Weekly frequencies are not available for Adrian et al. (2013) or Cochrane and Piazzesi (2005).

	Monthly	Quarterly	Annual
Expectations Hypothesis	0.080	0.107	0.217
	[0.031]	[0.098]	[0.046]
Survey of Prof Forecasters		0.294	0.284
		[0.007]	[0.049]
Adrian et al. (2013)	0.080	0.096	0.168
	[0.007]	[0.083]	[0.245]
Kim and Wright (2005)	0.803	0.596	0.489
	[0.000]	[0.000]	[0.008]
Bauer and Rudebusch (2020)		0.066	0.084
		[0.359]	[0.338]
Cochrane and Piazzesi (2005)	0.082	0.089	0.429
	[0.057]	[0.198]	[0.029]
Term Spread	0.043	0.063	0.169
	[0.042]	[0.108]	[0.124]
Random Walk	0.036	0.045	0.028
	[0.097]	[0.230]	[0.344]
Regression-rescaled risk-based	0.062	0.175	0.254
	[0.035]	[0.044]	[0.035]
Obs. frequency	Weekly	Weekly	Monthly

Note: Diebold-Mariano one-sided p-value in denoted in square brackets

Table 5: **Comparison of interest-rate risk premium estimates**

This table reports average risk premia from different models for 10-year swap rates, and their correlation with the risk-based estimates from this paper. The left two columns show the results for quarterly forecasts, and the right two columns for annual forecasts. "Realized" shows the ex-post change in rates vs forward-implied rates. "Risk-based" uses the variance risk premium method developed in this paper. Other models are as described in Section 3.2. Monthly data from 2002–2023. Means are expressed in percentage points. Kim & Wright forecasts are not available at a quarterly horizon.

	Quarterly		Annual	
	Mean	Corr(risk-based)	Mean	Corr(risk-based)
Realized	0.08		0.34	
Risk-based out-of-sample forecast	0.12		0.31	
Risk-based full-sample estimate	0.11		0.46	
(Std. error)	(0.03)		(0.20)	
Other forecasts:				
Survey of Prof Forecasters		0.37	-0.12	0.63
Adrian et al. (2013)	0.02	0.61	0.06	0.72
Kim and Wright (2005)	-0.34		-0.36	0.33
Bauer and Rudebusch (2020)	-0.06	-0.14	0.08	0.05
Cochrane and Piazzesi (2005)	-0.04	-0.23	-0.20	-0.03
Term spread	0.15	0.51	0.63	0.69
Random Walk	0.07	0.59	0.29	0.70

Table 6: **Determinants of marginal investor exposure**

This table reports coefficients from estimating exposure parameter λ_t as a linear function of different state variables. λ_t captures the marginal investor's sensitivity to interest-rate increases. PC1, PC2, and PC3 are the first three principal components of the yield curve, in *ppt*, demeaned. $Skew_t^*$ is the risk-neutral quarterly skewness of the 10-year rate. ΔDur is changes in the log of aggregate bond duration (from the Bloomberg Agg index) to GDP. Equity beta is rolling estimates of sensitivity of S&P 500 returns to 10-year rate changes. Point estimates use weekly data from 2002–2023, and standard errors are calculated using monthly data and Newey-West standard errors allowing for 12 lags.

	(1) Constant	(2) PCs only	(3) Main	(4) ΔAgg duration	(5) Equity β	(6): 4 & 5
Intercept	44.42***	35.05***	34.10*	43.46***	55.20***	57.93***
	(8.86)	(7.95)	(19.99)	(9.47)	(10.00)	(7.22)
PC 1		5.77***	5.50***			
		(1.22)	(1.12)			
PC 2		-5.75***	-5.80***			
		(2.17)	(2.01)			
PC 3		-2.08	-2.47			
		(6.06)	(6.71)			
σ_t^*			4.36			
			(19.42)			
$Skew_t^*$			-5.06			
			(7.20)			
ΔDur				32.34		67.12**
				(35.01)		(32.14)
Equity β					-1.72*	-2.49***
					(0.97)	(0.90)
Avg $\hat{\lambda}_t$	0.44***	0.35***	0.35***	0.45***	0.43***	0.43***
	(0.09)	(0.08)	(0.11)	(0.09)	(0.08)	(0.08)
Avg $\widehat{RP}_t$	0.12***	0.11***	0.11***	0.12***	0.12***	0.12***
	(0.02)	(0.02)	(0.03)	(0.02)	(0.02)	(0.02)
R2	0.41	0.55	0.55	0.42	0.44	0.46

Table 7: **Declines in risk-neutral variance during FOMC announcement windows**

This table reports the average proportional decline in risk-neutral variance during 3-day windows around FOMC meetings for different swap rate tenors (rows) and swaption maturities (columns). Maturity refers to the time of expiry of the swaption, tenor refers to the maturity of the underlying swap. Changes are calculated from close on day $t-1$ to close on day $t+1$ where day t is the FOMC announcement. Sample period is 2007–2023. Standard errors in parentheses are clustered by FOMC announcement date.

	Option Maturity		
Swap Tenor	Quarter	Year	5 Year
1 year	-2.5%	-4.1%	-1.6%
	(3.4%)	(1.3%)	(0.6%)
2 year	-4.3%	-3.4%	-1.8%
	(1.7%)	(1.3%)	(0.5%)
5 year	-3.2%	-2.8%	-1.5%
	(1.3%)	(0.8%)	(0.6%)
10 year	-2.6%	-2.3%	-1.4%
	(1.0%)	(0.7%)	(0.6%)
20 year	-2.6%	-2.5%	-1.4%
	(1.1%)	(0.7%)	(0.6%)
30 year	-2.9%	-2.9%	-1.2%
	(1.0%)	(0.8%)	(0.6%)

Table 8: **Stock-risk-premium correlation under different models**

The first column calculates the correlation of market returns with changes in 10y-in-1y risk premium calculated using the methodology described in Section 4. The second column uses the 1-year instantaneous out-of-sample term premium from Kim and Wright (2005). The horizon for changes and market returns is given by the rows. Market returns are the CRSP value-weighted return. Sample period is 2002–2018 for weekly and monthly data, 2008–2018 for daily data. Newey-West standard errors in parentheses.

Horizon	Risk-based	Kim and Wright (2005)
Daily	-0.08	0.30
	(0.03)	(0.03)
Weekly	-0.02	0.22
	(0.02)	(0.04)
Monthly	-0.19	0.18
	(0.06)	(0.05)
Note:		*p<0.1; **p<0.05; ***p<0.01

Table 9: **Effects of different sources of residual coskewness on estimated interest-rate risk premium**

This table summarizes the results from re-estimating the full-sample quarterly interest-rate risk premium, after accounting for different possible sources of residual coskewness, as described in Section 5. The first column shows the correlation with the main zero-residual-coskew results. The second column shows the average difference in quarterly interest-rate risk premium in ppt. The third column shows how much this difference implies the main results would over- or under-state the true risk premium by, in ppt. The rows allow for residual coskew from risk aversion, equity exposure, fixed income convexity, and three-factor interest rate exposures. Risk premium is estimated on the full sample, 2002-2023

	Correlation with main results	Avg difference (ppt)	Implied under- (over-) estimation
2nd order approx	0.976	0.044	29.1%
Equity exposure (CAPM)	1.000	-0.006	(5.5%)
Bond market convexity	0.999	-0.004	(3.6%)
CRRA $\gamma = 2$	0.995	0.060	34.5%
CRRA $\gamma = 4$	0.986	0.113	49.7%
Multiple interest-rate factors	0.962	-0.025	(27.3%)

Table 10: **Size of VRP implied by non-interest-rate risks**

The left column shows the average quarterly proportional variance risk premium, calculated as risk-neutral variance minus average realized variance, divided by average realized variance. The right three columns show the average variance risk premium that could be explained by different linear factor models of the SDF, after removing interest-rate risk. The methodology is described in Section 5. "FF3" and "FF5" denote the Fama & French 3 and 5 factor models, Corp denotes a factor consisting of corporate bond excess returns over treasuries, and "Mom" denotes the momentum factor. Standard errors are calculated from a 1-year block bootstrap.

Observed VRP	VRP implied by factor models			
	CAPM	FF 3	FF3 + Corp	FF 5 + Mom + Corp
16.7%	1.6%	1.8%	2.5%	1.7%
(4.3%)	(0.8%)	(0.9%)	(1.3%)	(1.6%)

Figures

Figure 3: **Conditional variance risk premium**

This figure plots the difference between risk-neutral variance and forecasted physical variance for 10-year swap rate at a quarterly horizon, January 2002 – June 2023. Positive values indicate that investors are willing to pay a premium for protection against large changes in interest rates. Units are percentage point squared.

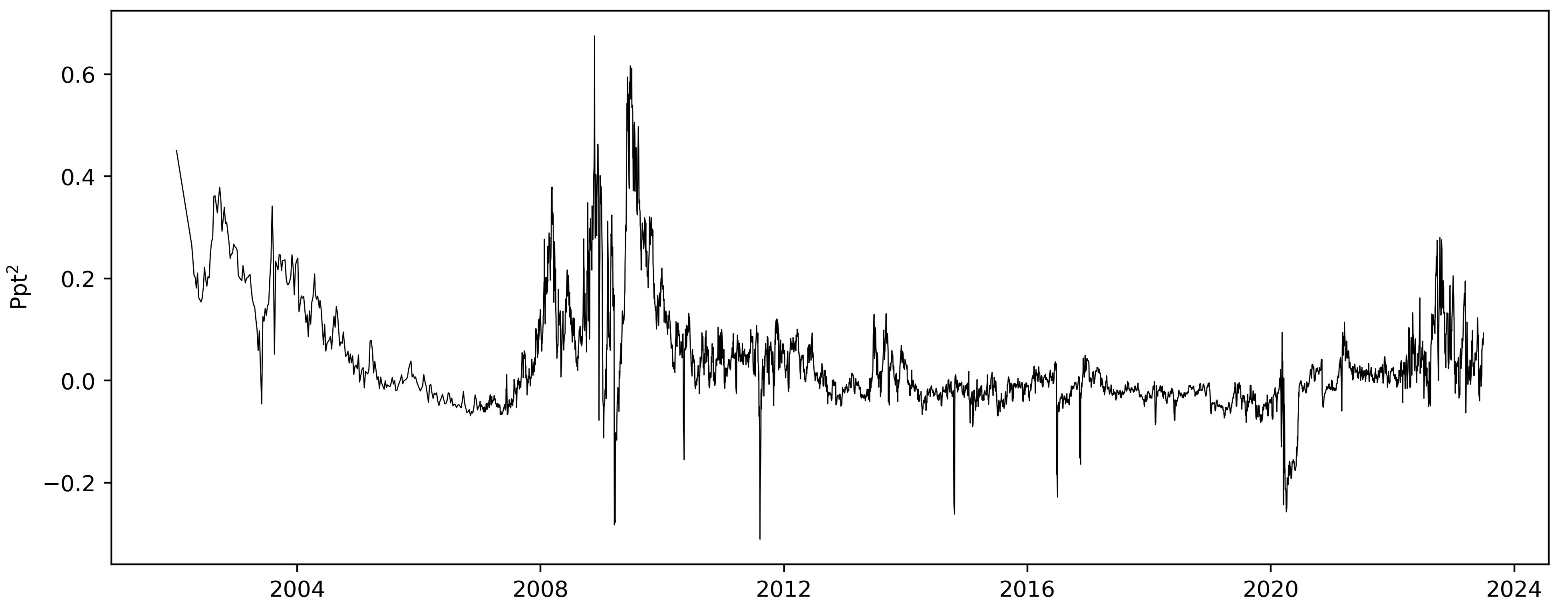

Figure 4: **Risk-based and term-structure-model risk premium estimates**

The blue line plots the risk-based out-of-sample 10y-in-1q interest-rate risk premium. The orange line plots the equivalent figure for the Adrian et al. (2013) dynamic term-structure model. Units are percentage points and represent the difference between 1-quarter-ahead forward rates and 1-quarter-ahead expected rates. Dates are 2002 to June 2023. Methodology is described in Section 3.

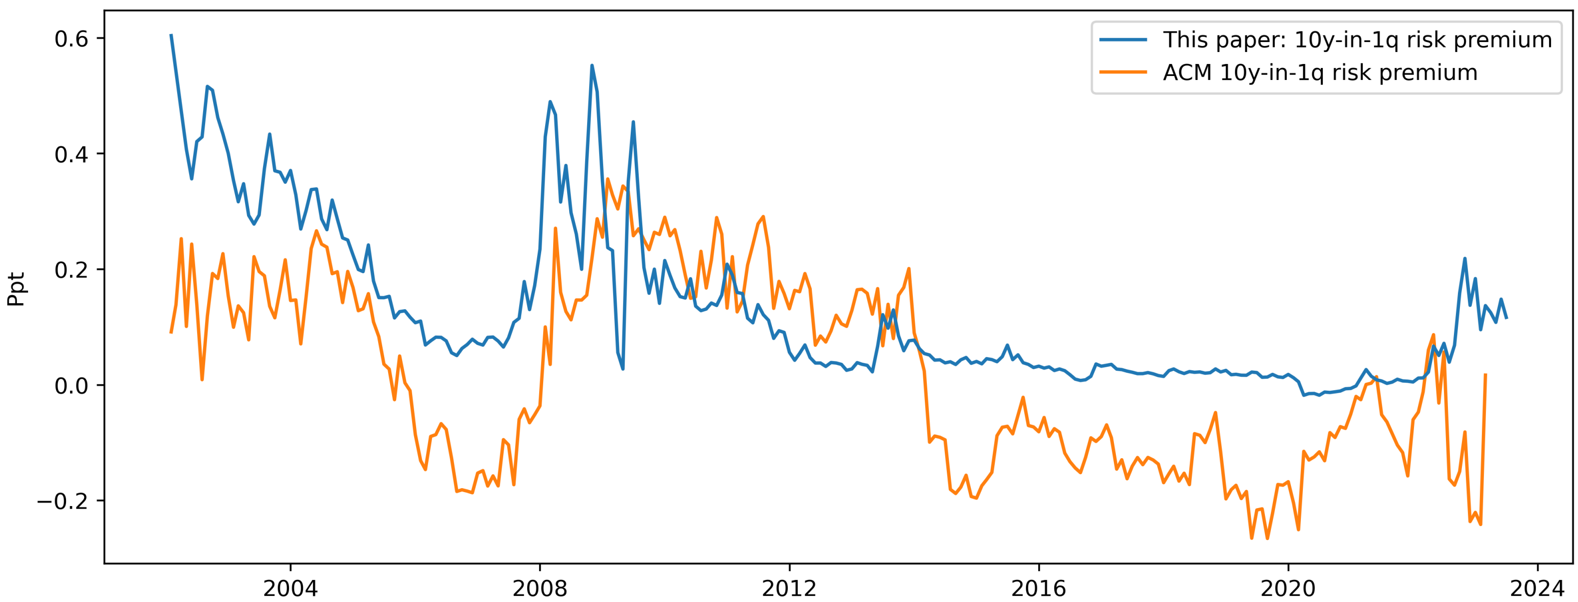

Figure 5: **Estimated investor exposure parameter,** λ_t

This figure plots the whole-sample estimates of exposure, λ_t, using the methodology described in Section 3. λ_t can be interpreted as the duration of a log-utility fixed income investor's portfolio or, approximately, as the $duration \times \gamma$ for a CRRA investor with risk aversion coefficient γ.

Figure 6: **Cumulative returns from interest-rate trading strategies**

This figure shows cumulative dollar returns from strategies that take positions in 1-month interest-rate swap forwards proportional to risk premia estimated from various models. Initial investment is normalized to \$1 in April 2002. Positions are scaled by $1/(2\sigma_t^{*2})$, in line with the implied portfolio of a CRRA investor with risk aversion coefficient $\gamma = 2$ who perceives the measured risk premium. Risk premium is estimated using this paper's risk-based measure, Survey of Professional Forecasters expectations, the Adrian et al. (2013) affine model, and a random walk assumption. Returns are calculated before transaction costs. The Survey of Professional Forecasters strategy uses quarterly forwards instead of monthly.

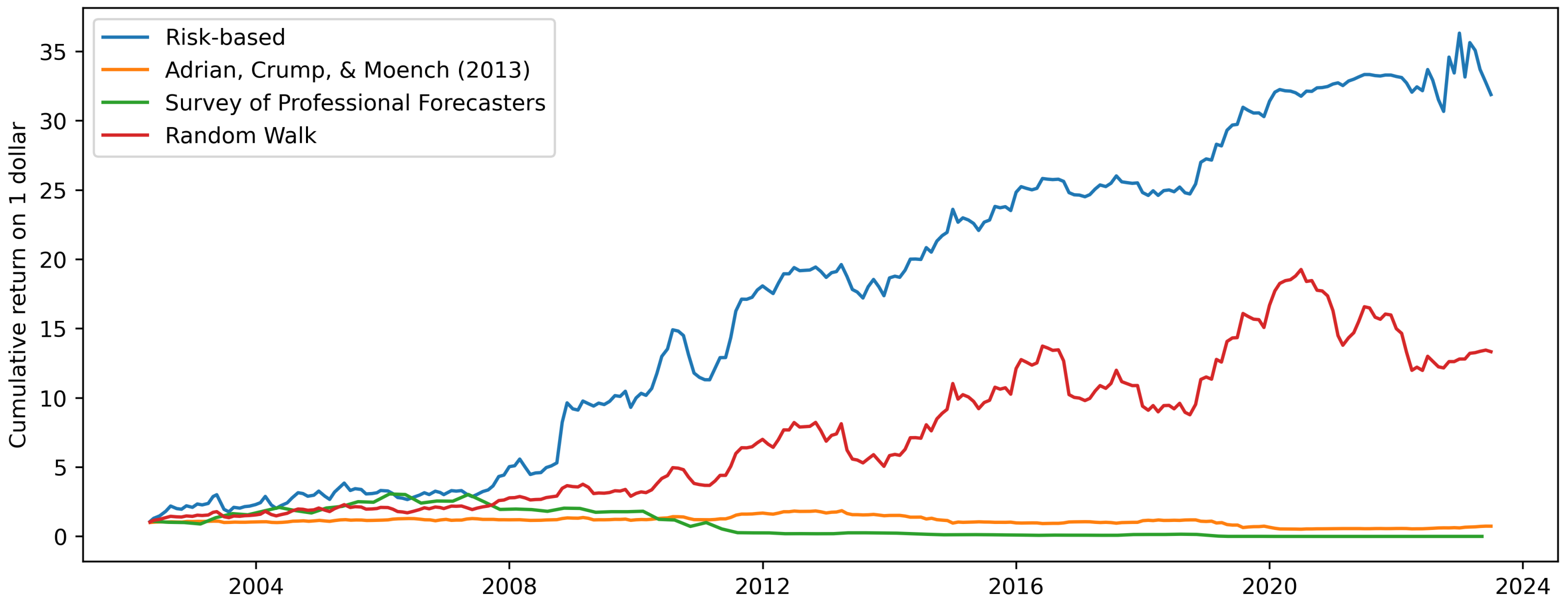

Figure 7: **Interest-rate risk premium after allowing for different sources of residual coskewness**

This figure plots the the results from re-estimating the full-sample conditional interest-rate risk premium, after accounting for different possible sources of residual coskewness, as described in Section 5. The different lines allow for residual coskew from risk aversion, equity exposure, fixed income convexity, and three-factor interest rate exposures.

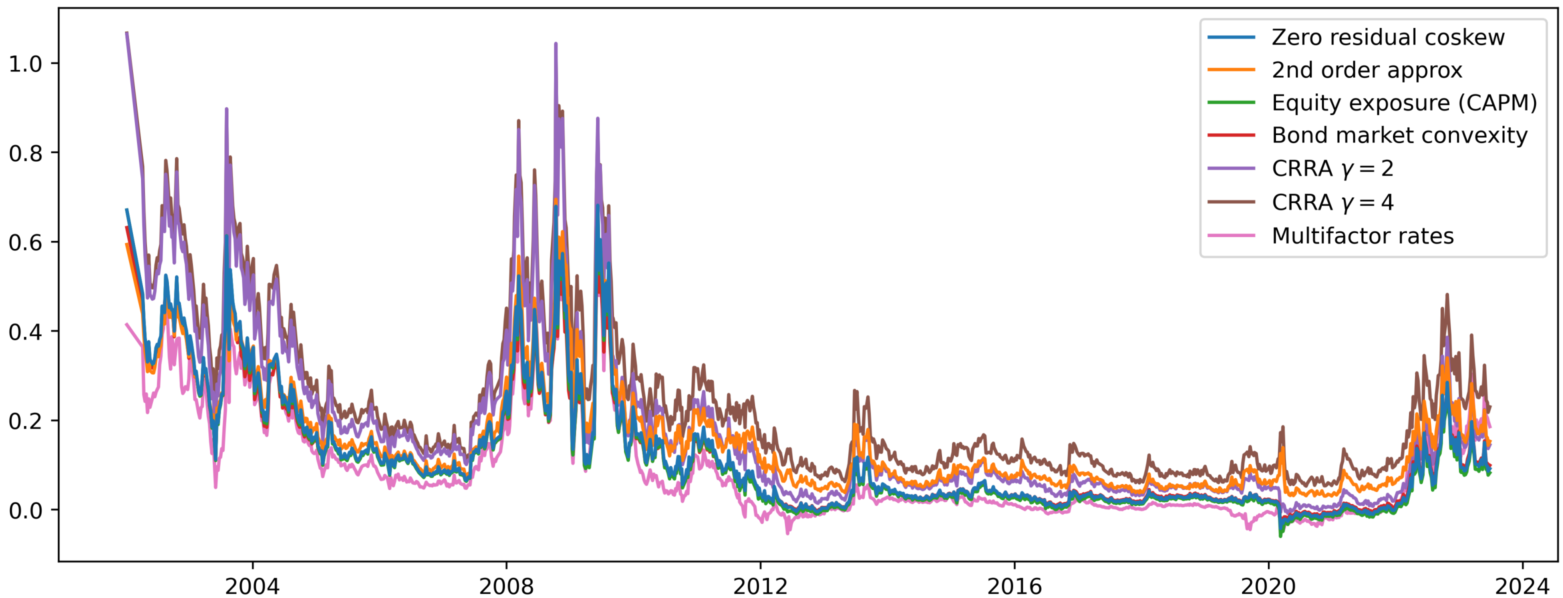

Appendices

A Forecasting other interest-rate tenors

Table 11 below reports the relative improvement in forecasting performance by applying this same methodology to the 1-year, 2-year, 5-year, 20-year, and 30-year swap yield. As a simple approximation I assume the value of λ_t for all tenors is the same as the values calculated in the main body of this paper.

Recalculating different values of λ_t for each tenor based on their respective variance risk premia can be done. However, there are some challenges in estimating physical variance for shorter tenors because the rates are at the zero lower bound for much of the sample and more subject to jumps around policy announcements.

Comparison model	Tenor	Monthly	Quarterly	Annual
Expectations Hypothesis	1	0.052	0.110	0.154
	2	0.063	0.131	0.185
	5	0.068	0.128	0.229
	10	0.065	0.102	0.230
	20	0.066	0.062	0.191
	30	0.072	0.042	0.174
Adrian et al. (2013)	1	0.126	0.228	0.153
	2	0.077	0.141	0.170
	5	0.060	0.096	0.162
	10	0.070	0.095	0.184

Table 11: Multi-tenor out-of-sample forecast performance. The table reports the improvement in R^2 from risk-based interest-rate expectation measure versus the expectations hypothesis and the Adrian et al. (2013) DTSM. Improvement in R^2 is defined as $1 - \sum \varepsilon_t^2 / \sum \nu_t^2$, where ε_t is the forecast error from the risk-based model and ν_t is the forecast error from the alternative. 2002-2023, monthly data.

B Interest-rate risk premium with non-parametric form for λ_t

Figure 8 compares interest-rate risk premium estimated from the parametric specification for λ_t, described in Section 3, with interest-rate risk premium estimated by solving for λ separately each period, without a functional form. Two roots are possible for lambda. I choose the root with the positive (or least negative) interest-rate risk premium. Confidence intervals are provided for each period for the non-parametric version. Where no value of λ_t was consistent with the measured variance risk premium, I take the closest value. If this implies a value for the variance risk premium that is outside of its 95% confidence interval of the variance-prediction regression, I plot this observation with a dotted line.

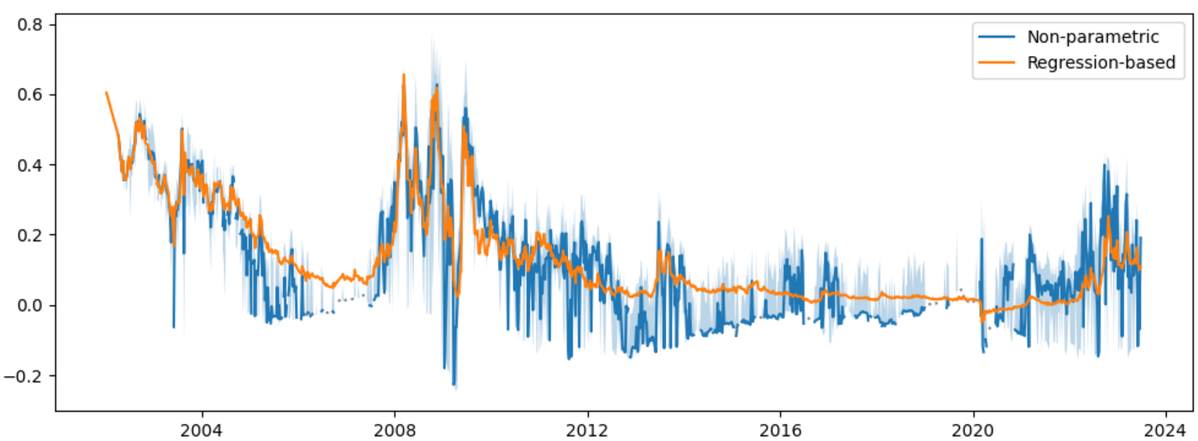

Figure 8: The orange line shows interest-rate risk premium estimated from the parametric specification for λ_t. The blue line shows interest-rate risk premium estimated by solving for λ separately each period. Quarterly forecasts in ppt, 2002-2023.

C Interest-rate risk premium with non-parametric form for σ_t^2

Proposition 1 tells us that the expectations hypothesis is equivalent to $\lambda_t = 0$ — i.e., there is only no risk-premium if there is no exposure. This section presents the results from a simple way to test this hypothesis by estimating the constant exposure parameter λ that best explains the realized variance risk premium.

The simplest way to estimate physical variance is to simply use realized variance measured over the subsequent period. Replacing the physical variance in Proposition 3 with the realized variance from daily data plus an error term and assuming zero residual coskew yields:

$$\sigma_t^{*2} - RV_{t\to t+1} = \lambda_t E_t^*(\Delta y_{t+1}^3) + \sigma^{*4}\lambda_t^2 + \eta_{t+1}$$

Now all terms in this equation are observable except for λ_t and the error term η_t. If the conditional variance is equal to the conditional expectation of realized variance, then this error term should be uncorrelated with the other terms in the equation. If we assume λ_t is constant then we can simply solve for the value of λ that minimizes the sum of squared errors $\sum \eta_t^2$ with no further assumptions.

I find $\hat{\lambda} = 41$, with a standard error of 6, using monthly observations and GMM Newey West standard errors with 3 lags. This amounts to a very strong rejection of the expectations hypothesis with a t-statistic of over 5. The statistical significance is high enough that a simple difference in priors could not account for the apparent risk premium, as suggested by Farmer et al. (2024) for the level-based forecasts.

Multiplying this exposure by the risk-neutral variance yields an average interest-rate risk premium of 11 bp, or 44 bp annualized. This is similar to the 34 bp realized risk premium. Markets seem to have been demanding compensation for bearing interest-rate risk, rather than making systematic mistakes.

This exposure estimate corresponds to a duration of 41 years for a log-utility fixed-income investor, or approximately $41/\gamma$ for a CRRA investor. To check how precise this approximation is, in Appendix F I re-estimate the parameters allowing for the nonlinearity in the SDF resulting from higher risk aversion and find that a γ of 4, gives a duration of approximately 16. This high sensitivity to interest-rate changes could be consistent with models where leveraged intermediaries are the marginal investors in fixed-income markets. For example, Kekre et al. (2024) suggest durations of 10 to 30 years for fixed-income arbitrageurs.

D Risk-neutral moment time series

Figure 9 plots the time series of risk-neutral variance and skewness 2002-2023.

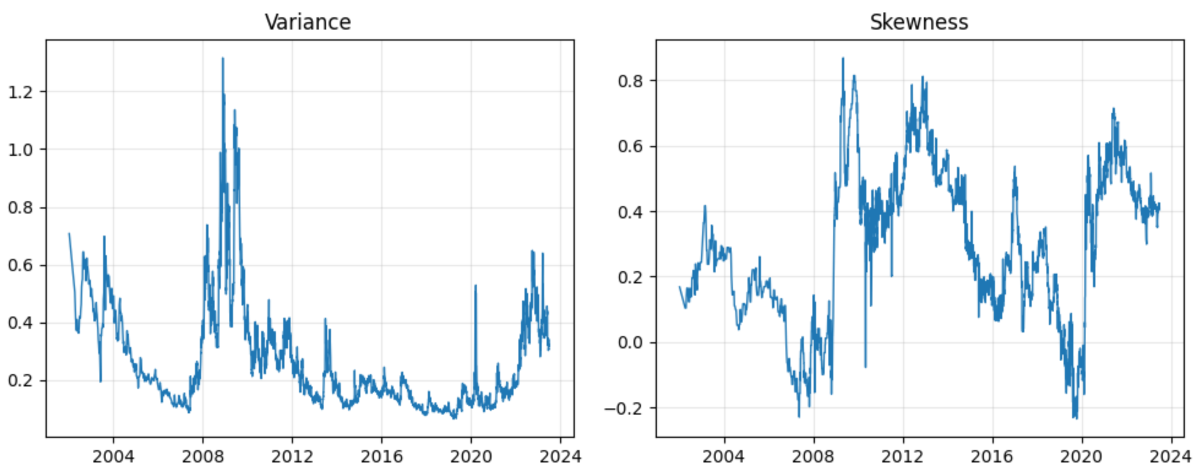

Figure 9: Quarterly risk-neutral variance (left) and skewness (right) of 10-year swap rates, 2002–2023. Variance and skewness are extracted from swaption prices, using the model free approach approach of Carr and Madan (1998). Variance is expressed in percentage points squared.

E Duration and exposure for a CRRA fixed-income investor

CRRA utility with coefficient γ implies:

$$\frac{1}{M} = R_f \frac{(1 - D\Delta y)^{\gamma}}{E^*((1 - D\Delta y)^{\gamma})}$$

Where I will ignoring the time subscripts throughout this section.

D is here the "duration" of the CRRA investor's portfolio. Assume that γ is a positive integer. Since λ is the projection coefficient of the inverse SDF on yields, by binomial expansion we have:

$$\lambda = -\frac{1}{R_f} \frac{cov^*(\frac{1}{M}, \Delta y)/\sigma^{*2}}{E^*((1 - D\Delta y)^{\gamma})} = -\frac{\sum_{k=1}^{\gamma} \binom{\gamma}{k}(-D)^k \mu_{k+1}^*/\sigma^{*2}}{1 + \sum_{k=2}^{\gamma} \binom{\gamma}{k}(-D)^k \mu_k^*}$$

Where μ_k^* is the kth risk-neutral central moment of y.

The first term in the expansion of the numerator is $\frac{\gamma D}{E^*((1-D\Delta y)^{\gamma})} \approx \gamma D$. To demonstrate that the other terms are small, Table 12 estimates the duration implied by a λ of 40 (or 0.4 if using ppt^2 units for variance) under different values of γ using the time series average levels of the risk-free rate and the first four quarterly risk-neutral moments after 2011. where I do not have risk-neutral moment data available (e.g. for the fifth moment), I use the values implied by a normal distribution. The approximation $D \approx \frac{\lambda}{\gamma}$ is never off by more than 2%. Repeating the exercise for every individual day from 2011 – 2023 gives a highest absolute approximation error of 5% for $\gamma = 4$.

F Additional content on the residual coskew

In this Appendix I report how I estimate the effects of residual coskew on estimated interest-rate risk premium after allowing for higher risk aversion, bond market convexity, and

CRRA Coefficient	Implied Duration	Approximation error
1	40.0	0.0%
2	20.1	0.7%
3	13.4	0.3%
4	10.0	0.0%

Table 12: Accuracy of the Duration = $\frac{\lambda}{\gamma}$ approximation. Each row provides the duration that would deliver $\lambda = 0.4$ for a CRRA investor with various levels of risk aversion and with linear exposure to the interest rate y_t. The first 4 quarterly risk-neutral moments are taken as their time series average 2011-2023, and the 5th moment is assumed to be 0. Approximation error shows the percentage difference between the calculated duration and the simple approximation $D \approx \frac{\lambda}{\gamma}$

multiple interest rate factors.

F.1 Effects of risk aversion greater than 1

F.1.1 Re-estimating risk premium with $\gamma = 2$

Higher relative risk aversion will lead the main methodology of this paper to underestimate the true risk aversion (proof at end of this subsection). We can calculate the size of this underestimation using the higher risk-neutral moments of interest rates. For example, if we consider the case of a CRRA investor with risk aversion coefficient $\gamma = 2$ and a linear exposure of duration D_t to interest rates, then the inverse SDF will be given by:

$$\frac{1}{M_{t+1}} = R_{f,t}\frac{(1 - D_t\Delta y_{t+1})^2}{1 + D_t^2\sigma_t^{*2}}$$

where the denominator ensures that $E_t^*(\frac{1}{M_{t+1}}) = R_{f,t+1}$.

Following some simple algebra the residual coskew becomes:

$$cov_t^*(\varepsilon, \Delta y_{t+1}^2) = D_t^2\frac{(\kappa_t^* - 1)\sigma_t^{*4} - E_t^*(\Delta y_t^3)^2/\sigma_t^{*2}}{1 + D_t^2\sigma_t^{*2}} \tag{10}$$

where κ_t^* is the conditional risk-neutral kurtosis.

The average quarterly risk-neutral kurtosis of interest rates 2011–2022 (when data is available) is 4.5. If duration is approximately 20 (as implied by $\lambda = 40$), other unconditional averages of risk-neutral moments from Table 1 suggests an unconditional average residual coskew of approximately 0.01 ppt^2. Since the variance risk premium is 0.05, this suggests a roughly 20% underestimation of interest-rate risk premium.

To test this scaling more thoroughly, I redo the same estimation described in Sections 2 and 3 on the full sample, but allowing for CRRA utility with $\gamma = 2$. Specifically, I follow the same steps described in the main body of the text, solving for the time varying D_t parameters that minimize the sum of the squared errors η_t from the equation:

$$VRP_t = \lambda_t(D_t)E_t^*(\Delta y_{t+1}^3) + \lambda_t(D_t)^2\sigma_t^{*4} + ResidCoskew_t(D_t) + \eta_t$$

where:

$$D_t = D_0 + D_1 f_{1,t} + D_2 f_{2,t} + D_3 f_{3,t} + D_4 \sigma_t^{*2} + D_5 skew_t^*$$

and $ResidCoskew_t(D_t)$ is given by Equation (10), and exposure as a function of duration, $\lambda_t(D_t)$ follows the relationship given in Appendix E (approximately $\lambda_t(D_t) \approx \gamma D_t$).

Since I only measure the fourth moment after 2011, I assume that the pre-2011 kurtosis is constant at its average post 2011 level of 4.5. This is likely to be an overestimate because of the very high kurtosis in 2020–2021 when variance was low, but with fat tails.

The resulting time series of interest-rate risk premium is 99% correlated with the main forecasts of this paper, as shown in Table 9. However, the λ and risk premium are higher. If the true model uses $\gamma = 2$, then the main estimates of this paper would underestimate exposure by 18, or about 30%, on average.

F.1.2 Allowing for $\gamma > 2$

More generally, if the investor has CRRA utility with a, then we can write the inverse SDF as:

$$\frac{1}{M} = R_f \frac{(1 - D\Delta y)^\gamma}{E^*((1 - D\Delta y)^\gamma)}$$

where I will ignoring the time subscripts throughout in this subsection. D is here the "duration" of the CRRA investor's portfolio.

With a positive integer risk aversion coefficient γ, applying the binomial expansion formula we can write the residual coskew as:

$$
\begin{aligned}
cov^*(\varepsilon, \Delta y^2) &= cov^*\left(\frac{1}{R_f M} + \lambda \Delta y, \Delta y^2\right) \\
&= \frac{\sum_{k=2}^{\gamma} \binom{\gamma}{k} (-D)^k \mu_{k+2}^* - \mu_k^* \sigma^{*2} - \mu_{k+1}^* E_t^*(\Delta y_{t+1}^3)/\sigma^{*2}}{1 + \sum_{k=2}^{\gamma} \binom{\gamma}{k} (-D)^k \mu_k^*}
\end{aligned}
$$

In general, this term will be positive and lead us to underestimate the exposure and risk premium of a CRRA investor if we ignore it. A brief proof of this statement is provided at the end of this section.

Calculating the size of this term for values of $\gamma > 2$ requires fifth and higher moments, which gather information from far out-of-the-money quotes that may not be easily executable. This is particularly problematic during the pre-2011 sample where I must extrapolate about the higher moments based on observed data on variances and skewnesses, as described in Appendix I.1. Nonetheless, to conduct a simple check of whether the residual coskew is likely to become much larger as risk aversion increases, I re-estimate the average value of risk premium as described the previous subsection for $\gamma = 2$. The results are shown below in Table 13. The level of bias introduced by the residual coskew increases only modestly beyond $\gamma = 2$.

F.1.3 Proof that a high risk aversion investor's SDF exhibits high residual coskew

Proof. We need to show that $cov_t^*(\varepsilon_{t+1}, (\Delta y_{t+1})^2) > 0$ where $\varepsilon_{t+1} = \frac{1}{M_{t+1}} - \lambda_t \Delta y_{t+1}$, with $\frac{1}{M_{t+1}} = (R_{f,t} + D\Delta y_{t+1})^\gamma$ for $\gamma > 1$ and $E_t^*[\Delta y_{t+1}] = 0$.

	Correlation with main results	Avg difference (ppt)	Implied under- (over-) estimation
CRRA $\gamma = 2$	0.994	0.061	34.0%
CRRA $\gamma = 3$	0.989	0.092	44.0%
CRRA $\gamma = 4$	0.982	0.109	48.2%

Table 13: Effects of different levels of CRRA risk aversion on estimated interest-rate risk premium. This table summarizes the results from re-estimating the full-sample quarterly interest-rate risk premium, after accounting for different possible sources of residual coskewness, as described in Section 5. The first column shows the correlation with the main zero-residual-coskew results. The second column shows the average difference in quarterly interest-rate risk premium in ppt. The third column shows how much this difference implies the main results would over- or under-state the true risk premium by, in ppt. Risk premium is estimated on the full sample, 2002-2023

First, expand $\frac{1}{M_{t+1}}$ around $\Delta y_{t+1} = 0$:

$$\begin{aligned}\frac{1}{M_{t+1}} =& R_{f,t} + \gamma R_{f,t}^{\gamma-1} D\Delta y_{t+1} + \frac{\gamma(\gamma-1)}{2} R_{f,t}^{\gamma-2}(D\Delta y_{t+1})^2 + \\ & \frac{\gamma(\gamma-1)(\gamma-2)}{6} R_{f,t}^{\gamma-2}(D\Delta y_{t+1})^3 + O((\Delta y_{t+1})^4)\end{aligned}$$

Since $E_t^*[\Delta y_{t+1}] = 0$ and $\lambda_t = \frac{cov_t^*\left(\frac{1}{M_{t+1}}, \Delta y_{t+1}\right)}{\sigma_t^{*2}}$, we have:

$$\lambda_t = \gamma R_{f,t}^{\gamma-1} D + \frac{\gamma(\gamma-1)}{2} R_{f,t}^{\gamma-2} D^2 \frac{E_t^*[\Delta y_{t+1}^3]}{\sigma_t^{*2}} + O(E_t^*[\Delta y_{t+1}^2])$$

The covariance of interest is:

$$cov_t^*\left(\varepsilon_{t+1}, \Delta y_{t+1}^2\right) = cov_t^*\left(\frac{1}{M_{t+1}}, \Delta y_{t+1}^2\right) - \lambda_t E_t^*[\Delta y_{t+1}^3]$$

Computing the first term using the Taylor expansion:

$$cov_t^*\left(\frac{1}{M_{t+1}}, \Delta y_{t+1}^2\right) = \gamma R_{f,t}^{\gamma-1} D E_t^*[\Delta y_{t+1}^3] + \frac{\gamma(\gamma-1)}{2} R_{f,t}^{\gamma-2} D^2 var_t^*[\Delta y_{t+1}^2] + O(E_t^*[\Delta y_{t+1}^4])$$

Substituting the expression for λ_t and simplifying, the leading-order term becomes:

$$cov_t^*(\varepsilon_{t+1}, (\Delta y_{t+1})^2) = \frac{\gamma(\gamma-1)}{2} R_{f,t}^{\gamma-2} D^2 \left(var_t^*[\Delta y_{t+1}^2] - \frac{(E_t^*[\Delta y_{t+1}^3])^2}{\sigma_t^{*2}} \right) + O(E_t^*[\Delta y_{t+1}^4]) \tag{11}$$

By the Cauchy-Schwarz inequality:

$$(E_t^*[\Delta y_{t+1}^3])^2 \leq \sigma_t^{*2} \cdot E_t^*[\Delta y_{t+1}^4] \tag{12}$$

Therefore, $var_t^*[\Delta y_{t+1}^2] - \frac{(E_t^*[\Delta y_{t+1}^3])^2}{\sigma_t^{*2}} > 0$ unless Δy_{t+1} is degenerate. Since $\gamma > 1$ implies $\gamma(\gamma-1) > 0$, we conclude that $cov_t^*(\varepsilon_{t+1}, \Delta y_{t+1}^2) > 0$. □

F.2 Effects of bond market convexity

Bonds with embedded options, particularly mortgage-backed securities, can exhibit negative convexity. When rates fall, prepayments accelerate and duration shrinks; when rates rise, prepayments slow and duration extends. Could this convexity explain the variance risk premium?

For a log investor holding a portfolio with duration D and convexity C, the inverse SDF becomes:

$$\frac{1}{M_{t+1}} = R_{f,t} - D\Delta y_{t+1} + \frac{1}{2}C(\Delta y_{t+1}^2 - \sigma_t^{*2})$$

The residual covariance term equals:

$$E_t^*(\Delta y_{t+1}^2 \varepsilon_{t+1}) = \frac{1}{2}C \times \left(\mathrm{var}_t^*(\Delta y_{t+1}^2) - \frac{E_t^*(\Delta y_{t+1}^3)^2}{\sigma_t^{*2}} \right)$$

Using the relationship $\mathrm{var}_t^*(\Delta y_{t+1}^2) = (\mathrm{kurt}_t^*(\Delta y_{t+1}) - 1)\sigma_t^{*4}$, and noting that risk-neutral kurtosis averages 4.5 from 2011-2023 while quarterly risk-neutral variance squared averages

0.1, the convexity needed to generate the observed variance risk premium would be:

$$C = 2\frac{-E(\text{VRP})}{E\left((\text{kurt}^*(\Delta y) - 1)\sigma^{*4} - -\frac{E_t^*(\Delta y_{t+1}^3)^2}{\sigma_t^{*2}}\right)} \approx 2\frac{-0.04}{3.5 \times 0.1 - 0.02} = -0.24$$

To put this in perspective, the Bloomberg Barclays Aggregate Bond Index recorded its most negative convexity at -0.005 in 2004 and most positive convexity at 0.006 (after converting to our percentage point units). A log investor would therefore need negative convexity equivalent to a 48-times levered position in the aggregate bond portfolio at its moment of peak negative convexity. This is implausible for a reasonably representative investor.

F.3 Multiple interest rate factors

I will assume that interest rates can be exactly decomposed into three principal components $\{f_{1,t}, f_{2,t}, f_{3,t}\}$ with constant loadings of each interest rate on these factors. This is a relatively innocuous approximation: the first three principal components explain 99.9% of quarterly variance in swap rates. If we allow an investor to have linear exposure to all swap rates, we can then write the inverse SDF as a loading on the changes in each factor:

$$\frac{1}{M_{t+1}} = R_{f,t}(1 - \lambda_t' \Delta f_{t+1})$$

where λ_t is now a size 3 vector of exposures to the 3 factors.

The same steps described in Section 1 then allows us to then write a matrix expression for the multi-factor variance risk premium:

$$\Sigma_t^* - \Sigma_t = \lambda_t' S_t^* + \Sigma_t^* \lambda_t \lambda_t' \Sigma_t^*$$

where Σ_t^* and Σ_t are the risk-neutral and physical factor covariance, and S_t^* is the $3 \times 3 \times 3$ risk-neutral "cube" of factor third moments, and hence $\lambda_t' S_t^*$ is a 3×3 matrix.

For any specific tenor of interest rates with factor loading ϕ_i such that $\Delta y_{i,t+1} = \phi_i \Delta f_{t+1}$, its variance risk premium can be easily calculated from this multi-factor variance risk premium expression as:

$$\sigma^{*2}_{i,t} - \sigma^2_{i,t} = \phi_i' \left(\Sigma^*_t - \Sigma_t\right) \phi_i$$

If we can measure the risk-neutral moments, we can therefore pursue the exact same estimation strategy in as in the main body of this paper. I parameterize λ_t as consisting of 3 constants λ_0 and a 3×3 matrix of loadings Λ of each exposure term on each principal component of rates:

$$\lambda_t = \lambda_0 + \Lambda_1 f_t$$

I then solve for the values of λ_0 and Λ_1 that minimize the sum of squared errors $\sum_{i=1}^{3} \sum_{t=0}^{T} \eta^2_{i,t}$ from predicting the variance risk premium of each tenor of interest rate:

$$\sigma^{*2}_{i,t} - \sigma^2_{i,t} = \phi_i' \left(\lambda_t' S^*_t + \Sigma^*_t \lambda_t \lambda_t' \Sigma^*_t\right) \phi_i + \eta_{i,t}$$

where physical variance is estimated from regressions n each interest rate tenor.

Swaptions data allows us to measure the risk-neutral moments of multiple interest rates at the same time. If we assume a constant factor loading, this allows us to measure the factor co-moments. In particular, I calculate for each period the risk-neutral covariance matrix Σ^*_t by finding the value such that for each tenor with factor loading ϕ_i

$$\sigma^*_{i,t} = \phi_i' \Sigma^*_t \phi_i$$

Since I observe six different tenors, and there are six independent elements of the covariance matrix, the matrix exactly identified in each period.

I follow the same procedure for the third moment cube. The level of each tenor's skewness is a function of the third moment cube and the factor loadings. The coskewness cube has 10 independent elements and so it is not exactly identified in each period. I choose the value that minimises the sum of squared coskewnesses. An alternative approach using a Kalman filter for identification assuming that coskewness change smoothly, yields similar results.

G Testing for information effects in risk-neutral variance

While the magnitudes suggest risk premium changes, I conduct three tests to evaluate whether information effects could explain the patterns.

G.1 Mechanical variance effects

FOMC meetings are high-variance events. Once a meeting passes, forward-looking variance mechanically declines by removing this event from the forecast window. Could this explain the observed patterns?

The mechanical effect is far too small. Three-day FOMC-window variance for daily 10-year rate changes is 37% higher than non-FOMC periods. Passing through an FOMC window should therefore reduce:

- Quarterly variance by $0.37 \times \frac{3}{61} = 1.8\%$
- Annual variance by $0.37 \times \frac{3}{250} = 0.4\%$
- Five-year variance by less than 0.1%

I observe declines of 2.9% for annual variance—more than seven times the mechanical effect.

Moreover, the cross-sectional pattern contradicts the mechanical story. FOMC-day variance is twice as high for 1-year rates as for 10-year rates, yet the proportional variance declines are roughly similar across tenors. The risk-neutral variance of rates five years forward falls substantially, which mechanical effects cannot explain.

G.2 Size of changes

For forward rates, the FOMC-window decline equals the total decline, consistent with gradual learning. But for variance, the cumulative FOMC effect dwarfs the total change.

For five-year risk-neutral variance on 10-year rates:

- Cumulative FOMC-window decline: -12.2 percentage points squared
- Net change 2007-2023: $+1.2$ percentage points squared
- Implied between-meeting increase: $+13.4$ percentage points squared

If markets were learning about future volatility, they would need to systematically "over-learn" during FOMC meetings that volatility will be low, then receive opposite information between every meeting. This asymmetric updating pattern seems implausible for rational learning but fits naturally with temporary risk premium compression.

G.3 Forecasting ability

If FOMC meetings reveal information about future volatility, then FOMC-window changes in risk-neutral variance should predict realized variance. I test:

$$\sigma^2_{RV,t} - \hat{\sigma}^2_{t-1} = \beta \Delta \sigma^{*2,\mathrm{FOMC}}_t + \varepsilon_t \tag{13}$$

where $\sigma^2_{RV,t}$ is subsequently realized variance and $\hat{\sigma}^2_{t-1}$ is the pre-FOMC forecast from Section 3.

Table 14 shows the results. The R^2 is below 0.01 for both quarterly and annual horizons. The coefficient is insignificantly different from zero and, for annual horizons, significantly different from one. FOMC-window variance changes contain essentially no information about future realized variance.

	Quarterly	Yearly
const	0.013	0.070
	(0.021)	(0.124)
$\Delta\sigma^{*2}$	0.232	0.239
	(0.617)	(0.311)
R-squared	0.006	0.003
N	137	137

Table 14: Information content of FOMC variance changes. The table reports regressions of realized variance forecast errors on FOMC-window changes in risk-neutral variance: $\sigma^2_{RV,t} - \hat{\sigma}^2_{t-1} = \alpha + \beta\Delta\sigma^{*2}_{FOMC,t} + \varepsilon_t$. If FOMC variance changes reflect information about future volatility, β should equal one. $\sigma^2_{RV,t}$ is subsequently realized variance, $\hat{\sigma}^2_{t-1}$ is the pre-FOMC regression forecast, and $\Delta\sigma^{*2}_{FOMC,t}$ is the change in risk-neutral variance during the FOMC window. Standard errors calculated using Newey-West with 2 lags (quarterly) and 8 lags (annual). Sample includes 137 FOMC meetings from 2007–2023.

H Constructing swap forward rates

Daily 1–10 year LIBOR swap rates data is taken from LSEG Workplace (formerly Refinitiv), with an example ticker being `USDSB3L10Y` for the 10-year rate. I construct 1-quarter and 1-year swap forward rates using the swap rates from 1 to 30 years and the risk-free yield curve calibrations from Gürkaynak, Sack, and Wright (2007).

I use the standard swap-rate formula as a function of risk-free rates and LIBOR forwards. I solve for the series of quarterly spreads between LIBOR forward rates and Treasury forward rates that successfully match the observed swap yields. This procedure assumes that the LIBOR forward spread is piecewise constant between the available swap tenors. For instance, given swap data for 1 and 2-year maturities, I assume a constant LIBOR spread between 3-months and 1-year, and a (different) constant spread between 1-year and 2-years. Once I have the full series of LIBOR forward rates and risk-free forwards, I construct the 3-month and 1-year swap forwards using standard formulas.

The Refinitiv swap data is available from 2002. To construct data prior to 2002, I use the following methods:

- For swap variance data (used to construct variance forecasts), I use the realized variance of daily par yields calculated from the Gurkayak, Sack, and Wright dataset. To this, I add a constant equal to the average difference between the LIBOR swap rate realized variance and the Treasury par yield realized variance over the 2002–2023 period.

- For swap rate levels (used only for illustrations, as in Figure 1), I take the Treasury par yield and add the average 2002–2023 LIBOR spread.

I Constructing Risk-Neutral Moments

I.1 Data

I use swaption implied volatility data from the Bloomberg "volatility cube." Quotes are provided as normal (Bachelier) implied volatilities at different tenors, maturities, and strike prices. Quotes are provided by strike relative to the at-the-money (ATM) rate, with 11 intervals available from ATM-200bp to ATM+200bp. An example ticker is `USSRAC10` for the ATM - 25bp, 3m×10y swaption.

To calculate the risk-neutral moments, I first apply the Bachelier option pricing formula to convert the provided implied volatilities into forward swaption prices. I interpolate option prices between the provided strikes by linearly interpolating the implied volatility. For strikes more than 200bp out of the money, I assume a constant implied volatility. Finally, I numerically integrate across the swaption prices to calculate the risk-neutral moments, using the formula provided in Appendix I.2.

Out-of-the-money (OTM) quotations are only available from 2011. To extend the risk-neutral moment data before this date, I use swaption-implied risk-neutral moment data calculated by Trolle and Schwartz (2014) from December 2001 to January 2010, augmented with other available data. Specifically, the data I use is:

- **Risk-Neutral Mean Proxy:** Swap forward rates.

- **Risk-Neutral Variance Proxy:**
  - 10y-in-1y swaption risk-neutral variance calculated by Trolle and Schwartz (2014) and shown in Figure 1
  - At-the-money (ATM) option implied variance (i.e., the Bloomberg implied volatility squared)
- **Risk-Neutral Skewness Proxies:**
  - 10y-in-1y swaption risk-neutral skewness from Trolle and Schwartz (2014)
  - The average spread between 10y-in-1y and 10y-in-1q risk-neutral skewness from Trolle and Schwartz (2014) Table 2
  - 10y treasury risk-neutral skewness from Bauer and Chernov (2024)

For the risk-neutral mean, I simply backfill the means using the coefficients from a regression of the 2011–2023 risk-neutral mean on the forward yield. This regression has a $> 99.9\%$ R^2.

For the risk-neutral variance, I extract the Trolle and Schwartz (2014) risk-neutral variance data by digitizing their Figure 1. Their paper calculates risk-neutral variance of swap rates under the annuity measure. This is on average around 1.5% lower than the risk-neutral variance I calculate using the methodology described in the next section. To estimate the true risk-neutral variance before 2011 I use the coefficients from a post-2011 regression of risk-neutral variance on ATM variance and the spread between the 10y-in-1y ATM variance and 10y-in-1y annuity-measure variance. To avoid any issues with date alignment of the digitized plot, I use 6-month rolling averages of the estimated spreads between risk-neutral variance and ATM variance. This spread is slow-moving during the periods where we observe both. This regression has a 99.7% R^2 for the 10y-in-1q risk-neutral variance. For the window between when the Trolle and Schwartz (2014) data ends and the Bloomberg OTM data starts, I simply use a regression of ATM variance on risk-neutral variance (R^2 of 99.6%).

For the risk-neutral skewness, I again use the Trolle and Schwartz (2014) data digitized from Figure 1. This figure represents skews calculated under the annuity measure which are approximately 10% lower on average than those using the risk-neutral measure approximation described in the next section. I therefore backfill the 10y-in-1y risk-neutral skewness using the coefficients from a regression of 10y-in-1y risk-neutral skewness on annuity-measure skewness. This regression has an R^2 of 99.8%. For the window between when the Trolle and Schwartz (2014) data ends and the Bloomberg OTM data starts I use the Bauer and Chernov (2024) treasury risk-neutral skewness in its place.

For the 10y-in-1q skewness, I simply subtract the average spread between the 10y-in-1q and 10y-in-1y skewness reported by Trolle and Schwartz (2014). During the period where I have the full dataset, a constant spread from the 10y-in-1y skewness describes the 10y-in-1q skewness with an R^2 (uncentered) of 90%, and the third moment with 94%.

I.2 Annuity approximation

Swaptions are not precisely options on the interest rate. I therefore need to employ a simple approximation. I assume that changes in the annuity yield are the same as changes in the swap yield. This approximation is unlikely to lead to substantive errors — changes in the 10-year annuity yield and swap yield are >99% correlated at the quarterly or annual horizon.

Swaps and swaption prices

By a standard result, we can write the value of a swap agreed at fixed rate K with tenor T as $(y - K)A$ where A is the price of the T-period annuity, and y is the swap rate (i.e., the rate that sets the value of the swap to 0).

The payoff of a 1 period swaption agreed at rate K is therefore

- Pay-fixed: $Max\left\{(y_{t+1} - K)\, A_{t+1}, 0\right\}$
- Receive-fixed: $Max\left\{(K - y_{t+1})\, A_{t+1}, 0\right\}$

And the forward price of the swaption can be written:

- Pay-fixed: $C\left(k\right)=E_t^*(Max\left\{\left(y_{t+1}-K\right)A_{t+1},0\right\})$
- Receive-fixed: $P\left(k\right)=E_t^*(Max\left\{\left(K-y_{t+1}\right)A_{t+1},0\right\})$

By standard results, we can change measure from the risk-neutral measure to the "forward annuity measure," defined such that for a random variable X, $E_t^*(A_{t+1}X_{t+1})=E_t^*(A_{t+1})E_t^A(X_{t+1})$, where $E_t^*(A_{t+1})$ is the observable forward annuity price. The forward swaption price divided by the forward annuity price is a linear option under the swap rate under this measure:

- Pay-fixed: $\frac{C(k)}{E_t^*(A_{t+1})}=E_t^A(Max\left\{\left(y_{t+1}-K\right),0\right\})$
- Receive-fixed: $\frac{P(k)}{E_t^*(A_{t+1})}=E_t^A(Max\left\{\left(K-y_{t+1}\right),0\right\})$

Applying Breeden & Litzenberger

We can apply Breeden and Litzenberger (1978) to write the expectation of any function of the swap rate in terms of the prices of linear options on the swap rate. Since the options are linear under the annuity measure, we can write the expectation under the annuity measure of any function of the swap rate as:

$$E^A\left(g\left(y\right)\right)=\left(g(E_t^A(y_{t+1}))+\int_{-\infty}^{E_t^A(y_{t+1})}g''\left(k\right)\frac{P\left(k\right)}{E_t^*(A_{t+1})}dk+\int_{E_t^A(y_{t+1})}^{\infty}g''\left(k\right)\frac{C\left(k\right)}{E_t^*(A_{t+1})}dk\right)$$

Where $E_t^A(y_{t+1})$ is the observable swap forward rate.

However, this paper uses the moments under the risk-neutral measure, not the annuity measure. So I will assume that changes in the annuity yield vs its forward yield are always identical to changes in the swap yield vs its forward yield. The annuity price then becomes a function of the swap yield:

$$A_{t+1}=\sum_{j=1}^{4T}\frac{1}{(1+y_{t+1+j})^{t/4}}=A(y_{t+1})$$

Now we can find the risk-neutral expectation of any arbitrary function $f(y_{t+1})$ by letting:

$$g(y_{t+1}) = \frac{E_t^*(A_{t+1})}{A(y_{t+1})} f(y_{t+1})$$

.

The annuity measure is defined such that for any X:

$$E_t^A \left(\frac{E^*(A_{t+1})}{A(y_{t+1})} X_{t+1} \right) = E_t^*(X_{t+1})$$

And hence we have

$$E_t^A(g(y_{t+1})) = E_t^*(f(y_{t+1}))$$

Calculating the moments

For this paper I need to estimate calculate the first three risk-neutral moments of y_{t+1}, that is: $E_t^*(f(y_{t+1}))$ for $f(y_{t+1}) = y_{t+1}$, $f(y_{t+1}) = (y_{t+1} - E_t^*(y_{t+1}))^2$, and $f(y_{t+1}) = (y_{t+1} - E_t^*(y_{t+1}))^3$.

In each case, I first solve the second derivative of the function:

$$g(y_{t+1}) = \frac{E_t^*(A_{t+1})}{A(y_{t+1})} f(y_{t+1})$$

I then use my data on the forward prices of payer and receiver swaptions to calculate, using the standard:

$$E^A\left(g\left(y\right)\right) = \left(g(E_t^A(y_{t+1})) + \int_{-\infty}^{E_t^A(y_{t+1})} g''\left(k\right) \frac{P\left(k\right)}{E_t^*(A_{t+1})} dk + \int_{E_t^A(y_{t+1})}^{\infty} g''\left(k\right) \frac{C\left(k\right)}{E_t^*(A_{t+1})} dk \right)$$

J Constructing Alternative Risk Premium Estimates

This section describes how I construct proxies for the interest-rate risk premium predictions from other prominent models in the literature.

Adrian et al. (2013)

I replicate the approach from the original Adrian et al. (2013) (ACM) paper, fitting the model using expanding windows starting from 1964. I then calculate the 10-year par yield expected in 3 months and the 10-year forward par yield. The risk premium is the difference between these two values. My replication uses as a starting point the code made available by Arnab Biswas.[15]

Kim and Wright (2005)

I use term premium estimates provided by the Federal Reserve Board. To ensure the estimates are out-of-sample, I use the "older" calibration provided by the Board up to 2019 and the "newer" calibration thereafter. For quarterly and monthly horizon forecasts I divide the annual forecast by 4 or 12. Since 10y-in-1y risk premium estimates are not available I use the estimated risk premium on the 1-year instantaneous forward rate ("THREEFFTP0100"). This is likely to be close to the 10y-in-1y rate. For the ACM model, for example, the 10y-in-1y and instantaneous 1y term premium have a 90% correlation and the coefficient from regressing one on the other is approximately one.

Bauer and Rudebusch (2020)

I use the replication code provided by the authors. I calculate the quarterly forecast and risk premium of 8-year spot yields from the "observed-shifting-endpoints" model up to 2018. I use 8-year yields because this maturity approximately matches the duration of a 10-year par coupon bond.

[15] https://github.com/arnab13061989

Cochrane and Piazzesi (2005)

As in the original article, I regress the average annual log excess returns of the 2–5 year zero coupon bonds onto the 5 annual log forward rates 1–5 years from the Fama & Bliss dataset. I scale this factor with a regression Δy_{t+1} onto the factor. I construct the factor out-of-sample using expanding window regressions from 1964.

Term Spread

To construct a term-spread-based forecast, I regress Δy_{t+1} onto the 10-year minus 3-month (10y–3m) term spread available from FRED, using expanding window regressions