The Natural Rate of Interest in the Euro Area: Evidence from Inflation-Indexed Bonds

Jens H. E. Christensen[†]

&

Sarah Mouabbi[‡]

Abstract

The natural rate of interest, widely known as r_t^*, is a key variable used to judge the stance of monetary policy. We offer a novel euro-area r_t^* estimate based on a dynamic term structure model estimated directly on the prices of bonds with cash flows indexed to the euro-area harmonized index of consumer prices with adjustments for bond-specific risk and real term premia. Despite a recent increase, our estimate indicates that the natural rate in the euro area has fallen about 2 percentage points on net since 2002. We also devise a related measure of the stance of monetary policy, which suggests that monetary policy in the euro area was not accommodative at the height of the COVID-19 pandemic.

JEL Classification: C32, E43, E52, G12

Keywords: affine arbitrage-free term structure model, financial market frictions, convenience premium, monetary policy, rstar

We thank participants at the Second ESM Workshop on Economic Modelling in Policy Institutions and the 13th Bundesbank Term Structure Workshop for helpful comments, including our discussants Peter Hördahl and Andreea Vladu. Furthermore, we thank seminar participants at the Swiss National Bank, the Banque de France, the European Banking Authority, the European Central Bank, and the Danish National Bank for helpful comments. Moreover, we thank Jean-Paul Renne and Timo Reinelt for helpful comments and suggestions. Finally, we thank Gus Kmetz and Caroline Paulson for excellent research assistance. The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of San Francisco or the Federal Reserve System, or those of the Banque de France or the Eurosystem.

[†]Corresponding author: Federal Reserve Bank of San Francisco, 101 Market Street MS 1130, San Francisco, CA 94105, USA; phone: 1-415-974-3115; e-mail: jens.christensen@sf.frb.org.

[‡]Banque de France; e-mail: sarah.mouabbi@banque-france.fr.

This version: January 28, 2025.

1 Introduction

The so-called natural rate of interest, widely known as r_t^*, is a key variable in finance and macroeconomic theory. For investors, the steady-state level of the real short-term interest rate serves as an anchor for projections of the future discount rates used in valuing assets (e.g., Clarida 2014). For policymakers and researchers, the natural rate of interest is a policy lodestar that provides a neutral benchmark to calibrate the stance of monetary policy: Monetary policy is expansionary if the real short-term interest rate lies below the natural rate and contractionary if it lies above. A good estimate of the natural rate is also necessary to operationalize popular monetary policy rules such as the Taylor rule. For fiscal policy, the natural rate of interest is instrumental to assessing the sustainability of public finances in the long run. More broadly, in the decades prior to the COVID-19 pandemic, the possibility of a lower new normal for interest rates was at the center of key policy debates about bond market conundrums, global saving gluts, and secular stagnation.[1] More recently, the post-pandemic spike in interest rates globally has given rise to intense policy debates about whether interest rates will hold steady at the new higher levels or revert back towards their pre-pandemic lows.[2] In short, the natural rate of interest is a variable of immense importance.

Unfortunately, despite its importance, the natural rate of interest is not directly observable. Instead, it has to be inferred from economic data. In the literature, most estimates of the natural rate are drawn from *macroeconomic* models and data, including the widely cited Laubach and Williams (2003) model. In this paper, we follow Christensen and Rudebusch (2019, henceforth CR) and use *financial* models. Specifically, we rely on bond prices denominated in euros and indexed with the harmonized index for consumer prices (HICP) for our analysis and therefore offer a euro-area perspective on recent trends in the natural rate of interest.

To further motivate our focus on the euro area, we note that euro-area yield data are unique in that the European Central Bank (ECB) is a major central bank that has gone far in exploring the true lower bound for its key policy rate. One relevant policy question is therefore whether this extreme policy choice has caused the natural rate to be lower in the euro area than in other advanced economies. Alternatively, the causation could run in the other direction, namely that the ECB was forced to pursue what might appear to be an extremely accommodative stance of monetary policy *because* the natural rate in the euro area was already really low. We will attempt to provide an answer to this important question, which is likely to also have major implications for what to expect going forward in the post-pandemic world.

[1] See, for example, Greenspan (2005), Bernanke (2005), and Summers (2014, 2015), respectively, on these three debates.

[2] See, for example, Blanchard (2023) and Summers (2023).

The bonds we consider have coupon and principal payments indexed to the HICP (ex tobacco) and provide compensation to investors for the erosion of purchasing power due to price inflation in the euro area as a whole.[3] Therefore, their prices can be expressed directly in terms of real yields. The basic premise of our analysis is that the longer-term expectations embedded in these bond prices reflect financial market participants' views about the steady state of the euro-area economy, including its natural rate of interest.

To provide the cleanest possible read on investors' expectations for the natural rate in the euro area, we limit our focus to bonds issued by the French government. In principle, we could have included bonds indexed to HICP (ex tobacco) issued by other euro-area countries such as Germany, Italy, or Spain,[4] but it would complicate the analysis in terms of accounting for differences in credit and liquidity risks across these different markets and with few apparent benefits, in particular it would *not* provide us with a longer sample for our analysis.

The French government first issued bonds indexed to the HICP (ex tobacco), known as OAT€, in October 2001. However, given that we need at least two bonds to be trading, we start our analysis in October 2002. This long sample allows us to provide a 20-year perspective on the components that have influenced euro-area real yields in recent decades. Besides its length, this sample choice offers additional advantages. First, France has deep and liquid markets for government debt. Second, with maturities of up to 33 years, the OAT€ market contains the farthest forward-looking information among all the inflation-indexed bond markets in the euro area and hence is likely to provide the clearest evidence for the question at hand. Third, by relying on inflation-indexed bonds, we avoid any issues related to the effective lower bound that applies to the ECB's policy rate and other nominal interest rates. Furthermore, as the underlying factors affecting long-term interest rates are likely global in nature—such as worldwide demographic shifts or changes in productivity trends—the euro-area government bond market in general, and the French government bond market specifically, may well be as informative as any other major sovereign bond market. Finally, the French government held a AA credit rating from all major rating agencies during our sample period, which ends in December 2022. Hence, there is a minimum of credit risk to account for in our French bond price data.

Despite all these advantages the use of inflation-indexed bonds for measuring the natural rate of interest entails its own empirical challenges. One problem is that inflation-indexed bond prices include a real term premium. Given the generally upward slope of the OAT€ yield curve, the real term premium is presumably usually positive. However, little is known with certainty about its size or variability. In addition, despite the fairly large notional amount

[3] HICP is the price index targeted by the ECB for monetary policy purposes, but for historical reasons the HICP-indexed bonds issued in the euro area reference HICP (ex tobacco); see Ejsing et al. (2007).

[4] See Christensen et al. (2025) for an analysis of the limited universe of German inflation-linked government bonds indexed to the HICP (ex tobacco).

of outstanding OAT€s, these securities face unique market risks due to high demand from institutional investors such as pension funds and life insurance companies.[5]

To estimate the natural rate of interest in the presence of market risk and real term premia, we use an arbitrage-free dynamic term structure model of real yields augmented with a bond-specific risk factor. The identification of the bond-specific risk factor comes from its unique loading for each individual bond security as in Andreasen et al. (2021, henceforth ACR). Our analysis uses prices of individual bonds rather than the more usual input of yields from fitted synthetic curves. The underlying mechanism assumes that, over time, an increasing proportion of the outstanding inventory is locked up in buy-and-hold investors' portfolios. Given forward-looking investor behavior, this lock-up effect means that a particular bond's sensitivity to the market-wide bond-specific risk factor will vary depending on how seasoned the bond is and how close to maturity it is. In a careful study of nominal U.S. Treasuries, Fontaine and Garcia (2012) find a pervasive bond-specific factor that affects all bond prices, with loadings that vary with the maturity and age of each bond. By observing a cross section of bond prices over time—each with a different time-since-issuance and time-to-maturity—we can identify the overall bond-specific risk factor and each bond's loading on that factor. This technique is particularly useful for analyzing inflation-indexed debt when only a limited sample of bonds may be available, for example early in our sample.[6]

The theoretical arbitrage-free formulation of the model also provides identification of a time-varying real term premium in the pricing of OAT€s. Identifying the bond-specific risk premium and real term premium allows us to estimate the underlying frictionless real rate term structure and the natural rate of interest, which we measure as the average expected real short rate over a five-year period starting five years ahead—consistent with the longer-run perspective emphasized by Laubach and Williams (2016). Our preferred estimate of the natural rate of interest, r_t^*, is shown in Figure 1 along with ten-year nominal and real yields.[7] Both nominal and real long-term yields in the euro area trended down together during the 2002-2021 period, and this concurrence suggests little net change in inflation expectations or the inflation risk premium during that 20-year period. The estimated natural rate fell from above 1.5 percent to below -1.5 percent by the end of 2021, before retracing some of that decline during 2022. Accordingly, our results show that more than 75 percent of the 4-percentage-point decline in longer-term yields by the end of 2021 represents a reduction in the natural rate of interest. Our model estimates also indicate that about 75 percent of

[5]OAT€s also provide protection against net deflation over the life of each bond. However, the value of this protection is likely to be low and is therefore not considered; see Christensen and Mouabbi (2023).

[6]Finlay and Wende (2012) examine prices from a limited number of Australian inflation-indexed bonds but do not account for bond-specific liquidity or convenience premia.

[7]These yields are constructed using a model of French standard nominal government bonds, known as OATs, and a separate model of French OAT€ prices, each estimated directly on the observed bond prices as advocated by Andreasen et al. (2019).

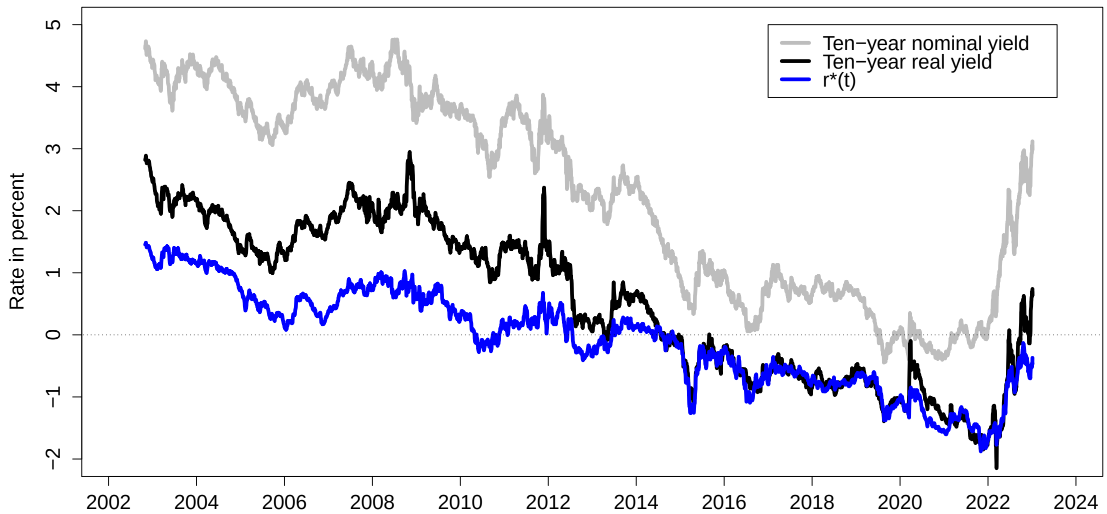

Figure 1: **Long-Term Nominal and Real Yields and an Estimate of r***
Ten-year nominal and real yields and our preferred AFNS-R model estimate of the equilibrium real short rate, r_t^*, i.e., the 5- to 10-year risk-neutral real rate.

the interest rate increases the last year of our sample reflect increases in the natural rate of interest. However, model projections suggest that the natural rate of interest is likely to revert only very gradually towards its old mean in the years ahead. Thus, policy rates in the euro area may return to levels close to the effective lower bound during economic downturns once the economy moves past the recent spell of high inflation. Finally, to evaluate the model more fully, we note that we perform our analysis using daily data. This could also be used to examine the impact of specific ECB policy announcements relying on established high-frequency event-study technology for identification, as in Christensen and Rudebusch (2012), but we leave that venue for future research.

As a separate contribution and to demonstrate the applicability of our model for economic analysis, we use it to devise market-based measures of the stance of monetary policy in the euro area. This is achieved by deducting our r_t^* estimate from observed measures of one-year real yields. We consider the latter to be a reasonable proxy for the theoretically ideal, but unobserved, instantaneous real short rate r_t appearing in textbook formulas of the stance of monetary policy measured as the gap between the current real short-term interest rate and its natural level. The results indicate that it took significant time for monetary policy in the euro area to reach an accommodative stance during both the Global Financial Crisis (GFC) and the COVID-19 pandemic.

We further validate our estimate of the stance of monetary policy by comparing it with a text-based measure introduced by Hubert and Portier (2024), who use machine learning to analyze the statement and transcript of the press conference following each ECB governing council meeting. Although similar most of the time, the market- and text-based measures of the policy stance deviate during three key periods, namely the European Sovereign Debt Crisis, the COVID-19 pandemic, and the post-pandemic economic reopening characterized by highly elevated inflation. During the first and last of these three episodes, text-based analysis points to a hawkish posture among policymakers, while our market-based measure suggests that monetary policy was in fact quite accommodative. In contrast, policymakers clearly tried to achieve an accommodative stance for policy through their communications in response to the COVID-19 pandemic. However, our market-based measure suggests that policy did not become accommodative until into 2021. Thus, our results underscore the challenges of central bank communication during times of crisis. At the same time, though, our market-based measure of the stance of monetary policy offers a way to examine in real time to what extent investors' and financial market participants' perceptions about the stance of monetary policy is aligned with the one communicated by policymakers. Hence, we see our measure as a potentially important policy tool going forward, but we leave it for future research to examine its usefulness for this purpose.

Our analysis focuses on a real term structure model that only includes the prices of inflation-indexed bonds. This methodology contrasts with previous term structure research in two ways. First, previous term structure models are almost universally estimated not on observed bond prices but on synthetic zero-coupon yields obtained from fitted yield curves. Fontaine and Garcia (2012) argue that the use of such synthetic yields can erase useful information on bond-specific price effects, and they provide a rare exception of the estimation of a term structure model with bond prices. More generally, the use of interpolated yield curves in term structure analysis can introduce arbitrary and unnecessary measurement error.[8] A second difference is that past analysis of inflation-indexed bonds has jointly modeled both the real and nominal yield curves, e.g., Christensen et al. (2010), Abrahams et al. (2016), and D'Amico et al. (2018) for the United States and Joyce et al. (2010) and Carriero et al. (2018) for the United Kingdom. Such joint specifications can also be used to estimate the steady-state real rate—though this earlier work has emphasized only the measurement of inflation expectations and risk premia.[9] Relative to our procedure of using just inflation-indexed bonds to estimate the natural rate, including both real and nominal yields has the advantage of being able to estimate a model on a much larger sample of bond yields. However, a joint

[8]Dai et al. (2004) found notable differences in empirical results across four different yield curve interpolation schemes. For further discussion of these issues; see Andreasen et al. (2019).

[9]Joyce et al. (2012) use dynamic term structure models of U.K. index-linked government bond yields to study long-term real rate expectations while accounting for real term premia though not bond-specific risk or liquidity premia.

specification also requires additional modeling structure—including specifying more pricing factors, an inflation risk premium, and inflation expectations. The greater number of modeling elements—along with the requirement that this more elaborate structure remains stable over the sample—raises the risk of model misspecification, which can contaminate estimates of the natural rate and model inference more generally. In particular, if the inflation components are misspecified, the whole dynamic system may be compromised, a valid concern in the recent high-inflation environment. Furthermore, during the 2009-2021 period when the ECB kept its policy rate close to its effective lower bound, the dynamic interactions of short- and medium-term *nominal* yields were likely affected. Such a constraint is very difficult to include in an empirical term structure model of nominal yields (see Swanson and Williams 2014 and Christensen and Rudebusch 2015 for discussions). By relying solely on real yields, which are not subject to a lower bound, we avoid this complication altogether.

The analysis in this paper relates to several important literatures. Most directly, it speaks to the burgeoning literature on measurement of the natural rate of interest. Second, our estimates of the real yield curve that would prevail without trading frictions have implications for asset pricing analysis on the true slope of the real yield curve. Furthermore, our results relate to research on financial market liquidity and convenience premia. Finally, the paper contributes to the rapidly growing literature on the economic consequences of the COVID-19 pandemic.

The remainder of the paper is organized as follows. Section 2 contains a description of the French OAT€ bond data, while Section 3 details the no-arbitrage term structure models we use and presents the empirical results. Section 4 describes the estimated real bond-specific premia, while Section 5 analyzes our OAT€-based estimate of the natural rate and compares it with other measures. Finally, Section 6 introduces our market-based measure of the stance of the ECB's monetary policy before Section 7 concludes.

2 The French OAT€ Bond Data

This section briefly describes the available data downloaded from Bloomberg for the market for French inflation-indexed bonds referencing the harmonized index for consumer prices (HICP) (ex tobacco) and known as OAT€s.

To give a sense of the size of the French government bond market, we note up front that, as of the end of December 2022, the total outstanding notional amount of marketable bonds issued by the French government was €2,28 trillion. In terms of medium- and long-term debt, the outstanding notional amount was €2,13 trillion of which €262 billion, or 12.3 percent, represented inflation-indexed securities, and out of this amount OAT€s represented €183.7

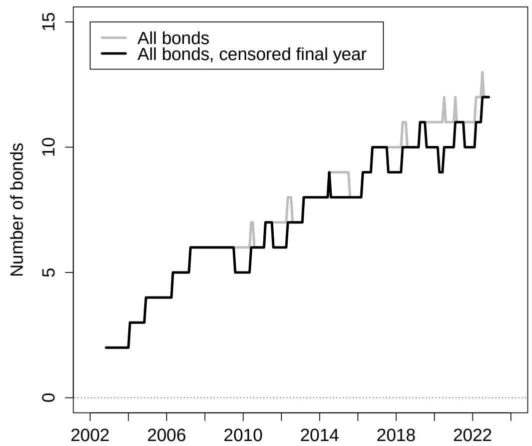

(a) Number of OAT€ bonds

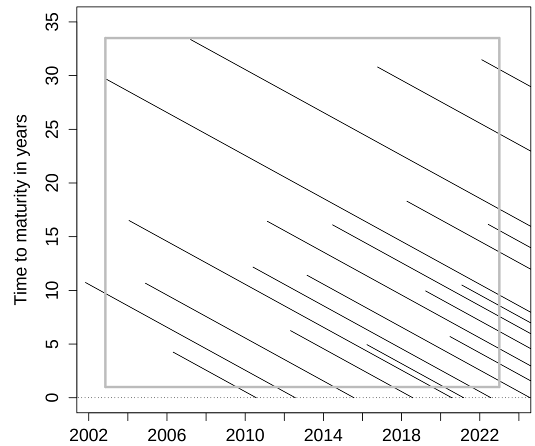

(b) Distribution of OAT€ bonds

Figure 2: **Overview of the French OAT€ Bond Data**
Panel (a) reports the number of outstanding OAT€ bonds at a given point in time. Panel (b) shows the maturity distribution of all French OAT€ bonds issued since October 2001. The solid gray rectangle indicates the sample used in our analysis, where the sample is restricted to start on October 31, 2002, and limited to bond prices with more than one year to maturity after issuance.

billion, or 70.1 percent.[10] Despite the large size of the French government bond market, the French government still held a AA rating from all major rating agencies during our sample period. As a consequence, there is essentially no credit risk to account for in our bond price data, as also suggested by measures of the credit risk premia of French government bond examined in Section 2.1.

The French government issued its first inflation-indexed bond referencing HICP on October 31, 2001. At the end of December 2022, the outstanding amount of French OAT€s was €184 billion as already noted. Thus, this is a large market in a European context. The total number of such bonds outstanding over time in our sample is shown as a solid gray line in Figure 2(a). At the end of our sample, 12 French OAT€s were outstanding. However, as noted by Gürkaynak et al. (2010) and ACR, prices of inflation-indexed bonds near their maturity tend to be somewhat erratic because of the indexation lag in their payouts. Therefore, to facilitate model estimation, we censor the prices of OAT€s from our sample when they have less than one year to maturity. Using this cutoff, the number of OAT€s in the sample is modestly reduced, as shown with a solid black line in Figure 2(a).

Figure 2(b) shows the distribution of the available universe of French OAT€s, where we note that a repeated, although somewhat infrequent, issuance of ten-, fifteen-, and thirty-year

[10]This information is available at
https://www.aft.gouv.fr/files/medias-aft/7_Publications/7.2_BM/392_Monthly%20bulletin%20january%202023.pdf

OAT€ bond	No. obs.	Issuance Date	Issuance amount	Total uplifted amount
(1) 3% 7/25/2012	2,278	10/31/2001	787	14,494
(2) 3.15% 7/25/2032	5,258	10/31/2002	587	12,098
(3) 2.25% 7/25/2020	4,045	1/22/2004	298	20,310
(4) 1.6% 7/25/2015	2,522	11/23/2004	3,527	14,052
(5) 1.25% 7/25/2010	849	4/25/2006	3,634	9,325
(6) 1.8% 7/25/2040	4,119	3/14/2007	347	12,929
(7) 1.1% 7/25/2022	2,910	5/25/2010	2,883	19,928
(8) 1.85% 7/25/2027	3,094	2/16/2011	418	23,433
(9) 0.25% 7/25/2018	1,370	2/16/2011	2,520	11,257
(10) 0.25% 7/25/2024	2,566	2/26/2013	2,320	14,644
(11) 0.7% 7/25/2030	2,225	6/18/2014	429	17,232
(12) 0.1% 3/1/2021	1,029	3/21/2016	2,290	7,566
(13) 0.1% 7/25/2047	1,629	10/5/2016	556	13,027
(14) 0.1% 7/25/2036	1,233	4/6/2018	416	12,747
(15) 0.1% 3/1/2029	984	3/21/2019	2,128	17,772
(16) 0.1% 3/1/2026	660	6/18/2020	3,044	12,736
(17) 0.1% 7/25/2031	505	1/24/2021	2,370	11,741
(18) 0.1% 7/25/2053	239	2/1/2022	217	6,447
(19) 0.1% 7/25/2038	153	6/1/2022	549	7,089

Table 1: **Sample of French OAT€ Bonds**

The table reports the characteristics, first issuance date and amount, and total amount issued in millions of euros either at maturity or as of December 31, 2022, for the sample of French OAT€ bonds. Also reported are the number of daily observation dates for each bond during the sample period from October 31, 2002, to December 31, 2022.

OAT€s implies that there is a fairly wide range of available maturities in the data going back to the start of our sample in October 2002. It is this cross-sectional dispersion that provides the econometric identification of the factors in our models, including the inflation-indexed bond-specific risk factor. Finally, Table 1 contains the contractual details of all 19 French OAT€s in our data as well as the number of daily observations for each in our sample.

Figure 3 shows the yields to maturity for all French OAT€ bonds in our sample at daily frequency from October 31, 2002, to December 30, 2022. Note the following regarding these yield series. First, the significant persistent decline in real yields over this 20-year period is clearly visible. Long-term real yields in the euro area were close to 3 percent in late 2002 and had dropped below -1 percent by late 2021 before retracing some of that decline during 2022. The empirical question we are interested in is to what extent these persistent fluctuations represent changes in the natural rate or are driven by other factors such as term or other bond-specific risk premia. Second, business cycle variation in the shape of the yield curve is pronounced around the lower trend. The yield curve tends to flatten ahead of recessions and steepen during the initial phase of economic recoveries. These characteristics are the

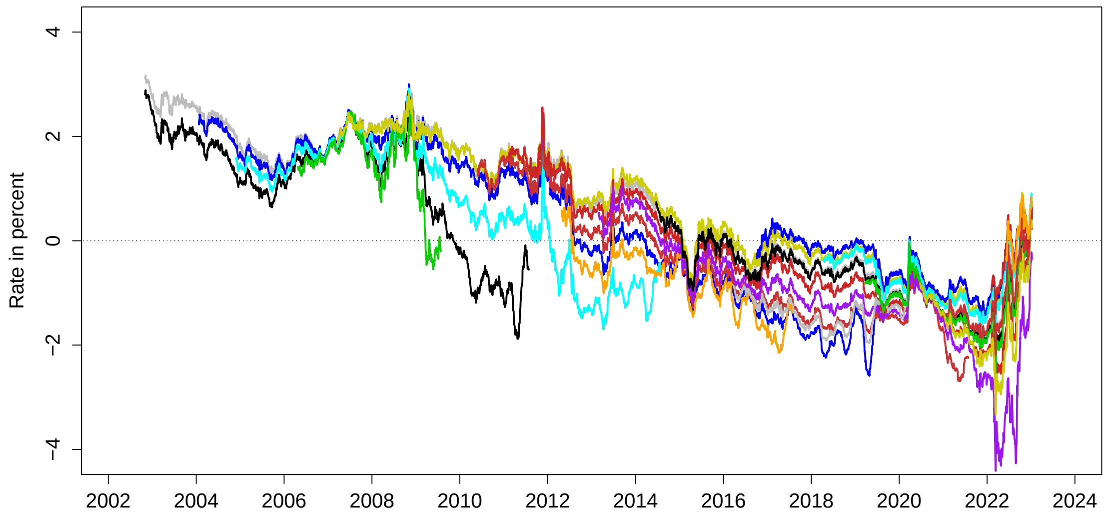

Figure 3: **Yield to Maturity of French OAT€ Bonds**

practical motivation behind our choice of using a three-factor model for the frictionless part of the euro-area real yield curve, adopting an approach similar to what is standard for U.S. and U.K. nominal yield data; see Christensen and Rudebusch (2012).

Figure 4 shows the inflation index ratios for all 19 French OAT€s in our sample. We note that none of the bonds have been exposed to any prolonged period of deflation, defined as periods with inflation index ratios below one. Indeed, thanks to the generally positive inflation environment in the euro area, the ratios tend to relatively quickly become significantly positive. This suggests that their offered deflation protection is likely to be of modest value, similar to what Christensen and Mouabbi (2023, henceforth CM) find for French government bonds indexed using the French CPI and known as OATi's. We therefore disregard this component in our analysis and leave it for future research to assess its value.

2.1 The Credit Risk of French Government Bonds

In this section, we assess whether there are any material credit risk issues to consider in modeling French OAT€ bond prices.

First, we examine rates on so-called credit default swap (CDS) contracts. They reflect the annual rate investors are willing to pay to buy protection against default-related losses on these bonds over a fixed period of time stipulated in the contract. Such derivatives have been used to price the credit risk of many countries, including France, since the early 2000s.

In Figure 5, we plot the available series downloaded from Bloomberg for the five-year

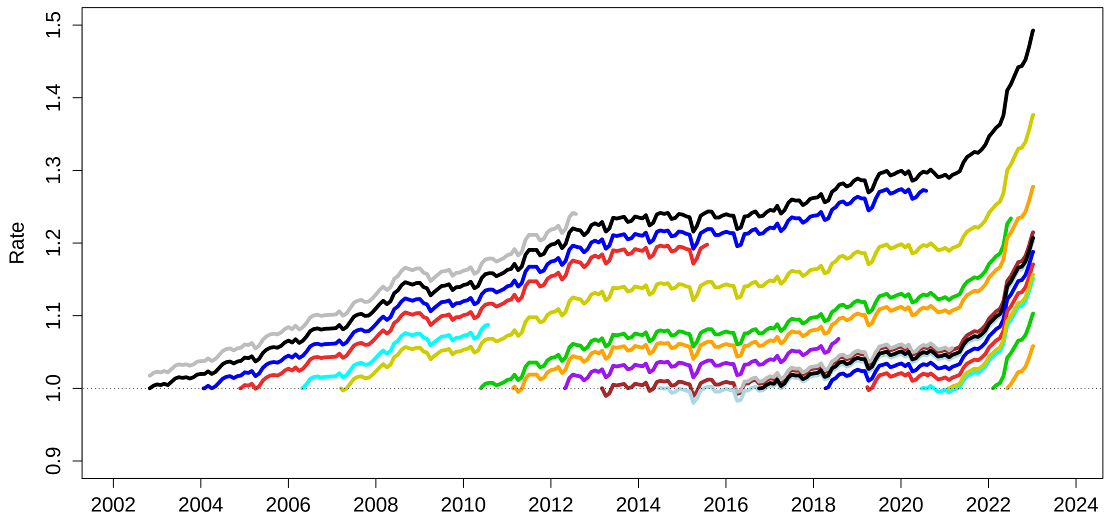

Figure 4: **Inflation Index Ratios of French OAT€ Bonds**

French CDS rate shown with a solid blue line. In addition to occasionally missing values, the five-year CDS rate is missing entirely between September 24, 2018, and March 11, 2020. As a consequence of the missing data, we consider an alternative measure of the credit risk of French government bonds. Specifically, we include the five-year forward yield spread between French and German inflation-indexed bond yields for a period starting five years ahead. Due to the late launch of the German inflation-indexed government bond program, we can only construct the 5yr5yr Franco-German real yield spread starting June 12, 2009, when the third such German bond was issued; see Christensen et al. (2025) for details. The available series since then through the end of our sample is shown with a solid grey line in Figure 5. Similar to regular German bunds, German inflation-indexed government bonds trade at a convenience premium as documented by Christensen et al. (2025). However, given their lower liquidity, we refer to these premia as safety premia; see Christensen and Mirkov (2022). For the same reason we interpret the 5yr5yr Franco-German real yield spread as mainly reflecting differences in credit risk premia rather than differences in liquidity risk premia.

In light of the incomplete sample histories for both credit risk measures, we construct a composite measure of the credit risk of French government bonds by averaging the two measures. To further smooth out idiosyncratic noise, we calculate the four-week moving average of this composite series, which is shown with a solid black line in Figure 5. This series is available daily from April 1, 2003, to December 30, 2022, and we take it to be a representative proxy for the credit risk premium of French government bonds.

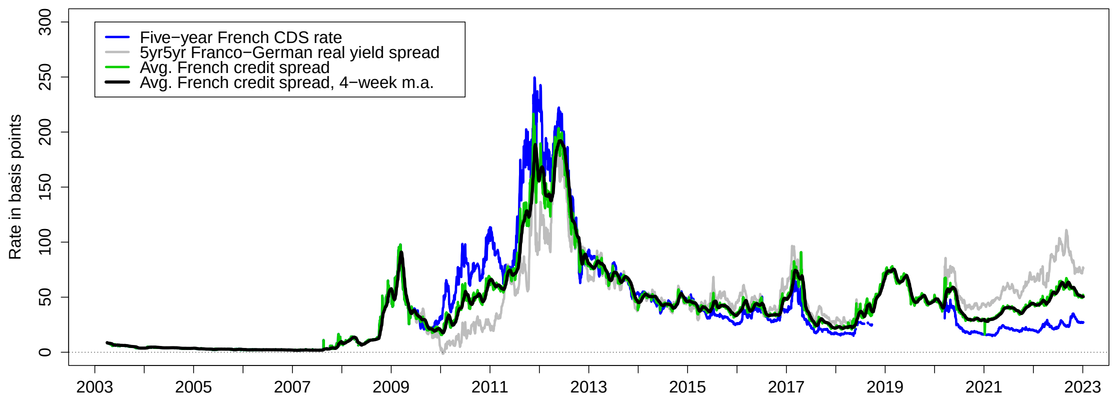

Figure 5: **Measures of the Credit Risk of French Government Bonds**

We note that, beyond being elevated around the European sovereign debt crisis in the 2010-2013 period, our composite measure of credit risk has remained fairly stable fluctuating around 50 basis points. Thus, the main takeaway for our analysis is that changes in credit risk premia cannot account for the persistent trends in the OAT€ bond yields during our sample period. Furthermore and importantly, if anything, given the modest positive net change in the composite credit risk measure from April 2003 to December 2022, the credit risk component should have *pushed up* French real yields. Instead, French real yields of all maturities have experienced a persistent significant net *decline* since 2003. Thus, we feel that we can rule out with great confidence credit risk components as an important driver of French real yields during our sample period.

2.2 Bid-Ask Spreads of OAT€ Bonds

In this section, to shed light on the trading frictions in the market for French OAT€ bonds, we examine their bid-ask spreads.

To begin, we note that reliable bid-ask spreads for individual OAT€ bonds are available from Bloomberg starting in 2011.[11] In Figure 6, we show the smallest and largest observed bid-ask spread for each observation date as well as the median bid-ask spread. Although elevated during the European sovereign debt crisis in 2011 and the first half of 2012, the median bid-ask spread since then has followed a stable and declining trend that has left it close to 1 basis point by the end of our sample. This points to high, and even improving, liquidity in the market for OAT€s during the last decade of our sample. That said, we still

[11]Speck (2021) reports bid-ask spreads for French inflation-indexed bonds back to 2006, but his data comes form a different source.

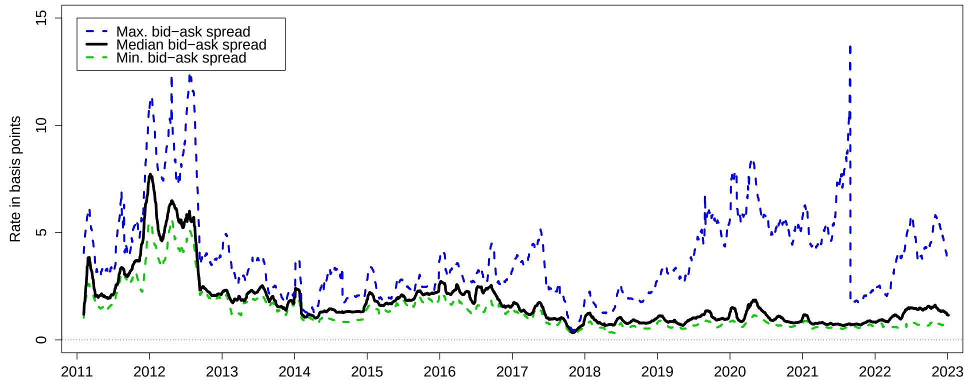

Figure 6: **Bid-Ask Spreads of OAT€ Bonds**

want to account for any bond-specific effects tied to liquidity in our analysis, in particular in light of the fact that there is one or more OAT€s that face challenging trading conditions on an on-going basis as evidenced by the elevated maximum bid-ask spreads observed for extended periods in Figure 6.

A key purpose of the remainder of the paper is to quantify the importance of these bond-specific risk premia in the pricing of OAT€ bonds and what adjustments for them may imply about bond investors' underlying real short-rate expectations and associated real term premia.

3 Model Estimation and Results

In this section, we first describe how we model yields in a world without any frictions to trading. This model of frictionless dynamics is fundamental to our analysis. We then detail the augmented model that accounts for the bond-specific premia in inflation-indexed yields. This is followed by a description of the restrictions imposed to achieve econometric identification of this model and its estimation. We end the section with a brief summary of our estimation results.

3.1 A Frictionless Arbitrage-Free Model of Real Yields

To capture the fundamental or frictionless factors operating the OAT€ real yield curve, we choose to focus on the tractable affine dynamic term structure model introduced in Christensen et al. (2011).[12]

[12] Although the model is not formulated using the canonical form of affine term structure models introduced by Dai and Singleton (2000), it can be viewed as a restricted version of the canonical Gaussian model; see Christensen et al. (2011) for details.

In this arbitrage-free Nelson-Siegel (AFNS) model, the state vector is denoted by $X_t = (L_t, S_t, C_t)$, where L_t is a level factor, S_t is a slope factor, and C_t is a curvature factor. The instantaneous risk-free real rate is defined as

$$r_t = L_t + S_t. \tag{1}$$

The risk-neutral (or $\mathbb{Q}$-) dynamics of the state variables are given by the stochastic differential equations[13]

$$\begin{pmatrix} dL_t \\ dS_t \\ dC_t \end{pmatrix} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & -\lambda & \lambda \\ 0 & 0 & -\lambda \end{pmatrix} \begin{pmatrix} L_t \\ S_t \\ C_t \end{pmatrix} dt + \Sigma \begin{pmatrix} dW_t^{L,\mathbb{Q}} \\ dW_t^{S,\mathbb{Q}} \\ dW_t^{C,\mathbb{Q}} \end{pmatrix}, \tag{2}$$

where Σ is the constant covariance (or volatility) matrix that is assumed to be diagonal, as recommended by Christensen et al. (2011).[14] Based on this specification of the $\mathbb{Q}$-dynamics, real zero-coupon bond yields preserve the Nelson-Siegel factor loading structure as

$$y_t(\tau) = L_t + \left(\frac{1 - e^{-\lambda\tau}}{\lambda\tau}\right) S_t + \left(\frac{1 - e^{-\lambda\tau}}{\lambda\tau} - e^{-\lambda\tau}\right) C_t - \frac{A(\tau)}{\tau}, \tag{3}$$

where $A(\tau)$ is a convexity term that adjusts the functional form in Nelson and Siegel (1987) to ensure absence of arbitrage (see Christensen et al. (2011)).

To complete the description of the model and to implement it empirically, we will need to specify the risk premia that connect these factor dynamics under the $\mathbb{Q}$-measure to the dynamics under the real-world (or physical) $\mathbb{P}$-measure. It is important to note that there are no restrictions on the dynamic drift components under the empirical $\mathbb{P}$-measure beyond the requirement of constant volatility. To facilitate empirical implementation, we use the essentially affine risk premium specification introduced in Duffee (2002). In the Gaussian framework, this specification implies that the risk premia Γ_t depend on the state variables; that is,

$$\Gamma_t = \gamma^0 + \gamma^1 X_t,$$

where $\gamma^0 \in \mathbf{R}^3$ and $\gamma^1 \in \mathbf{R}^{3\times 3}$ contain unrestricted parameters.

Thus, the resulting unrestricted three-factor AFNS model has $\mathbb{P}$-dynamics given by

$$\begin{pmatrix} dL_t \\ dS_t \\ dC_t \end{pmatrix} = \begin{pmatrix} \kappa_{11}^{\mathbb{P}} & \kappa_{12}^{\mathbb{P}} & \kappa_{13}^{\mathbb{P}} \\ \kappa_{21}^{\mathbb{P}} & \kappa_{22}^{\mathbb{P}} & \kappa_{23}^{\mathbb{P}} \\ \kappa_{31}^{\mathbb{P}} & \kappa_{32}^{\mathbb{P}} & \kappa_{33}^{\mathbb{P}} \end{pmatrix} \left(\begin{pmatrix} \theta_1^{\mathbb{P}} \\ \theta_2^{\mathbb{P}} \\ \theta_3^{\mathbb{P}} \end{pmatrix} - \begin{pmatrix} L_t \\ S_t \\ C_t \end{pmatrix} \right) dt + \Sigma \begin{pmatrix} dW_t^{L,\mathbb{P}} \\ dW_t^{S,\mathbb{P}} \\ dW_t^{C,\mathbb{P}} \end{pmatrix}.$$

[13]As discussed in Christensen et al. (2011), with a unit root in the level factor, the model is not arbitrage-free with an unbounded horizon; therefore, as is often done in theoretical discussions, we impose an arbitrary maximum horizon.

[14]As per Christensen et al. (2011), $\theta^{\mathbb{Q}}$ is set to zero without loss of generality.

This is the transition equation in the Kalman filter estimation.

3.2 An Arbitrage-Free Model of Real Yields with Bond-Specific Risk

In this section, we augment the frictionless AFNS model introduced above to account for any bond-specific risk premia embedded in the OAT€ prices. To do so, let $X_t = (L_t, S_t, C_t, X_t^R)$ denote the state vector of the four-factor AFNS-R model with bond-specific risk premium adjustment. As in the non-augmented model, we let the frictionless instantaneous real risk-free rate be defined by equation (1), while the risk-neutral dynamics of the state variables used for pricing are given by

$$\begin{pmatrix} dL_t \\ dS_t \\ dC_t \\ dX_t^R \end{pmatrix} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & \lambda & -\lambda & 0 \\ 0 & 0 & \lambda & 0 \\ 0 & 0 & 0 & \kappa_R^{\mathbb{Q}} \end{pmatrix} \left[\begin{pmatrix} 0 \\ 0 \\ 0 \\ \theta_R^{\mathbb{Q}} \end{pmatrix} - \begin{pmatrix} L_t \\ S_t \\ C_t \\ X_t^R \end{pmatrix} \right] dt + \Sigma \begin{pmatrix} dW_t^{L,\mathbb{Q}} \\ dW_t^{S,\mathbb{Q}} \\ dW_t^{C,\mathbb{Q}} \\ dW_t^{R,\mathbb{Q}} \end{pmatrix},$$

where Σ continues to be a diagonal matrix.

In the augmented model, OAT€ yields are sensitive to bond-specific risks because the net present value of their future cash flow is calculated using the following discount function:

$$\overline{r}^i(t, t_0^i) = r_t + \beta^i(1 - e^{-\lambda^{R,i}(t-t_0^i)})X_t^R = L_t + S_t + \beta^i(1 - e^{-\lambda^{R,i}(t-t_0^i)})X_t^R. \quad (4)$$

CR show that the net present value of one unit of consumption paid by OAT€ i at time $t+\tau$ has the following exponential-affine form

$$\begin{aligned} P_t(t_0^i, \tau) &= E^{\mathbb{Q}}\left[e^{-\int_t^{t+\tau} \overline{r}^i(s,t_0^i)ds}\right] \\ &= \exp\Big(B_1(\tau)L_t + B_2(\tau)S_t + B_3(\tau)C_t + B_4(t, t_0^i, \tau)X_t^R + A(t, t_0^i, \tau)\Big). \end{aligned}$$

This result implies that the model belongs to the class of Gaussian affine term structure models. Note also that, by fixing $\beta^i = 0$ for all i, we recover the AFNS model.

Now, consider the whole value of OAT€ i issued at time t_0^i with maturity at $t + \tau^i$ that pays an annual coupon C^i. Its price is given by[15]

$$\begin{aligned} \overline{P}_t(t_0^i, \tau^i, C^i) &= C^i(t_1 - t)E^{\mathbb{Q}}\left[e^{-\int_t^{t_1} \overline{r}^i(s,t_0^i)ds}\right] + \sum_{j=2}^{N} C^i E^{\mathbb{Q}}\left[e^{-\int_t^{t_j} \overline{r}^i(s,t_0^i)ds}\right] \\ &\quad + E^{\mathbb{Q}}\left[e^{-\int_t^{t+\tau^i} \overline{r}^i(s,t_0^i)ds}\right]. \end{aligned}$$

[15] This is the clean price that does not account for any accrued interest and maps to our observed bond prices.

There are only two minor omissions in this bond pricing formula. First, it does not account for the lag in the inflation indexation of the OAT€ bond payoff. The potential error from this omission should be modest (see Grishchenko and Huang 2013), especially as we exclude bonds from our sample when they have less than one year of maturity remaining. Second, we do not account for the value of deflation protection offered by OAT€s, as already noted. However, CM find these values to be very small for French OATi bonds indexed to the French consumer price index, and, given that HICP inflation has run quite a bit above French CPI inflation during our sample, the value of this protection for OAT€ bonds is likely to be entirely negligible.

Finally, to complete the description of the AFNS-R model, we again specify an essentially affine risk premium structure, which implies that the risk premia Γ_t take the form

$$\Gamma_t = \gamma^0 + \gamma^1 X_t,$$

where $\gamma^0 \in \mathbf{R}^4$ and $\gamma^1 \in \mathbf{R}^{4\times 4}$ contain unrestricted parameters. Thus, the resulting unrestricted four-factor AFNS-R model has $\mathbb{P}$-dynamics given by

$$\begin{pmatrix} dL_t \\ dS_t \\ dC_t \\ dX_t^R \end{pmatrix} = \begin{pmatrix} \kappa_{11}^{\mathbb{P}} & \kappa_{12}^{\mathbb{P}} & \kappa_{13}^{\mathbb{P}} & \kappa_{14}^{\mathbb{P}} \\ \kappa_{21}^{\mathbb{P}} & \kappa_{22}^{\mathbb{P}} & \kappa_{23}^{\mathbb{P}} & \kappa_{24}^{\mathbb{P}} \\ \kappa_{31}^{\mathbb{P}} & \kappa_{32}^{\mathbb{P}} & \kappa_{33}^{\mathbb{P}} & \kappa_{34}^{\mathbb{P}} \\ \kappa_{41}^{\mathbb{P}} & \kappa_{42}^{\mathbb{P}} & \kappa_{43}^{\mathbb{P}} & \kappa_{44}^{\mathbb{P}} \end{pmatrix} \left(\begin{pmatrix} \theta_1^{\mathbb{P}} \\ \theta_2^{\mathbb{P}} \\ \theta_3^{\mathbb{P}} \\ \theta_4^{\mathbb{P}} \end{pmatrix} - \begin{pmatrix} L_t \\ S_t \\ C_t \\ X_t^R \end{pmatrix} \right) dt + \Sigma \begin{pmatrix} dW_t^{L,\mathbb{P}} \\ dW_t^{S,\mathbb{P}} \\ dW_t^{C,\mathbb{P}} \\ dW_t^{R,\mathbb{P}} \end{pmatrix}.$$

This is the transition equation in the Kalman filter estimation.

3.3 Model Estimation and Econometric Identification

Due to the nonlinear relationship between the state variables and the bond prices, the model cannot be estimated with the standard Kalman filter. Instead, we use the extended Kalman filter as in Kim and Singleton (2012); see CR for details. Furthermore, to make the fitted errors comparable across bonds of various maturities, we scale each bond price by its duration. Thus, the measurement equation for the bond prices takes the following form

$$\frac{P_t^i(t_0^i, \tau^i)}{D_t^i(t_0^i, \tau^i)} = \frac{\widehat{P}_t^i(t_0^i, \tau^i)}{D_t^i(t_0^i, \tau^i)} + \varepsilon_t^i,$$

where $\widehat{P}_t^i(t_0^i, \tau^i)$ is the model-implied price of bond i and $D_t^i(t_0^i, \tau^i)$ is its duration, which is calculated before estimation. See Andreasen et al. (2019) for evidence supporting this formulation of the measurement equation.

Furthermore, since the bond-specific risk factor is a latent factor that we do not observe, its level is not identified without additional restrictions. As a consequence, we let the second

OAT€ bond, which was issued right at the start of our sample, have a unit loading on this factor, that is, the 30-year OAT€ bond issued on October 31, 2002, and maturing on July 25, 2032, with 3.15 percent coupon has $\beta^i = 1$. This choice implies that the β^i sensitivity parameters measure bond-specific risk sensitivity relative to that of the 30-year 2032 OAT€ bond.

Finally, we note that the $\lambda^{R,i}$ parameters can be hard to identify if their values are too large or too small. As a consequence, we follow ACR and impose the restriction that they fall within the range from 0.0001 to 10, which is without practical consequences, as demonstrated by CM. Also, for numerical stability during model optimization, we impose the restriction that the β^i parameters fall within the range from 0 to 250, which turns out to be a binding constraint for two of the 19 bonds in our sample, but it is again the case that these two constraints are without practical consequences.

3.4 Estimation Results

This section presents our benchmark estimation results. In the interest of simplicity, in this section we focus on a version of the AFNS-R model where $K^{\mathbb{P}}$ and Σ are diagonal matrices. As shown in ACR, these restrictions have hardly any effects on the estimated bond-specific risk premium for each inflation-indexed bond, because it is identified from the model's $\mathbb{Q}$-dynamics, which are independent of $K^{\mathbb{P}}$ and only display a weak link to Σ through the small convexity adjustment in the bond yields. Furthermore, we stress that we relax this assumption in Section 5 when we analyze estimates of r_t^*, which are indeed sensitive to the specification of the models' $\mathbb{P}$-dynamics.

Table 2 reports the summary statistics for the fitted errors of individual OAT€s as well as for all OAT€s combined. With the single exception of OAT€ number 4 in our sample, there is otherwise uniform improvement in model fit from incorporating the bond-specific risk factor into the AFNS model. Still, it is worth noting that the AFNS model is able to deliver a root mean-squared fitted error of 5.6 basis points across all bonds combined, which in general could be characterized as a satisfactory fit, but obviously not as good as the RMSE of 4.3 basis points for all bonds combined achieved by the AFNS-R model, which represents a really good fit to the entire cross section of yields. Note also that neither the 15- nor 30-year bonds pose any particular challenges for the two models. Thus, both the AFNS and AFNS-R models are clearly able to fit those long-term bond yields to a satisfactory level of accuracy.

Table 3 contains the estimated dynamic parameters. Note that the dynamics of the first three factors are qualitatively very similar across the two estimations. Hence, the frictionless dynamics of the state variables within the AFNS-R model are essentially statistically indistinguishable from the corresponding dynamics in the simpler AFNS model. We take this as a sign of the robustness of our results. Furthermore, λ is smaller in the AFNS-R model.

OAT€ bond	Pricing errors				Estimated parameters			
	AFNS		AFNS-R		AFNS-R			
	Mean	RMSE	Mean	RMSE	β^i	SE	$\lambda^{R,i}$	SE
(1) 3% 7/25/2012	0.32	4.29	0.55	3.00	249.9962	1.3687	0.0022	0.0001
(2) 3.15% 7/25/2032	1.09	4.24	0.85	2.62	1	n.a.	9.9999	1.3562
(3) 2.25% 7/25/2020	-0.88	4.81	0.58	2.90	45.4290	1.2766	0.0024	0.0001
(4) 1.6% 7/25/2015	-4.88	9.16	-5.75	12.86	58.6938	0.8397	0.7940	0.0444
(5) 1.25% 7/25/2010	1.14	4.33	0.94	2.59	0.5500	0.1740	9.9941	1.3542
(6) 1.8% 7/25/2040	-1.19	4.70	0.72	2.81	0.9419	0.0669	9.9945	1.3535
(7) 1.1% 7/25/2022	-0.94	4.23	-0.58	3.13	2.9259	0.3679	0.1050	0.0222
(8) 1.85% 7/25/2027	2.10	4.23	1.53	2.91	0.8847	0.0254	10.0000	1.3522
(9) 0.25% 7/25/2018	-2.19	4.99	0.45	2.06	4.6753	0.1780	1.9798	0.6843
(10) 0.25% 7/25/2024	0.26	5.24	0.65	2.59	1.4392	0.0439	9.2613	1.3507
(11) 0.7% 7/25/2030	-1.76	4.73	-0.16	2.28	3.7724	1.0018	0.0481	0.0155
(12) 0.1% 3/1/2021	8.07	9.44	2.17	3.45	1.2439	0.0401	1.0239	0.1341
(13) 0.1% 7/25/2047	3.11	5.01	0.11	2.17	249.9910	1.3611	0.0028	0.0001
(14) 0.1% 7/25/2036	-0.30	2.98	0.31	2.14	1.0284	0.0538	10.0000	1.3318
(15) 0.1% 3/1/2029	2.65	3.65	1.32	2.48	143.6258	1.3572	0.0014	0.0000
(16) 0.1% 3/1/2026	14.87	16.24	1.30	3.25	35.0249	1.3548	0.0100	0.0006
(17) 0.1% 7/25/2031	-4.07	7.19	0.47	2.14	1.8264	0.1884	0.6665	0.1139
(18) 0.1% 7/25/2053	1.98	7.38	0.41	3.89	29.8711	1.3287	0.2621	0.0169
(19) 0.1% 7/25/2038	3.39	4.95	0.04	2.99	1.3281	0.0781	9.9998	1.1221
All yields	0.23	5.61	0.20	4.25	-	-	-	-
Max $\mathcal{L}^{EKF}$	217,238.6		234,570.8		-		-	

Table 2: **Pricing Errors and Estimated Bond-Specific Risk Parameters**
This table reports the mean pricing errors (Mean) and the root mean-squared pricing errors (RMSE) of French OAT€ bonds in the AFNS and AFNS-R models estimated with a diagonal specification of $K^{\mathbb{P}}$ and Σ. The errors are computed as the difference between the French OAT€ bonds market price expressed as yield to maturity and the corresponding model-implied yield. All errors are reported in basis points. Standard errors (SE) are not available (n.a.) for the normalized value of β^2.

This implies that the yield loadings of the slope factor decays toward zero more slowly as the maturity increases. At the same time, the peak of the curvature yield loadings is located at a later maturity compared with its loading in the AFNS model. As a consequence, slope and curvature matter more for longer-term yields in the AFNS-R model. This helps explain part of the better fit to the entire cross section of bonds within that model.

The estimated paths of the level, slope, and curvature factors from the two models are shown in Figure 7. While the two models' slope factors are close to each other most of the time, their level factors have a wedge between them. However, they generally move in tandem, as both exhibit a persistent decline from 2002 through the end of 2021 that is partially offset by a sharp reversal during the last year of our sample. The lower path of the level factor in the AFNS model is offset by a mostly higher path of the curvature factor in that model compared to the AFNS-R model. Accordingly, the main impact of accounting for bond-specific risk premia in the pricing of the OAT€s is on the level and curvature factors of the frictionless

Parameter	AFNS		AFNS-R	
	Est.	SE	Est.	SE
$\kappa_{11}^{\mathbb{P}}$	0.0194	0.0473	0.0441	0.0767
$\kappa_{22}^{\mathbb{P}}$	0.3754	0.2020	0.2522	0.1952
$\kappa_{33}^{\mathbb{P}}$	0.4188	0.2578	0.4964	0.2697
$\kappa_{44}^{\mathbb{P}}$	-	-	0.0876	0.1432
σ_{11}	0.0036	0.0000	0.0054	0.0000
σ_{22}	0.0129	0.0002	0.0117	0.0002
σ_{33}	0.0183	0.0003	0.0184	0.0003
σ_{44}	-	-	0.0189	0.0025
$\theta_1^{\mathbb{P}}$	0.0340	0.0322	0.0383	0.0248
$\theta_2^{\mathbb{P}}$	-0.0235	0.0120	-0.0211	0.0156
$\theta_3^{\mathbb{P}}$	-0.0096	0.0139	-0.0209	0.0126
$\theta_4^{\mathbb{P}}$	-	-	-0.0290	0.0426
λ	0.3860	0.0012	0.3245	0.0013
$\kappa_R^{\mathbb{Q}}$	-	-	7.5059	0.9816
$\theta_R^{\mathbb{Q}}$	-	-	0.0002	0.0000
σ_y	0.0006	7.4×10^{-7}	0.0003	1.14×10^{-6}

Table 3: **Estimated Dynamic Parameters**
The table shows the estimated dynamic parameters for the AFNS and AFNS-R models estimated with a diagonal specification of $K^{\mathbb{P}}$ and Σ.

real yield curve. As we demonstrate later, this affects the models' longer-run projections of real rates and hence the estimates of the natural rate. The fourth factor in the AFNS-R model, the bond-specific risk factor, is shown in Figure 7(d). It follows a persistent process with a very stable path near zero for the first 15 years before it experiences a pronounced downward trend during the last 7 years of the sample that leaves it with a significantly negative value at the end of our sample.

4 The OAT€ Bond-Specific Risk Premium

In this section, we analyze the French OAT€ bond-specific risk premia implied by the estimated AFNS-R model described in the previous section. First, we formally define the bond-specific risk premium and study its historical evolution. We then briefly assess its robustness, including its sensitivity to the high-frequency daily data we use. We end the section with an examination of the determinants of the average estimated bond-specific risk premium using regression analysis.

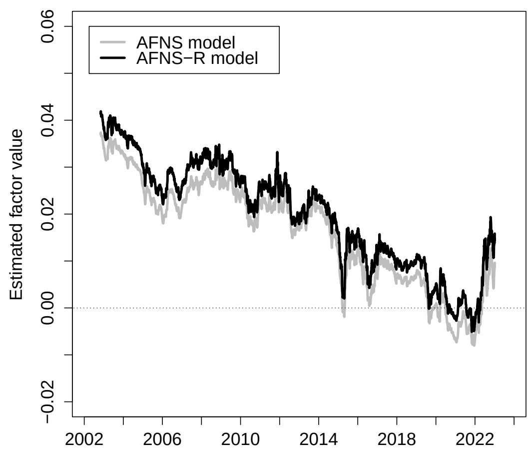

(a) L_t

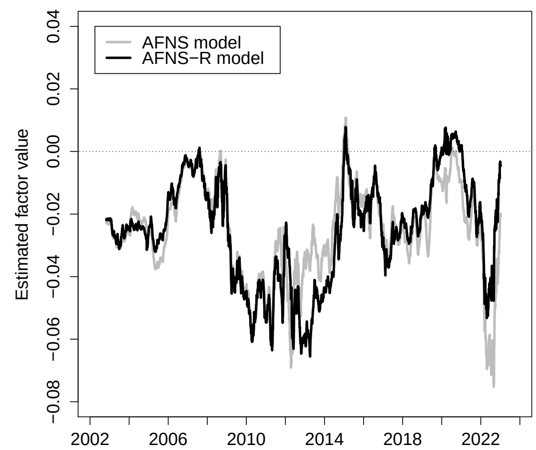

(b) S_t

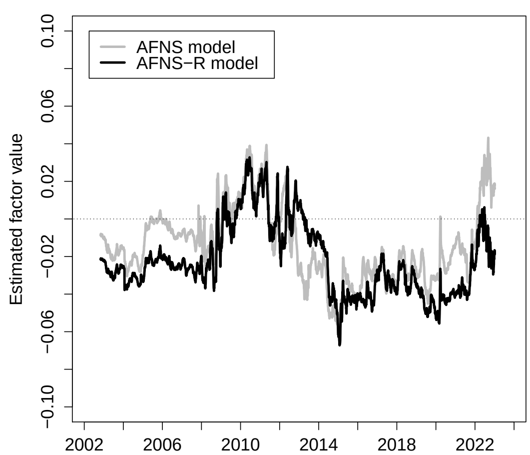

(c) C_t

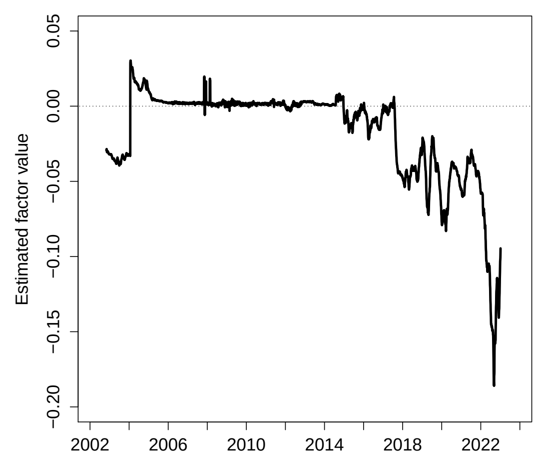

(d) X_t^R

Figure 7: **Estimated State Variables**
Illustration of the estimated state variables from the AFNS and AFNS-R models.

4.1 The Estimated OAT€ Bond-Specific Risk Premia

We now use the estimated AFNS-R model to extract the bond-specific risk premia in the OAT€ market. To compute these premia, we first use the estimated parameters and the filtered states $\left\{X_{t|t}\right\}_{t=1}^{T}$ to calculate the fitted OAT€ prices $\left\{\hat{P}_t^i\right\}_{t=1}^{T}$ for all outstanding OAT€ securities in our sample. These bond prices are then converted into yields to maturity

$\left\{\hat{y}_t^{c,i}\right\}_{t=1}^{T}$ by solving the fixed-point problem

$$\begin{aligned}\hat{P}_t^i &= C(t_1 - t)\exp\left\{-(t_1 - t)\hat{y}_t^{c,i}\right\} + \sum_{k=2}^{n} C\exp\left\{-(t_k - t)\hat{y}_t^{c,i}\right\} \\ &\quad + \exp\left\{-(T - t)\hat{y}_t^{c,i}\right\},\end{aligned} \tag{5}$$

for $i = 1, 2, ..., n_{OATe}$, meaning that $\left\{\hat{y}_t^{c,i}\right\}_{t=1}^{T}$ is approximately the real rate of return on the ith OAT€ if held until maturity (see Sack and Elsasser 2004). To obtain the corresponding yields with correction for the bond-specific risk premia, we compute a new set of model-implied bond prices from the estimated AFNS-R model using only its frictionless part, i.e., using the constraints that $X_{t|t}^R = 0$ for all t as well as $\sigma_{44} = 0$ and $\theta_R^Q = 0$. These prices are denoted $\left\{\tilde{P}_t^i\right\}_{t=1}^{T}$ and converted into yields to maturity $\tilde{y}_t^{c,i}$ using equation (5). They represent estimates of the prices that would prevail in a world without any financial frictions or special demands for certain bonds. The bond-specific risk premium for the ith OAT€ is then defined as

$$\Psi_t^i \equiv \hat{y}_t^{c,i} - \tilde{y}_t^{c,i}. \tag{6}$$

Figure 8 shows the average estimated OAT€ bond-specific risk premium $\bar{\Psi}_t$ across the outstanding OAT€s at each point in time. Note that a negative value means that the fitted OAT€ price is *above* the model-implied frictionless price, i.e., OAT€ prices are higher than they should be in a world without any frictions. Importantly, though, the mean of the series is -0.56 basis point, that is, less than 0.0001 in absolute size. Thus, on average, OAT€ prices are *not* biased by bond-specific risk premia unlike French OATi's, whose prices contain a large convenience premium as documented by CM. That said, there are clearly still some trends and time variation in the series, which explains the standard variation of 9.46 basis points. Furthermore, toward the end of our sample, the average bond-specific premium dropped significantly into negative territory, reaching a historic low of -41.94 basis points on August 31, 2022. Hence, at that point in time, the average OAT€ bond was trading at a significant price or convenience premium. When HICP inflation spiked sharply in 2022, one implication was that bonds like OAT€s, whose principal and cash flows adjust with the changes in the HICP, became very desirable and convenient assets to hold—so much so that investors were willing to give up 0.42 percent in annual return, or equivalently overpay a corresponding amount, to hold these securities. In contrast, it reached its maximum of 37.25 basis points in late 2007, coinciding with a few single-day large spikes. Notably, a large *positive* premium here means that the average OAT€ was trading at a liquidity discount, or at low prices. This is the typical pattern in fixed-income markets when investors are concerned about liquidity and their ability to sell a bond back to the market, and such spells of illiquidity tend to be

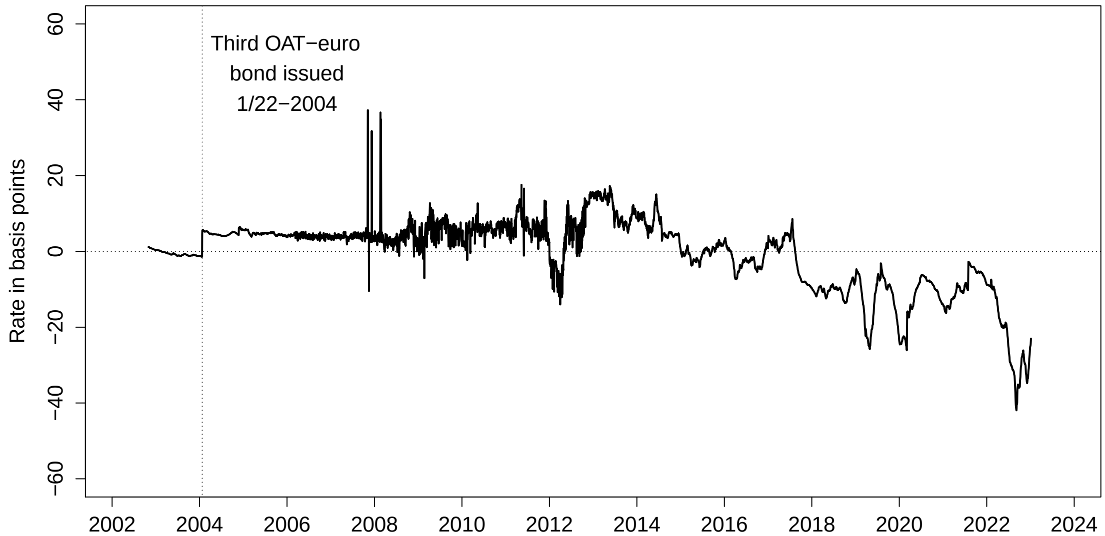

Figure 8: **Average Estimated OAT€ Bond-Specific Risk Premium**
Illustration of the average estimated bond-specific risk premium of French OAT€s for each observation date implied by the AFNS-R model. The bond-specific risk premia are measured as the estimated yield difference between the fitted yield to maturity of individual OAT€s and the corresponding frictionless yield to maturity with the bond-specific risk factor turned off. The data are daily and cover the period from October 31, 2002, to December 30, 2022.

fairly short lived. Thus, the single-day spikes driven by illiquidity events fit that historical pattern well.

Finally, we note the abrupt uptick on January 22, 2004, when the third OAT€ bond was issued and entered our sample. By having pricing information from three bonds instead of two, the model learns that the bond-specific risk premia in the early years of this market most likely were modestly positive. Hence, the estimated bond-specific risk premia prior to January 22, 2004, should be interpreted with caution. This contrasts with the later years in our sample, when our AFNS-R model has sufficient pricing information to identify all four state variables. This makes the bond-specific risk premia very robustly estimated as we demonstrate in Section 4.2.

In Figure 9, we show the individual estimated bond-specific risk premium series for each OAT€ bond. In general, these bonds start out with bond-specific risk premia very close to zero in the first many years of trading. However, as the bonds become seasoned and less traded because the majority of their outstanding notional amount is locked up in buy-and-hold investors' portfolios, their pricing starts to become rather sensitive to market conditions. To demonstrate this, two bonds are highlighted in Figure 9: the OAT€ 1.6% 7/25/2015

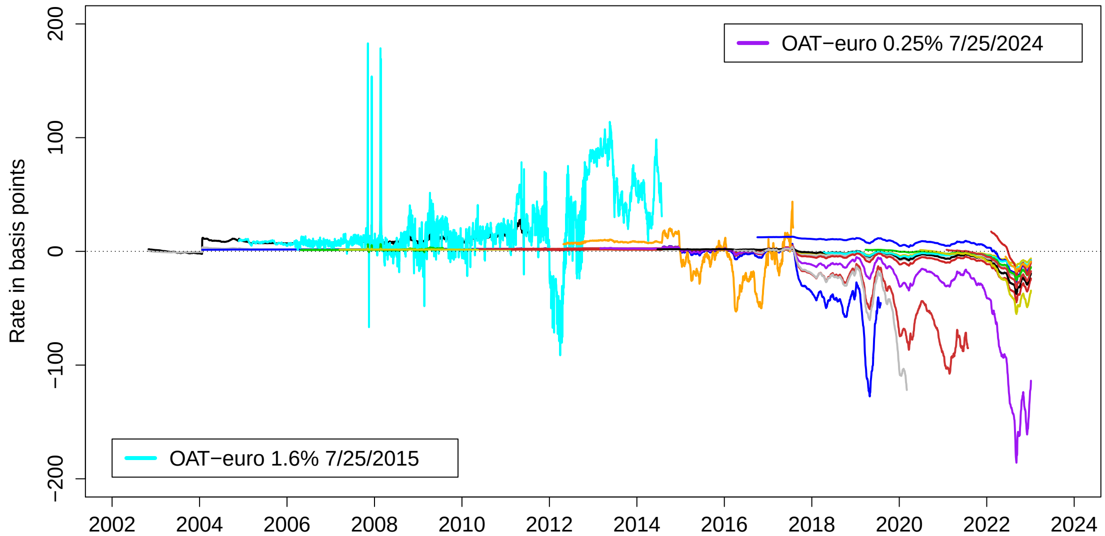

Figure 9: **Individual Estimated OAT€ Bond-Specific Risk Premia**
Illustration of the individual estimated bond-specific risk premia of French OAT€s for each observation date implied by the AFNS-R model. The bond-specific risk premia are measured as the estimated yield difference between the fitted yield to maturity of individual OAT€s and the corresponding frictionless yield to maturity with the bond-specific risk factor turned off. The data are daily and cover the period from October 31, 2002, to December 30, 2022.

that reached maturity during our sample and OAT€ 0.25% 7/25/2024 that was approaching maturity by the end of our sample.

For the OAT€ 1.6% 7/25/2015, which reached this critical phase during the European Sovereign Debt Crisis, we see a sizable and volatile *positive* bond-specific risk premium, meaning there was a material liquidity discount in its pricing in 2012 and 2013. As shown in Figure 6, market conditions for OAT€s as measured by bid-ask spreads were indeed challenging in 2012. Under those circumstances OAT€s approaching maturity are likely to trade at a liquidity discount similar to what the OAT€ 1.6% 7/25/2015 did at the time. Importantly, though, the remaining universe of OAT€s continued to trade with close to zero bond-specific risk premia even during this challenging period.

For the OAT€ 0.25% 7/25/2024, which was approaching maturity towards the end of our sample, we see the opposite outcome, namely a sizable and volatile *negative* bond-specific risk premium, meaning it was trading at a material price premium. That happened in the context of highly elevated inflation well above the ECB's 2 percent target. Under those conditions, inflation-indexed bonds become very convenient assets to hold. As a consequence, the entire outstanding universe of OAT€s was trading at a price premium towards the end of our sample.

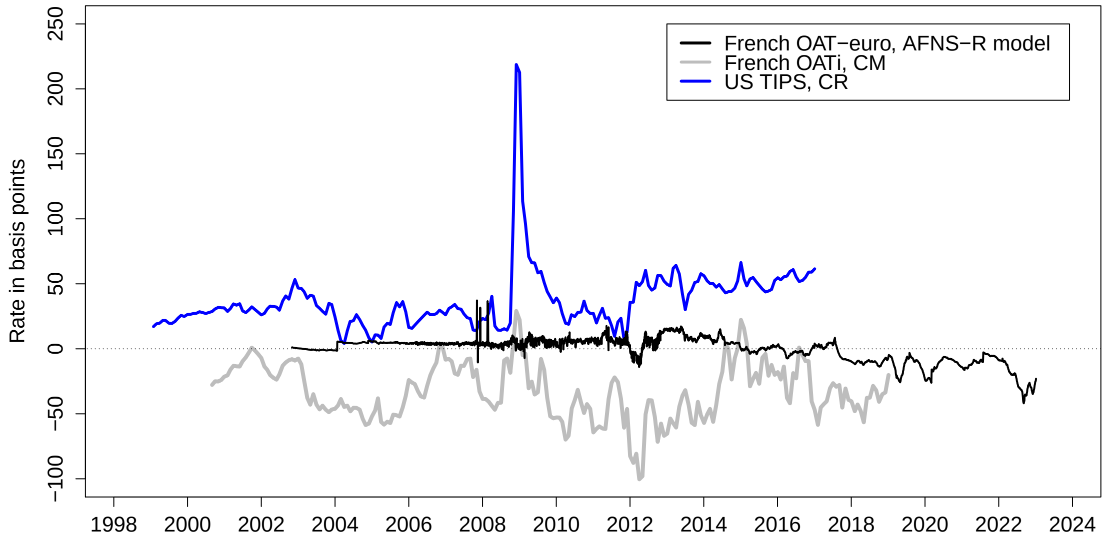

Figure 10: **Comparison of Average Estimated Bond-Specific Risk Premia**
Illustration of the average estimated bond-specific risk premium of French OAT€s implied by the AFNS-R model. Also shown are the average estimated bond-specific risk premium in French OATi yields reported by CM and the average estimated bond-specific risk premium in U.S. TIPS yields reported by CR.

As a final exercise and to put our average estimated bond-specific risk premium from the market for French OAT€s into an international context, we compare it to similar estimates from two other major inflation-indexed bond markets, specifically the market for French OATi's with cash flows adjusted to the French consumer price index examined by CM and the much larger market for U.S. TIPS with cash flows adjusted to the U.S. consumer price index examined by CR. Figure 10 shows the respective average estimated bond-specific risk premium series from all three markets.

We note that U.S. TIPS prices contain a sizable liquidity premium discount, which is well documented in the literature; see ACR, D'Amico et al. (2018), and Pflueger and Viceira (2016), among many others. Cardozo and Christensen (2024) offer a rationale for the illiquidity of inflation-indexed securities like TIPS. By being protected against inflation, indexed securities are inherently less traded than nominal securities. In addition, foreigners not exposed to the domestic price index do not benefit from owning them. Combined this significantly reduces their trading volumes and makes the market for these securities be dominated by patient domestic buy-and-hold investors. This drives up the search frictions in the over-the-counter market for these bonds and leads to a steady-state outcome with their prices containing a large liquidity discount.

In contrast, CM document that French OATi's prices contain a sizable convenience premium averaging close to 0.40 percent. They explain this with the fact that French banks are obliged by law to offer their customers a special type of savings account, known as livret A, the interest of which is tied mechanically to French CPI through a somewhat complicated formula. This creates a regulatory-driven natural demand for OATi bonds as French banks need them to hedge the promised interest payments on these savings accounts.

Based on our average estimated bond-specific risk premium for French OAT€s, this market falls in between these two extremes. On the one hand, there does not seem to be any regulatory-driven benefits of holding OAT€s. As a consequence, there is little reason for them to trade at a convenience premium outside of unusual circumstances with highly elevated inflation when they are obviously convenient assets to hold. On the other hand, by being the largest safe euro-denominated market for bonds indexed to the HICP (ex tobacco), these bonds may be able to attract sufficient demand from non-French investors in the euro area to offset the otherwise negative price dynamics implied by the inherent illiquidity of inflation-indexed bonds.

Overall, our results suggest that the French OAT€ market is a rich and relatively unbiased source of information about bond investors' real rate expectations in the euro area that is not overly influenced by either liquidity discounts or flight-to-safety convenience premia. This makes it an ideal source for our purposes of understanding the trends in the natural rate in the euro area. Moreover, it makes it an ideal input for the construction of breakeven inflation for the euro area, but we leave that task for future research.

To summarize, we feel that the average estimated OAT€ bond-specific risk premium broadly follows a reasonable time series pattern. More importantly, these premia only constitute a minor distortion in the observed OAT€ prices. This provides support for our approach in which we rely on these bond prices for evidence on bond investors' outlook for future real rates in the euro area.

4.2 Robustness Analysis

This section examines the robustness of the average bond-specific risk premium reported in the previous section to some of the main assumptions imposed so far. Throughout the section, the AFNS-R model with diagonal $K^{\mathbb{P}}$ and Σ matrices serves as the benchmark.

First, we assess whether the specification of the dynamics within the AFNS-R model matters for the estimated OAT€ bond-specific risk premium. To do so, we estimate the AFNS-R model with unconstrained dynamics, that is, the AFNS-R model with unrestricted $K^{\mathbb{P}}$ and lower triangular Σ matrix. Figure 11 shows the estimated OAT€ bond-specific risk premium from this estimation and compares it to the series produced by our benchmark

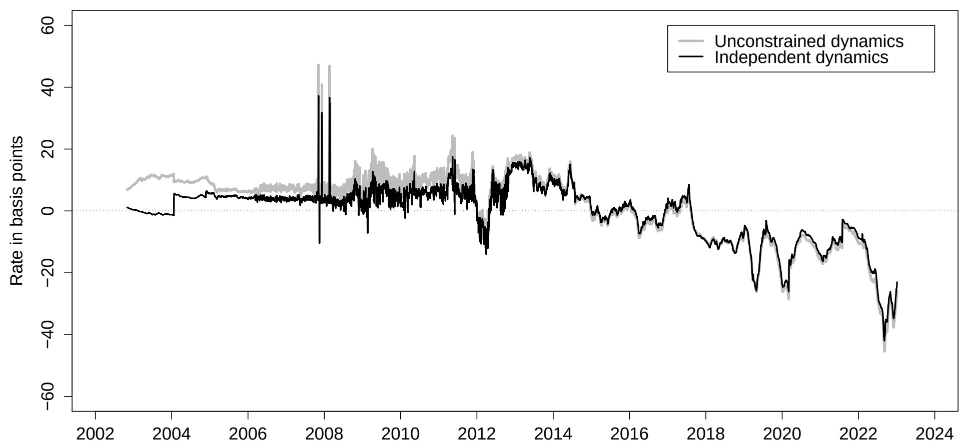

Figure 11: **Average Estimated OAT€ Bond-Specific Risk Premium: Alternative $\mathbb{P}$ Dynamics**
Illustration of the average estimated bond-specific risk premium of French OAT€s for each observation date implied by the AFNS-R model when estimated with unconstrained dynamics as detailed in the text instead of independent factor dynamics. In both cases, the bond-specific risk premia are measured as the estimated yield difference between the fitted yield to maturity of individual OAT€s and the corresponding frictionless yield to maturity with the bond-specific risk factor turned off.

model. Note that they are barely distinguishable. Thus, we conclude that the specification of the dynamics within the AFNS-R model only play a very modest role for the estimated bond-specific risk premia, which is consistent with the findings of ACR in the context of U.S. TIPS.

Second, we assess whether the data frequency plays any role for our results. To do so, we estimate the AFNS-R model using daily, weekly, monthly, and quarterly data, and based on the results above it suffices to focus on the most parsimonious AFNS-R model with diagonal $K^{\mathbb{P}}$ and Σ matrices. Figure 12 shows the average estimated OAT€ bond-specific risk premium series from all four estimations. Note that they are barely distinguishable during the last decade of our sample, while there are some notable discrepancies during the first decade of our sample between the high-frequency daily and weekly series, on one hand, and the lower-frequency monthly and quarterly series, on the other.

As to the importance of these early discrepancies, we stress that, in explaining the large swings in OAT€ yields observed in Figure 3, the relatively minor differences between the high- and low-frequency series during the first 10 years of the sample are clearly *not* the source of those declines.

At a technical level, the issue is that, at low frequency, some variation in the OAT€ yields

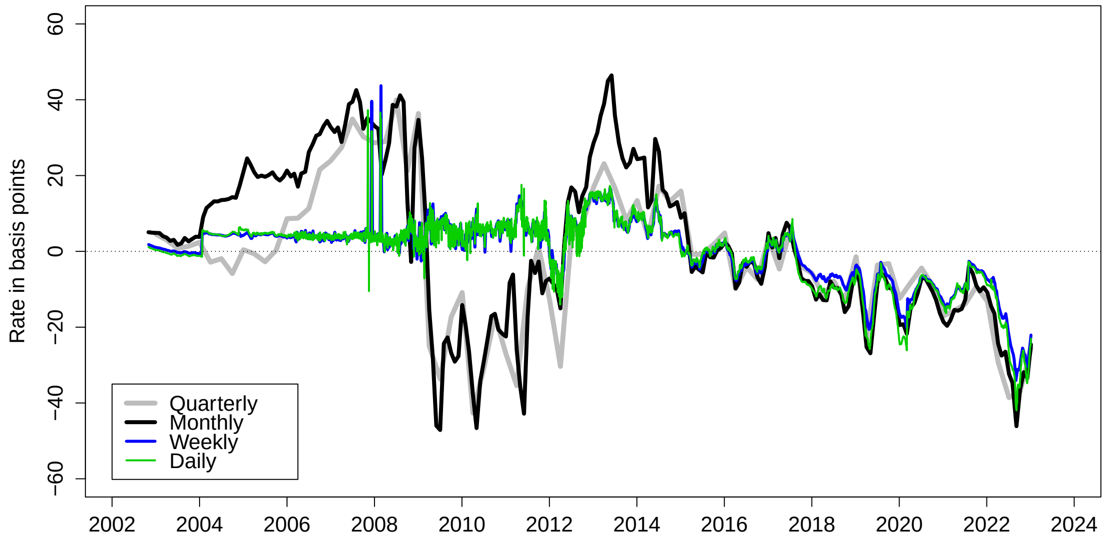

Figure 12: **Average Estimated OAT€ Bond-Specific Risk Premium: Data Frequency**
Illustration of the average estimated bond-specific risk premium of French OAT€s for each observation date implied by the AFNS-R model when estimated using daily, weekly, monthly, and quarterly data. In all cases, the bond-specific risk premia are measured as the estimated yield difference between the fitted yield to maturity of individual OAT€s and the corresponding frictionless yield to maturity with the bond-specific risk factor turned off.

gets ascribed to the nonfundamental bond-specific risk premia that, at higher daily or weekly frequency, the AFNS-R model is able to tell should go into the fundamental frictionless level, slope, and curvature factors. Given that the ideal is to have as much of the bond yield variation explained by the fundamental level, slope, and curvature factors rather than bond-specific risks, these findings provide one justification for us to prefer the implementation based on high-frequency daily data over the more conventional monthly data frequently considered in the literature, despite the significantly higher computational costs.

Third, we assess whether the data censoring choice matters for our results. To do so, we estimate the AFNS-R model using alternative data cutoffs: No cutoff (i.e. 0 months), 6 months, 18 months, 24 months, and 30 months in addition to our benchmark choice of using 12 months as the censoring point for OAT€ bonds approaching maturity. We note that we perform this exercise for our preferred AFNS-R model to be described in Section 5.2, but we stress that the results are not sensitive to this choice as demonstrated by the results above. Figure 13 shows the average estimated OAT€ bond-specific risk premium series from all six estimations. In general, except when we do not impose any cutoff, the estimated risk premium

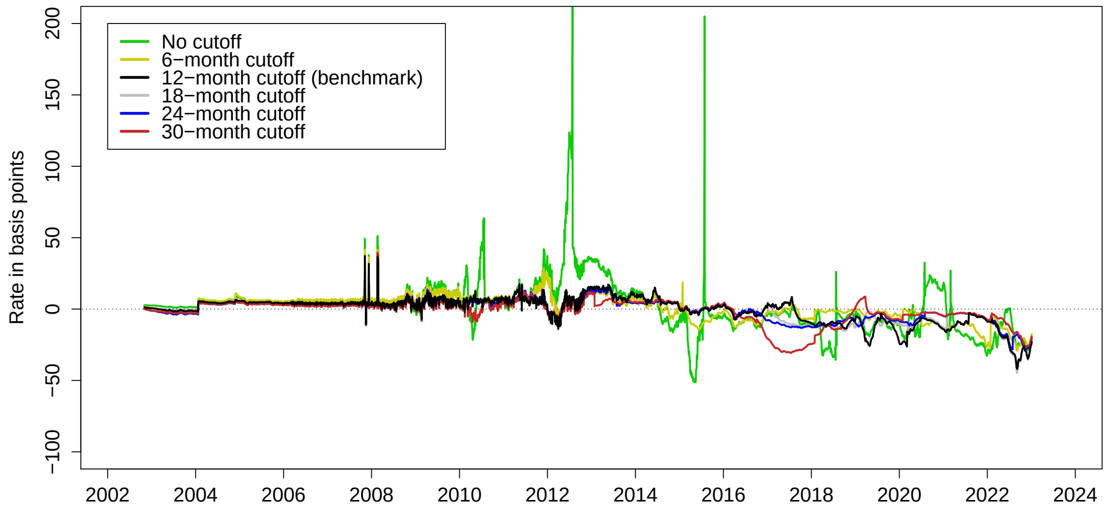

Figure 13: **Average Estimated OAT€ Bond-Specific Risk Premium: Data Cutoff**
Illustration of the average estimated bond-specific risk premium of French OAT€s for each observation date implied by the AFNS-R model when estimated using daily, weekly, monthly, and quarterly data. In all cases, the bond-specific risk premia are measured as the estimated yield difference between the fitted yield to maturity of individual OAT€s and the corresponding frictionless yield to maturity with the bond-specific risk factor turned off.

series are very similar and close to each other. Furthermore, the cutoff choice matters little during the first 6-7 years of our sample as no bonds are approaching maturity early on in our sample. In choosing a preferred cutoff point, there is a tension between, on the one hand, keeping as much information as possible, and at the same time limit the impact of noisy observations on the estimation results. We think of our benchmark choice to use a 12-month cutoff similar to ACR as striking a sensible balance between these two considerations for our specific data sample.

4.3 Determinants of the Bond-Specific Risk Premia

In this section, we explore which factors matter for the size of the bond-specific risk premia in the OAT€ prices. To explain the variation of the average estimated bond-specific risk premium series, we run regressions with it as the dependent variable and a wide set of explanatory variables that are thought to play a role for the bond-specific risk premia as explained in the following.

To begin, we are interested in the role of factors that are believed to matter for OAT€ market liquidity specifically or bond market liquidity more broadly as they could matter for the estimated bond-specific risk premia. First, we include the average OAT€ bond age and the one-month realized volatility of the 10-year OAT€ bond yield as proxies for OAT€ bond liquidity following the work of Houweling et al. (2005).[16] Inspired by the analysis of Hu et al. (2013), we also include a noise measure of OAT€ bond prices to control for variation in the amount of arbitrage capital available in this market.[17] Finally, we add the euro overnight interbank rate to proxy for the opportunity cost of holding money and the associated liquidity premia of French government bonds, including OAT€ bonds, as explained in Nagel (2016). Combining these four explanatory variables tied to market liquidity and functioning produces the results reported in regression (1) in Table 4. We note a relatively modest adjusted R^2 of 0.30. The average OAT€ bond age, the one-month realized volatility of the ten-year OAT€ bond yield, and the overnight rate all have statistically significant negative coefficients. This implies that an increase in the liquidity risk of OAT€ bonds is associated with lower average estimated bond-specific risk premia. Moreover, the noise measure, which serves as a proxy for financial frictions in the market for OAT€s, has a positive, but insignificant coefficient in this regression. Hence, we take these results to show that our average estimated bond-specific risk premia in the OAT€ prices behave more like convenience premia than liquidity premia.

After having explored the role of liquidity factors, we examine the effects of factors reflecting risk sentiment domestically and globally on the average estimated bond-specific risk premia. This set of variables includes the VIX, which represents near-term uncertainty about the general stock market as reflected in options on the Standard & Poor's 500 stock price index and is widely used as a gauge of investor fear and risk aversion. The set also contains the yield difference between seasoned (off-the-run) U.S. Treasury securities and the most recently issued (on-the-run) U.S. Treasury security of the same ten-year maturity. This on-the-run (OTR) premium is a frequently used measure of financial frictions in the U.S. Treasury market. To control for factors related to the uncertainty about the interest rate environment, we include the MOVE index. The fourth variable is the U.S. TED spread, which is calculated as the difference between the three-month U.S. LIBOR and the three-month U.S. T-bill interest rate. This spread represents a measure of the perceived general credit risk in global financial markets. As an additional indicator of credit risk and credit risk sentiment, we use the composite measure of the credit risk of French government bonds shown in Figure 5. The next variable in the set is the ten-year U.S. Treasury yield from the Federal Reserve's H.15 database, which is included to control for reach-for-yield effects in advanced economies. This may be particularly relevant for our sample during the period between December 2008 and

[16] The ten-year OAT€ bond yield is the ten-year fitted real yield implied by the AFNS model estimated using our sample of daily OAT€ prices.

[17] The noise measure is the mean absolute fitted error from the estimated daily AFNS model, where each error is calculated as the difference between the observed OAT€ price converted into yield to maturity and the fitted OAT€ price also converted into yield to maturity.

Explanatory variables	(1)	(2)	(3)
Avg. bond age (yrs)	-3.039***		-4.367***
	(0.379)		(0.391)
One-month realized volatility of ten-year real yield (bps)	-0.203***		-0.233***
	(0.068)		(0.054)
Bond noise measure (bps)	0.385		-0.780***
	(0.267)		(0.284)
Overnight rate (%)	-0.708*		-1.421**
	(0.420)		(0.565)
VIX (%)		-0.322***	-0.070
		(0.068)	(0.059)
Ten-year OTR premium (bps)		0.698***	0.525***
		(0.116)	(0.070)
MOVE Index (bps)		-0.141***	-0.114***
		(0.039)	(0.022)
TED spread (bps)		0.024*	0.036***
		(0.013)	(0.011)
Composite credit risk measure (bps)		0.068***	0.055***
		(0.013)	(0.016)
Ten-year US Treasury yield (%)		1.960***	-2.738***
		(0.422)	(0.609)
WTI ($)		0.054***	0.159***
		(0.018)	(0.019)
Constant	17.081***	-1.753	29.015***
	(2.567)	(3.189)	(3.542)
N	4838	4838	4838
Adjusted R^2	0.30	0.27	0.53

Table 4: **Regression Results for Average Estimated OAT€ Bond-Specific Risk Premium**
The table reports the results of regressions with the average estimated bond-specific risk premium of French OAT€s as the dependent variable and 11 explanatory variables. Standard errors computed by the Newey-West estimator (with 13 lags) are reported in parentheses. Asterisks *, ** and *** indicate significance at the 10 percent, 5 percent, and 1 percent levels, respectively.

December 2015 and again in the 2020-2021 period when U.S. short-term interest rates were constrained by the zero lower bound. Finally, we include the West Texas Intermediate (WTI) Cushing crude oil price to proxy for energy prices, which represent a significant risk to the inflation outlook in many countries around the world, including many euro area member states. The results of the regression with these seven explanatory variables is reported in regression (2) in Table 4. This produces a modest adjusted R^2 of 0.27. We note that all seven variables have some explanatory power as their estimated coefficients are all statistically significant.

To assess the robustness of the results from the first two regressions, we include all variables with the results reported in column (3) in Table 4. This joint regression produces a high adjusted R^2 of 0.53. The significant increase in the adjusted R^2 suggests that there is little overlap between the two sets of explanatory variables. The first set is squarely focused on the liquidity in the OAT€ market, while the second set represents global risk sentiment and flight-to-safety effects.

With the systematic negative coefficients on the liquidity risk variables—and on the VIX and the MOVE Index in the full regression model—we feel that we can confidently reject the conjecture that our average estimated bond-specific risk premia in the OAT€ prices should represent liquidity risk premia. This is also consistent with the interpretation we offered in Section 4.1, when we contrasted our estimates with the estimated U.S. TIPS liquidity premia reported by CR. Hence, the trading dynamics in the OAT€ market seem to be fundamentally different from those prevailing in the U.S. TIPS market.

Finally, changes in perceived credit risk as reflected in either the TED spread or our composite credit risk measure are significantly positively correlated with changes in the bond-specific risk premia. Based on these results we conclude that some part of the bond-specific risk premia seems to reflect compensation for credit risk.

5 A New Normal for Euro-Area Interest Rates?

In this section, we first go through a careful model selection process to find a preferred specification of the AFNS-R model's objective $\mathbb{P}$-dynamics. We then use this AFNS-R model to account for bond-specific risk and standard term premia in the OAT€ prices and obtain expected real short rates and the associated measure of the natural rate. Finally, we compare this estimate to other market- and macro-based estimates from the literature and consider model projections to assess its likely path going forward.

5.1 Definition of the Natural Rate

Our working definition of the natural rate of interest r_t^* is

$$r_t^* = \frac{1}{5}\int_{t+5}^{t+10} E_t^{\mathbb{P}}[r_s^R]ds, \tag{7}$$

that is, the average expected real short rate over a five-year period starting five years ahead, where the expectation is with respect to the objective $\mathbb{P}$-probability measure. As noted in the introduction, this 5yr5yr forward average expected real short rate should be little affected by short-term transitory shocks. Alternatively, r_t^* could be defined as the expected real short-term interest rate at an infinite horizon. However, this quantity will depend crucially on whether the factor dynamics exhibit a unit root. As is well known, the typical spans of time series data that are available do not distinguish strongly between highly persistent stationary processes and nonstationary ones. Our model follows the finance literature and adopts the former structure, so strictly speaking, our infinite-horizon steady-state expected real rate is constant. However, we view our data sample as having insufficient information in the ten-year to infinite horizon range to definitively pin down that steady state, so we prefer our definition with a medium- to long-run horizon. Moreover, we examine the sensitivity of our results to using alternative integration intervals in the definition r_t^* and find them to be robust.

5.2 Model Selection

For estimation of the natural rate and associated real term premia, the specification of the mean-reversion matrix $K^{\mathbb{P}}$ is crucial as noted earlier. To select the best-fitting specification of the model's real-world dynamics, we use a general-to-specific modeling strategy in which the least significant off-diagonal parameter of $K^{\mathbb{P}}$ is restricted to zero and the model is re-estimated. This strategy of eliminating the least significant coefficient is carried out down to the most parsimonious specification, which has a diagonal $K^{\mathbb{P}}$ matrix. The final specification choice is based on the value of the Bayesian information criterion (BIC), as in Christensen et al. (2014).[18]

The summary statistics of the model selection process are reported in Table 5. The BIC is minimized by specification (12), which has a $K^{\mathbb{P}}$-matrix given by

$$K_{BIC}^{\mathbb{P}} = \begin{pmatrix} \kappa_{11}^{\mathbb{P}} & 0 & 0 & 0 \\ 0 & \kappa_{22}^{\mathbb{P}} & \kappa_{23}^{\mathbb{P}} & 0 \\ 0 & 0 & \kappa_{33}^{\mathbb{P}} & 0 \\ 0 & 0 & 0 & \kappa_{44}^{\mathbb{P}} \end{pmatrix}.$$

[18]The Bayesian information criterion is defined as $\text{BIC} = -2\log L + k\log T$, where k is the number of model parameters and $T = 5{,}258$ is the number of daily data observations.

Alternative specifications	Goodness of fit statistics			
	$\log L$	k	p-value	BIC
(1) Unrestricted $K^{\mathbb{P}}$	234,604.3	65	n.a.	-468,651.7
(2) $\kappa_{34}^{\mathbb{P}} = 0$	234,603.0	64	0.11	-468,657.7
(3) $\kappa_{34}^{\mathbb{P}} = \kappa_{12}^{\mathbb{P}} = 0$	234,602.3	63	0.24	-468,664.8
(4) $\kappa_{34}^{\mathbb{P}} = \kappa_{12}^{\mathbb{P}} = \kappa_{31}^{\mathbb{P}} = 0$	234,601.7	62	0.27	-468,671.6
(5) $\kappa_{34}^{\mathbb{P}} = \ldots = \kappa_{43}^{\mathbb{P}} = 0$	234,600.7	61	0.16	-468,678.6
(6) $\kappa_{34}^{\mathbb{P}} = \ldots = \kappa_{42}^{\mathbb{P}} = 0$	234,599.5	60	0.12	-468,684.9
(7) $\kappa_{34}^{\mathbb{P}} = \ldots = \kappa_{32}^{\mathbb{P}} = 0$	234,598.8	59	0.24	-468,692.1
(8) $\kappa_{34}^{\mathbb{P}} = \ldots = \kappa_{41}^{\mathbb{P}} = 0$	234,593.8	58	< 0.01	-468,690.7
(9) $\kappa_{34}^{\mathbb{P}} = \ldots = \kappa_{14}^{\mathbb{P}} = 0$	234,589.9	57	< 0.01	-468,691.5
(10) $\kappa_{34}^{\mathbb{P}} = \ldots = \kappa_{13}^{\mathbb{P}} = 0$	234,584.9	56	< 0.01	-468,690.0
(11) $\kappa_{34}^{\mathbb{P}} = \ldots = \kappa_{21}^{\mathbb{P}} = 0$	234,579.1	55	< 0.01	-468,687.0
(12) $\kappa_{34}^{\mathbb{P}} = \ldots = \kappa_{24}^{\mathbb{P}} = 0$	234,577.4	54	0.07	**-468,692.2**
(13) $\kappa_{34}^{\mathbb{P}} = \ldots = \kappa_{23}^{\mathbb{P}} = 0$	234,570.8	53	< 0.01	-468,687.5

Table 5: **Evaluation of Alternative Specifications of the AFNS-R Model**
There are 13 alternative estimated specifications of the AFNS-R model. Each specification is listed with its maximum log likelihood ($\log L$), number of parameters (k), the p-value from a likelihood ratio test of the hypothesis that it differs from the specification above with one more free parameter, and the Bayesian information criterion (BIC). The period analyzed covers daily data from October 31, 2002, to December 30, 2022.

$K^{\mathbb{P}}$	$K_{\cdot,1}^{\mathbb{P}}$	$K_{\cdot,2}^{\mathbb{P}}$	$K_{\cdot,3}^{\mathbb{P}}$	$K_{\cdot,4}^{\mathbb{P}}$	$\theta^{\mathbb{P}}$		Σ
$K_{1,\cdot}^{\mathbb{P}}$	0.0448	0	0	0	0.0388	σ_{11}	0.0054
	(0.0785)				(0.0249)		(0.0000)
$K_{2,\cdot}^{\mathbb{P}}$	0	1.0132	0.8448	0	-0.0260	σ_{22}	0.0117
		(0.3200)	(0.2718)		(0.0105)		(0.0002)
$K_{3,\cdot}^{\mathbb{P}}$	0	0	0.5067	0	-0.0200	σ_{33}	0.0184
			(0.2678)		(0.0123)		(0.0003)
$K_{4,\cdot}^{\mathbb{P}}$	0	0	0	0.0817	-0.0287	σ_{44}	0.0189
				(0.1420)	(0.0441)		(0.0025)

Table 6: **Estimated Dynamic Parameters of the Preferred AFNS-R Model**
The table shows the estimated parameters of the $K^{\mathbb{P}}$ matrix, $\theta^{\mathbb{P}}$ vector, and diagonal Σ matrix for the preferred AFNS-R model according to the BIC. The estimated value of λ is 0.3245 (0.0013), while $\kappa_R^{\mathbb{Q}} = 7.5228$ (0.9785), and $\theta_R^{\mathbb{Q}} = 0.0002$ (0.0000). The maximum log likelihood value is 234,577.4. The numbers in parentheses are the estimated parameter standard deviations.

This specification shows that the model's $\mathbb{P}$-dynamics preferred by the data have a structure similar to the one assumed under the risk-neutral $\mathbb{Q}$-dynamics used for pricing to achieve the Nelson-Siegel factor loading structure, which is comforting.

The estimated parameters of the preferred specification are reported in Table 6. The estimated $\mathbb{Q}$-dynamics used for pricing and determined by (Σ, λ, $\kappa_R^{\mathbb{Q}}$, $\theta_R^{\mathbb{Q}}$) are very close to those reported in Table 3 for the AFNS-R model with diagonal $K^{\mathbb{P}}$. This implies that both model fit and the estimated OAT€ bond-specific risk premia from the preferred AFNS-R model are very similar to those already reported and therefore not shown. Furthermore, the estimated objective $\mathbb{P}$-dynamics in terms of $\theta^{\mathbb{P}}$ and Σ are also qualitatively similar to those reported in Table 3.

Still, to understand the role played by the mean-reversion matrix $K^{\mathbb{P}}$ for estimates of the natural rate, we will later analyze the most flexible model with unrestricted mean-reversion matrix $K^{\mathbb{P}}$ and the most parsimonious model with diagonal $K^{\mathbb{P}}$, in addition to our preferred specification described above.

5.3 Estimates of the Natural Rate

Our market-based measure of the natural rate is the average expected real short rate over a five-year period starting five years ahead. This 5yr5yr forward average expected real short rate should be little affected by short-term transitory shocks and well positioned to capture the persistent trends in the natural rate.

To illustrate the decomposition underlying our definition of r_t^*, recall that the real term premium is defined as

$$TP_t(\tau) = y_t(\tau) - \frac{1}{\tau}\int_t^{t+\tau} E_t^{\mathbb{P}}[r_s]ds.$$

That is, the real term premium is the difference in expected real returns between a buy-and-hold strategy for a τ-year real bond and an instantaneous rollover strategy at the risk-free real rate r_t. Note that $y_t(\tau)$ in this definition is the *frictionless* yield clean of any bond-specific risk premia. Figure 14 shows the preferred AFNS-R model decomposition of the 5yr5yr forward frictionless real yield based on this definition. The solid gray line is the 5yr5yr forward real term premium, which, although volatile, has fluctuated around a fairly stable level since the early 2000s. As suggested by theory, this premium is countercyclical and elevated during economic recessions. In contrast, the estimate of the natural rate of interest implied by the AFNS-R model—the black line—shows a gradual decline from above 1.5 percent in the early 2000s to well below -1.5 percent by late 2021, with a partial retracing of that decline during the last year of our sample. Importantly, the vast majority of the persistent trends in the 5yr5yr forward real yield is driven by similar trends in this measure of r_t^*.

To examine the sensitivity of our r_t^* estimate to our choice to define r_t^* as the average expected real short rate over a five-year period starting five years ahead, we consider three alternative definitions that all embed a longer view about the time it takes for the euro area economy to reach steady state. The first assumes that this takes seven years and then measures the neutral real rate as the average expected real short over the following three years. It is referred to as the 7yr3yr r_t^* estimate. The second alternative takes an even longer view and assumes that it takes nine years to reach steady state and then measures the neutral

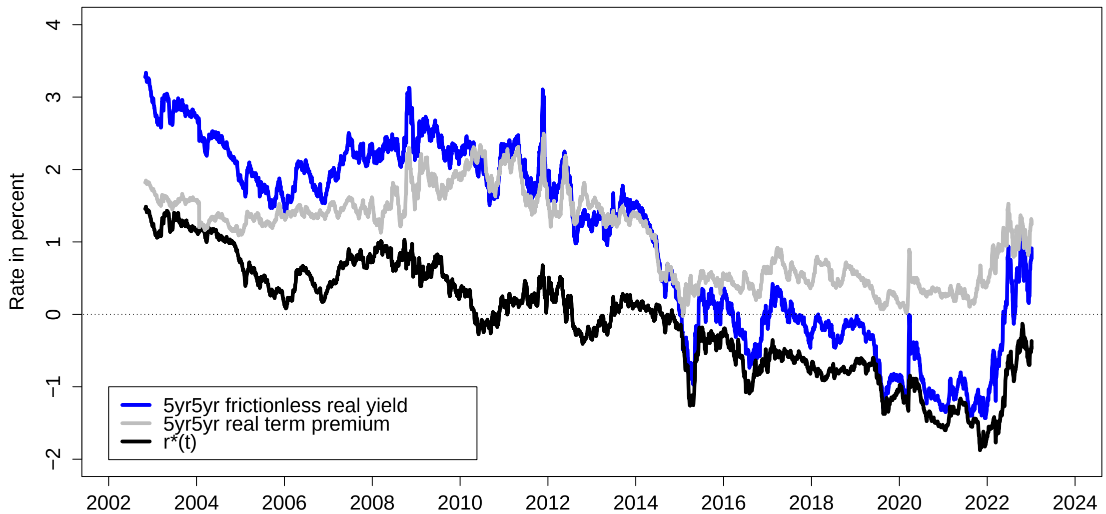

Figure 14: **AFNS-R Model 5yr5yr Real Yield Decomposition**

real rate as the average expected real short rate over a short one-year period. It is referred to as the 9yr1yr r_t^* estimate. Finally, the third alternative takes the longest view and assumes that it takes ten years for the economy to reach steady state and then measures r_t^* as the average expected real short rate over a five-year horizon as in our benchmark definition. It is referred to as the 10yr5yr r_t^* estimate. Figure 15 shows all four r_t^* estimates, which are very close to each other thanks to the high estimated persistence of the state variables within our preferred AFNS-R model. Hence, our r_t^* estimate of the neutral real rate for the euro area has very little sensitivity to the specific definition of r_t^* used. Thus, our reported results are very robust from that perspective.

To assess the sensitivity of our r_t^* estimate to the specification of the mean-reversion matrix $K^{\mathbb{P}}$, we compare it in Figure 16 to the estimates from the AFNS-R models with unrestricted and diagonal $K^{\mathbb{P}}$ matrix, respectively. As noted in Figure 16, our r_t^* estimate is indeed very sensitive to this model choice, but parsimonious specifications like our preferred AFNS-R model specification favored by the data tend to give fairly similar r_t^* estimates. Still, these results demonstrate how insignificant off-diagonal parameters in the specification of the mean-reversion $K^{\mathbb{P}}$ matrix can materially distort estimates of r_t^*. Hence, the results underscore the importance of our careful model selection procedure needed to identify appropriate specifications of $K^{\mathbb{P}}$ supported by the bond price data.

The effect on the estimated natural rate from accounting for the bond-specific risk premia in OAT€ prices is the subject of Figure 17. The black line is the estimate of r_t^* from the

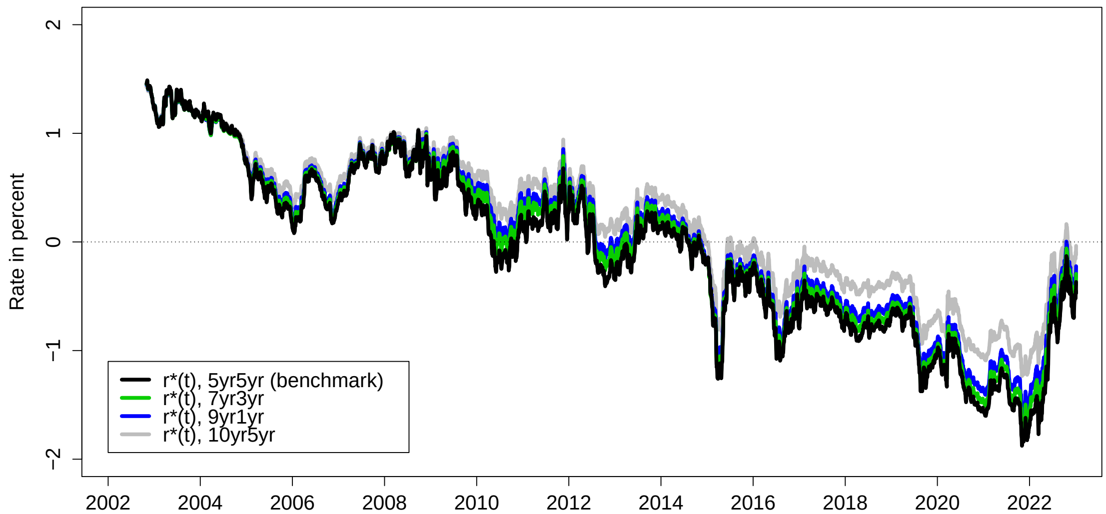

Figure 15: **The Sensitivity of r* Estimate to Alternative Definitions**

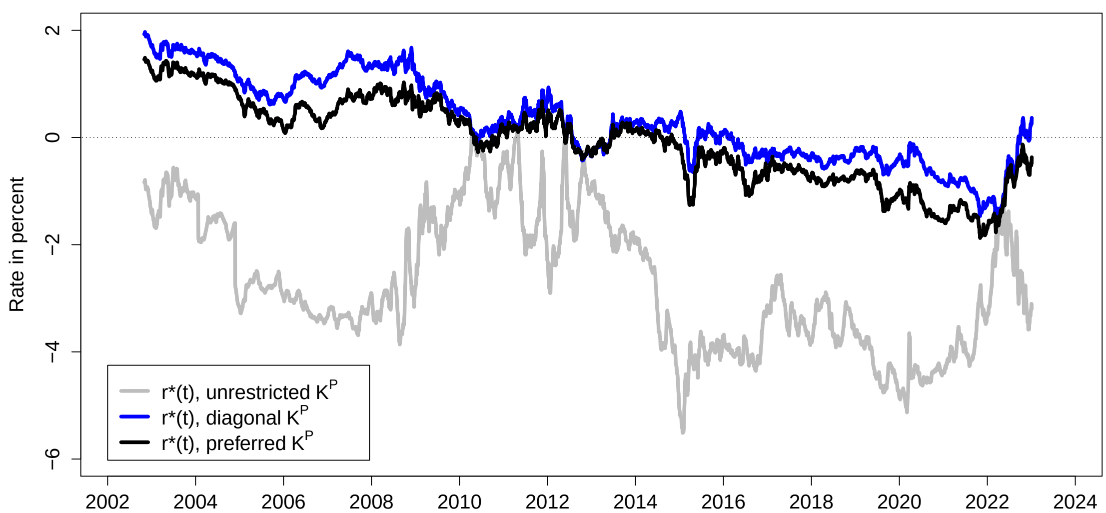

Figure 16: **The Sensitivity of r* Estimate to $K^{\mathbb{P}}$ Specification**

AFNS-R model, while the gray line is the estimate from the AFNS model, which does not account for time-varying bond-specific risk premium effects in OAT€ prices.[19] Accounting

[19]For the AFNS model, we also go through a careful model selection process and use the BIC to determine a preferred specification, as described in online Appendix A.

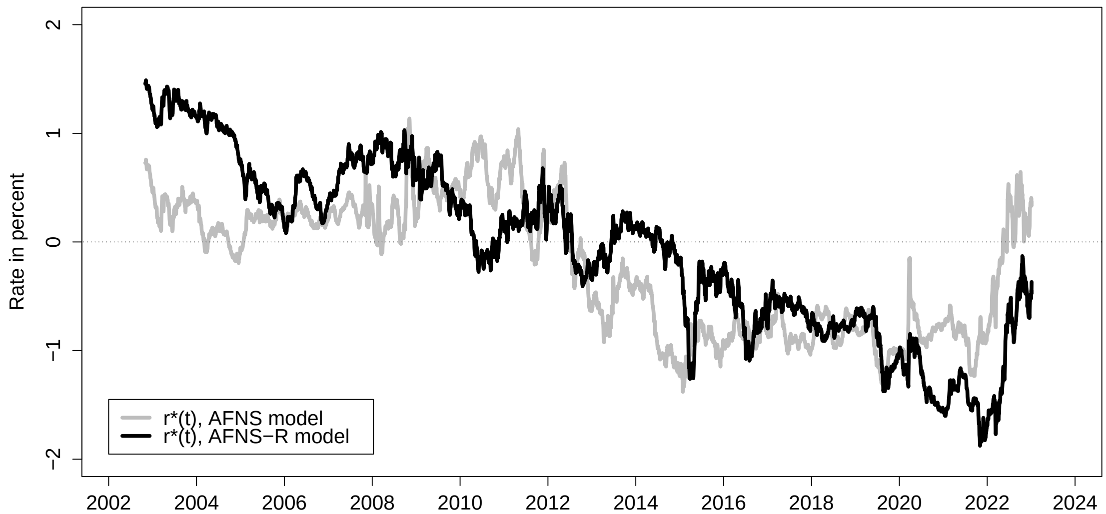

Figure 17: **Effect of the Bond-Specific Risk Adjustment on Estimates of r***

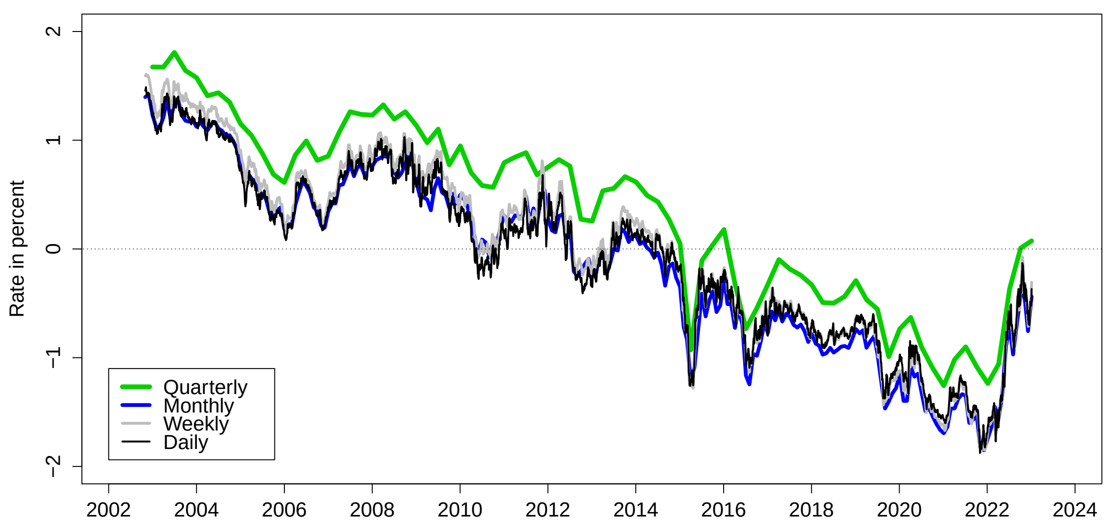

Figure 18: **The Sensitivity of r* Estimate to Data Frequency**

for the bond-specific risk premia in OAT€ prices leads to a persistent and diverging difference in the two natural rate estimates. Thus, even though both average close to zero during our sample period, it is crucial to account for the bond-specific risk premia to produce reliable estimates of the natural rate of interest.

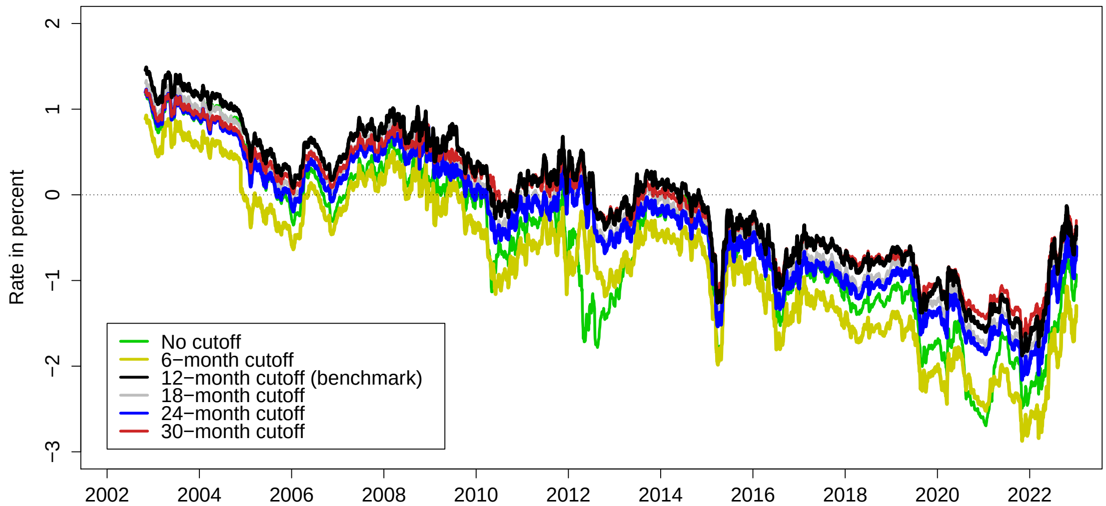

Figure 19: **The Sensitivity of r* Estimate to Data Cutoff**

The role of the data frequency is examined in Figure 18, which shows the r_t^* estimates implied by our preferred AFNS-R model estimated at daily, weekly, monthly, and quarterly frequency. The results show that our estimate has little sensitivity to our choice to focus on high-frequency daily data. This also underscores the usefulness of our model for real-time analysis as we also demonstrate later on in Section 6.2.

As a final robustness check, we vary the data cutoff used to censor the data for each OAT€ that approaches maturity from zero months (i.e. no censoring of any observations) up to 30 months in six month increments with the 12-month cutoff being our benchmark. Figure 19 shows the resulting six r_t^* estimates. We note that our r_t^* estimate is sensitive to the cutoff choice to some extent. In general, there is a pattern whereby a later cutoff—meaning more data is kept and included in the estimation—leads to lower r_t^* estimates. This is explained by the fact that the OAT€ yield curve is generally upward sloping most of the time. As a consequence, keeping short-term OAT€s in the sample implies that the estimated real short rate r_t, which is the launch point for the projections underlying our definition of r_t^*, will tend to be lower all else being equal. In addition, there is a mild tendency for the r_t^* estimates to be slightly more volatile as we lower the cutoff point towards zero. Overall, we consider our choice to use a 12-month cutoff as recommended by ACR to strike a sensible balance between including as much data as possible and the stability of our r_t^* estimate.

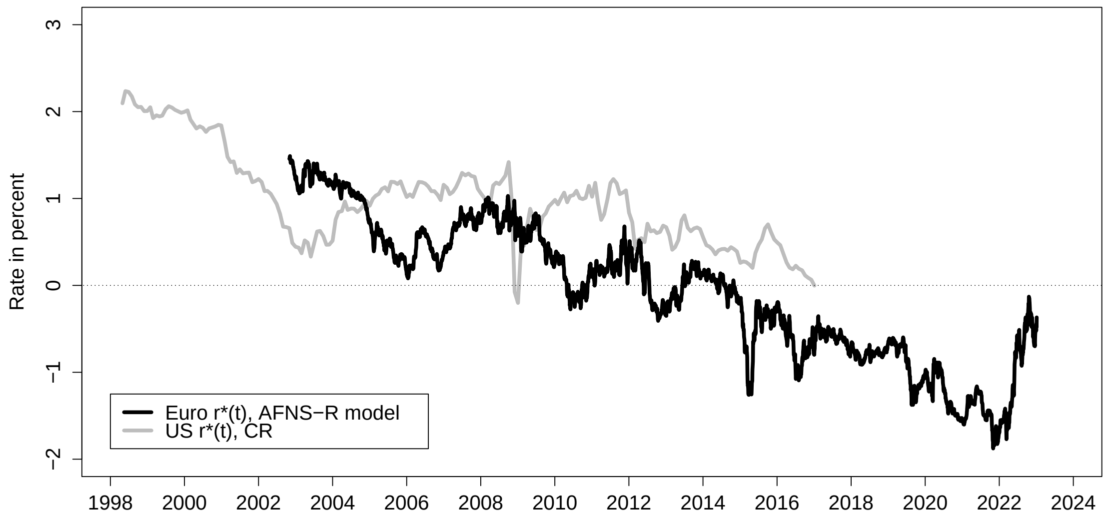

Figure 20: **Comparison with Foreign Market-Based Estimate of r***

5.4 Comparison of Estimates of the Natural Rate

In this section, we compare our estimate of the natural rate to other existing estimates of the natural interest rate in the literature. To start, we compare our r_t^* estimate from the preferred AFNS-R model to the U.S. market-based estimate reported by CR using solely the prices of U.S. TIPS. These two market-based estimates of the natural rate are shown in Figure 20. Their high positive correlation and similar downward trend are both evident. Also, they share the common feature, that their most pronounced declines over the past two decades happened before and after, but not *during* the GFC. These observations combined suggest that the factors depressing U.S. and euro-area interest rates are likely to be global in nature and are not particularly tied to developments surrounding the GFC.

Now, we turn to the crucial comparison of our finance-based estimate of r_t^* with estimates based on macroeconomic data. Figure 21 shows the r_t^* estimate from our preferred AFNS-R model, along with the macro-based estimate of r_t^* from Holston et al. (2017, henceforth HLW), which is the filtered estimate generated by applying the approach described in Laubach and Williams (2003) to euro-area macroeconomic series. The r_t^* estimate from HLW starts in 1972. However, until the onset of the GFC, this macro-based estimate appears to be stationary and remains close to 2.5 percent the whole time. This is consistent with the received wisdom of that era in monetary economics that viewed the natural rate as effectively constant—for example, as assumed in the large Taylor rule literature. It is only in the aftermath of the GFC that we see a persistent large downward movement in the macro-based r_t^* estimate, which is much later and smaller than the sizable drop in our market-based estimate. Importantly, at the end of our sample, this macro-based estimate is -0.68 percent and hence close to our market-based estimate.

The second macro-based estimate of r_t^* is taken from Del Negro et al. (2019, henceforth DGGT). They estimate a flexible vector autoregression model with common trends for a sample of data from 7 advanced economies, including Germany, France, and Italy, covering the period from 1870 to 2016 and here extended through the end of 2020.[20] They use annual data on short- and long-term government bond yields, consumer prices, and real consumption per capita in addition to Moody's Baa corporate bond yields. In their analysis, it is an assumption of no arbitrage in the long run that implies a factor structure for the trend of real interest rates across countries. They find that real interest rates across these 7 countries share a global common trend that has been particularly pronounced since the 1970s. Moreover, using regression analysis, they find that declining consumption growth and increasing convenience yields from the safety and liquidity offered by government bonds from these countries are the main drivers of declining real rates the past 40 years. We calculate the average of their r_t^* estimates for Germany, France, and Italy to get a representative estimate for the euro area. This series is shown with a solid blue line in Figure 21. Note that their r_t^* estimate for the euro area was increasing back in the 1960s and 1970s before starting a pronounced secular decline in the early 1980s. The trend lower continues through the end of the shown sample and leaves it below zero by 2020 not much above our market-based r_t^* estimate.

The third macro-based estimate of r_t^* is taken from Ferreira and Shousha (2023, henceforth FS) and shown with a solid green line in Figure 21. They consider a panel of 11 advanced economies and estimate the longer-run neutral real interest rates while accounting for changes in the global supply of safe assets and their convenience yields in addition to productivity and demographics and global spillovers from their developments. Their r_t^* estimate for the euro area starts in 1960 and fell steadily until the mid-1970s. It reversed some of the decline in the early 1980s and remained fairly stable until the late 1990s. It then steadily declined until 2008 when it reached a historic low of -0.89 percent. Since then it has gradually trended higher and stood at 0.1 percent by the end of 2023. This upward trend the past 15 years with a net increase of about 1 percentage point sets it apart from the other estimates, which mostly trend lower during this period.

The fourth and final macro-based estimate of r_t^* comes from Davis et al. (2024, henceforth DFHMT). They introduce a unified no-arbitrage macro-finance model with two trend factors used to estimate the natural rate r_t^* for 10 advanced economies, including Germany, France, and Spain. Using a multitude of data sources on trend growth and inflation in addition to risk premium series, DFHMT also underscore the need for a coherent model approach like

[20]We thank Marco Del Negro for sharing the updated data.

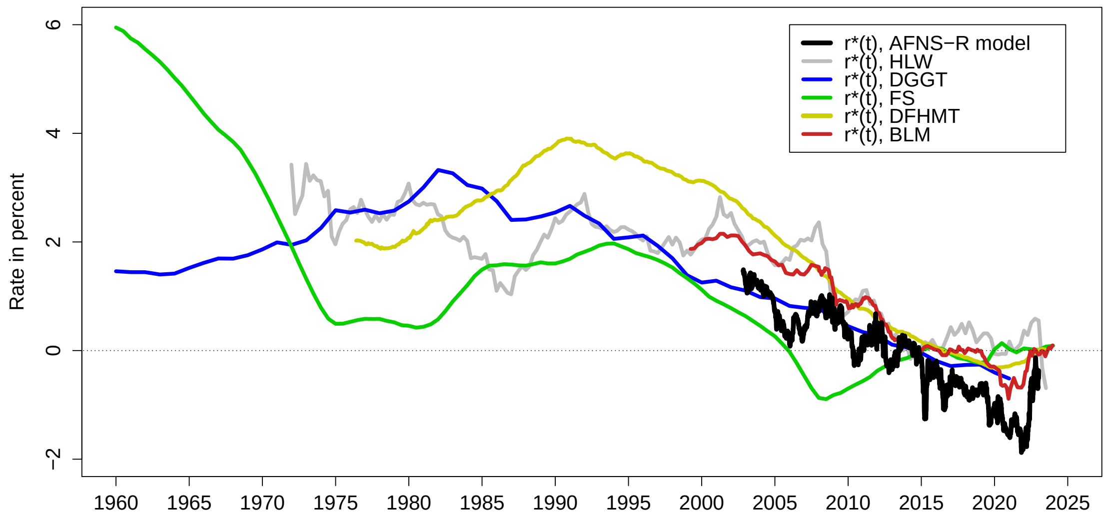

Figure 21: **Comparison with Macro-Based Estimates of r***

ours. Importantly, the interpretation of their natural rate r_t^* is consistent with Laubach and Williams (2003) of representing a medium-run real rate anchor for monetary policy. Our finance-based definition taken from CR is intended to capture the same concept. Hence, the DFHMT r_t^* estimate should be comparable to ours. One notable difference, though, is that, by relying on a single average yield equation, their estimation is not fully exploiting all the information in the yield curve unlike our approach.

To get a representative estimate for the euro area, we calculate the average of their r_t^* estimates for Germany, France, and Italy. The resulting r_t^* series is shown with a solid yellow line in Figure 21. While increasing in the 1970s and 1980s, the DFHMT r_t^* estimate peaked in late 1989. In the subsequent more than 30 years, r_t^* fell more than 4 percentage points according to their estimate and ends the sample slightly below zero. For the overlapping period this entails a close similarity between their r_t^* estimate for the euro area and our market-based r_t^* estimate. Overall, this pattern aligns well with the observed OAT€ yields shown in Figure 3. Moreover, using panel regressions, DFHMT relate their r_t^* estimates to economic growth and demographic variables and find that slowing growth and population aging have been significant factors in driving down the natural rate r_t^* globally, and in their three European countries in particular. Given the similarity of our market-based r_t^* to their estimate, we speculate that our estimate are likely influenced by those same factors.

The final series shown in Figure 21 is the median of a variety of r_t^* estimates reported by Brand et al. (2024, henceforth BLM). They include both macro- and market- as well as survey-based estimates of r_t^* for the euro area.[21] The similarities in both the declining trend and the general level of their median r_t^* estimate and our market-based r_t^* estimate are striking. In particular, they both suggest that the natural rate experienced a significant decline early on during the COVID-19 pandemic and a fairly sharp recovery of that decline in early 2022. As a result, both series suggest that r_t^* in the euro area has changed little on net since before the pandemic. Still, all six considered measures suggest that r_t^* in the euro area has declined notably the past 20-30 years and remain close to zero at the end of our sample despite the recent sharp increases in long-term interest rates in the euro area and other major advanced economies. This obviously matters for judgments about the stance of monetary policy, as we will discuss later on.

5.5 Projections of the Natural Rate

In light of the intense debate among researchers, investors, and policymakers about whether there is a new lower normal for interest rates, we end our analysis in this section by presenting the outlook for the natural rate based on our preferred AFNS-R model. We follow the approach of Christensen et al. (2015) and simulate 10,000 factor paths over a ten-year horizon conditioned on the shape of the OAT€ yield curve and investors' embedded forward-looking expectations as of the end of our sample (that is, using estimated state variables and factor dynamics as of December 30, 2022). The simulated factor paths are then converted into forecasts of r_t^*. Figure 22 shows the median projection and the 5th and 95th percentile values for the simulated natural rate over a ten-year forecast horizon.[22]

First, we note that our r_t^* estimate experienced some reversal of the declines from the past two decades during the last year of our sample, which left it at -0.37 percent at its end. The median r_t^* projection shows a persistent, but very gradual further reversal throughout the ten-year projection period that would put it close to 0.2 percent by 2032. The upper 95th percentile rises more rapidly and moves slightly above 2 percent by the end of the projection period, while the lower 5th percentile represents outcomes with the natural rate trending persistently lower into negative territory and remaining there over the entire forecast horizon. Although stationary, these results show that a highly persistent model like our preferred AFNS-R model can deviate from the estimated mean for several decades. Thus, nonstationary dynamics such as unit roots or trending shifting end points are not necessary to satisfactorily model the secular persistent decline of interest rates observed in the OAT€ market the past two decades. Of course, like most estimates of persistent dynamics, the model may still suffer from some finite-sample bias in the estimated parameters of its mean-reversion matrix $K^{\mathbb{P}}$,

[21] We thank Claus Brand for sharing this series.

[22] Note that the lines do not represent short rate paths from a single simulation run over the forecast horizon; instead, they delineate the distribution of all simulation outcomes at a given point in time.

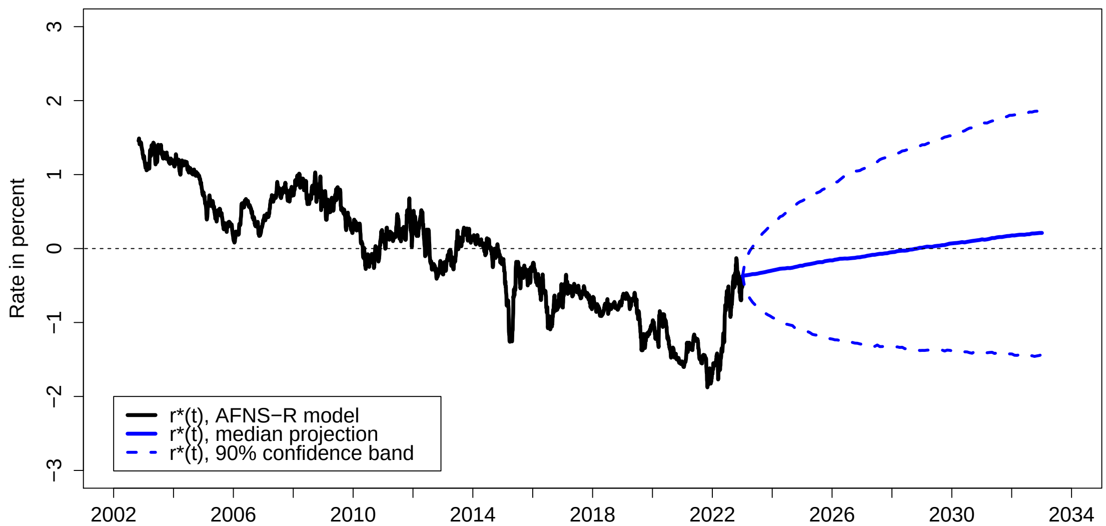

Figure 22: **Ten-Year Projections of r* from AFNS-R Model**

which would imply that it does not exhibit a sufficient amount of persistence—as described in Bauer et al. (2012). In turn, this would suggest (all else being equal) that the outcomes below the median are more likely than a straight read of the simulated probabilities indicate, and correspondingly those above the median are less likely than indicated. As a consequence, we view the projections in Figure 22 as an upper bound estimate of the true probability distribution of the future path for the natural rate.

Finally, our OAT€-based estimate of r_t^* appears relevant to the debate about the source of the decline in the natural rate. In particular, our measure of the natural rate did not fluctuate much in response to the GFC. This relative stability suggests that flight-to-safety and safety premium explanations of the lower natural rate, which have been put forward to explain low U.S. interest rates, are unlikely to be key drivers of the downtrend in euro-area interest rates. Instead, our estimates appear more broadly consistent with many of the explanations that attribute the decline in the natural rate to real-side fundamentals such as changing demographics (e.g., Carvalho et al. 2016, Favero et al. 2016, and Gagnon et al. 2016).

6 The Stance of ECB Monetary Policy

In this section, as a final application of our market-based estimate of r_t^*, we use it to construct measures of the stance of the ECB's monetary policy.

In theory, the stance of monetary policy would be given by the difference between the

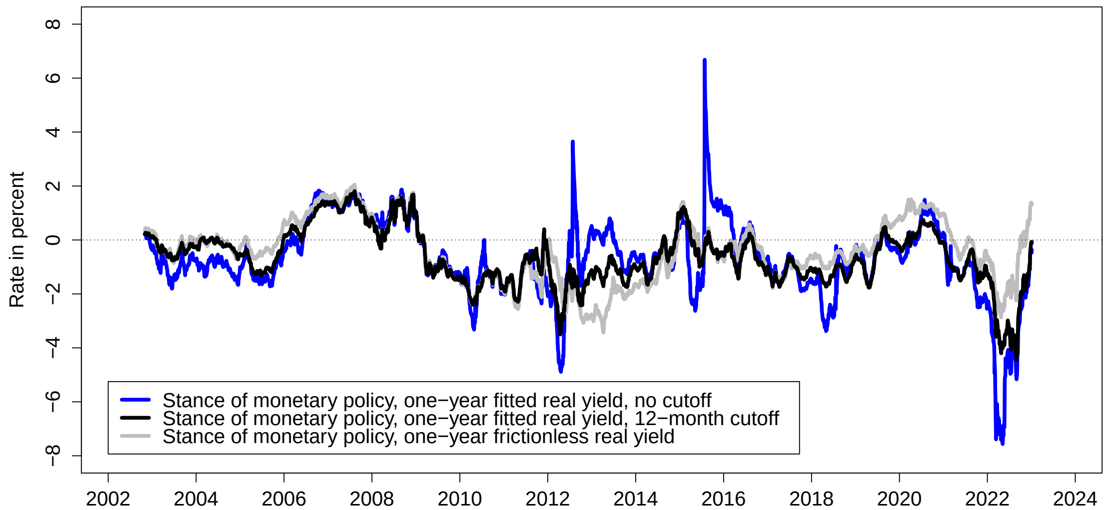

Figure 23: **Market-Based Measures of the Stance of Monetary Policy**

current real instantaneous short rate and its natural level as reflected in r_t^*, i.e., it would be defined as

$$\zeta_t = r_t - r_t^*.$$

The intuition behind this definition is straightforward. When the current real short rate is above its natural level, interest rates of all kinds are likely to be above their steady-state level and will provide some headwind for economic activity though higher borrowing costs and help slowdown the economy. And vice versa, when the current real short rate is below its natural level, the general interest rate level is likely to be below what is needed to maintain trend growth, and businesses and households may take advantage of that by making investments in new projects or housing at cheap financing rates, which will help boost economic activity.

Unfortunately, the instantaneous real short rate is not directly observable because we do not have a continuous measure of the very short end of the OAT€ yield curve, given that individual OAT€s reach maturity infrequently as noted in Figure 2(b). Furthermore, as explained earlier, OAT€s, like other inflation-indexed bonds, tend to have rather erratic prices close to maturity thanks to both low liquidity and the unpredictability of the final inflation adjustments to be earned—the sudden and very sharp spike in HICP inflation in 2022 is very illustrative in this regard.[23] Thus, to make the definition above operational, we

[23]For comparison, a standard fixed-coupon bond pays a principal of 1 and fixed coupons C. Thus, there is no uncertainty about its final cash flow in the months leading up to its maturity date, which helps maintain the liquidity of such securities.

consider instead three proxies that we think of as reasonable substitutes for r_t. The first is given by the one-year fitted real OAT€ yield from an estimation of the AFNS model without censoring any bond price information, that is, OAT€ prices remain in the sample until they mature. This provides the best possible coverage around the one-year maturity point but comes at the cost of adding significant noise from the prices of OAT€s close to maturity. Still, one can argue that this yield measures the full actual real yields observed in financial markets—including noise and frictions—and hence represents the most realistic real-world equivalent to the textbook short-term real rate embedded in the definition of ζ_t. To limit the noise and erratic behavior while preserving the desirable economic interpretation, we consider a second proxy for the stance of monetary policy calculated using instead the one-year fitted real yield from an estimation of the AFNS model imposing our baseline censoring of price information when bonds have less than one year to maturity. The third and final proxy is the one-year frictionless real yield implied by our preferred AFNS-R model. This is a cleaner and more stable measure of the one-year real yield as it adjusts for the noise from the bond-specific risk premia. However, in doing so, it may be different from the textbook concept of the real short rate r_t appearing in the original definition of ζ_t. Moreover, as OAT€ bond prices with less than one year to maturity are censored in the estimation of our preferred AFNS-R model, it may capture the short end of the OAT€ real yield curve less accurately similar to our second proxy.

The resulting three empirical measures of the ECB's stance of monetary policy are shown in Figure 23. In general, the three measures are quite similar and highly positively correlated, although there are important differences to note. Allowing for no data cutoff in the estimation that produces the short-term real yields, provides a stance measure with sharp spikes up or down whenever a bond in the sample approaches maturity. Crucially, these sharp short-lived gyrations are uncorrelated with the stance of monetary policy, which leads us to reject this measure. In comparison, using frictionless real yields to construct our measure of the stance of monetary policy, provides a more smooth and stable estimate. Unfortunately, as we demonstrate in Figure 24 below, this measure of the stance of monetary policy is very sensitive in the first decade of our sample to the data frequency used in the model estimation, which is an undesirable feature. In contrast, using fitted real yields based on our baseline approach with censoring of the bond prices with less than one year to maturity, provides a stance measure that is both relatively stable *and* robust to the data frequency used in the model estimation. As a consequence, this is our preferred measure of the stance of monetary policy in the euro area.

Comfortingly, there are several important commonalities across the three measures worth highlighting. First, monetary policy in the euro area was tight going into the GFC in 2007 and remained above neutral into 2009 before finally reaching an accommodative level. Second, in the 2015-2018 period, quantitative easing and other unconventional measures along with forward guidance managed to push the stance of monetary policy into accommodative territory and keep it there for several years according to all three measures. Third, at the peak of the COVID-19 pandemic in spring 2020, monetary policy reached a tightening stance and did not become accommodative until early 2021. Finally and similar to the United States, the ECB response to the spike in inflation following the global economic reopening after the pandemic was delayed, which had the implication that monetary policy remained very accommodative for an extended period of time and did not reach a tightening posture until the very end of our sample, and only according to one of our three measures. This may have contributed to prolonging the spell of high inflation in the euro area during this period, but it falls well outside the focus of this paper to make any determinations to that effect, so we leave it for future research to explore that question further.

Based on these observations we think of our empirical market-based measures of the ECB's monetary policy stance as realistic and representative. Moreover, as demonstrated by our analysis, they can be estimated at daily frequency and hence used for truly real-time policy analysis. This represents a major advantage relative to existing macro-based estimates, which are only available with a lag and may be subject to significant data revisions.

6.1 Comparison with a Text-Based Policy Stance Measure

To validate our market-based measures of the stance of ECB's monetary policy, we focus on Hubert and Portier (2024, henceforth HP), who construct a text-based measure of the ECB's stance of monetary policy. Specifically, they use textual analysis techniques to identify words that are either dovish or hawkish in the policy statement and during the press conference following each ECB governing council meeting. By subtracting the dovish count from the hawkish count and divide by the total word count, they obtain a measure of the net hawkish signal or stance conveyed after each policy meeting since 2001.

For this comparison, we use our preferred measure of the stance of ECB monetary policy based on the one-year fitted real yield using 12-month censoring in the model estimation. In Figure 25, we compare our chosen market-based measure of the stance of monetary policy in the euro area to the text-based ECB stance measure reported by HP. We note that, although broadly similar in the 2003-2008 period and again in the 2015-2019 period, the two measures imply sharply different assessments of ECB's monetary policy stance during the European Sovereign Debt Crisis in 2010-2013, during the COVID-19 pandemic in 2020-2021, *and* during the post-pandemic spike in inflation. Interestingly, the text-based measure suggests that ECB policy was neutral-to-net hawkish during the sovereign debt crisis, while our market-based measure suggests that monetary policy in the euro area was accommodative during this period by historical standards. During the COVID-19 pandemic we see the opposite pattern whereby

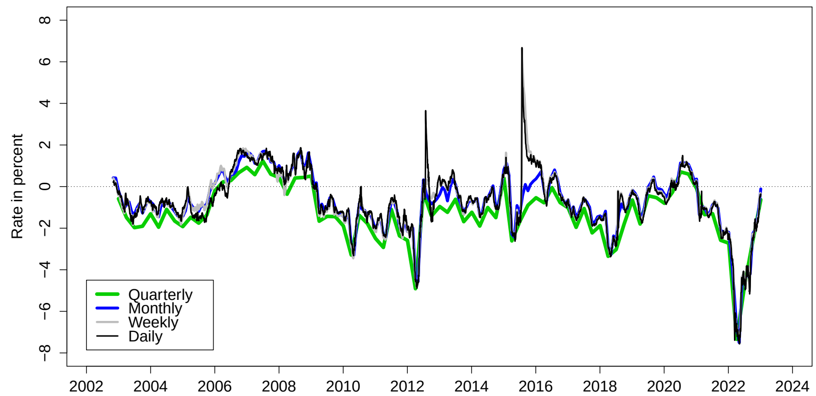

(a) Stance of monetary policy, one-year fitted real yield, no cutoff

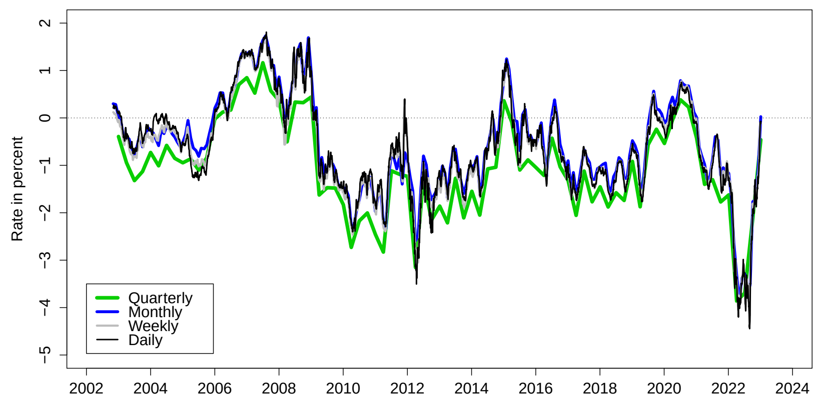

(b) Stance of monetary policy, one-year fitted real yield, 12-month cutoff

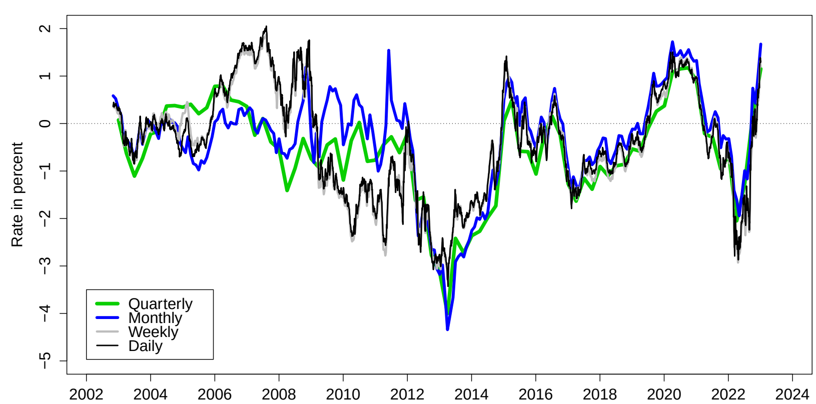

(c) Stance of monetary policy, one-year frictionless real yield, 12-month cutoff

Figure 24: **Sensitivity of Measures of the Stance of Monetary Policy to Data Frequency**

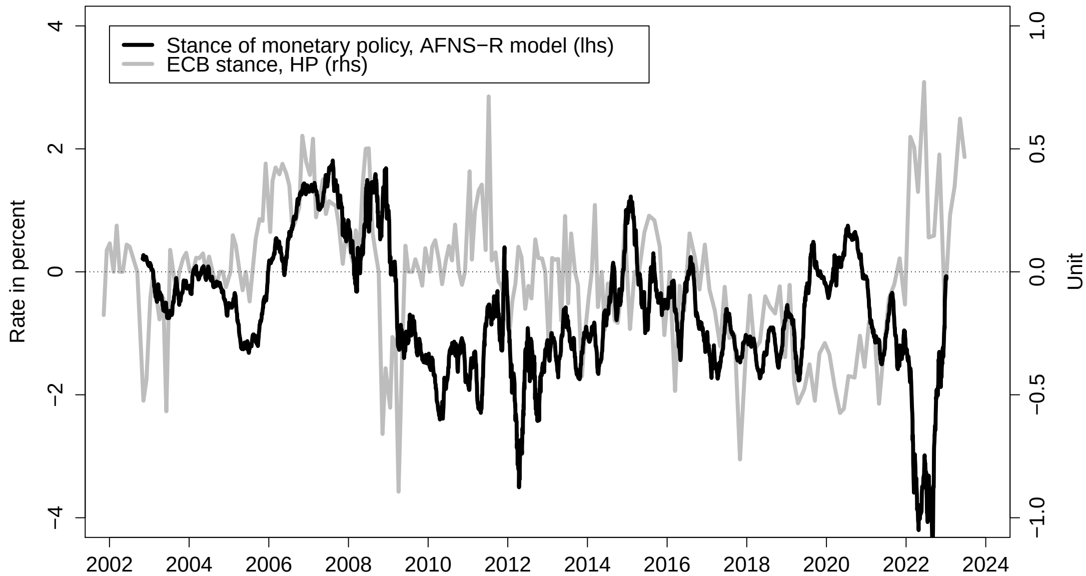

Figure 25: **Comparison with a Text-Based Measure of the Stance of Monetary Policy**

ECB policymakers tried to convey an accommodative stance, but the market-based measure indicates that this was not achieved until after the pandemic when inflation spiked up. During this latter, equally interesting period the measures again flip sign. While ECB policymakers were trying to send very strong hawkish signals, the market-based measure suggests that monetary policy was very accommodative initially as short-term real rates were very negative due to the high inflation. As a consequence, monetary policy only reached a restrictive stance by late 2022 according to our market-based measure.

What explains this very different pattern for the text-based measure during these crucial periods? To offer an answer, we note that HP's measure reflects—in a very direct way—the monetary policy stance communicated by ECB officials in the statement and through the answers to questions during the press conference. However, by design, it fails to capture to what extent the messaging is actually registered by financial market participants. In contrast, our market-based measure is designed to exactly capture the information investors have priced into the deep and liquid OAT€ bond market. Under the assumption that investors are forward looking and have every monetary incentive to use what they deem to be the best available information in devising their trading strategies, this "best available" information gets reflected in the bond prices. Under the additional assumption that our model is well specified, it should extract this information and the embedded investor expectations in a reliable manner. This is the theoretical and econometric argument for why our market-based measure of the stance of monetary policy should be preferable to the text-based measure produced by HP. This also explains why they may be different and not necessarily highly positively correlated.

6.2 Real-Time Analysis

As a final exercise, we examine the real-time behavior of both our r_t^* estimate and our proposed measure of the stance of ECB's monetary policy. Given the documented robustness of both measures to the data frequency used, we choose to perform the exercise at the monthly frequency, in part to save on computing time and in part because this is the frequency conventionally used in macro-based policy analysis. Practically, we start the model estimation on January 31, 2015, add one month of data to our sample, re-estimate the models, and continue this process until we reach our full sample that ends on December 30, 2022. This allows us to study the real-time model performance before, during, and after the COVID-19 pandemic.

Figure 26 shows the resulting real-time estimates with a comparison with the corresponding full-sample estimates. We note that the real-time estimates of both r_t^* and the policy stance measure are very close to their full-sample counterparts. We take this evidence to demonstrate that our model can be reliably used for real-time analysis.

7 Conclusion

Given the historic downtrend in yields in recent decades, many researchers have investigated the factors pushing down the steady-state level of the safe short-term real interest rate. However, all of this empirical work has been based on *macroeconomic* models and data, and uncertainty about the correct macroeconomic specification has led some to question the resulting macro-based estimates of the natural rate. We avoid this debate by introducing a market-based measure of the natural rate derived from an empirical dynamic term structure model estimated solely on the prices of bonds issued by the French government and indexed to the HICP—known as OAT€s. By adjusting for both OAT€ bond-specific risk premia and real term premia, we uncover investors' expectations for the underlying frictionless real short rate for the five-year period starting five years ahead. This measure of the natural rate of interest exhibits a gradual decline over the past two decades that accounts for about 75 percent of the general decline in euro-area bond yields. Specifically, as of the end of December 2022, the AFNS-R model estimate of r_t^* is -0.37 percent, with a net decline of slightly less than 2 percentage points since the early 2000s.

Given that our measure of the natural rate of interest is based on the forward-looking information priced into the active inflation-indexed OAT€ market and can be updated at a daily frequency as we demonstrate, it could serve as an important input for real-time

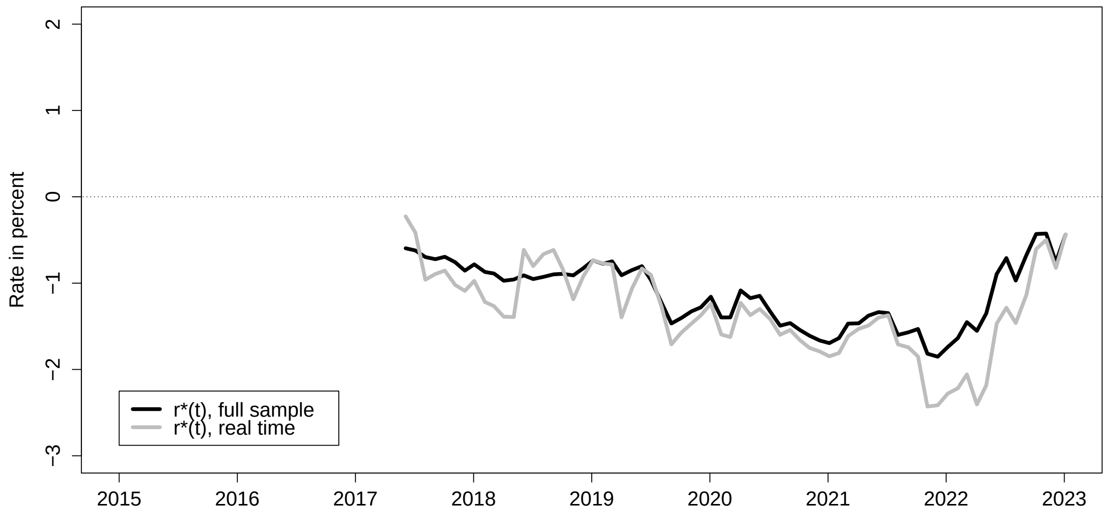

(a) r_t^*

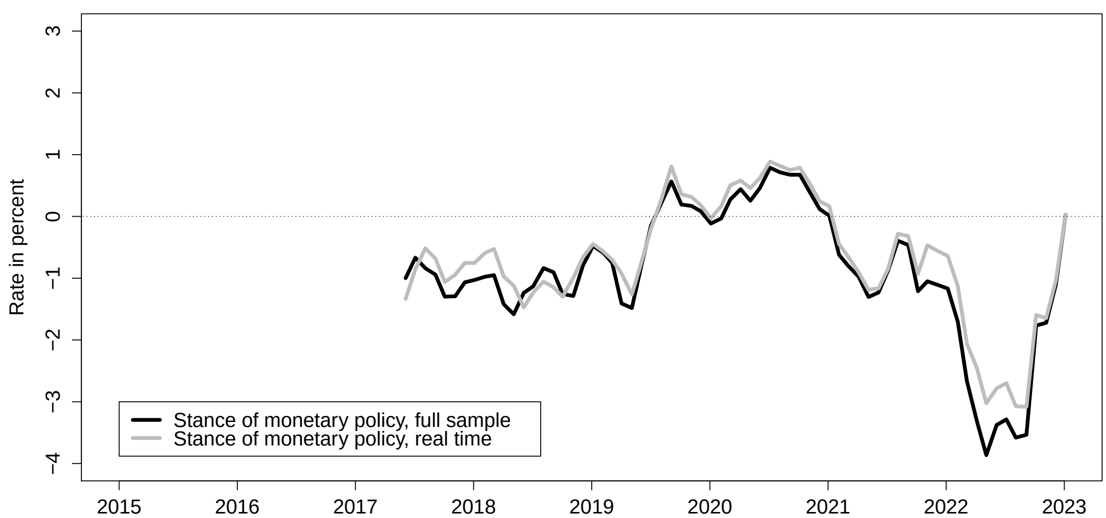

(b) Stance of monetary policy

Figure 26: **Real-Time Estimates of r_t^* and the Stance of Monetary Policy**

monetary policy analysis. Our related empirical measures of the stance of monetary policy would seem to be particularly relevant to examine further in this regard. For future research, our methods could also be expanded along an international dimension. With a significant degree of capital mobility, the natural rate will depend on global saving and investment, so the joint modeling of inflation-indexed bonds in several countries could be informative (see

HLW for an international discussion of the natural rate). Finally, our measure could be incorporated into an expanded joint macroeconomic and finance analysis—particularly with an eye towards further understanding the determinants of persistent changes in the natural rate.

References

Abrahams, Michael, Tobias Adrian, Richard K. Crump, Emanuel Moench, and Rui Yu, 2016, "Decomposing Real and Nominal Yield Curves," *Journal of Monetary Economics*, Vol. 84, 182-200.

Andreasen, Martin M., Jens H. E. Christensen, and Simon Riddell, 2021, "The TIPS Liquidity Premium," *Review of Finance*, Vol. 25, No. 6, 1639-1675.

Andreasen, Martin M., Jens H. E. Christensen, and Glenn D. Rudebusch, 2019, "Term Structure Analysis with Big Data: One-Step Estimation Using Bond Prices," *Journal of Econometrics*, Vol. 212, 26-46.

Bauer, Michael D., Glenn D. Rudebusch, and Jing (Cynthia) Wu, 2012, "Correcting Estimation Bias in Dynamic Term Structure Models," *Journal of Business and Economic Statistics*, Vol. 30, No. 3, 454-467.

Bernanke, Ben, 2005, "The Global Saving Glut and the U.S. Current Account Deficit," Sandridge Lecture, Richmond VA, March 10.

Blanchard, Olivier, 2023, *Fiscal Policy under Low Interest Rates.* MIT Press: Cambridge, Massachusetts.

Brand, Claus, Noëmie Lisack, and Falk Mazelis, 2024, "Estimates of the Natural Interest Rate for the Euro Area: An Update," European Central Bank, Economic Bulletin Issue 1.

Cardozo, Cristhian H. R. and Jens H. E. Christensen, 2024, "The Benefit of Inflation-Indexed Debt: Evidence from an Emerging Market," Working Paper 2023-04, Federal Reserve Bank of San Francisco.

Carriero, Andrea, Sarah Mouabbi, and Elisabetta Vangelista, 2018, "U.K. Term Structure Decompositions at the Zero Lower Bound," *Journal of Applied Econometrics*, Vol. 33, 643-661.

Carvalho, Carlos, Andrea Ferrero, and Fernanda Nechio, 2016, "Demographics and Real Interest Rates: Inspecting the Mechanism," *European Economic Review*, Vol. 88, 208-226.

Christensen, Jens H. E., Francis X. Diebold, and Glenn D. Rudebusch, 2011, "The Affine Arbitrage-Free Class of Nelson-Siegel Term Structure Models," *Journal of Econometrics*, Vol. 164, No. 1, 4-20.

Christensen, Jens H. E., Jose A. Lopez, and Glenn D. Rudebusch, 2010, "Inflation Expectations and Risk Premiums in an Arbitrage-Free Model of Nominal and Real Bond Yields," *Journal of Money, Credit and Banking*, Supplement to Vol. 42, No. 6, 143-178.

Christensen, Jens H. E., Jose A. Lopez, and Glenn D. Rudebusch, 2014, "Do Central Bank Liquidity Facilities Affect Interbank Lending Rates?," *Journal of Business and Economic Statistics*, Vol. 32, No. 1, 136-151.

Christensen, Jens H. E., Jose A. Lopez, and Glenn D. Rudebusch, 2015, "A Probability-Based Stress Test of Federal Reserve Assets and Income," *Journal of Monetary Economics*, Vol. 73, 26-43.

Christensen, Jens H. E. and Nikola Mirkov, 2022, "The Safety Premium of Safe Assets," Working Paper 2019-28, Federal Reserve Bank of San Francisco.

Christensen, Jens H. E. and Sarah Mouabbi, 2023, "Pre- and Post-Pandemic Inflation Expectations in France: A Bond Market Perspective," Manuscript, Federal Reserve Bank of San Francisco.

Christensen, Jens H. E., Sarah Mouabbi, and Caroline Paulson, 2025, "German Inflation-Linked Bonds: Overpriced, yet Undervalued," Manuscript, Federal Reserve Bank of San Francisco.

Christensen, Jens H. E. and Glenn D. Rudebusch, 2012, "The Response of Interest Rates to U.S. and U.K. Quantitative Easing," *Economic Journal*, Vol. 122, F385-F414.

Christensen, Jens H. E. and Glenn D. Rudebusch, 2015, "Estimating Shadow-Rate Term Structure Models with Near-Zero Yields," *Journal of Financial Econometrics*, Vol. 13, No. 2, 226-259.

Christensen, Jens H. E. and Glenn D. Rudebusch, 2019, "A New Normal for Interest Rates? Evidence from Inflation-Indexed Debt," *Review of Economics and Statistics*, Vol. 101, No. 5, 933-949.

Clarida, Richard, 2014, "Navigating the New Neutral, " *Economic Outlook,* PIMCO, November.

Dai, Qiang and Kenneth J. Singleton, 2000, "Specification Analysis of Affine Term Structure Models," *Journal of Finance*, Vol. 55, No. 5, 1943-1978.

Dai, Qiang, Kenneth J. Singleton, and Wei Yang, 2004, "Predictability of Bond Risk Premia and Affine Term Structure Models," manuscript, Stanford University.

D'Amico, Stefania, Don H. Kim, and Min Wei, 2018, "Tips from TIPS: The Informational Content of Treasury Inflation-Protected Security Prices," *Journal of Financial and Quantitative Analysis* Vol. 53, 243-268.

Davis, Josh, Cristian Feunzalida, Leon Huetsch, Benjamin Mills, and Alan M. Taylor, 2024, "Global Natural Rates in the Long Run: Postwar Macro Trends and the Market-Implied r^* in 10 Advanced Economies," *Journal of International Economics*, Vol. 149, 103919.

Del Negro, Marco, Domenico Giannone, Marc P. Giannoni, and Andrea Tambalotti, 2019, "Global Trends in Interest Rates," *Journal of International Economics*, Vol. 118, 248-262.

Duffee, Gregory R., 2002, "Term Premia and Interest Rate Forecasts in Affine Models," *Journal of Finance*, Vol. 57, No. 1, 405-443.

Ejsing, Jacob, Juan Angel Garcia, and Thomas Werner, 2007, "The Term Structure of Euro Area Break-Even Inflation Rates: The Impact of Seasonality," European Central Bank Working Paper Series No. 830.

Favero, Carlo A., Arie E. Gozluklu, and Haoxi Yang, 2016, "Demographics and the Behavior of Interest Rates," *IMF Economic Review* Vol. 64, No. 4, 732-776.

Ferreira, Thiago R. T. and Samer Shousha, 2023, "Determinants of Global Neutral Interest Rates," *Journal of International Economics*, Vol. 145, 103833.

Finlay, Richard and Sebastian Wende, 2012, "Estimating Inflation Expectations with a Limited Number of Inflation-Indexed Bonds," *International Journal of Central Banking*, Vol. 8, No. 2, 111-142.

Fontaine, Jean-Sébastien and René Garcia, 2012, "Bond Liquidity Premia," *Review of Financial Studies*, Vol. 25, No. 4, 1207-1254.

Gagnon, Etienne, Benjamin K. Johannsen, and David Lopez-Salido, 2016, "Understanding the New Normal: The Role of Demographics," Finance and Economics Discussion Series 2016-080. Washington: Board of Governors of the Federal Reserve System, http://dx.doi.org/10.17016/FEDS.2016.080.

Greenspan, Alan, 2005, Federal Reserve Board's Semiannual Monetary Policy Report to the Congress, Testimony to the U.S. Senate, February 16.

Grishchenko, Olesya V. and Jing-Zhi Huang, 2013, "Inflation Risk Premium: Evidence from the TIPS Market," *Journal of Fixed Income*, Vol. 22, No. 4, 5-30.

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

Holston, Kathryn, Thomas Laubach, and John C. Williams, 2017, "Measuring the Natural Rate of Interest: International Trends and Determinants," *Journal of International Economics*, Vol. 108, 559-575.

Houweling, Patrick, Albert Mentink, and Ton Vorst, 2005, "Comparing Possible Proxies of Corporate Bond Liquidity," *Journal of Banking and Finance*, Vol. 29, No. 6, 1331-1358.

Hu, Grace Xing, Jun Pan, and Jiang Wang, 2013, "Noise as Information for Illiquidity," *Journal of Finance*, Vol. 68, No. 6, 2341-2382.

Hubert, Paul and Rose Portier, 2024, "Monetary Policy and the Slope of the Yield Curve," Manuscript, Banque de France.

Joyce, Michael A. S., Iryna Kaminska, and Peter Lildholdt, 2012, "Understanding the Real Rate Conundrum: An Application of No-Arbitrage Models to the U.K. Real Yield Curve," *Review of Finance*, Vol. 16, 837-866.

Joyce, Michael, Peter Lildholdt, and Steffen Sørensen, 2010, "Extracting Inflation Expectations and Inflation Risk Premia from the Term Structure: A Joint Model of the U.K. Nominal and Real Yield Curves, " *Journal of Banking and Finance*, Vol. 34, 281-294.

Kim, Don H. and Kenneth J. Singleton, 2012, "Term Structure Models and the Zero Bound: An Empirical Investigation of Japanese Yields," *Journal of Econometrics*, Vol. 170, No. 1, 32-49.

Laubach, Thomas and John C. Williams, 2003, "Measuring the Natural Rate of Interest," *Review of Economics and Statistics*, Vol. 85, No. 4, 1063-1070.

Laubach, Thomas and John C. Williams, 2016, "Measuring the Natural Rate of Interest Redux," *Business Economics*, Vol. 51, No. 2, 57-67.

Nagel, Stefan, 2016, "The Liquidity Premium of Near-Money Assets," *Quarterly Journal of Economics*, Vol. 131, No. 4, 1927-1971.

Nelson, Charles R. and Andrew F. Siegel, 1987, "Parsimonious Modeling of Yield Curves," *Journal of Business*, Vol. 60, No. 4, 473-489.

Pflueger, Carolin E. and Luis M. Viceira, 2016, "Return Predictability in the Treasury Market: Real Rates, Inflation, and Liquidity," in Handbook of Fixed-Income Securities, edited by P. Veronesi, Wiley, Hoboken, N.J., pp. 191-209 (Chapter 10).

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

Speck, Christian, 2021, "Market Microstructure of Inflation-Indexed Bonds and Technical Biases in Market-Based Inflation Indicators," Deutsche Bundesbank Internal Note.

Summers, Lawrence H., 2014, "U.S. Economic Prospects: Secular Stagnation, Hysteresis, and the Zero Lower Bound," *Business Economics*, Vol. 49, No. 2, 65-73.

Summers, Lawrence H., 2015, "Demand Side Secular Stagnation," *American Economic Review, Papers and Proceedings*, Vol. 105, No. 5, 60-65.

Summers, Lawrence H., 2023, "Back to Secular Stagnation?," Speech at the Annual Meeting of the American Economic Association, January 7, 2023.

Swanson, Eric T. and John C. Williams, 2014, "Measuring the Effect of the Zero Lower Bound on Medium- and Longer-Term Interest Rates," *American Economic Review*, Vol. 104, No. 10, 3154-3185.

Online Appendix

The Natural Rate of Interest in the Euro Area: Evidence from Inflation-Indexed Bonds

Jens H. E. Christensen
Federal Reserve Bank of San Francisco
jens.christensen@sf.frb.org

and

Sarah Mouabbi
Banque de France
sarah.mouabbi@banque-france.fr

The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of San Francisco or the Federal Reserve System, or those of the Banque de France or the Eurosystem.

This version: January 28, 2025.

Contents

A Model Selection in the Daily AFNS Model

In this appendix, we go through a careful model selection procedure for the AFNS model estimated at daily frequency similar to the one described in the main text for the AFNS-R model.

Alternative Specifications	Goodness of fit statistics			
	$\log L$	k	p-value	BIC
(1) Unrestricted $K^{\mathbb{P}}$	217,252.5	17	n.a.	-434,359.4
(2) $\kappa_{31}^{\mathbb{P}} = 0$	217,252.5	16	1.00	-434,367.9
(3) $\kappa_{31}^{\mathbb{P}} = \kappa_{21}^{\mathbb{P}} = 0$	217,251.9	15	0.27	-434,375.3
(4) $\kappa_{31}^{\mathbb{P}} = \kappa_{21}^{\mathbb{P}} = \kappa_{32}^{\mathbb{P}} = 0$	217,250.8	14	0.14	-434,381.7
(5) $\kappa_{31}^{\mathbb{P}} = \ldots = \kappa_{23}^{\mathbb{P}} = 0$	217,248.6	13	0.04	-434,385.8
(6) $\kappa_{31}^{\mathbb{P}} = \ldots = \kappa_{12}^{\mathbb{P}} = 0$	217,247.3	12	0.11	**-434,391.8**
(7) $\kappa_{31}^{\mathbb{P}} = \ldots = \kappa_{13}^{\mathbb{P}} = 0$	217,238.6	11	< 0.01	-434,383.0

Table 1: **Evaluation of Alternative Specifications of the AFNS Model**
There are seven alternative estimated specifications of the AFNS model. Each specification is listed with its maximum log likelihood ($\log L$), number of parameters (k), the p-value from a likelihood ratio test of the hypothesis that it differs from the specification above with one more free parameter, and the Bayesian information criterion (BIC). The period analyzed covers daily data from October 31, 2002, to December 30, 2022.

For estimates of r_t^* based on our definition, the specification of the mean-reversion matrix $K^{\mathbb{P}}$ is critical. To select the best fitting specification of the AFNS model's real-world dynamics, we use a general-to-specific modeling strategy in which the least significant off-diagonal parameter of $K^{\mathbb{P}}$ is restricted to zero and the model is re-estimated. This strategy of eliminating the least significant coefficient is carried out down to the most parsimonious specification, which has a diagonal $K^{\mathbb{P}}$ matrix. As in the main text, the final specification choice is based on the value of the Bayesian information criterion (BIC).

The summary statistics of the model selection process are reported in Table 1. The BIC is minimized by specification (6), which has a $K^{\mathbb{P}}$ matrix given by

$$K_{BIC}^{\mathbb{P}} = \begin{pmatrix} \kappa_{11}^{\mathbb{P}} & 0 & \kappa_{13}^{\mathbb{P}} \\ 0 & \kappa_{22}^{\mathbb{P}} & 0 \\ 0 & 0 & \kappa_{33}^{\mathbb{P}} \end{pmatrix}.$$

The estimated parameters of this preferred specification are reported in Table 2. We note that most of the parameters are very close to those reported in the main text for the AFNS

$K^{\mathbb{P}}$	$K^{\mathbb{P}}_{\cdot,1}$	$K^{\mathbb{P}}_{\cdot,2}$	$K^{\mathbb{P}}_{\cdot,3}$	$\theta^{\mathbb{P}}$		Σ
$K^{\mathbb{P}}_{1,\cdot}$	0.1709	0	-0.1702	0.0318	$\Sigma_{1,1}$	0.0036
	(0.0732)		(0.0392)	(0.0169)		(0.0000)
$K^{\mathbb{P}}_{2,\cdot}$	0	0.3863	0	-0.0242	$\Sigma_{2,2}$	0.0129
		(0.2018)		(0.0115)		(0.0002)
$K^{\mathbb{P}}_{3,\cdot}$	0	0	0.2717	0.0073	$\Sigma_{3,3}$	0.0183
			(0.2361)	(0.0177)		(0.0003)

Table 2: **Estimated Parameters in the Preferred AFNS Model**
The estimated parameters for the mean-reversion matrix $K^{\mathbb{P}}$, the mean vector $\theta^{\mathbb{P}}$, and the volatility matrix Σ in the AFNS model preferred according to the BIC. The $\mathbb{Q}$-related parameter is estimated at $\lambda = 0.3861$ (0.0012). The maximum log likelihood value is 217,247.3. The numbers in parentheses are the estimated standard deviations.

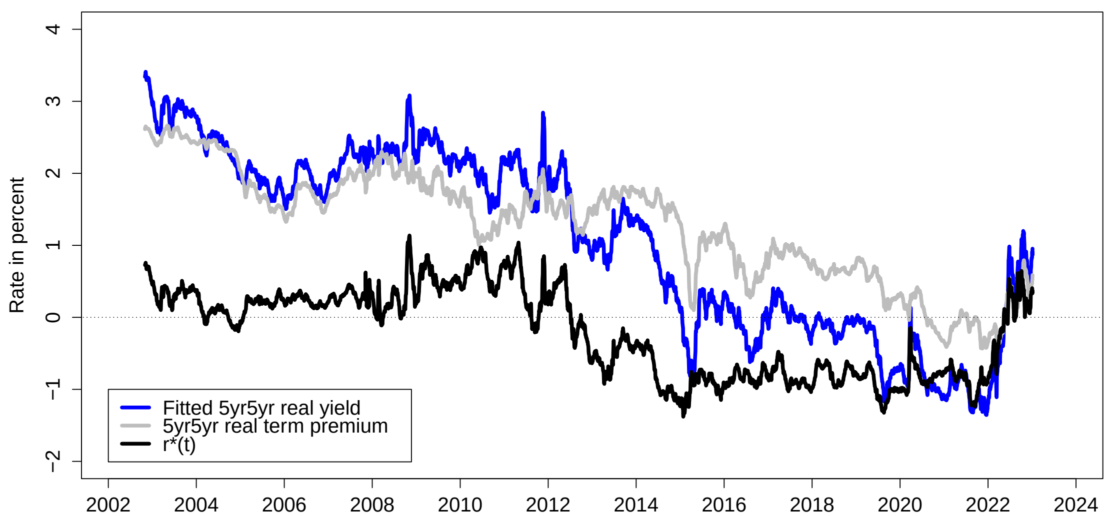

Figure 1: **5yr5yr Real Yield Decomposition**

model with diagonal $K^{\mathbb{P}}$ matrix, which seems reasonable given that only the off-diagonal $\kappa^{\mathbb{P}}_{13}$ parameter separates the two models.

Figure 1 shows the 5yr5yr real yield decomposition implied by the preferred AFNS model. Its estimate of the natural real rate r_t^* is stable with persistent fluctuations around zero. As a result, the model implies that the lower trend in the 5yr5yr real yield is driven by declines in the 5yr5yr real term premium.

To examine the sensitivity of the estimated r_t^* from the preferred AFNS model to the specification of the $K^{\mathbb{P}}$ matrix, we consider the AFNS models with unrestricted and diagonal

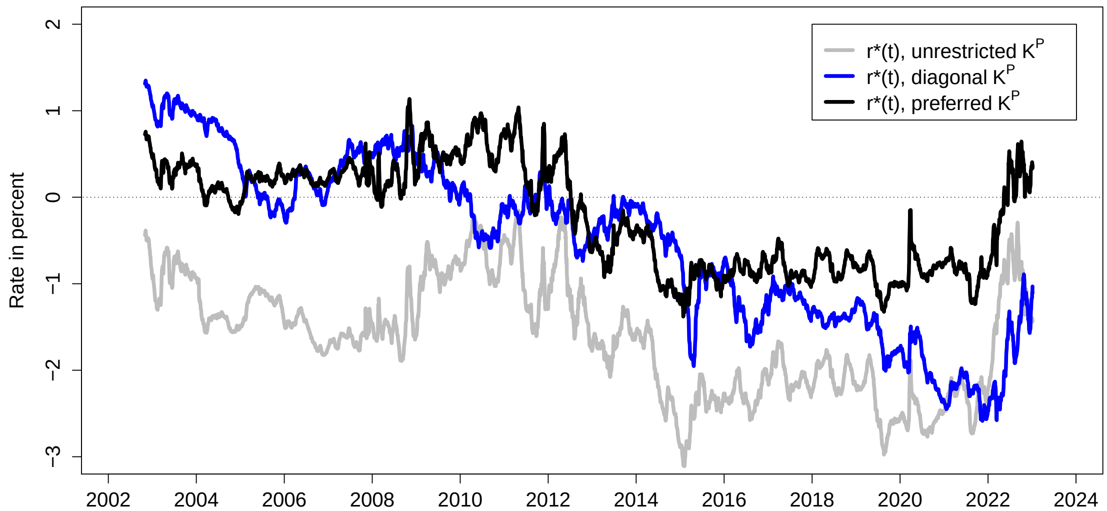

Figure 2: **Sensitivity of r* Estimate to $K^{\mathbb{P}}$ Specification**

$K^{\mathbb{P}}$ matrix. The resulting r_t^* estimates are shown in Figure 2 where we note that the estimates are indeed very sensitive to this choice. This underscores the importance of going through a careful model selection procedure like the one described above.

B Sensitivity of Estimated State Variables to Data Frequency

In this appendix, we examine the sensitivity of the estimated state variables within the AFNS-R model to the data frequency. To do so, we focus on the most parsimonious specification of the model with diagonal $K^{\mathbb{P}}$ mean-reversion matrix and diagonal Σ volatility matrix estimated at daily, weekly, monthly, and quarterly frequency, respectively.

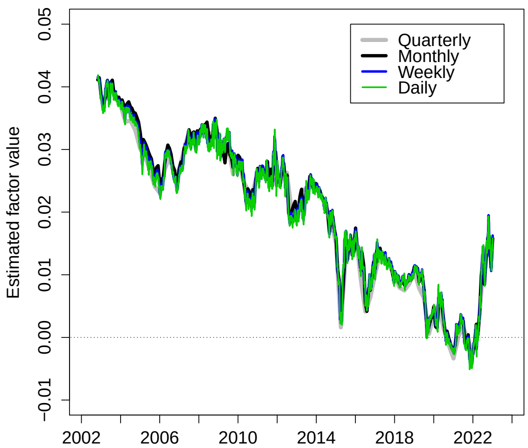

(a) L_t

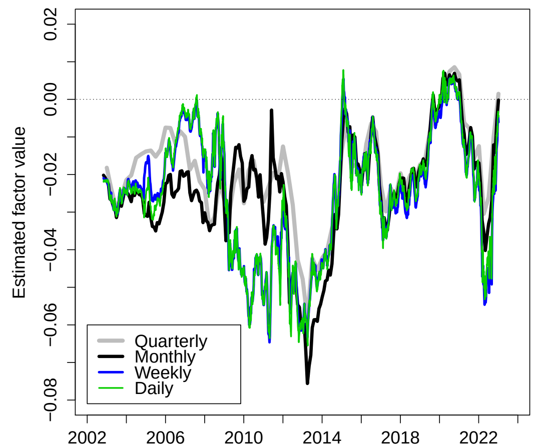

(b) S_t

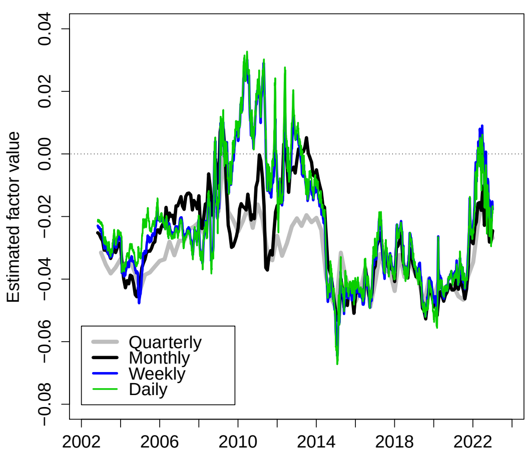

(c) C_t

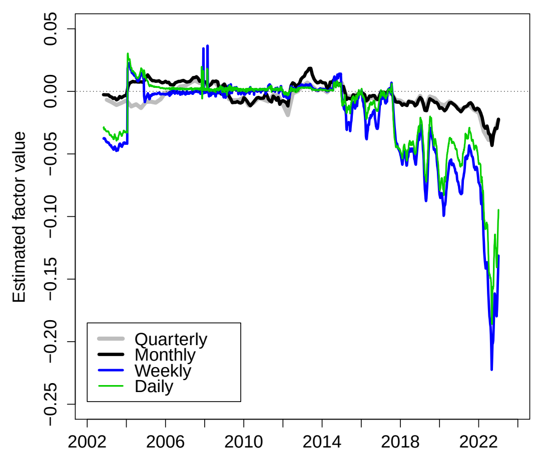

(d) X_t^R

Figure 3: **Estimated State Variables: Data Frequency**
Illustration of the estimated state variables from the AFNS-R model when estimated using daily, weekly, monthly, and quarterly data.

Figure 3 shows the estimated paths for all four state variables from the four estimations.

We note that, due to the limited number of observed bond prices in the early years of our sample, all state variables are not fully identified during that period. As a consequence, we do see some differences in the filtered state variables depending on the frequency of the data used in the model estimation. Importantly, though, the dominating level factor *is* well identified and its filtered path is insensitive to the data frequency. Moreover, roughly starting in 2012 the filtered paths for the frictionless level, slope, and curvature factors become insensitive to the data frequency thanks to the sufficiently large number of observed bond prices during the remaining part of the sample. Furthermore, we do see some differences in the estimated bond-specific risk factor X_t^R depending on the data frequency even after 2012. However, these differences do not translate into differences in the average estimated bond-specific risk premium series during the last ten years of our sample as demonstrated in Figure 12 in the main text. Finally and most importantly, we stress that it follows from Figure 18 in the main text that the r_t^* estimates from our preferred AFNS-R model estimated at different data frequencies are very similar and all exhibit the same trending patterns. Hence, the crucial r_t^* output for our analysis has little sensitivity to the data frequency used, which supports our choice to focus on the highest possible daily data frequency for our analysis.

C Sensitivity of Estimated State Variables to Data Cutoff

In this appendix, we examine the sensitivity of the estimated state variables within the AFNS-R model to the data cutoff used in the model estimation. To do so, we focus on the preferred AFNS-R model identified in Section 5.2 in the main text. We estimate this model with no cutoff (i.e. 0 months) as well as a cutoff of 6 months, 18 months, 24 months, and 30 months as an alternative to our benchmark choice of using a 12-month cutoff.

Figure 4 shows the estimated paths for all four state variables from the six estimations. In general, the state variables have relatively little sensitivity to the cutoff choice. As a consequence, the sensitivity of our r_t^* estimate to this implementation choice is also relatively modest as demonstrated in Section 5.3 in the main text.

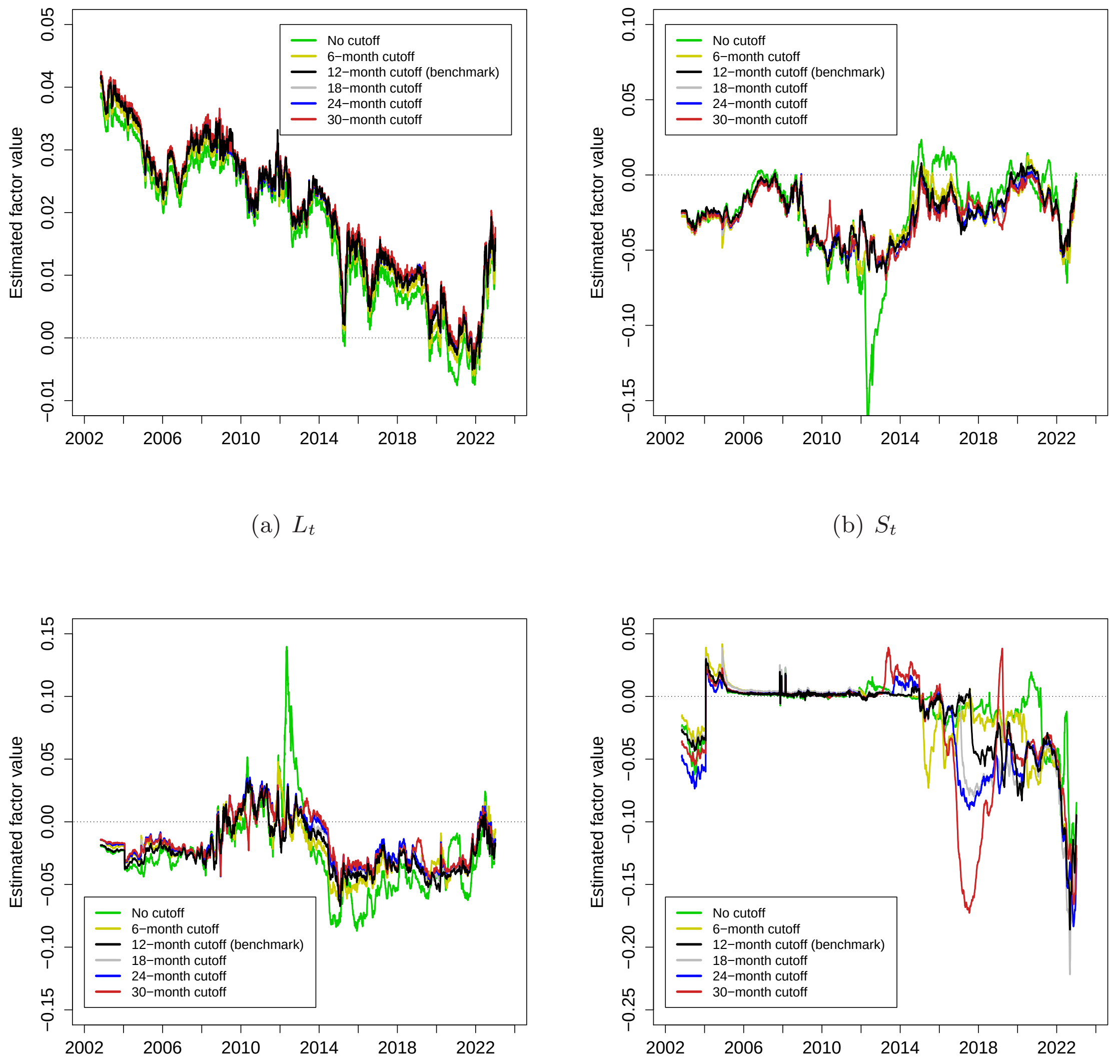

(a) L_t (b) S_t

(c) C_t (d) X_t^R

Figure 4: **Estimated State Variables: Data Cutoff**
Illustration of the estimated state variables from the AFNS-R model when estimated with varying data cutoff choices ranging from censoring with zero months to maturity up to censoring with 30 months to maturity.