Macro Takes Time: Term Structure of Risk Premia in Bonds and Currencies

Gustavo Amarante* Gustavo Soares†

March 15, 2026

Abstract

We apply the term-structure-of-risk-premia framework of Bryzgalova, Huang, and Julliard (2024, BHJ) to a cross section of 29 foreign exchange forwards and government bonds. Using a hierarchical Bayesian model that extracts latent priced shocks from returns and traces their propagation into macroeconomic variables, we estimate horizon-specific risk premia for seven US macro and financial factors. Four of them bear statistically significant premia: CPI inflation, the Fed Funds Rate, the yield curve slope, and the US–Europe yield differential. In each case, premia are small at short horizons and grow with the investment horizon, consistent with priced shocks that markets absorb immediately but that map into observable macro series only gradually. GDP, the real yield, and the VIX are not significantly priced. This pattern contrasts sharply with BHJ's equity results, where real-activity variables bear the dominant long-horizon premia: in our bond-and-currency cross section, the priced disturbance is concentrated in the nominal-monetary block rather than in real activity. We derive a closed-form decomposition showing that the model-implied mimicking portfolio preserves the same asset-space direction across all horizons, with only its leverage changing. The Bayesian implementation delivers coherent inference over the entire term structure and over portfolio weights in a setting where non-parametric alternatives would be underpowered.

1 Introduction

Macroeconomic risk lies at the center of modern asset pricing theory, yet empirical detection of macro risk premia remains challenging because macro variables are persistent, noisy, and may react with delay to the shocks that are immediately reflected in asset prices. Bryzgalova, Huang, and Julliard (2024, BHJ) address this identification problem by turning the canonical approach "macro factors price returns" on its head: they use a large cross section of returns to extract the priced shocks spanned by financial markets and then quantify how these shocks propagate through macroeconomic time series across horizons, thereby delivering a coherent, horizon-specific term structure of macro risk premia. This paper is an application of that priced-macro-shocks extraction framework. The key difference is the asset universe and economic setting: BHJ focus on identifying macro risk premia in equities and using mainly

*gustavoca2@insper.edu.br
†gustavobs3@insper.edu.br

real-activity-related macro factors, emphasizing that conventional evidence for priced macro risks in stock returns is limited and that macro premia emerge primarily at longer horizons.

Our setting implements the same methodology in foreign exchange and government bond returns. This application matters because the role of macro shocks is both more central and more directly interpretable in rates and currency markets than in equities, due to the mechanics of how these contracts work. However, there are distinct empirical traditions in these asset classes. In term structure research, macro factors and shocks are foundational ingredients in no-arbitrage yield curve models and in the empirical study of bond risk premia, including work that embeds macro dynamics into the term structure, like Ang and Piazzesi (2003), documents strong return predictability and time variation in bond risk premia, like Cochrane and Piazzesi (2005), and connects broad macro conditions to variation in expected bond returns, like Ludvigson and Ng (2009), alongside influential econometric formulations of dynamic term structure models like Joslin, Singleton, and Zhu (2011) and Adrian, Crump, and Moench (2012).

In currency markets, a large literature likewise links excess returns to macro risks and global risk factors: the cross section of currency risk premia has been tied to consumption-based risk, like Lustig and Verdelhan (2007), and habit-based mechanisms, like Verdelhan (2010), while common components in currency returns and their associated premia have been emphasized in multi-currency factor structures Lustig, N. Roussanov, and Verdelhan (2011). Recent work further highlights that currency carry premia vary systematically across investment horizons, suggesting a genuinely term-structured object in FX Lustig, Stathopoulos, and Verdelhan (2019). Our contribution is to bring the unified "priced macro shocks" framework from BHJ to these markets and to quantify, within a single internally consistent methodology, the horizon-specific premia of macro factors spanned by FX and bonds in a way that is directly comparable across macro variables.

Empirically, applying the BHJ machinery to FX and government bonds shows clear evidence of priced macro components associated with inflation, short-term nominal interest rates, yield curve slope and yield differential: the estimated premia for these four macro variables are statistically significant. We find no evidence of priced factors in GDP, the TIPS real 10-year yield and the VIX. Importantly, the estimated premia exhibit a pronounced term structure: average macro premia are small at short horizons and become larger at longer investment horizons, mirroring one of the central results of the BHJ framework, but now in markets where macroeconomic shocks are widely believed to be first-order pricing determinants. By placing rates and currencies into the same priced-shocks-and-propagation lens, we provide a disciplined and comparable map of macro risk compensation across horizons in two cornerstone asset classes beyond equities.

A central technical contribution of BHJ is to re-express macro risk compensation as a horizon-by-horizon object tied to the propagation of priced return shocks into a macroeconomic variable, rather than to the contemporaneous correlation between the macro series and a single-period factor-mimicking portfolio. In their framework, the risk premium of a (possibly non-traded) factor is defined through its covariance with the stochastic discount factor (SDF) over the relevant investment horizon. This definition is designed to capture precisely the mechanism that motivates the "term structure" in BHJ paper: when the macro series reacts sluggishly to shocks that are immediately priced in returns, short-horizon risk premia can be small even though longer-horizon premia are economically meaningful.

What makes this definition especially reassuring is that BHJ connect it directly to the more familiar non-parametric linear projection logic underlying "mimicking portfolios". Consider the conventional approach: to price a non-traded factor one projects it onto the span of traded returns and then reads its risk premium from the average excess return of the resulting portfolio. BHJ show that, when one formulates the linear projection at the same horizon as the object being priced — i.e., using multi-period returns and the macro variable's multi-period growth — the resulting horizon-specific mimicking portfolio has a risk premium that converges, in a large cross section, to the premium implied by the SDF-covariance definition. The underlying asymptotics are cross-sectional: as the number of test assets becomes large, and provided the return space is sufficiently rich to span the relevant priced shocks (a "strong factor" condition), the idiosyncratic component of returns remains well-behaved, and pricing errors do not aggregate into asymptotic arbitrage, the linear-projection-based construction becomes equivalent to pricing via the parametric SDF. Put differently, the horizon-specific mimicking portfolio effectively isolates the same priced component that the BHJ model attributes to the macro variable through its loading on the common return shocks, so the two routes to risk premia coincide asymptotically. This equivalence is not merely conceptual: it provides a bridge between (i) a non-parametric, covariance-driven approach that relies on the geometry of the return space and (ii) a parametric, SDF-based approach that relies on an explicit factor structure and a propagation model for the macro series. Hence, in BHJ framework is particularly well-suited to estimating identifying macro risk premia in FX and government bond returns where cross section of test assets is modest.

Two further implications are worth emphasising for our purposes. First, the equivalence is inherently horizon-specific. BHJ show that the canonical, single-period mimicking portfolio can be largely uninformative about long-horizon premia precisely because it ignores the macro variable's lagged response to priced shocks. Second, even if one wished to pursue horizon-specific linear projections non-parametrically, doing so comes at a steep statistical cost: multi-period returns and multi-period macro growth induce overlapping observations, eroding effective sample size and power just where the term-structure question becomes most interesting. The BHJ parametric propagation structure — estimating how priced shocks map into the macro variable over time — addresses this power problem by pooling information across horizons through an economically interpretable dynamic restriction, rather than forcing each horizon to be estimated as a separate overlapping-sample linear projection.

These considerations are particularly salient in our FX and government bonds application. Unlike equities, where the cross section can be made very large via characteristics sorts or extensive portfolio formation, our test-asset universe is necessarily more limited: the set of liquid currency excess returns and investable government bonds delivers a cross section that is informative but not asymptotically "large" in the sense required for the strongest non-parametric equivalence arguments. In such a setting, the covariance matrix underpinning mimicking-portfolio weights can be noisy, and small-sample instability in linear projection weights becomes a first-order concern, especially once one moves to longer horizons. Consequently, we do not treat non-parametric mimicking portfolios as a primary identification strategy. Instead, we lean on the parametric content of the BHJ model—the joint factor structure in returns and the distributed-lag propagation of priced shocks into macro variables—and we use the Bayesian implementation to obtain coherent inference for the entire term structure of macro premia in a setting where frequentist delta-method uncertainty would be cumbersome and non-parametric overlap would be underpowered.

Within this parametric structure, our central findings mirror the key economic mechanism highlighted by BHJ, albeit in a different asset universe and transmission environment. We identify priced components associated with monetary variables (inflation and short-term nominal interest rates) and some financial variables (yield curve slope and yield differential), but not for real variables like GDP and the TIPS 10-year real yield. Among the financial variables, the VIX also did not show any priced components for our set of returns. Moreover, the estimated premia display a pronounced term structure: they are small at short horizons but grow with the investment horizon. Interpreted through the BHJ lens, this pattern is the natural consequence of slow macro propagation of shocks that markets price immediately, so that the relevant covariance with the SDF accumulates over time even when the contemporaneous relation is weak.

The remainder of this study is organized as follows. Section 2 describes the model from and elaborates on its main objects, the term structure of risk premia and the mimicking portfolios. Section 3 describes our dataset of currencies, bond futures and cash bonds returns, and briefly comments on their comovements. Section 4 shows our main results, the composition of the mimicking portfolios and their term structure of interest rate for our seven macro factors, which are grouped in three categories: monetary, real and financial. Section 5 concludes.

2 Methodology

2.1 The Model

We use the model developed by Bryzgalova, Huang, and Julliard (2024). Following their notation, r_t is the vector of excess returns on N assets. The factor g_t, tradable or not, is a covariance-stationary observable variable. The cumulative returns and cumulative change of the factor between $t-1$ and $t+S$ are

$$r_{t-1\to t+S} = \sum_{t=-1}^{S} r_{t+i} \qquad g_{t-1\to t+S} = g_{t+S} - g_{t-1}$$

Returns are modeled with a linear latent factor model

$$r_t = \mu_r + \beta_\upsilon \upsilon_t + w_{r,t} \tag{1}$$

$$\upsilon_t \overset{iid}{\sim} N\left(0, I_K\right) \qquad w_{r,t} \overset{iid}{\sim} N\left(0, \Sigma_{wr}\right) \qquad \upsilon_t \perp w_{r,t}$$

where υ_t is a vector of priced components with K uncorrelated latent factors, β_υ are the factor loadings, μ_r is the expected log-returns, and $w_{r,t}$ are unpriced idiosyncratic errors. The assumption that υ_t is serially uncorrelated has important consequences for our application of mimicking portfolios, as we will see in section 2.3. We interpret this assumption as more of a simplifying assumption than an assumption on which the results heavily depend. BHJ relax this assumption on their conditional version of the model and, in a future version of this study, we intend to relax this assumption as well.

Given the structure in equation (1), there exists an equivalent representation of a log-SDF that is linear in the latent factors:

$$m_t - \kappa_m = -\lambda_\upsilon' \upsilon_t$$

and since the factor covariance is the identity matrix, their risk prices are risk premia. As we will show later, some of the contracts that we use in our sample are not denominated in US dollars, but all the total return indexes that we built have been converted to US dollars.

This means that the SDF m_t is the only one we need to keep track of, there is no need for a foreign currency denominated SDF like in Lustig, Stathopoulos, and Verdelhan (2019).

To model the observable factor g_t we use moving average representation of a combination of priced and unpriced components.

$$g_t = \mu_g + \sum_{s=0}^{\bar{S}} \rho_s \underbrace{\eta_g' v_{t-s}}_{f_{t-s}} + w_{g,t} \tag{2}$$

where μ_g is the uncoditional mean of g_t, η_g are coefficients proportional to risk premia and normalized to unit norm, f_t is the spanned component that drives both returns r_t and the factor g_t, $\{\rho_s\}_{s=0}^{\bar{S}}$ is a square-summable sequence of scalars, and $w_{g,t}$ is a pottentially autocorrelated shock unrelated to v_t and $w_{r,t}$. The number of lags in the process may be infinite, in which case $\bar{S}$ would be a high-enough finite lag to generate a good approximation of the process.

A notable feature of this model is that g_t can react to both current and lagged shocks to asset returns, which allows for the interpretations of returns reacting immediately to innovations and allowing non-tradable macroeconomic factors to have a slow and/or delayed response. And given its MA structure we can even interpret $\{\rho_s\}_{s=0}^{\bar{S}}$ as the impulse response function of g_t to a return shock f_t.

Another intended extension of this structure is to change the one-dimensional object g_t in equation 2 into a multivariate one.

2.2 Term Structure of Risk Premia

Another measure of interest is the risk premium coming from the factor g_t, which is defined by $-Cov(m_t, g_t)$. If g_t is a tradeable factor, this interpretation is direct. If g_t is not tradeable, we can interpret $-Cov(m_t, g_t)$ as the risk premium on an asset that delivers a payoff that follows the moves of g_t. And since g_t may have a persistent behavior, the risk premium may change for different horizons S. BHJ shows we can expand the risk premium definition and look at the per-period average covariance over different horizons:

$$\lambda_g^S = -\frac{Cov\left(m_{t-1\to t+S}, g_{t-1\to t+S}\right)}{S+1} = \left(\frac{\sum_{\tau=0}^{S}\sum_{s=0}^{\tau}\rho_s}{S+1}\right)\underbrace{\eta_g'\lambda_v}_{\lambda_f} \tag{3}$$

We can interpret the elements of equation (3). λ_f is the risk premium of the spanned component f_t, and $\frac{\sum_{\tau=0}^{S}\sum_{s=0}^{\tau}\rho_s}{S+1}$ as the per-period loading of g_t on return shocks.

2.3 Mimicking Portfolios

In the original article of Bryzgalova, Huang, and Julliard (2024), the primary object of interest is the term structure of risk premia λ_g^S. Our interest here extends beyond this measure to the composition of the portfolios w_{MP}^S that deliver this estimated premia. In rates and currency markets, practitioners and policymakers care not only about whether macro risk is compensated, but about where that compensation resides in the cross section. BHJ work with 275 Fama-French characteristic-sorted portfolios, where individual portfolio weights carry limited economic content. Our cross section of 29 instruments makes the weights directly interpretable as trading positions.

We can build the linear projection based portfolio w^S_{MP} as the projection of the macro factor onto the span of asset returns.

$$w^S_{MP} = Var\left(r_{t-1\to t+S}\right)^{-1} Cov\left(r_{t-1\to t+S}, g_{t-1\to t+S}\right) \tag{4}$$

which has risk premium λ^S_{MP}. As $N \to \infty$, BHJ shows that $\lambda^S_{MP} \to \lambda^S_g$. This means that as the number assets in the cross section increases, the risk premium of the linear-projection-based mimicking portfolio converges to that of the horizon-specific portfolio of priced factors. We can also compute the w^S_{MP} implied by the model:

Proposition 1. *Under the distributional assumptions of equations (1) and (2), the horizon-specific mimicking portfolio weights are given by*

$$w^S_{MP} = \varphi\left(S\right) \bar{w}_{MP} \tag{5}$$

where

$$\varphi\left(S\right) = \frac{\left[\sum_{l=0}^{\min\left(\bar{S}, S+1\right)} \rho_l \left(S+2-l\right)\right]}{S+2} \qquad and \qquad \bar{w}_{MP} = \Sigma^{-1}_{rr}\beta_v \Sigma_v \eta_g$$

The proof of this proposition is shown in appendix A.1. Two features of this decomposition are worth highlighting. First, the only term that depends on the horizon S is $\varphi\left(S\right)$ and it is a scalar. The direction of the portfolio in asset space is pinned down entirely by the base portfolio $\bar{w}_{MP}$, which is invariant to S. What changes across horizons is only the leverage of this base portfolio: longer horizons scale the same portfolio as the cumulative impulse response grows. In economic terms, the identity of which bonds and currencies hedge, say, inflation risk is the same whether one is a short-run or a long-run investor, only the size of the position differs. This horizon-invariance of direction is a strong and testable restriction of the model and it follows from the assumption that the latent factors v_t are serially uncorrelated. If one were to allow these to have a dependency structure, the base portfolio would itself become horizon-dependent. This is an intended extension of a future version of this study. That being said, it is unlikely that the mimicking portfolio will be strikingly different depending on the investment horizon. Hence, we see this assumption as more of a simplifying assumption than an assumption on which the results heavily depend.

Second, the base portfolio $\bar{w}_{MP}$ has a transparent economic interpretation. It identifies the portfolio of traded assets that best tracks the priced component of g_t, given the model's structural parameters rather than noisy sample covariances. In our setting of 29 assets with a relatively short sample, this model-implied construction works as an ideintification of an object that would be unfeasible with the direct non-parametric linear projection, which would require inverting a 29×29 sample covariance matrix of multi-period overlapping returns.

As shown by BHJ, the latent factors v_t are only identified up to a linear rotation, but the objects of the risk premia term structure are invariant to these rotations. The same happens for the $\bar{w}_{MP}$ portfolio. The intuition behind this result is that object in equation (4) only depends on second-order moments of observable variables. Rotating the base of unobservable factors does not change the object.

2.4 Estimation

We now describe the hierarchical Bayesian framework from Bryzgalova, Huang, and Julliard (2024), and show how the factor-mimicking portfolios from Proposition 1 are embedded as an additional step of the sampler.

We start with the following distributional assumptions

$$g_t = \mu_g + \sum_{s=0}^{\bar{S}} \rho_s \eta_g' \left(\upsilon_{t-s} - \mu_\upsilon\right) + w_{g,t} \tag{6}$$

$$\upsilon_t \overset{iid}{\sim} N\left(\mu_\upsilon, \Sigma_\upsilon\right) \qquad w_{g,t} \overset{iid}{\sim} N\left(0_N, \sigma_{wg}^2\right)$$

$$r_t = \mu_r + \beta_\upsilon \left(\upsilon_t - \mu_\upsilon\right) + w_{r,t} \tag{7}$$

$$w_{r,t} \overset{iid}{\sim} N\left(0_N, \Sigma_{wr}\right)$$

$$\Sigma_{wr} = \operatorname{diag}\left(\sigma_{1,wr}^2, \ldots, \sigma_{N,wr}^2\right)$$

$$\upsilon_t \perp w_{g,t} \perp w_{r,t}$$

Define the vector of parameters of the macro equation as

$$\rho_g = \begin{pmatrix} \mu_g \\ \rho_0 \\ \rho_1 \\ \vdots \\ \rho_{\bar{S}} \end{pmatrix}_{\left(\bar{S}+2\right)\times 1}$$

We assign uninformative prior distributions to parameters of the timeseries dimension

$$\pi\left(\rho_g, \eta_g, \sigma_{wg}^2\right) \propto \left(\sigma_{wg}^2\right)^{-1}$$

$$\pi\left(\upsilon\right) \propto 1 \qquad \pi\left(\mu_\upsilon, \Sigma_\upsilon\right) \propto \left|\Sigma_\upsilon\right|^{-\frac{K+1}{2}}$$

$$\pi\left(\beta_\upsilon\right) \propto 1 \qquad \pi\left(\mu_r, \Sigma_{wr}\right) \propto \left|\Sigma_{wr}\right|^{-\frac{N+1}{2}}$$

Given these priors we can build the Gibbs sampler of the model. We initialize the chains with estimates from principal components for υ_t, factor loadings for β_υ, and frequentist analogous estimators for the factor MA process.

1. **Factor Equation**: Conditional on the data $\{g_t\}_{t=\bar{S}+1}^T$ and latent factor $\{\upsilon_t\}_{t=1}^T$, the parameters from the factor equation (σ_{wg}^2, ρ_g and η_g) follow a normal-inverse-gamma distribution. For point identification, after each draw we normalize $\eta_g'\eta_g = 1$.

2. **Return Equation**: Conditional on asset returns $\{r_t\}_{t=1}^T$ and latent factors $\{\upsilon_t\}_{t=1}^T$, we update model parameters from the return equation (Σ_{wr}, μ_r and β_υ) using a normal-inverse-Wishart distribution.

3. **Latent Factors**: Conditional on asset returns and μ_r, β_v and Σ_{wr}, we update the latent factors υ_t, their mean μ_v, and covariance parameters Σ_v using a normal-inverse-Wishart distribution.

4. **Risk Premia**: Conditional on all the draws from steps 1-3, the posterior distribution of λ_v is a Dirac distribution at $\lambda_v = (\beta_v'\beta_v)^{-1}\beta_v'\tilde{\mu}_r$ and the posterior distribution for the term structure of the factor risk premia for S periods ahead is a Dirac distribution at

$$\lambda_g^S = \left(\frac{\sum_{\tau=0}^{S}\sum_{s=0}^{\tau}\rho_s}{S+1}\right)\eta_g'\lambda_v$$

5. **Mimicking Portfolios**: Conditional on all the draws from steps 1-3, the posterior distribution of w_{MP}^S is a Dirac distribution at

$$w_{MP}^S = \frac{\left[\sum_{l=0}^{\min\left(\bar{S},S+1\right)}\rho_l\,(S+2-l)\right]}{S+2}\Sigma_{rr}^{-1}\beta_v\Sigma_v\eta_g$$

A practical consequence is that we obtain full Bayesian posterior distributions over portfolio weights at a near zero additional computational cost[1].

3 Data

3.1 Asset Returns

Our cross section consists of 29 global assets: 9 currency forwards, 13 government bond futures across several countries and maturities, and 7 cash government bonds of 10-year maturity. The sample runs from January 2000 to December 2025. In a future development of this study we intend to introduce more maturities and longer series of returns, increasing our sample both in the cross-section and in the timeseries dimensions. For each asset, we construct a daily excess return index, which is then resampled to the observation frequency of the factor g_t. Since the derivatives in our sample are unfunded positions, their returns are already measured in excess of the appropriate financing rate. We assume trades are executed at observed settlement prices, implying zero slippage and no trading costs. Given that the instruments in our sample are among the most liquid in global financial markets, this approximation is unlikely to materially affect the results.

Bond futures are settled daily in local currency, and we convert these daily settlements to US dollars, so that the resulting return series reflects no residual currency risk. Currency forwards, by construction, isolate the exchange rate component: their excess returns capture the interest rate differential plus the change in the spot exchange rate. Cash government bonds, by contrast, are converted at the end-of-period exchange rate, so their returns embed both duration risk and currency risk. We include cash bonds precisely because this joint exposure enriches the cross section, by adding assets whose covariation with macro factors reflects both rate and exchange rate channels simultaneously.

[1] In fact, steps 4 and 5 don't even need to be implemented inside the Gibbs sampler. These are computed after the full MCMC chain for the other parameters is generated. This allows steps 4 and 5 to run in a vectorized manner only on the draws after the burnin of the MCMC chain, speeding up the estimation process.

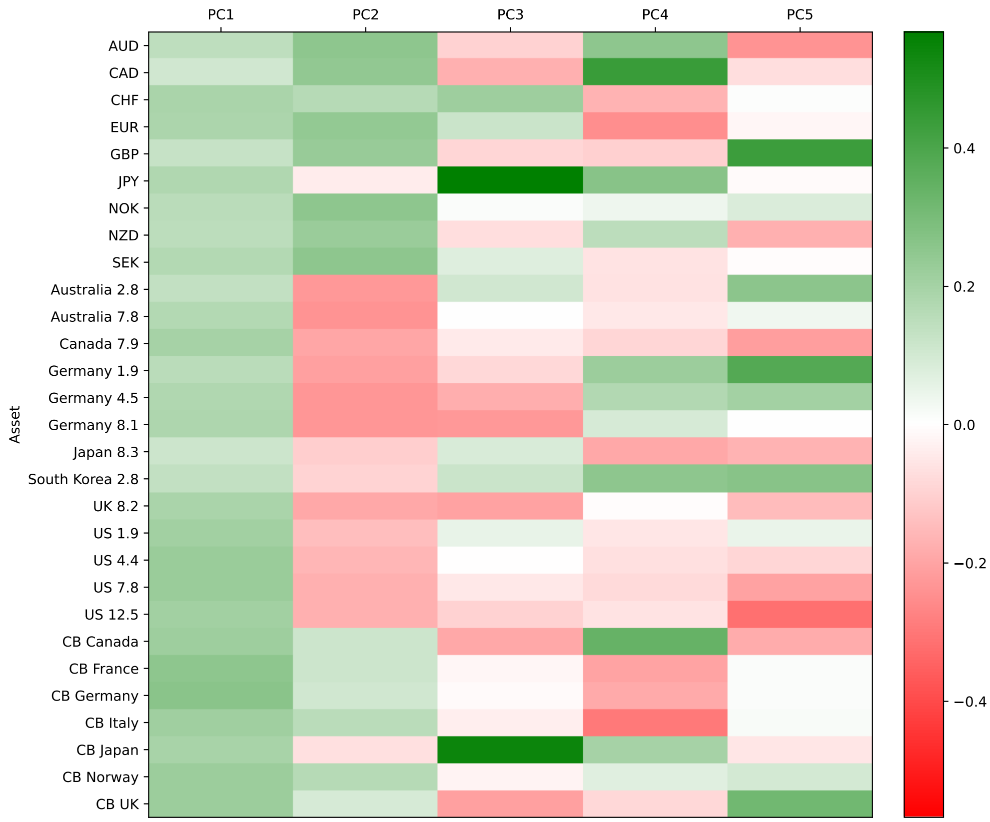

Figure 1: Factor Loadings from PCA on standardized monthly returns. The number next to the bond futures' names are their average duration. "CB" stands for cash bonds.

We run a principal component analysis on the monthly returns of our 29 global assets. The factor loadings in Figure 1 confirm that the distinction of yield and exchange rate exposure is economically meaningful. The first principal component acts as a market or global factor, with uniformly signed loadings across all assets, while the second principal component, which we interpret as a currency risk factor, loads positively on cash bonds and currency forwards while loading negatively on currency-hedged bond futures. This pattern is economically coherent and provides reassurance that the cross section, while modest in size, spans economically distinct sources of risk that are relevant for identifying macro risk premia. Figure 2 shows the explained variance of each factor. The first two principal components account for 68.4% of the total variance, with the marginal contribution of subsequent components being substantially smaller, bringing evidence that we have 2 main common driving forces behind these assets. This supports our baseline choice for the number of latent factors in the return equation, as will be described in Section 4.1.

3.2 Macro Factors

We apply the model to seven macro and financial variables, each observed for the United States. The focus on US variables is motivated by the large literature documenting the outsized role of US monetary conditions in driving global asset prices, what Miranda-Agrippino and Rey (2021) term the "global financial cycle". Since our asset universe spans government bonds and currencies of developed economies that are tightly linked to US monetary policy, using

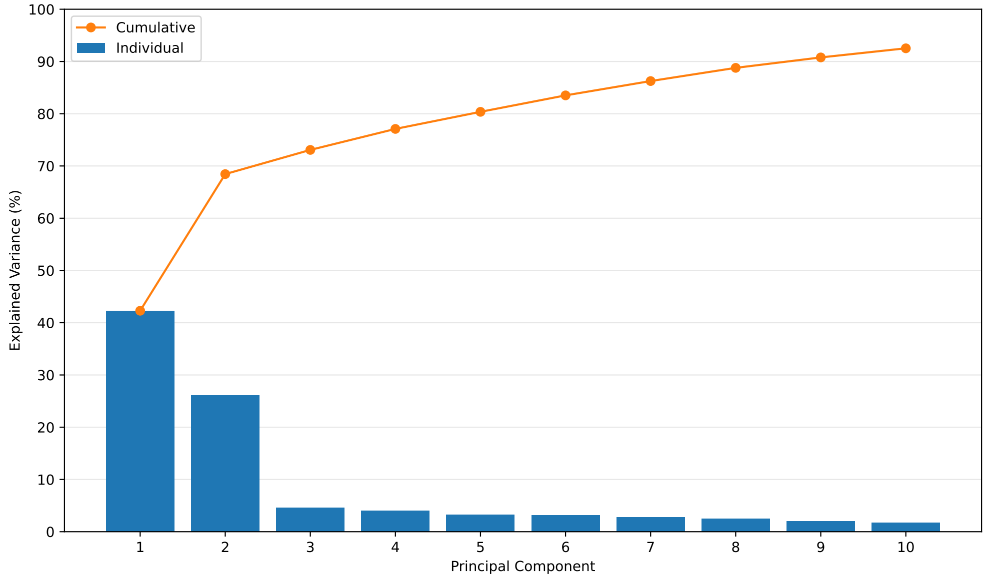

Figure 2: Explained variance decomposition from PCA on standardized monthly returns.

US factors allows us to ask a focused question: is US macro risk priced in global rates and currency markets?

We organize the factors into three groups based on economic content. The monetary factors are US CPI inflation (monthly percent changes) and the Fed Funds Rate (monthly average level). The real factors are US GDP growth (quarterly percent change) and the 10-year TIPS real yield (monthly average level). The financial factors are the slope of the US yield curve (10-year minus 2-year US Treasury zero-coupon yields), the difference between US and European 3-month government bond yields, and the VIX index. The yield curve slope and the US-Europe yield differential are commonly used as characteristics in the style factor literature, and therefore are relevant for the identification of risk premia. The VIX captures the price of aggregate equity volatility risk and serves as a proxy for shocks to the SDF.

4 Results

In this section, we show how we selected the number of latent factors, show the resulting mimicking portfolios and then comment on the risk premia term structure for each group of factors.

For each factor g_t, we estimate the model with the Gibbs sampler with 2000 burn-in draws, followed by 2000 posterior draws used in the analysis. Factors are standardized to unit variance so that estimates of λ_g^S are comparable across factors. For monthly factors we use $\bar{S} = 24$ lags and for quarterly GDP we use $\bar{S} = 8$ lags, corresponding to two years of propagation dynamics in both cases.

Figures 3 and 4 show the mimicking portfolios and their gross risk decomposition, respectively. Figures 5 through 7 show the estimated term structures of horizon-specific risk premia for each of the seven factors. Out of seven, four display posterior distributions whose 90% credible intervals exclude zero at some horizon: CPI, the Fed Funds Rate, the yield curve slope, and the US-Europe yield differential.

The common pattern across these significant factors is that the estimated premia are small at short horizons and become larger in magnitude as the horizon extends. In this application, the evidence points to priced macro shocks whose effects on returns are immediate, but whose mapping into observable macro variables is gradual enough that the economically relevant compensation is most visible at longer horizons. There are two complementary ways to interpret this pattern. The first is that macro factors are riskier for long-horizon investors than for short-horizon investors. The second is that time clears up the short-term noise and reveals the underlying risk premium that is difficult to see in the noisier short-term horizon.

An important interpretive caveat is that the model identifies the term structure of priced co-movement between returns and the observable factor g_t, but not, by itself, the primitive direction of causality. Since g_t is explicitly allowed to react to current and lagged return shocks, the evidence is consistent not only with nominal macro shocks being priced directly, but also with a reverse ordering in which a priced financial disturbance appears first in returns and only then passes through gradually into measured inflation and, with further delay, into the policy rate. This alternative is economically plausible given the large literature showing that exchange-rate movements pass through, albeit incompletely and heterogeneously, into import and consumer prices, and can also have broader macroeconomic consequences. In more mathematical terms, from the model's perspective what we have is a latent shock, and its origin is not directly attached to either equations (1) or (2), it has its nature in its assumed distribution.

4.1 Number of Latent Factors

Before estimating the model, we need to select the number of factors K in v_t for the returns equation (1). We follow the information criterion from Giglio and Xiu (2021), which balances the explanatory power of additional factors against the risk of overfitting in finite samples. The selected number of factors is

$$\hat{K} = \arg \min_{1 \leq j \leq K_{\max}} \left[N^{-1} T^{-1} \gamma_j (\bar{\boldsymbol{R}}^\top \bar{\boldsymbol{R}}) + j \times \frac{\hat{\gamma}}{2} (\log N + \log T) \left(N^{-\frac{1}{2}} + T^{-\frac{1}{2}} \right) \right] - 1$$

where $\bar{\boldsymbol{R}}$ is a $T \times N$ matrix of demeaned asset returns, $\gamma_j(\bar{\boldsymbol{R}}^\top \bar{\boldsymbol{R}})$ is the j-th eigenvalue of $\bar{\boldsymbol{R}}^\top \bar{\boldsymbol{R}}$, and $\hat{\gamma}$ is the median of the first $K_{\max}$ normalized eigenvalues $N^{-1}T^{-1}\gamma_j(\bar{\boldsymbol{R}}^\top \bar{\boldsymbol{R}})$. We set $K_{\max}$ to 15. For our cross section of 29 assets and 312 monthly observations, the criterion suggests $K = 2$ factors, which resonates with the evidence from the explained variance of each principal components from figure 2.

With only two latent factors driving returns, the model is parsimonious but potentially restrictive. But given our much smaller and more homogeneous cross section, choosing only two factors is not surprising. This number may appear low relative to the tradition of term-structure of interest rates. In yield-curve modeling, it is common to work with three latent yield factors, often interpreted as level, slope, and curvature, as emphasized early on by Litterman and Scheinkman (1991). Moreover, more flexible no-arbitrage term-structure specifications often employ additional factors to improve fit and capture term-premium dynamics. For instance, Adrian, Crump, and Moench (2012) report specification tests favoring models with five principal components of yields as pricing factors. At the same time, an important nuance is that "how many factors drive yields" and "how many factors drive time-variation in expected excess returns" need not coincide. Cochrane and Piazzesi (2005)

Figure 3: **Notional weights** for each factor g_t. Weights are normalized to 10% volatility of the portfolios. The number next to the bond futures' names are their average duration. "CB" stands for cash bonds.

show that a single return-forecasting factor made up from a particular linear combination of forward rates, predicts expected excess bond returns across maturities, highlighting that the relevant dimension for risk premia can be smaller than the dimension required for a high-fidelity representation of the entire yield curve.

4.2 Mimicking Portfolios

We now exploit the closed-form structure of Proposition 1 to examine which assets bear the macro risk premia. Recall that the model-implied mimicking portfolio weights decompose into a horizon-invariant base portfolio $\bar{w}_{MP}$ and a horizon-specific scalar $\varphi(S)$. To keep weights comparable across factors, the figures below focus on the $\bar{w}_{MP}$ direction, normalizing all the portfolios to a 10% annualized volatility. Different horizons would simply change the leverage of this portfolio, not its composition. Figure 3 reports the normalized notional weights of the base portfolio for each factor.

To quantify how individual assets contribute to the risk of the mimicking portfolio, we define the gross risk share (GRS) of asset i as

$$GRS_i = \frac{w_i \sigma_i}{\sum |w_i \sigma_i|}$$

where w_i is the portfolio weight and σ_i is the return volatility of asset i. Figure 4 reports the gross risk shares. Unlike notional weights, which can be misleading when assets have very different volatilities, gross risk shares measure each asset's contribution to the portfolio's total absolute risk exposure. Figure 4 reports the gross risk shares.

An important interpretive caveat applies before reading the portfolios. The model identifies the priced co-movement between returns and the observable factor g_t, but not the primitive

Figure 4: **Gross Risk Shares** for each factor g_t. Each column is $\frac{w_i\sigma_i}{\sum|w_i\sigma_i|}$. The number next to the bond futures' names are their average duration. "CB" stands for cash bonds.

direction of causality. Since g_t is explicitly allowed to react to current and lagged return shocks, the portfolio that hedges a given macro factor may reflect either a macro shock that moves asset prices, or a priced financial disturbance that subsequently passes through into the macro variable. We discuss both readings below.

4.3 Monetary Factors

Figure 5 shows the monetary factors. For CPI, the term structure of risk premia is negative and statistically significant at all horizons, with the posterior median at 24 months being roughly three times larger in magnitude than at horizon zero. This means that an asset whose payoff covaries with inflation commands increasingly negative risk compensation as the investment horizon lengthens; put differently, investors accept lower expected returns on inflation-hedging positions, and this willingness grows with the horizon. In our setting, that negative premium is economically intuitive. The mimicking portfolio appears to behave less like a conventional return-seeking portfolio and more like a hedge, or insurance contract, against adverse nominal states, in which inflation is high and real consumption opportunities deteriorate. The portfolio composition in Figures 3 and 4 is consistent with that interpretation: qualitatively, the hedge looks like a position that is long EUR relative to USD and short duration, especially through sizeable euro exposure and negative positions in longer-duration government-bond instruments. That is the sort of payoff one would expect to be valuable when inflation is adverse for nominal bonds and, in this sample, unfavourable for the dollar. This reading is also consistent with the broader literature. N. L. Roussanov, Liu, and Fang (2021) show that core inflation risk commands a negative premium across asset classes, while Cieslak and Pflueger (2023) emphasise that the sign of inflation compensation depends on the underlying inflation process rather than on inflation mechanically per se. On the bond side,

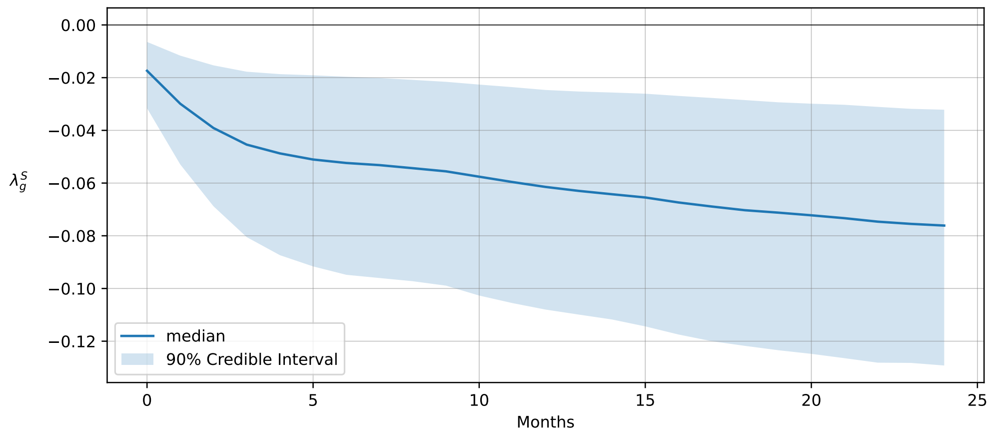

(a) US CPI

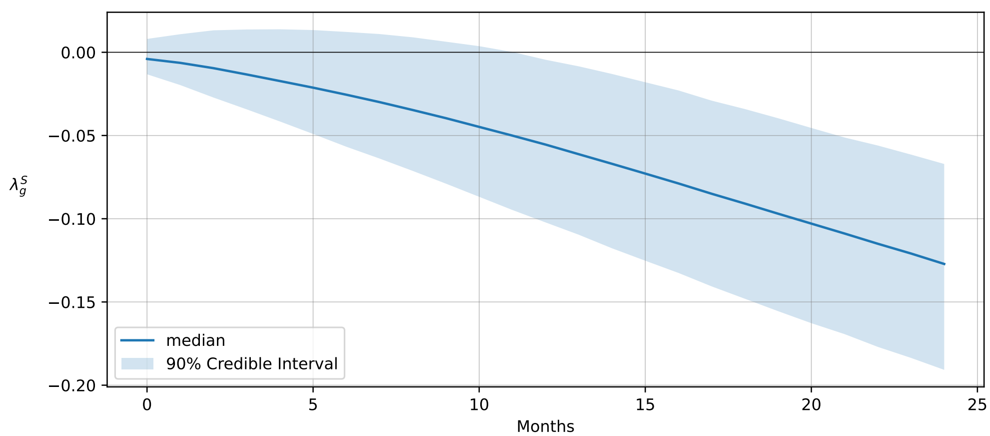

(b) Fed Funds Rate

Figure 5: Estimated term structures of risk premia for monetary variables as factors g_t.

Campbell, Sunderam, and Viceira (2009) show that nominal bonds are exposed to inflation risk in a way that can make them poor hedges in inflationary states. Accordingly, we interpret the CPI mimicking portfolio as evidence of an inflation hedge in our framework and sample, not as a universal claim that every inflation shock mechanically depreciates the dollar against the euro.

That said, the direction of interpretation need not run only from inflation to returns. The portfolio composition in Figures 3 and 4 is equally consistent with a priced FX disturbance, or broader nominal-financial shock, that subsequently passes through into import prices and consumer prices. The exchange-rate pass-through literature makes this interpretation natural.

Campa and Linda S. Goldberg (2005) document partial pass-through into import prices across OECD economies; Linda S Goldberg and Campa (2010) show that CPI responsiveness depends importantly on distribution margins and imported-input use; and Forbes, Hjortsoe, and Nenova (2018) show that pass-through into prices depends on the shock that moved the exchange rate in the first place. On this alternative reading, what we label as the CPI term structure may partly reflect the gradual pass-through of a priced financial shock into observed inflation, rather than inflation being the primitive shock itself.

For the Fed Funds Rate, the term structure is also negative but becomes significant only for horizons beyond 12 months. At short horizons, the credible interval includes zero, consistent with the idea that the policy rate's impact on asset returns accumulates gradually as the transmission mechanism of monetary policy works through the economy. A natural interpretation is that inflation is closer to the underlying nominal disturbance that markets price on impact, whereas the policy rate is an endogenous instrument that responds to that disturbance with inertia, smoothing, and a dependence on the expected future macro outlook. That distinction helps explain why CPI carries a short-horizon premium while the funds rate does not, even though both display increasingly negative premia at longer horizons. The close similarity between the inflation and Fed Funds Rate mimicking portfolios further suggests that the two are loading on a common nominal shock, with the policy rate reflecting that shock only more gradually. This interpretation also sits well with the monetary-policy and FX literature. Eichenbaum and Evans (1995) document large exchange-rate responses to US monetary shocks; Jarocinski and Karadi (2020) and Miranda-Agrippino and Ricco (2021) show that measured policy surprises mix genuine policy innovations with information effects unless they are carefully disentangled; and Müller, Wolf, and Hettig (2024) show that exchange-rate expectations adjust sluggishly after monetary shocks, generating delayed overshooting. In that sense, the long-horizon significance of the Fed Funds premium is not a contradiction of the CPI result, but a refinement of it: the inflation shock is the more immediate nominal object, while the policy-rate factor is the slower-moving policy manifestation of the same disturbance. As with CPI, the negative premium is intuitive if the mimicking portfolio is interpreted as insurance against adverse increases in the policy rate.

These two monetary factors therefore illustrate the central insight of the BHJ framework in a particularly transparent way. A conventional single-period exercise would detect that inflation is priced, but would understate how much larger the compensation becomes at business-cycle horizons; it would likely conclude that the Fed Funds Rate is not priced at all. Here, by contrast, the parametric propagation structure reveals that both objects carry economically meaningful long-horizon premia, with the difference between them stemming from the fact that markets price nominal disturbances immediately but observed policy rates absorb those disturbances only with delay.

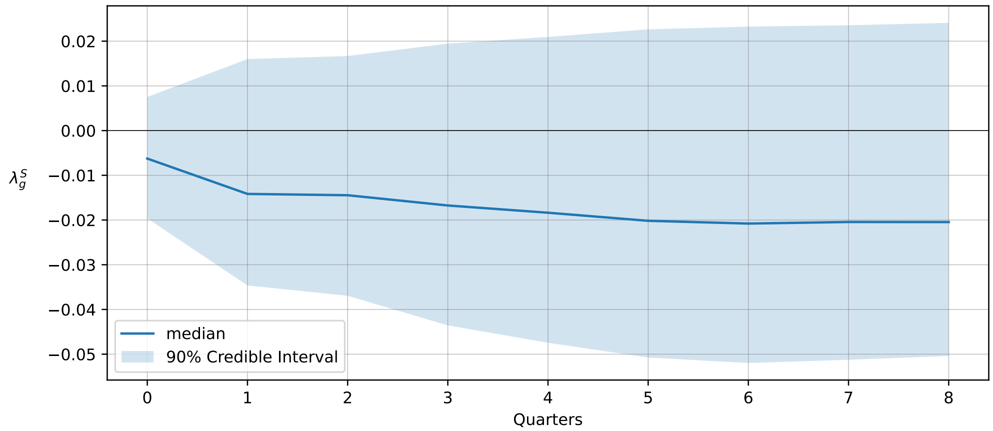

(a) US GDP

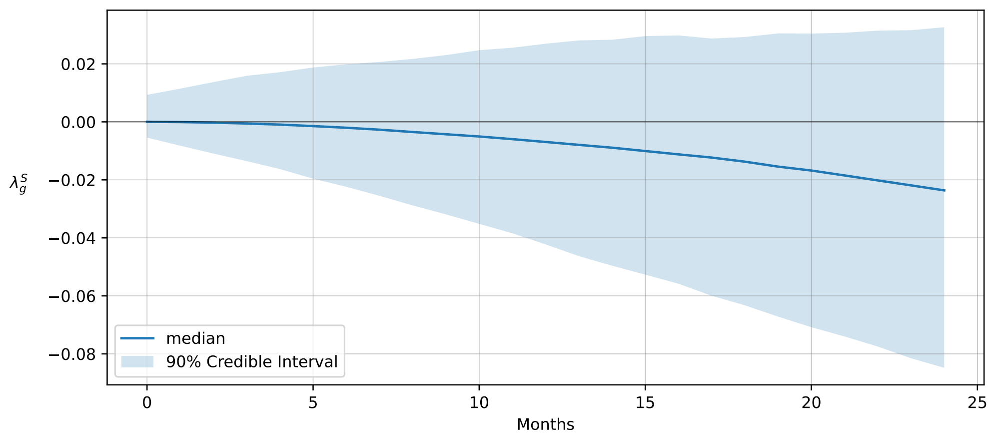

(b) 10-year Real Yield

Figure 6: Estimated term structures of risk premia for real variables as factors g_t.

4.4 Real Variables

Figure 6 turns to the real variables. In contrast to BHJ's equity results, where GDP carries significant and increasing risk premia at business-cycle frequencies, GDP is not significantly priced in our bond-and-currency cross section at any horizon. The same holds for the 10-year TIPS real yield, whose credible intervals include zero throughout. These null results are striking, but they should be framed carefully. For FX, they are consistent with the long exchange-rate-disconnect literature, beginning with Meese and Rogoff (1983) and formalised more sharply by Engel and West (2005), in which standard macro quantities have limited success in accounting for exchange-rate movements. For nominal bonds, however, the literature is more mixed: Ludvigson and Ng (2009) find that macro factors help forecast bond excess returns, while Bauer and Hamilton (2018) argue that some evidence for predictors beyond the standard yield-curve factors is less robust than initially believed. Our results should therefore not be read as showing that real variables never matter for bond premia or for macro-finance more generally. The narrower conclusion is that, in this particular FX-and-nominal-rates cross section, the directly priced component of expected returns appears more tightly linked to nominal monetary disturbances than to real shocks. Real activity may still matter, but mainly through its effect on the expected policy path rather than as a separately priced source of return variation within a parsimonious two-factor return space. That dominance result is, in our view, the right way to state the difference between our findings and those of BHJ.

4.5 Financial Variables

Figure 7 shows the three financial factors. The yield curve slope displays a positive and increasing term structure of risk premia, significant at all horizons. The US–Europe yield differential shows a negative and steepening term structure, also significant across all horizons. Yield curve slopes and yield differentials are used as carry measures in the style factors literature. Both results connect naturally to the carry and international term-structure literatures. Koijen et al. (2018) show that carry predicts returns across asset classes, including global bonds and currencies. Lustig, Stathopoulos, and Lustig, Stathopoulos, and Verdelhan (2019) show that currency carry premia are inherently term-structured and closely related to local bond term premia, while Chernov, Dahlquist, and Lochstoer (2023) stress that multi-horizon currency premia are tightly linked to sovereign yield differentials and bond risk premia. More broadly, the slope of the yield curve has long been interpreted as a summary statistic for future interest-rate dynamics and macro conditions in both finance and macroeconomics, as in Estrella and Hardouvelis (1991). In our setting, these two factors are probably best viewed less as structural shocks than as bridge variables: they connect the BHJ-style priced-shock framework to the literature that sorts currencies and rates on interest-rate differentials and curve information. Once the macro object g_t is generalised from a scalar to a vector of country-specific macro metrics, that bridge should become direct, allowing the framework to quantify how local versus US macro conditions propagate into USD returns across horizons and across markets.

The VIX, despite its prominent role as a measure of global risk appetite, does not show statistically significant risk premia in our setting: the credible intervals include zero at all horizons, though the point estimates are positive and mildly increasing. This is a notable contrast with BHJ's equity evidence, where the VIX exhibits a significant and downward-sloping term structure. The right interpretation, however, is not that VIX-related forces are irrelevant for FX or bond markets in general. On the contrary, the broader currency literature often treats global volatility and risk appetite as central state variables: Brunnermeier, Nagel, and Pedersen (2008) link carry returns to crash risk and funding conditions; Menkhoff, Sarno, Menkhoff et al. (2012) show that global FX volatility prices the cross section of carry portfolios; and Caballero and Doyle (2012) show that carry returns are closely related to VIX-rolldown strategies. Our result is therefore more specific: once nominal monetary factors are isolated, the VIX does not appear as a separately priced factor in this cross section. That strongly suggests that the macro forces driving currencies and nominal rates in our sample are distinct from the shocks that dominate time variation in equity risk premia.

5 Conclusion

This paper applies the priced-macro-shocks framework of Bryzgalova, Huang, and Julliard (2024) to a cross section of foreign exchange forwards and government bonds, delivering a term structure of macro risk premia in two cornerstone asset classes beyond equities. We derive a closed-form decomposition of the model-implied mimicking portfolios into a horizon-invariant

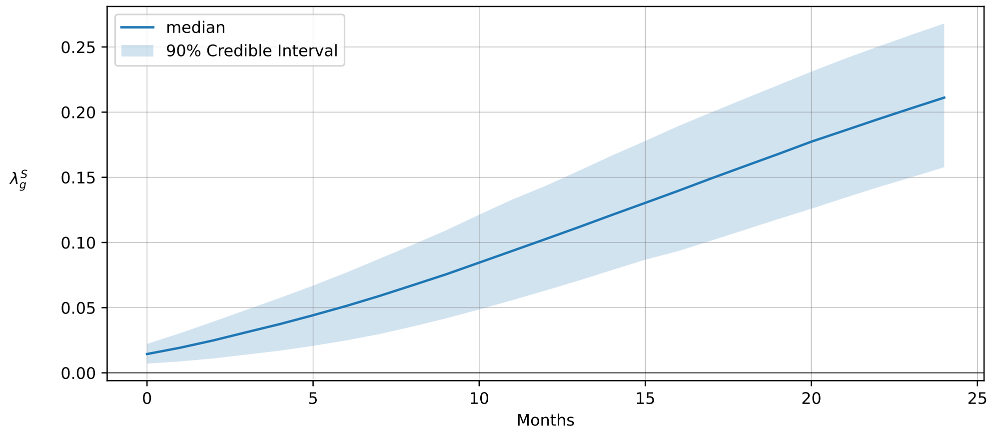

(a) US Yield Curve Slope

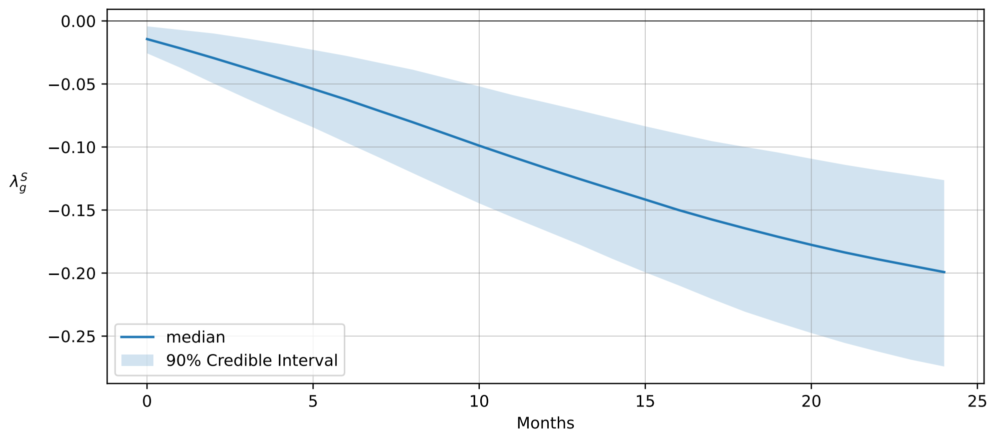

(b) Yield Differential

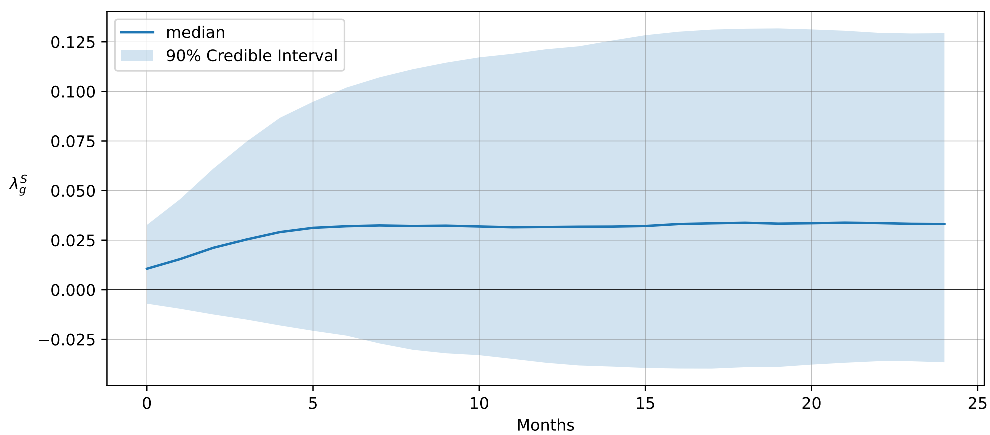

(c) VIX

Figure 7: Estimated term structures of risk premia for financial variables as factors g_t.

base portfolio and a horizon-specific scalar, which makes the composition of macro risk directly interpretable in a cross section of 29 tradable instruments.

Across the significant factors, the common qualitative feature is that premia are economically small or statistically weak at very short horizons and grow in magnitude as the horizon extends. There are two nearby ways to read this pattern. One is that nominal macro shocks are priced directly, but their economically relevant compensation becomes more visible only at business-cycle horizons as the propagation mechanism unfolds. The other is that priced financial shocks are absorbed first by FX and rates markets and only gradually pass through into inflation and, with further delay, into policy rates. Our framework does not force a sharp distinction between these narratives; what it establishes is that the priced disturbance in this cross section is concentrated in the nominal-financial block rather than in real activity when working with currencies and bonds.

The contrast with BHJ's equity results is informative. In equities, real-activity variables, such as GDP, bear the dominant long-horizon premia, while the VIX displays a significant and downward-sloping term structure. In our cross-section of currencies and bonds, real variables are not significantly priced, and the VIX is likewise insignificant. Instead, the significant premia belong to inflation, the policy rate, the yield curve slope, and the US–Europe 3-month yield differential. This difference is economically coherent: equities are claims on real cash flows and should reward exposure to real shocks, whereas nominal rates and exchange rates are more mechanicaly linked to the monetary transmission mechanism.

A methodological contribution of this paper is to show that the BHJ parametric structure is particularly well suited to settings where the cross section is informative but not asymptotically large. With 29 assets, non-parametric mimicking portfolios based on multi-period overlapping returns would be noisy and potentially unreliable. The Bayesian implementation pools information across horizons through an economically interpretable propagation structure, delivering coherent posterior distributions over the entire term structure and over portfolio weights at near-zero additional computational cost. The closed-form result for the mimicking portfolios extends the original framework and provides practitioners with directly tradable portfolios whose weights have transparent economic content.

Several extensions are natural. First, allowing the latent factors to be serially correlated would make the base portfolio itself state- and horizon-dependent, relaxing the strong restriction that only leverage changes across horizons. Second, extending the macro factor from a scalar to a vector of country-specific variables would allow the framework to disentangle how local versus US macro conditions propagate into global bond and currency returns, and identify the premia of common drivers of these factors. Third, incorporating persistent dynamics for the latent factors in the spirit of the time-varying risk premia extension in BHJ would permit the study of whether the term structure of macro premia in rates and FX markets is countercyclical, as it is in equities. Finally, expanding the cross section with additional assets would strengthen the spanning condition and sharpen identification of the priced shocks.

References

Adrian, Tobias, Richard K. Crump, and Emanuel Moench (2012). "Pricing the Term Structure with Linear Regressions". In: *SSRN Electronic Journal*.

Ang, Andrew and Monika Piazzesi (May 2003). "A no-arbitrage vector autoregression of term structure dynamics with macroeconomic and latent variables". In: *Journal of Monetary Economics* 50.4, pp. 745–787.

Bauer, Michael D. and James D. Hamilton (Feb. 1, 2018). "Robust Bond Risk Premia". In: *The Review of Financial Studies* 31.2, pp. 399–448.

Brunnermeier, Markus, Stefan Nagel, and Lasse Pedersen (Nov. 2008). *Carry Trades and Currency Crashes*. w14473. Cambridge, MA: National Bureau of Economic Research, w14473.

Bryzgalova, Svetlana, Jiantao Huang, and Christian Julliard (2024). "Macro Strikes Back: Term Structure of Risk Premia and Market Segmentation". In: *SSRN Electronic Journal*.

Caballero, Ricardo and Joseph Doyle (Dec. 2012). *Carry Trade and Systemic Risk: Why are FX Options so Cheap?* w18644. Cambridge, MA: National Bureau of Economic Research, w18644.

Campa, José Manuel and Linda S. Goldberg (Nov. 2005). "Exchange Rate Pass-Through into Import Prices". In: *Review of Economics and Statistics* 87.4, pp. 679–690.

Campbell, John, Adi Sunderam, and Luis Viceira (Feb. 2009). *Inflation Bets or Deflation Hedges? The Changing Risks of Nominal Bonds*. w14701. Cambridge, MA: National Bureau of Economic Research, w14701.

Chernov, Mikhail, Magnus Dahlquist, and Lars Lochstoer (Apr. 2023). "Pricing Currency Risks". In: *The Journal of Finance* 78.2, pp. 693–730.

Cieslak, Anna and Carolin Pflueger (Nov. 1, 2023). "Inflation and Asset Returns". In: *Annual Review of Financial Economics* 15.1, pp. 433–448.

Cochrane, John H and Monika Piazzesi (Feb. 1, 2005). "Bond Risk Premia". In: *American Economic Review* 95.1, pp. 138–160.

Eichenbaum, M. and C. L. Evans (Nov. 1, 1995). "Some Empirical Evidence on the Effects of Shocks to Monetary Policy on Exchange Rates". In: *The Quarterly Journal of Economics* 110.4, pp. 975–1009.

Engel, Charles and Kenneth D. West (June 2005). "Exchange Rates and Fundamentals". In: *Journal of Political Economy* 113.3, pp. 485–517.

Estrella, Arturo and Gikas A. Hardouvelis (June 1991). "The Term Structure as a Predictor of Real Economic Activity". In: *The Journal of Finance* 46.2, pp. 555–576.

Forbes, Kristin, Ida Hjortsoe, and Tsvetelina Nenova (Sept. 2018). "The shocks matter: Improving our estimates of exchange rate pass-through". In: *Journal of International Economics* 114, pp. 255–275.

Giglio, Stefano and Dacheng Xiu (July 1, 2021). "Asset Pricing with Omitted Factors". In: *Journal of Political Economy* 129.7, pp. 1947–1990.

Goldberg, Linda S and José Manuel Campa (May 2010). "The Sensitivity of the CPI to Exchange Rates: Distribution Margins, Imported Inputs, and Trade Exposure". In: *Review of Economics and Statistics* 92.2, pp. 392–407.

Jarocinski, Marek and Peter Karadi (Apr. 1, 2020). "Deconstructing Monetary Policy Surprises: The Role of Information Shocks". In: *American Economic Journal: Macroeconomics* 12.2, pp. 1–43.

Joslin, Scott, Kenneth J. Singleton, and Haoxiang Zhu (Mar. 2011). "A New Perspective on Gaussian Dynamic Term Structure Models". In: *Review of Financial Studies* 24.3, pp. 926–970.
Koijen, Ralph S.J. et al. (Feb. 2018). "Carry". In: *Journal of Financial Economics* 127.2, pp. 197–225.
Litterman, Robert B and Josè Scheinkman (June 30, 1991). "Common Factors Affecting Bond Returns". In: *The Journal of Fixed Income* 1.1, pp. 54–61.
Ludvigson, Sydney C. and Serena Ng (Dec. 2009). "Macro Factors in Bond Risk Premia". In: *Review of Financial Studies* 22.12, pp. 5027–5067.
Lustig, Hanno, Nikolai Roussanov, and Adrien Verdelhan (Nov. 2011). "Common Risk Factors in Currency Markets". In: *Review of Financial Studies* 24.11, pp. 3731–3777.
Lustig, Hanno, Andreas Stathopoulos, and Adrien Verdelhan (2019). "The Term Structure of Currency Carry Trade Risk Premia". In: *American Economic Review* 109.12, pp. 4142–4177.
Lustig, Hanno and Adrien Verdelhan (Feb. 1, 2007). "The Cross Section of Foreign Currency Risk Premia and Consumption Growth Risk". In: *American Economic Review* 97.1, pp. 89–117.
Meese, Richard A. and Kenneth Rogoff (Feb. 1983). "Empirical exchange rate models of the seventies". In: *Journal of International Economics* 14.1, pp. 3–24.
Menkhoff, Lukas et al. (Apr. 2012). "Carry Trades and Global Foreign Exchange Volatility". In: *The Journal of Finance* 67.2, pp. 681–718.
Miranda-Agrippino, Silvia and Hélène Rey (Oct. 2021). *The Global Financial Cycle*. w29327. Cambridge, MA: National Bureau of Economic Research, w29327.
Miranda-Agrippino, Silvia and Giovanni Ricco (July 1, 2021). "The Transmission of Monetary Policy Shocks". In: *American Economic Journal: Macroeconomics* 13.3, pp. 74–107.
Müller, Gernot J., Martin Wolf, and Thomas Hettig (July 1, 2024). "Delayed Overshooting: The Case for Information Rigidities". In: *American Economic Journal: Macroeconomics* 16.3, pp. 310–342.
Roussanov, Nikolai L., Yang Liu, and Xiang Fang (2021). "Getting to the Core: Inflation Risks Within and Across Asset Classes". In: *SSRN Electronic Journal*.
Verdelhan, Adrien (Feb. 2010). "A Habit-Based Explanation of the Exchange Rate Risk Premium". In: *The Journal of Finance* 65.1, pp. 123–146.

A Proofs

A.1 Horizon-Specific Mimicking Portfolios

To prove this proposition, we transform the 1-period moments of the model into their S-horizon analogous.

Proof. The weights of the horizon-specific factor-mimicking portfolio are given by

$$w_S^{MP} = cov\left(r_{t-1\to t+S}\right)^{-1} cov\left(r_{t-1\to t+S}, g_{t-1\to t+S}\right) = \Sigma_{rr,S}^{-1}\Sigma_{rg,S}$$

where $\Sigma_{rg,S}$ is a column vector with the correlations of the asset returns with the factors.

Define the cumulative horizon objects as in the mimicking-portfolio definition

$$r_{t-1\to t+S} = \sum_{i=-1}^{S} r_{t+i} \qquad g_{t-1\to t+S} = \sum_{i=-1}^{S} g_{t+i}$$

Using the fact that, in equation 1, both υ_t and $w_{r,t}$ are independet, then the S-horizon covariance matrix of returns is given by

$$\Sigma_{rr,S} = Var\left(\sum_{i=-1}^{S} r_{t+i}\right) = \sum_{i=-1}^{S} Var\left(r_{t+i}\right) = \left(S+2\right)\Sigma_{rr} \tag{8}$$

where the term Σ_{rr} is already part of the Gibbs sampling procedure.

We now show that Σ_{rg} is implied by β_υ, Σ_υ, η_g and ρ. The S-horizon cross-covariances of returns and the factor is given by

$$\Sigma_{rg,S} = Cov\left(\sum_{i=-1}^{S} r_{t+i}, \sum_{j=-1}^{S} g_{t+j}\right) = \sum_{i=-1}^{S}\sum_{j=-1}^{S} Cov\left(r_{t+i}, g_{t+j}\right)$$

$$\Sigma_{rg,S} = \sum_{i=-1}^{S}\sum_{j=-1}^{S} Cov\left(\beta_\upsilon\left(\upsilon_{t+i}-\mu_\upsilon\right) + w_{r,t+i}, \sum_{s=0}^{\bar{S}} \rho_s \eta_g'\left(\upsilon_{t+j-s}-\mu_\upsilon\right) + w_{g,t+j}\right)$$

Since $w_{r,t}$, $w_{g,t}$ and υ_t are orthogonal

$$\Sigma_{rg,S} = \sum_{i=-1}^{S}\sum_{j=-1}^{S} Cov\left(\beta_\upsilon \upsilon_{t+i}, \sum_{s=0}^{\bar{S}} \rho_s \eta_g' \upsilon_{t+j-s}\right)$$

$$\Sigma_{rg,S} = \sum_{i=-1}^{S}\sum_{j=-1}^{S}\left[\sum_{s=0}^{\bar{S}} \rho_s \beta_\upsilon Cov\left(\upsilon_{t+i}, \upsilon_{t+j-s}\right)\eta_g\right]$$

Because of the independence of υ_t

$$Cov\left(\upsilon_{t+i}, \upsilon_{t+j-s}\right) = \begin{cases} \Sigma_\upsilon & \text{if } s = j-i \\ 0 & \text{otherwise} \end{cases}$$

so for a given pair (i,j), there is a single term of the third sum that is different from zero. So we can rewrite

$$\Sigma_{rg,S} = \sum_{i=-1}^{S}\sum_{j=-1}^{S} \mathbb{I}_{\left\{0\le j-i\le \bar{S}\right\}} \rho_{j-i}\beta_\upsilon \Sigma_\upsilon \eta_g$$

$$\Sigma_{rg,S} = \beta_\upsilon \Sigma_\upsilon \eta_g \sum_{i=-1}^{S}\sum_{j=-1}^{S} \mathbb{I}_{\left\{0\le j-i\le \bar{S}\right\}} \rho_{j-i}$$

Define the index $l = j-i$

$$\Sigma_{rg,S} = \beta_\upsilon \Sigma_\upsilon \eta_g \sum_{l=0}^{\bar{S}} \rho_l \# \left\{(i,j) \in \{-1,\dots,S\}^2 \,|\, j-i = l\right\} \tag{9}$$

The number of elements per term is

$$-1 \leq i \leq S$$

$$-1 \leq j \leq S \qquad \Leftrightarrow \qquad -1 \leq i + l \leq S \qquad \Leftrightarrow \qquad -1 - l \leq i \leq S - l$$

which implies that $i \in \{-1, \ldots, S - l\}$, and the number of elements for a given l is

$$\# \left\{ (i,j) \in \{-1, \ldots, S\}^2 \,|\, j - i = l \right\} = \begin{cases} S + 2 - l & \text{if } l \in \{0, \ldots, S+1\} \\ 0 & \text{otherwise} \end{cases}$$

Substitute this back in equation (9)

$$\Sigma_{rg,S} = \beta_v \Sigma_v \eta_g \sum_{l=0}^{\bar{S}} \rho_l \left(S + 2 - l\right) \mathbb{I}_{\{0 \leq l \leq S+1\}}$$

$$\Sigma_{rg,S} = \beta_v \Sigma_v \eta_g \sum_{l=0}^{\min\left(\bar{S}, S+1\right)} \rho_l \left(S + 2 - l\right) \tag{10}$$

Putting together equations (8) and (10)

$$w_S^{MP} = \frac{1}{S+2} \left[\sum_{l=0}^{\min\left(\bar{S}, S+1\right)} \rho_l \left(S + 2 - l\right) \right] \Sigma_{rr}^{-1} \beta_v \Sigma_v \eta_g$$

□