Macro Strikes Back: Term Structure of Risk Premia and Market Segmentation*

Svetlana Bryzgalova† Jiantao Huang‡ Christian Julliard§

March 2024

Abstract

We develop a unified framework to study the term structure of risk premia of nontradable factors. Our method delivers level and time variation of risk premia, uncovers their propagation mechanism, is robust to misspecification and weak identification, and allows for segmented markets. Most macroeconomic factors are weakly identified at quarterly frequency, but have increasing (unconditional) term structures with large risk premia at business cycle horizons. Moreover, the slopes of their term structures are strongly procyclical. Most macroeconomic and intermediary-based factors command similar risk premia in equity and corporate bond markets, while we find strong evidence of segmentation for other factors.

Keywords: Macro-finance, asset pricing, risk premia, linear factor models, Bayesian inference, term structures, market segmentation.

JEL Classification Codes: C11, C58, E27, E44, G12, G17.

---

*Any errors or omissions are the responsibility of the authors. For helpful comments, discussions and suggestions, we thank Mikhail Chernov, Magnus Dahlquist, Xiang Fang, Marcelo Fernandes, Bernardo Guimaraes, Ian Martin, Lasse Pedersen, and seminar and conference participants at Copenhagen Business School, CUHK, FGV São Paulo, FinEML, London School of Economics.

†Department of Finance, London Business School; sbryzgalova@london.edu

‡Faculty of Business and Economics, The University of Hong Kong, huangjt@hku.hk

§Department of Finance, FMG, and SRC, London School of Economics, and CEPR; c.julliard@lse.ac.uk.

1 Introduction

Macroeconomic risk, often related to technology, consumption, or intermediary capital, is at the heart of most equilibrium-based asset pricing models. Yet reliable detection of macroeconomic risk premia remains elusive: 1) different time horizons often provide drastically different estimates of the priced risk, 2) most empirical models are widely known to be misspecified, calling for methods robust to the nature and number of risk factors, and 3) the weak contemporaneous link between macroeconomic factors and asset returns often leads to model parameters being weakly identified at best, causing a fundamental inference problem. All of these issues contribute to the empirical macrofinance disconnect.

We propose a new estimation framework that addresses all of the above. Unlike any existing approach, it produces not only reliable risk premia estimates but also their whole term structure in an internally consistent framework. Our method, which leverages the fact that many nontradable factors are persistent, relies on three key ingredients: 1) the moving average (MA) representation of the persistent component of the factor, driven by either priced or non-priced shocks, 2) an approximate factor structure for a wide cross-section of asset returns, which recovers priced shocks and is robust to model misspecification (Chamberlain and Rothschild (1983) and Giglio and Xiu (2021)), and 3) the hierarchical Bayesian inference method of Bryzgalova, Huang, and Julliard (2023a), which recovers both time series and cross-sectional properties of risk factors, and is by design robust to weak identification.

Our framework accurately identifies not only the joint comovement between nontradable factors and asset returns but also their propagation mechanism, and hence recovers the whole *term structure of risk premia*. As we show, the latter is crucial in assessing the role of macroeconomic risks in asset returns. We find that many macroeconomic variables (e.g., industrial production, consumption, and GDP growth) have increasing unconditional term structures and carry large and significant risk premia at business cycle frequencies (two–three years). Furthermore, their term structure is particularly steep during expansions and inverted during recessions.

Our findings are not a simple byproduct of factor persistence. We find that similarly persistent risk factors can have increasing (liquidity of Pástor and Stambaugh (2003)), flat (intermediary factor of He et al. (2017)), or decreasing term structures (VIX), or no significant risk premia at all (capital share growth of Lettau et al. (2019)). As we show, risk premia over different horizons can be directly mapped into per-period Sharpe ratios of factor-mimicking portfolios hedging multi-horizon priced innovations. The economic magnitude of our findings is striking: At business cycle frequencies, risk premia carried by, for example, industrial production, GDP, and consumption, are as large as that of the market. Crucially, all of our results are not based on ad-hoc frequency-based procedures, but are instead fully determined by the structural parameters of the model.

Our framework provides sharp identification of risk premia and reliable inference and is rooted in economic theory: Equilibrium asset prices are jump variables. Hence, news about current and *future* priced states should immediately be reflected into prices, albeit they might manifest in nontradable variables only with delay. Similar to Giglio and Xiu (2021), we leverage the fact that, while the actual drivers of asset returns are identifiable only up to a rotation, conditional and unconditional risk premia of observable factors are not affected by this issue and can, therefore, be reliably recovered from the data. As a result, our estimator is robust to the omitted variable bias, measurement error, and weak identification. Contrary to the existing literature, our method allows for the joint modeling of factor and return dynamics over different horizons, providing coherent insights into the whole term structure of risk premia. Tackling an inference problem in this setup would be challenging, if not infeasible, in frequentist estimation. Instead, we develop a simple Gibbs algorithm for Bayesian posterior sampling, with all the conditional posterior distributions available in closed form. Thus, we deliver not only point estimates of risk premia and deep model parameters but also valid credible intervals for all the objects of interest.

It is widely known that some nontradable factors have higher exposure to asset returns at longer horizons. (See, e.g., Jagannathan and Wang (2007), Cohen et al. (2009), and Hansen et al. (2008)) We uncover the mechanism generating this phenomenon and show that, in these cases, risk premia at different horizons are driven by the *same* priced innovations that slowly propagate through the nontradable risk factor. As a result, we also explain why many risk factors are statistically weak at quarterly frequency, yet become strongly identified at longer horizons. Consider GDP growth, for example. Although contemporaneous asset return shocks account for only 4% of the variation in GDP growth, they contribute to over 20% of its time series variation at business cycle frequencies. In such a case, the term structure of risk premia effectively boosts the signal-to-noise ratio of priced shocks in nontradable factors and sharply identifies the common priced component (which would normally be weak at best).

The empirical asset pricing literature has long recognized the persistent nature of many nontradable risk factors. As a result, researchers would usually first extract the AR(1) innovation from a factor and then proceed with measuring its risk premia via Fama-MacBeth (FM) regressions or the Generalized Method of Moments (GMM).[1] However, as we show, this common procedure fails to recover the true sources of priced risk. First, the conditional mean of the macroeconomic variable could follow a process different from AR(1). Second, the persistent component of the variable does not need to be driven only by priced shocks. As we show, AR(1) residuals do not recover actual priced innovations in many factors, leading to a significant bias in risk premia estimates. Our approach, rooted in the Wold decomposition, relies on the flexible MA representation of the risk factor. It efficiently separates priced and unpriced innovations and restores reliable inference on risk premia.

We use a large cross-section of 275 equity portfolios to estimate the term structures of risk premia of many nontradable risk factors. Contrary to the standard one-period inference, we find that a large part of the factors' conditional mean is driven by priced shocks, slowly propagating through the time series. Their overall dynamics display clear business cycle patterns and are common across many different macroeconomic factors.

Many risk factors are characterized by increasing term structures of risk premia. For example, while the risk premium of GDP growth is only 0.02 at the quarterly horizon, it increases to 0.14 at the two-year horizon and is strongly significant (while being spanned by the same shocks). Furthermore, term structures of macro risk exhibit strong commonality in their business cycle behavior: The average level is strongly procyclical, with larger risk premia during expansion and significantly reduced, or even becoming negative, during recessions; the slope of the term structures is also strongly procyclical, with longer (two- to three-year) maturities characterized by high conditional Sharpe ratios immediately before economic contractions.

We also observe factors commanding flat or downward-sloping unconditional term structures of risk premia. For example, the VIX risk premium is -0.11 at the monthly frequency, but

[1] See He et al. (2017), Pástor and Stambaugh (2003), and Giglio and Xiu (2021), among others.

its two-year counterpart is only -0.03. This observation is reassuring since the sign of the VIX risk premium, and its term structure, are mostly consistent with previous findings based on VIX derivatives (Eraker and Wu (2014), Dew-Becker et al. (2017), and Johnson (2017)). Intermediary factors (Adrian et al. (2014) and He et al. (2017)) carry significantly positive unconditional risk premia with a flat term structure. Furthermore, our estimate of risk premia for the intermediary factor of He et al. (2017) is rather close to the average return of its tradable version, providing further validation to our findings.

When asset markets are fully integrated, a risk factor commands the same risk premia across all of them. Recent empirical literature, however, provides growing support for partial segmentation in financial markets. (See Brandt et al. (2006), Choi and Kim (2018), Sandulescu (2022), and Patton and Weller (2022)) Motivated by this, we extend our framework to allow for partially or completely segmented markets. Intuitively, systematic priced shocks that drive asset returns could be (partially) different across market segments. As a result, the same risk factor could carry different risk premia when estimated in different asset classes. Our framework allows us not only to estimate the market-specific term structure of risk premia but also to formally test for its equality across market segments.

We compare the term structures of risk premia of nontradable factors, spanned by a large cross-section of equity returns, with those recovered from the cross-section of 40 corporate bond portfolios of Elkamhi, Jo, and Nozawa (2023). We find that many factors indeed command different risk premia in these asset classes. For example, since bonds typically have fixed nominal payoffs, oil price shocks (leading to inflation) command a large and positive risk premium. At the same time, equity markets provide a natural hedge against these shocks and produce a negative risk premium. The TED spread, widely known to increase with lower funding liquidity (Brunnermeier and Pedersen (2009) and Asness et al. (2013)), has a flat term structure of negative risk premia in corporate bonds but is not priced in stocks.

While there are many examples of partial market segmentation, many factors carry very similar risk premia across markets (both economically and statistically). The most striking example is the VIX index, which has an almost identical term structure of risk premia in both stock and bond markets. Other examples with nearly homogeneous risk premia estimates across the two asset classes include GDP growth, industrial production growth, durable and nondurable consumption growth, and the intermediary factor of He et al. (2017). Our paper indicates that despite growing evidence of market segmentation, there are many common fundamental risks driving both asset classes.

To further highlight the strength and robustness of our estimation approach, we conduct extensive simulations and study the empirical size and power of the procedure in detecting the persistent priced component of nontradable factors. We find that the Bayesian credible intervals provide proper posterior coverages of the pseudo-true risk premia and the risk premia difference in the case of market segmentation. Moreover, we show that the MA-based approach is essential in identifying the priced component in persistent factors and leads to a significant boost in the power of detecting segmented risk premia. Finally, our Bayesian inference remains valid even in the presence of persistent, yet weak, risk factors.

The remainder of the paper is organized as follows. In the next subsection we review the most closely related literature and our contribution to it. Section 2 outlines our estimation framework and its properties, while Section 3 provides simulation evidence on the power of the method in realistically small samples. Section 4 presents our empirical findings, and Section 5 concludes. Additional results, proofs, derivations, and a detailed description of the data sources, are reported in the Internet Appendix.

1.1 Closely Related Literature

Our paper naturally relates to the inference on risk premia in linear factor models. As shown by the past literature (e.g., Kan and Zhang (1999a,b), Kleibergen (2009), and Kleibergen and Zhan (2015)), weak factors invalidate risk premia estimates and cross-sectional fit of traditional FM and GMM estimators. Several studies (e.g., Kan et al. (2013), Gospodinov et al. (2014, 2019), Bryzgalova (2015), Kleibergen and Zhan (2020), and Bryzgalova et al. (2023a)) propose methods that are robust to weak factors and misspecification. Giglio and Xiu (2021) further emphasize that standard estimators of risk premia are biased if some priced factors are omitted and propose a three-pass method to resolve the issue. Likewise, Giglio et al. (2023) propose a supervised principal component analysis method to recover the risk premia of weak factors. Similarly, our method aligns with the literature that relates to principal component analysis in asset pricing (e.g., Chamberlain and Rothschild (1983), Connor and Korajczyk (1986, 1988), Kozak et al.

(2018, 2020), and Kelly et al. (2019)). Unlike them, we incorporate the dynamics of factors and returns to elicit the *entire* term structure of risk premia in an internally consistent manner. This additional dimension is economically meaningful since many variables, particularly macro variables, are significantly priced only at particular horizons.

Equilibrium macro-finance models have sharp and salient predictions for the term structures of risk premia of macro factors. For instance, as shown in Figure A1, the habit model of Campbell and Cochrane (1999) predicts flat term structures of risk premia for consumption and dividend growth.[2] Yet these same factors command upward-sloping term structures in the long-run risk model of Bansal and Yaron (2004). However, these predictions rely on ad hoc assumptions on cash flow dynamics and investors' preferences. To obtain model-free estimates, van Binsbergen et al. (2012) and van Binsbergen and Koijen (2017)) analyze traded dividend claims and observe a downward-sloping term structure of dividend risk, which contradicts the predictions of leading macro-finance models. Consequently, several equilibrium models (e.g., Belo et al. (2015), Hasler and Marfe (2016), Ai et al. (2018), and Kragt et al. (2020)) have been developed to explain this phenomenon. However, traded dividend strips data suffer from a short time series sample and liquidity concerns. Recent papers tackle these shortcomings by estimating either a regime-switching model (Bansal et al. (2021)) or an affine term structure model of expected returns and dividend growth (Giglio et al. (2023)). Both papers suggest an unconditionally upward-sloping, and conditionally procyclical, term structure of dividend risk – as we uncover for dividend growth risk premia. But crucially, our new method has much broader applicability than solely dividends, as it delivers the term structure of risk premia for *all* factors (traded and nontraded) of equilibrium models.

Our paper also connects to the large body of literature that emphasizes horizon-dependent risk premia (see, e.g., Chernov et al. (2021)). Extensive empirical evidence shows that consumption growth carries more significant premia at long horizons (Daniel and Marshall (1997), Parker and Julliard (2005), Jagannathan and Wang (2007), Hansen et al. (2008), Malloy et al. (2009), Ortu et al. (2013), Dew-Becker and Giglio (2016), and Bandi and Tamoni (2023)). In contrast, VIX (Eraker and Wu (2014), Dew-Becker et al. (2017), and Johnson (2017)) carry more sizable risk premia at short horizons. Our paper is motivated by these empirical facts and

[2]Calibration and derivation details can be found in Internet Appendix IA.4.

provides a much more extensive, and robust, investigation of the risk premia of more than 20 economic variables.

Furthermore, this paper contributes to the literature on market segmentation. Existing works either study the comovement between stochastic discount factors (SDFs) across markets (e.g., Chen and Knez (1995), Brandt et al. (2006), Kelly et al. (2023), Sandulescu et al. (2021), and Sandulescu (2022)) or tests heterogeneous risk premia for the same factor in different markets (e.g., Choi and Kim (2018), Chaieb et al. (2021), and Patton and Weller (2022)). Our paper is more closely related to the latter strand of literature. However, our econometric method accounts for both weak identification and omitted variable bias, which is essential to avoid spurious rejection of homogeneous risk premia. Furthermore, we test heterogeneous risk premia in the entire term structure, greatly increasing the power of the statistical test for serially dependent factors.

Finally, our paper is related to the recent developments of Bayesian econometrics in asset pricing (e.g., Barillas and Shanken (2018), Chib et al. (2020), Bryzgalova et al. (2023a), and Avramov et al. (2023)). Unlike most papers, which emphasize Bayesian model selection and/or aggregation, we estimate the posterior credible intervals of the term structure of risk premia.

2 Theory and Method

This section describes our Bayesian framework for estimating factors' risk premia. We aim to test whether a (covariance-stationary) factor g_t, either tradable or nontradable, is priced in a large cross-section of test assets. Throughout our analysis, we consider log variables; that is, g_t is the log growth rate of G_t between time $t-1$ and t, where G_t can be, for example, the portfolio value, GDP, or industrial production.

We denote the vector of log returns on N assets, in excess of the log risk-free rate (r_f), by $\boldsymbol{r_t} = (r_{1t}, \ldots, r_{Nt})^\top$. We further define the cumulative variable: $g_{t-1\to t+S} = \log(G_{t+S}) - \log(G_{t-1})$, which measures the multiperiod growth rate of G_t. Similarly, $\boldsymbol{r}_{t-1\to t+S}$ denote the cumulative log returns between time $t-1$ and $t+S$.

We assume a linear latent factor model for asset returns driven by K systematic factors, as follows:

$$\boldsymbol{r}_t = \boldsymbol{\mu}_r + \boldsymbol{\beta}_{\tilde{v}}\tilde{\boldsymbol{v}}_t + \boldsymbol{w}_{rt}, \quad \tilde{\boldsymbol{v}}_t \overset{\text{iid}}{\sim} \mathcal{N}(\mathbf{0}_K, \boldsymbol{I}_K), \quad \boldsymbol{w}_{rt} \overset{\text{iid}}{\sim} \mathcal{N}(\mathbf{0}_N, \boldsymbol{\Sigma}_{wr}), \quad \tilde{\boldsymbol{v}}_t \perp \boldsymbol{w}_{rt}, \tag{1}$$

where $\tilde{\boldsymbol{v}}_t$ are K uncorrelated latent factors with loadings $\boldsymbol{\beta}_{\tilde{v}}$, $\boldsymbol{w}_{rt}$ are unpriced idiosyncratic errors, and $\boldsymbol{\mu}_r$ denote expected returns of test assets. We relax the assumption of uncorrelated $\tilde{\boldsymbol{v}}_t$ in Section 2.2. We impose an approximate factor structure among asset returns, following Chamberlain and Rothschild (1983). Mathematically, the largest K eigenvalues of $\boldsymbol{r}_t$'s covariance matrix will explode as the number of assets goes to infinity (equivalently, the eigenvalues of $\boldsymbol{\beta}_{\tilde{v}}\boldsymbol{\beta}_{\tilde{v}}^\top$ will explode), while those of $\boldsymbol{\Sigma}_{wr}$ remain bounded. We allow for a certain degree of cross-sectional dependence of $\boldsymbol{w}_{rt}$, as discussed later in the simulation study. The number of latent factors, K, is assumed to be known in this section.

We further assume that factors' loadings, $\boldsymbol{\beta}_{\tilde{v}}$, can partially explain expected returns,

$$\boldsymbol{\mu}_r = \boldsymbol{\beta}_{\tilde{v}}\boldsymbol{\lambda}_{\tilde{v}} + \boldsymbol{\alpha}, \tag{2}$$

where $\boldsymbol{\lambda}_{\tilde{v}}$ denote risk premia associated with $\tilde{\boldsymbol{v}}_t$, and $\boldsymbol{\alpha}$ is a vector of pricing errors. In addition, we assume that each asset's pricing error, α_i, is independently and identically distributed (IID) and cross-sectionally independent of factor loadings. This form of model misspecification has been commonly used in the past literature (e.g., Kan et al. (2013), Gospodinov et al. (2014), Giglio and Xiu (2021), and Bryzgalova et al. (2023a)) and has a clear economic interpretation. Equation (2) is equivalent to a log SDF that is linear in latent factors $\tilde{\boldsymbol{v}}_t$, described as follows:

$$m_t = 1 - \boldsymbol{\lambda}_{\tilde{v}}^\top \tilde{\boldsymbol{v}}_t. \tag{3}$$

Since $\tilde{\boldsymbol{v}}_t$ have an identity covariance matrix, their risk prices are identical to risk premia.

The Wold decomposition implies that g_t has an MA representation,[3] as follows:

$$g_t = \sum_{s=0}^{\infty} \tilde{\rho}_s \epsilon_{g,t-s} + \mu_g, \quad \sum_{s=0}^{\infty} \tilde{\rho}_s^2 < \infty \ (\tilde{\rho}_0 = 1), \tag{4}$$

where $\epsilon_{g,t-s}$ is a white noise process with bounded second moment. Furthermore, we model the white noise innovation of g_t as (potentially, since parameters will be estimated) partially

[3] The only requirement for this to hold is that g_t be a covariance-stationary time series. Furthermore, μ_g could also vary in time as long as it is a deterministic process.

spanned by the contemporaneous asset return shocks $\tilde{\boldsymbol{v}}_t$, as follows:

$$\epsilon_{g,t-s} = \tilde{\boldsymbol{\eta}}_g^\top \tilde{\boldsymbol{v}}_{t-s} + e_{g,t-s}, \quad \tilde{\boldsymbol{v}}_{t-s} \perp e_{g,t-s}. \tag{5}$$

Taken together, equations (1)–(2) and (4)–(5) imply that the factor g might potentially react over multiple periods, or with delay, to the priced shocks to asset returns.

Plugging equation (5) into (4), we can rewrite the data-generating process of g_t as follows: $g_t = \mu_g + \sum_{s=0}^{\infty} \tilde{\rho}_s \tilde{\boldsymbol{\eta}}_g^\top \tilde{\boldsymbol{v}}_{t-s} + \tilde{w}_{gt}$, where $\tilde{w}_{gt} = \sum_{s=0}^{\infty} \tilde{\rho}_s e_{g,t-s}$. Since estimating MA coefficients with infinite lags is unrealistic, and the vector norms of $\tilde{\boldsymbol{\eta}}$ and $\tilde{\boldsymbol{\rho}}$ are unidentified, we truncate the number of lags at $\bar{S}$ and normalize $\tilde{\boldsymbol{\eta}}_g$ such that

$$g_t = \mu_g + \sum_{s=0}^{\bar{S}} \tilde{\rho}_s \underbrace{\tilde{\boldsymbol{\eta}}_g^\top \tilde{\boldsymbol{v}}_{t-s}}_{f_{t-s}} + w_{gt}, \quad \tilde{\boldsymbol{\eta}}_g^\top \tilde{\boldsymbol{\eta}}_g = 1, \tag{6}$$

so f_t, which we call the spanned component, has a unit variance. Since f_t is a white noise innovation, we can interpret $\{\tilde{\rho}_s\}_{s=0}^{\bar{S}}$ as g_t's impulse responses to the asset returns' shock f_t.[4]

Several features of equation (6) are noteworthy. First, g_t can react to both current and lagged asset return shocks $\tilde{\boldsymbol{v}}_t$. This assumption is motivated by past literature showing that asset returns can predict macro variables (e.g., Liew and Vassalou (2000), Ang et al. (2006), and Bryzgalova et al. (2023b)). Second, when g_t is a white noise process ($\tilde{\rho}_s = 0$ for $s > 0$), the model reduces to the setting studied in Giglio and Xiu (2021).[5] However, many factors, particularly macro variables, strongly deviate from the assumption of white noise processes. Past research often extracts the AR(1) innovations of macro variables and estimates the innovations' risk premia. As we show in the empirical analysis, this standard practice often fails to identify macro factors' risk premia, whereas the MA representation in equation (6) does. Finally, we allow for the measurement error of g_t, denoted by w_{gt} (which is unrelated to $\tilde{\boldsymbol{v}}_t$ and $\boldsymbol{w}_{rt}$), to be potentially autocorrelated.

We next define the risk premium of g_t by extending the study of Giglio and Xiu (2021). In their framework, g_t's risk premium is defined as the negative of the covariance between g_t

[4] We do not interpret f_t as a structural shock, so the impulse responses of g_t to f_t purely quantify the lead-lag correlations rather than the causal relationship between asset returns and g_t.

[5] Since their paper uses original rather than log returns, this statement is precise with the exception of the log-linearization approximation error.

and the SDF, $\lambda_g = -\text{cov}(g_t, m_t)$.[6] When g_t is a traded excess return, the fundamental asset pricing equation, $\mathbb{E}[m_t \cdot g_t] = 0$, directly implies that $\mathbb{E}[g_t] = -\text{cov}(g_t, m_t)$ (since we normalize $\mathbb{E}[m_t] = 1$). For a nontradable factor, one can interpret $-\text{cov}(g_t, m_t)$ as the pseudo expected excess return of g_t as if it were tradable. In other words, $-\text{cov}(g_t, m_t)$ is the risk premium on an asset that delivers a payoff that grows at the rate of g_t. We expand their definition by allowing for an entire term structure of risk premia. Specifically, the (average per-period) risk premium of g from $t-1$ to $t+S$ ($0 \le S \le \bar{S}$) is defined as the multiperiod covariance between the factor and the SDF, divided by the number of holding periods, as follows:

$$\lambda_g^S = -\frac{\text{cov}(m_{t-1\to t+S}, g_{t-1\to t+S})}{1+S} = \frac{\sum_{\tau=0}^{S}\sum_{s=0}^{\tau}\tilde{\rho}_s}{1+S} \cdot \underbrace{\tilde{\boldsymbol{\eta}}_g^\top \boldsymbol{\lambda}_{\tilde{\boldsymbol{v}}}}_{\lambda_f}. \tag{7}$$

There are two ways to interpret the definition in equation (7). First, λ_f is the risk premium of the spanned component ($f_t = \tilde{\boldsymbol{\eta}}_g^\top \tilde{\boldsymbol{v}}_t$) driving both asset returns and g_t, and $\frac{\sum_{\tau=0}^{S}\sum_{s=0}^{\tau}\tilde{\rho}_s}{1+S}$ is the per-period loading of $g_{t-1\to t+S}$ on multiperiod asset return shocks $f_{t-1\to t+S}$. Hence, λ_g^S, the risk premium of g over an investment horizon of $(1+S)$ periods, equals its loadings on f multiplied by f's risk premium.

Second, as established below, we can interpret λ_g^S as the risk premium of $g_{t-1\to t+S}$'s mimicking portfolio, with portfolio weights $\boldsymbol{w}^{MP} = \text{cov}(\boldsymbol{r}_{t-1\to t+S})^{-1}\text{cov}(\boldsymbol{r}_{t-1\to t+S}, g_{t-1\to t+S})$. The risk premium of this portfolio, normalized by the number of holding periods, is

$$\begin{aligned}
\lambda_g^{MP} = \frac{\mathbb{E}[\boldsymbol{r}_{t-1\to t+S}]^\top \boldsymbol{w}^{MP}}{1+S} &= \mathbb{E}[\boldsymbol{r}_{t-1\to t+S}]^\top \text{cov}(\boldsymbol{r}_{t-1\to t+S})^{-1}\frac{\text{cov}(\boldsymbol{r}_{t-1\to t+S}, g_{t-1\to t+S})}{1+S} \\
&= \mathbb{E}[\boldsymbol{r}_t]^\top \text{cov}(\boldsymbol{r}_t)^{-1}\boldsymbol{\beta}_{\tilde{v}}\frac{\text{cov}(\boldsymbol{v}_{t-1\to t+S}, g_{t-1\to t+S})}{1+S},
\end{aligned}$$

where the last equality uses the assumption that $\boldsymbol{v}_t$ are serially uncorrelated and $\boldsymbol{w}_{r,t-1\to t+S}$ are orthogonal to g. We relax the assumption of uncorrelated $\tilde{\boldsymbol{v}}_t$ in Section 2.2.

Using Proposition A.1 in the Appendix, we can simplify the risk premium of g_t's mimicking portfolio and show that, as the number of test assets goes to infinity, $\lambda_g^{MP} \to \frac{\boldsymbol{\lambda}_{\tilde{v}}^\top \text{cov}(\tilde{\boldsymbol{v}}_{t-1\to t+S}, g_{t-1\to t+S})}{1+S} = -\frac{\text{cov}(m_{t-1\to t+S}, g_{t-1\to t+S})}{1+S}$, where $m_{t-1\to t+S} = \sum_{\tau=0}^{S} m_{t+\tau-1,t+\tau} = 1 + S - \boldsymbol{\lambda}_{\tilde{v}}^\top \tilde{\boldsymbol{v}}_{t-1\to t+S}$. Therefore, our definition of g_t's risk premium in equation (7) is asymptotically equivalent to the risk premium of the mimicking portfolio in a large cross-section.

[6] This definition is consistent with Cochrane (2009, Chapter 6).

In the data, asset return factors, $\tilde{\boldsymbol{v}}_t$, are unidentified. That is, one can only estimate a linear rotation of $\tilde{\boldsymbol{v}}_t$, denoted by $\boldsymbol{v}_t = \boldsymbol{H}\tilde{\boldsymbol{v}}_t$, where $\boldsymbol{H}$ is a $K \times K$ nonsingular matrix. Since $\tilde{\boldsymbol{v}}_t \overset{\text{iid}}{\sim} \mathcal{N}(\boldsymbol{0}_K, \boldsymbol{I}_K)$, we have that $\boldsymbol{\Sigma}_v \equiv \text{cov}(\boldsymbol{v}_t) = \boldsymbol{H}\boldsymbol{H}^\top$. Even though $\tilde{\boldsymbol{v}}_t$ cannot be identified, g_t's risk premium is well-defined. In particular, the identification of λ_g^S builds upon the rotation invariance property emphasized in Giglio and Xiu (2021). The rotation invariance can be easily seen by rewriting the model as follows:

$$
\begin{aligned}
&\boldsymbol{r}_t = \boldsymbol{\alpha} + \underbrace{\boldsymbol{\beta}_{\tilde{v}}\boldsymbol{H}^{-1}}_{\boldsymbol{\beta}_v}\underbrace{\boldsymbol{H}\boldsymbol{\lambda}_{\tilde{v}}}_{\boldsymbol{\lambda}_v} + \underbrace{\boldsymbol{\beta}_{\tilde{v}}\boldsymbol{H}^{-1}}_{\boldsymbol{\beta}_v}\underbrace{\boldsymbol{H}\tilde{\boldsymbol{v}}_t}_{\boldsymbol{v}_t} + \boldsymbol{w}_{rt}, \quad g_t = \mu_g + \sum_{s=0}^{\bar{S}} \tilde{\rho}_s \underbrace{\tilde{\boldsymbol{\eta}}_g^\top \boldsymbol{H}^{-1}}_{\boldsymbol{\eta}_g^\top}\underbrace{\boldsymbol{H}\tilde{\boldsymbol{v}}_{t-s}}_{\boldsymbol{v}_{t-s}} + w_{gt}, \text{ and} \\
&m_t = 1 - \boldsymbol{\lambda}_v^\top(\boldsymbol{H}^{-1})^\top\boldsymbol{H}^{-1}\boldsymbol{v}_t = 1 - \boldsymbol{\lambda}_v^\top\boldsymbol{\Sigma}_v^{-1}\boldsymbol{v}_t, \quad \lambda_g^S = \frac{\sum_{\tau=0}^{S}\sum_{s=0}^{\tau}\tilde{\rho}_s}{1+S} \cdot \underbrace{\tilde{\boldsymbol{\eta}}_g^\top\boldsymbol{H}^{-1}}_{\boldsymbol{\eta}_g}\underbrace{\boldsymbol{H}\boldsymbol{\lambda}_{\tilde{v}}}_{\boldsymbol{\lambda}_v}.
\end{aligned}
\tag{8}
$$

Therefore, the most important quantity in our paper, λ_g^S, is well identified.

Estimating the confidence bands – or better, the statistical uncertainty – of λ_g^S is challenging in the frequentist framework. Specifically, λ_g^S is a function of $\boldsymbol{\rho}_g$, $\boldsymbol{\eta}_g$, and $\boldsymbol{\lambda}_v$, where the first two parameters depend on each other. Hence, the frequentist asymptotic covariance matrix of λ_g^S is quite complex despite its closed-form expression outlined above. Consequently, we adopt a Bayesian framework to provide valid inference for all model parameters and present it in the next subsection.

2.1 Bayesian Estimation of Risk Premia

This subsection describes our hierarchical Bayesian framework. We first consider the *time series* dimension, which is needed to estimate the joint posterior distribution of asset returns' latent factors and their loadings, expected asset returns, g_t's loadings on the latent factors, and the precision matrices of error terms. We make the following distributional assumptions:

$$
g_t = \mu_g + \sum_{s=0}^{\bar{S}} \rho_s \boldsymbol{\eta}_g^\top(\boldsymbol{v}_{t-s} - \boldsymbol{\mu}_{\boldsymbol{v}}) + w_{gt}, \quad w_{gt} \overset{\text{iid}}{\sim} \mathcal{N}(0, \sigma_{wg}^2), \quad \boldsymbol{v}_t \overset{\text{iid}}{\sim} \mathcal{N}(\boldsymbol{\mu}_{\boldsymbol{v}}, \boldsymbol{\Sigma}_{\boldsymbol{v}}), \tag{9}
$$

$$
\boldsymbol{r}_t = \boldsymbol{\mu}_r + \boldsymbol{\beta}_v(\boldsymbol{v}_t - \boldsymbol{\mu}_v) + \boldsymbol{w}_{rt}, \quad \boldsymbol{w}_{rt} \overset{\text{iid}}{\sim} \mathcal{N}(\boldsymbol{0}_N, \boldsymbol{\Sigma}_{wr}), \quad \boldsymbol{\Sigma}_{wr} = \text{diag}\{\sigma_{1,wr}^2, \ldots, \sigma_{N,wr}^2\}, \text{ and} \tag{10}
$$

$$
\boldsymbol{v}_t \perp w_{gt} \perp \boldsymbol{w}_{rt}, \text{ and let } \boldsymbol{\rho}_g = (\mu_g, \rho_0, \ldots, \rho_{\bar{S}})^\top, \tag{11}
$$

where $\boldsymbol{v}_t$ are linear and nonsingular rotations of the true K latent factors $\tilde{\boldsymbol{v}}_t$. Since these rotations are arbitrary, we need to estimate their unconditional means ($\boldsymbol{\mu}_{\boldsymbol{v}}$) and covariance matrix ($\boldsymbol{\Sigma}_{\boldsymbol{v}}$). Direct modeling of $\boldsymbol{\mu}_v$ is critical for obtaining a proper posterior distribution of expected excess returns $\boldsymbol{\mu}_r$.[7] According to equation (11), the error terms, w_{gt} and $\boldsymbol{w}_{rt}$, are orthogonal, which implies that we can estimate the model parameters in g_t and $\boldsymbol{r}_t$ separately.

The systems in (9) and (10) introduces a potential degree of misspecification relative to the true data-generating processes described in equations (2) and (6). First, the error w_{gt} could be serially correlated. As Müller (2013) shows, posteriors are still asymptotically normal and centered at the maximum likelihood estimate under this assumption, although the canonical posterior covariance matrix of the model parameters is incorrect and should be replaced with a sandwich covariance matrix. Therefore, we incorporate this correction within our method.

Second, $\boldsymbol{\Sigma}_{\boldsymbol{wr}}$ is assumed to be diagonal. Our posterior characterization below does not require this assumption, and indeed, we impose it only to avoid numerical problems when considering very large cross-sectional dimensions (i.e., when the number of assets approaches or exceeds the time series dimension) and relax it in all other instances. However, as we will show through simulations, the diagonal assumption does not have material effects on the posterior distributions. Hence, this assumption is harmless. This robustness result is not surprising since, in a frequentist setting, this type of misspecification would affect only efficiency but not consistency.

We assign the standard uninformative prior distributions to the time series parameters

$$\begin{aligned}&\pi(\boldsymbol{\rho}_g, \boldsymbol{\eta}_g, \sigma^2_{wg}) \propto (\sigma^2_{wg})^{-1}, \quad \pi(\boldsymbol{v}) \propto 1, \quad \pi(\boldsymbol{\mu}_v, \boldsymbol{\Sigma}_v) \propto |\boldsymbol{\Sigma}_v|^{-\frac{K+1}{2}}, \text{ and}\\ &\pi(\boldsymbol{\beta}_v) \propto 1, \quad \pi(\boldsymbol{\mu}_r, \boldsymbol{\Sigma}_{wr}) \propto |\boldsymbol{\Sigma}_{wr}|^{-\frac{N+1}{2}}.\end{aligned} \tag{12}$$

In the *cross-sectional* dimension, conditional on the recovered sources of risk v_t in the time series dimension, the SDF and its risk prices, λ_v, can then be recovered using the Bayesian-SDF estimator (B-SDF) in Definition 1 of Bryzgalova et al. (2023a). That is, conditional on the recovered v_t being the sources of risk driving the cross-section, we have the SDF

$$m_t = 1 - \boldsymbol{\lambda}_v^\top \boldsymbol{\Sigma}_v^{-1} \boldsymbol{v}_t \;\Rightarrow\; \boldsymbol{\mu}_r = \boldsymbol{\beta}_v \boldsymbol{\lambda}_v. \tag{13}$$

[7] The sample average of $\boldsymbol{r}_t$ is $\boldsymbol{\mu}_r + \boldsymbol{\beta}_v \frac{1}{T}\sum_{t=1}^T(\boldsymbol{v}_t - \boldsymbol{\mu}_v) + \frac{1}{T}\sum_{t=1}^T \boldsymbol{w}_{rt}$. If we always demean the latent factors to have zero sample averages, the first source of uncertainty about $\boldsymbol{\mu}_r$, originated from $\frac{1}{T}\sum_{t=1}^T(\boldsymbol{v}_t - \boldsymbol{\mu}_v)$, will disappear. Consequently, the credible intervals for $\boldsymbol{\mu}_r$ will be too tight if we do not directly model $\boldsymbol{\mu}_v$.

Recall that we nevertheless allow for pricing errors as outlined in (2). With extensive simulation studies, we show in Section 3 that this approach delivers valid posterior distributions.

Within the frequentist paradigm, constructing proper inference for the system in equations (9)–(13) is, if not infeasible, at least a daunting task. As we are about to show in Proposition 1 below, this is both simple and transparent within the Bayesian paradigm.

There are two reasons for this. First, a joint distribution, say $p(x, y)$, can be traced by generating a Markov chain that sequentially samples from $p(x|y)$ and $p(y|x)$ – the so-called Gibbs sampling.

Second, the hierarchical structure of the time series and cross-sectional layers of the estimation problem yields well-defined and well-understood conditional posterior distributions. Specifically, if $\boldsymbol{v}_t$ were known (i.e., conditioning on it), equation (9) would simply be an ordinary linear regression problem with well-known properties: in a Bayesian setting, under diffuse and/or conjugate priors, a normal-inverse-gamma posterior distribution (i.e., the analogous of the t-distribution that would arise for frequentist inference in this case).

Similarly, if $\boldsymbol{v}_t$ were known, equation (10) would simply be a canonical multivariate linear regression, thereby yielding (under diffuse and/or conjugate priors) a well-known posterior distribution: a normal-inverse-Whishart (the Bayesian analogous to the frequentist multivariate t-distribution result).

Furthermore, conditional on knowing both the parameters in equation (10) and the data, the distribution of the latent factors $\boldsymbol{v}_t$ can be obtained by inverting its relationship with asset returns. Finally, conditional on the parameters and latent factors in the time series layer, the distribution of the risk prices, $\boldsymbol{\lambda}_v$, simply follows from Definition 2 of Bryzgalova et al. (2023a). Note that this layer is fundamental since it de facto selects which of (and how) the latent drivers $\boldsymbol{v}_t$ are actually sources of priced risk – the crucial stage for measuring the risk premia associated with g_t.

We formalize this hierarchal characterization of the posterior in the proposition below and derive it in Internet Appendix IA.1.1.

Proposition 1 (Gibbs sampler of the baseline model)**.** *Under the assumptions in equations (9)–(13), the posterior distribution of the model parameters can be sampled from the following conditional distributions:*

(1) Conditional on the data, $\{g_t\}_{t=1+\bar{S}}^T$, and latent factors, $\{\boldsymbol{v}_t\}_{t=1}^T$, the parameters of the g_t process (σ_{wg}^2, $\boldsymbol{\rho}_g$, and $\boldsymbol{\eta}_g$) follow the normal-inverse-gamma distribution in equations (IA.1)–(IA.3) of Internet Appendix IA.1.1. For point identification purposes, draws of $\boldsymbol{\rho}_g$ and $\boldsymbol{\eta}_g$ are normalized such that $\boldsymbol{\eta}_g^\top \boldsymbol{\eta}_g = 1$.

(2) Conditional on asset returns, $\{\boldsymbol{r}_t\}_{t=1}^T$, and latent factors, the parameters of the $\boldsymbol{r}_t$ process ($\boldsymbol{\Sigma}_{wr}$ and $\boldsymbol{B}_r^\top = (\boldsymbol{\mu}_r, \boldsymbol{\beta}_v)$) follow the normal-inverse-Wishart distribution in equations (IA.4)–(IA.5) of Internet Appendix IA.1.1.

(3) Conditional on asset returns and $(\boldsymbol{\mu}_r, \boldsymbol{\beta}_v, \boldsymbol{\Sigma}_{wr})$, the latent factors, $\boldsymbol{v}_t$, their mean, and covariance matrix can be sampled from

$$\begin{aligned} &\boldsymbol{v}_t \mid \boldsymbol{r}_t, \boldsymbol{\mu}_r, \boldsymbol{\beta}_v, \boldsymbol{\Sigma}_{wr}, \boldsymbol{\mu}_v, \boldsymbol{\Sigma}_v \sim \\ &\mathcal{N}\left(\left(\boldsymbol{\beta}_{\boldsymbol{v}}^\top \boldsymbol{\Sigma}_{\boldsymbol{wr}}^{-\mathbf{1}} \boldsymbol{\beta}_{\boldsymbol{v}}\right)^{-1} \left[\boldsymbol{\beta}_{\boldsymbol{v}}^\top \boldsymbol{\Sigma}_{\boldsymbol{wr}}^{-\mathbf{1}} (\boldsymbol{r}_t - \boldsymbol{\mu}_r + \boldsymbol{\beta}_{\boldsymbol{v}} \boldsymbol{\mu}_v)\right],\ \left(\boldsymbol{\beta}_{\boldsymbol{v}}^\top \boldsymbol{\Sigma}_{\boldsymbol{wr}}^{-\mathbf{1}} \boldsymbol{\beta}_{\boldsymbol{v}}\right)^{-1} \right), \end{aligned} \tag{14}$$

$$\boldsymbol{\Sigma}_v \mid \{\boldsymbol{v}_t\}_{t=1}^T \sim \mathcal{W}^{-1}\left(T-1,\ \sum_{t=1}^T (\boldsymbol{v}_t - \bar{\boldsymbol{v}})(\boldsymbol{v}_t - \bar{\boldsymbol{v}})^\top\right),\ \textit{and} \tag{15}$$

$$\boldsymbol{\mu}_v \mid \boldsymbol{\Sigma}_v, \{\boldsymbol{v}_t\}_{t=1}^T \sim \mathcal{N}\left(\bar{\boldsymbol{v}},\ \boldsymbol{\Sigma}_v / T\right), \tag{16}$$

where $\mathcal{N}(\cdot)$ and $\mathcal{W}^{-1}(\cdot)$ denote, respectively, the normal and inverse-Wishart distributions.

(4) Conditional on the posterior draws from the time series steps (1)–(3), the posterior distribution of $\boldsymbol{\lambda}_v$ is a Dirac distribution at $(\boldsymbol{\beta}_v^\top \boldsymbol{\beta}_v)^{-1} \boldsymbol{\beta}_v^\top \boldsymbol{\mu}_r$, yielding a Dirac conditional posterior for the term structure of g_t's risk premia at $\lambda_g^S = \frac{\sum_{\tau=0}^S \sum_{s=0}^\tau \rho_s}{1+S} \cdot \boldsymbol{\eta}_g^\top \boldsymbol{\lambda}_{\boldsymbol{v}}$, where $0 \le S \le \bar{S}$.

Several features of our Bayesian Gibbs sampler are noteworthy. First, although we do not know in closed-form the joint distribution of all parameters, all conditional distributions, such as inverse-gamma, multivariate normal, and inverse-Wishart distributions, are well-defined and standard.

Second, we follow Müller (2013) and adjust the posterior covariance matrix of $\boldsymbol{\rho}_g$ and $\boldsymbol{\eta}_g$ for the autocorrelation in the residuals, w_{gt} and $\boldsymbol{w}_{rt}$, using the Newey and West (1987) type of sandwich estimator.[8]

[8] The number of lags is set to be $\bar{S}$ since $w_{gt}\boldsymbol{x}_t$ and $w_{g,t-l}\boldsymbol{x}_{t-l}$ become serially uncorrelated for $l > \bar{S}$, where

Third, the posterior distribution of $\boldsymbol{v}_t$ in Step 3 of Proposition 1 ignores the information embedded in g_t, balancing the trade-off between model simplicity and estimation efficiency. Since g_t depends on many lags of the latent factors, incorporating its information in estimating $\boldsymbol{v}_t$ is feasible but requires a more computationally demanding approach, such as the Kalman filter. More importantly, we consider large cross-sections of test assets; hence, the discarded information is negligible as $N \to \infty$. Finally, in empirical applications, not conditioning on g_t in the extraction of $\boldsymbol{v}_t$ provides a level playing field when comparing the estimated risk premia of different variables g_t.

Fourth, Proposition 1 does *not* require a diagonal $\boldsymbol{\Sigma}_{wr}$. Nevertheless, for empirical applications where N is close to the time series sample size, we impose diagonality to avoid numerical difficulties. Our simulation studies confirm that the assumption of a diagonal $\boldsymbol{\Sigma}_{wr}$ does not result in invalid confidence intervals, even though $\boldsymbol{w}_{rt}$ is cross-sectionally correlated in the hypothetical true data-generating process. In contrast, in empirical applications where the number of test assets is relatively small (i.e., $N \leq 50$, such as in the cross-section of corporate bonds), we use a nondiagonal $\boldsymbol{\Sigma}_{wr}$ in estimation.

Fifth, the cross-sectional dimension (Step 4 in Proposition 1) defines latent factors' risk premia as $(\boldsymbol{\beta}_v^\top \boldsymbol{\beta}_v)^{-1}\boldsymbol{\beta}_v^\top \boldsymbol{\mu}_r$ and, via the sequential resampling, accounts for the uncertainty about the expected returns, the factor loadings, and the latent factors' means $\boldsymbol{\mu}_v$.

In addition to risk premia estimates, our Bayesian framework can produce valid posterior distributions for other economic quantities of interest, including, but not limited to, the time series fit in g_t's equation (R_g^2), cumulative impulse responses of g_t to the asset return shocks ($\{\tilde{\rho}_s\}_{s=0}^{\bar{S}}$), and the cross-sectional fit in explaining average returns.

Past literature often adopts the Fama-MacBeth regression to estimate factors' risk premia. In Proposition 1, steps 2–4 echo the time series and cross-sectional steps of the Bayesian Fama-MacBeth in Bryzgalova et al. (2023a) for principal components of asset returns. Step 1 is the additional step that models the joint dynamics of asset returns and g_t. As Giglio and Xiu (2021) argue, estimating factors' risk premia using principal components of asset returns can avoid the omitted variable bias and attenuation bias from measurement errors.

Finally, the traditional Fama-MacBeth regression suffers from weak identification (see, e.g.,

x_t denote the regressors in g_t's equation and is the linear transformation of latent factors $\{v_{t-s}\}_{s=0}^{\bar{S}}$.

Kan and Zhang (1999a,b)), particularly for macro factors. One contribution of our paper is to use the factors' cumulative loadings on asset returns, proxied by $\{\tilde{\rho}_s\}_{s=0}^{\bar{S}}$, to identify their risk premia. In short, we will show in both simulation studies and real-world data that our Bayesian estimates are not only robust to the weak identification but, more importantly, help recover the risk premia of persistent macro factors.

2.2 Time-Varying Risk Premia and Their Term Structures

From an economic standpoint, a salient feature of macro-finance equilibrium models is the time variation in risk premia. In this section, we extend our Bayesian framework for estimating time-varying term structures and show that our unconditional formulation in the previous sections is consistent even in the presence of time-varying risk premia.

We now require the SDF to price assets *conditionally*; that is,

$$\mathbb{E}_t[m_{t+1} \cdot r_{i,t+1}] = 0, \quad i = 1 \dots N, \tag{17}$$

where $\mathbb{E}_t$ denotes the conditional expectation at time t. Leveraging Hansen and Jagannathan (1991), we focus on the conditional SDF projections on the space of returns as follows:

$$m_{t+1} = 1 - \boldsymbol{b}_t^\top \big(\boldsymbol{r}_{t+1} - \mathbb{E}_t[\boldsymbol{r}_{t+1}]\big), \quad \text{where} \;\; \boldsymbol{b}_t = \text{cov}_t(\boldsymbol{r}_{t+1})^{-1}\mathbb{E}_t[\boldsymbol{r}_{t+1}]. \tag{18}$$

The return process, as before, follows an approximate factor structure,

$$\boldsymbol{r}_t = \boldsymbol{\mu}_r + \boldsymbol{\beta}_{\tilde{v}}\tilde{\boldsymbol{v}}_t + \boldsymbol{w}_{rt}, \quad \tilde{\boldsymbol{v}}_t \perp \boldsymbol{w}_{rt}, \quad \mathbb{E}_{t-1}[\boldsymbol{w}_{rt}] = \boldsymbol{0}_N, \quad \mathbb{E}[\tilde{\boldsymbol{v}}_t] = \boldsymbol{0}_K, \tag{19}$$

where, importantly, the priced systematic factors $\tilde{\boldsymbol{v}}_t$ are potentially predictable. That is, $\tilde{\boldsymbol{v}}_t = \boldsymbol{\mu}_{\tilde{v},t-1} + \boldsymbol{\epsilon}_{\tilde{v}t}$, where $\boldsymbol{\mu}_{\tilde{v},t-1} \equiv \mathbb{E}_{t-1}[\tilde{\boldsymbol{v}}_t]$; hence $\boldsymbol{\mu}_{\tilde{v},t-1} \perp \boldsymbol{\epsilon}_{\tilde{v}t}$. We normalize the innovations to the latent factors such that $\text{cov}(\boldsymbol{\epsilon}_{\tilde{v}t}) = \boldsymbol{I}_K$.

As previously, unconditional mean returns are partially explained by $\boldsymbol{\beta}_{\tilde{v}}$ in equation (2). The only additional assumption that we require is that the eigenvalues of $\text{cov}(\boldsymbol{\mu}_{\tilde{v},t-1})$ are bounded. This formulation yields the SDF[9]

$$m_{t+1} = 1 - \boldsymbol{\lambda}_{\tilde{v}}^\top \boldsymbol{\epsilon}_{\tilde{v},t+1} - \boldsymbol{\mu}_{\tilde{v}t}^\top \boldsymbol{\epsilon}_{\tilde{v},t+1}, \tag{20}$$

[9] Using equations (18) and (19), we can show that $\boldsymbol{b}_t = \big(\boldsymbol{\beta}_{\tilde{v}}\boldsymbol{\beta}_{\tilde{v}}^\top + \boldsymbol{\Sigma}_{wr}\big)^{-1}\big(\boldsymbol{\alpha} + \boldsymbol{\beta}_{\tilde{v}}\boldsymbol{\lambda}_{\tilde{v}} + \boldsymbol{\beta}_{\tilde{v}}\boldsymbol{\mu}_{\tilde{v},t}\big)$ and $\boldsymbol{r}_{t+1} - \mathbb{E}_t[\boldsymbol{r}_{t+1}] = \boldsymbol{\beta}_{\tilde{v}}\boldsymbol{\epsilon}_{\tilde{v},t+1} + \boldsymbol{w}_{rt}$. Ignoring the unpriced idiosyncratic shocks $\boldsymbol{w}_{rt}$, we can represent the linear SDF as $m_{t+1} = 1 - \boldsymbol{\alpha}^\top\big(\boldsymbol{\beta}_{\tilde{v}}\boldsymbol{\beta}_{\tilde{v}}^\top + \boldsymbol{\Sigma}_{wr}\big)^{-1}\boldsymbol{\beta}_{\tilde{v}}\boldsymbol{\epsilon}_{\tilde{v},t+1} - (\boldsymbol{\lambda}_{\tilde{v}} + \boldsymbol{\mu}_{\tilde{v},t})^\top\boldsymbol{\beta}_{\tilde{v}}^\top\big(\boldsymbol{\beta}_{\tilde{v}}\boldsymbol{\beta}_{\tilde{v}}^\top + \boldsymbol{\Sigma}_{wr}\big)^{-1}\boldsymbol{\beta}_{\tilde{v}}\boldsymbol{\epsilon}_{\tilde{v},t+1}$. Following similar derivations as in Appendix A.1, we can derive that $m_{t+1} \to 1 - (\boldsymbol{\lambda}_{\tilde{v}} + \boldsymbol{\mu}_{\tilde{v},t})^\top\boldsymbol{\epsilon}_{\tilde{v},t+1}$ as $N \to \infty$.

where $\boldsymbol{\mu}_{\tilde{v}t}^{\top}\boldsymbol{\epsilon}_{\tilde{v},t+1}$ captures time-varying risk premia of asset return shocks.

Since the Wold representation requires the MA formulation to depend only on innovations, the process for g is modified as follows:

$$g_t = \mu_g + \sum_{s=0}^{\bar{S}} \tilde{\rho}_s \underbrace{\tilde{\boldsymbol{\eta}}_g^{\top}\boldsymbol{\epsilon}_{\tilde{v},t-s}}_{f_{t-s}} + w_{gt}, \quad \tilde{\boldsymbol{\eta}}_g^{\top}\tilde{\boldsymbol{\eta}}_g = 1. \tag{21}$$

That is, g is potentially driven by the innovations of the priced systematic factors $\tilde{\boldsymbol{v}}_t$. Hence, defining the *conditional* risk premia analogously as the unconditional ones, we have that the time-varying term structure of risk premia is given by

$$\lambda_{g,t-1}^{S} = -\frac{\mathrm{cov}_{t-1}(m_{t-1\to t+S}, g_{t-1\to t+S})}{1+S} = \sum_{\tau=0}^{S}\sum_{s=0}^{\tau} \frac{\tilde{\rho}_s \boldsymbol{\eta}_g^{\top}(\boldsymbol{\lambda}_{\tilde{v}} + \boldsymbol{\mu}_{\tilde{v},t-1})}{1+S}. \tag{22}$$

Four important observations are in order. First, the dynamics of the conditional mean of the priced systematic risks, $\boldsymbol{\mu}_{\tilde{v},t-1}$, drive the time variation of the term structure of risk premia. Second, since by construction $\mathbb{E}[\boldsymbol{\mu}_{\tilde{v},t-1}] = 0$, the implied unconditional term structure is the same as that of equation (7), which was obtained with uncorrelated sources of systematic risk. That is, the estimator derived in Section 2.1 is consistent even in the presence of time-varying risk premia. Third, despite the added generality, the risk premia of g remain point-identified due to the rotation invariance property of our setting.[10] Fourth, to elicit the time-variation of the term structure, we need to explicitly model the conditional mean process, that is, the dynamics of $\tilde{\boldsymbol{v}}$.

We assume that $\tilde{\boldsymbol{v}}_t$ are driven by some predictors, such as $\tilde{\boldsymbol{v}}_t$'s lags and p external variables $\boldsymbol{z}_t$. Let $\boldsymbol{x}_t = (\tilde{\boldsymbol{v}}_t^{\top}, \boldsymbol{z}_t^{\top})^{\top}$, which follows a vector autoregressive (VAR) model of order q:[11]

$$\boldsymbol{x}_t = \boldsymbol{\phi}_0 + \boldsymbol{\phi}_1\boldsymbol{x}_{t-1} + \cdots + \boldsymbol{\phi}_q\boldsymbol{x}_{t-q} + \boldsymbol{\epsilon}_{xt}, \quad \boldsymbol{\epsilon}_{xt} \overset{\text{iid}}{\sim} \mathcal{N}(\mathbf{0}_{K+p}, \boldsymbol{\Sigma}_{\epsilon x}). \tag{23}$$

The additional layer in equation (23) requires a minimal change to our Gibbs sampler to characterize the posterior distribution. The only deviation from Section 2.1 is that $\boldsymbol{v}_t$ follows a VAR process rather than an IID normal distribution. In particular, using the canonical diffuse prior $\pi(\boldsymbol{\phi}_0, \dots, \boldsymbol{\phi}_q, \boldsymbol{\Sigma}_{\epsilon x}) \propto |\boldsymbol{\Sigma}_{\epsilon x}|^{-\frac{K+p+1}{2}}$, the conditional posterior of the parameters in

[10] See Appendix A.2 for details.

[11] The VAR assumption is often adopted in past literature studying return predictability (e.g., Campbell and Shiller (1988), Campbell and Vuolteenaho (2004), and Campbell et al. (2013)).

this additional layer follows the usual normal-inverse-Wishart distribution and can be sampled accordingly. We summarize the Gibbs sampler in Proposition A2 of Appendix A.2 and derive it in Internet Appendix IA.1.2.

2.3 Risk Premia in Potentially Segmented Markets

The econometric framework in the previous subsection imposes market integration, which implies that risk premia estimates are homogeneous across different asset markets. However, market integration has been challenged by both theoretical and empirical work. Theoretically, arbitrageurs are constrained by financial frictions (see, e.g., Shleifer and Vishny (1997), Garleanu and Pedersen (2011), Gromb and Vayanos (2002, 2018), Duffie (2010), and Greenwood et al. (2018)), so the law of one price can be violated across market segments. Empirically, we have seen evidence suggesting market segmentation in (i) international stock and currency markets and (ii) US equity and corporate bond markets.[12]

Motivated by past research, we extend the framework in Section 2.1 by allowing segmented risk premia in unconditional models. We aim to test whether the factor g_t carries the same risk premium in two potentially segmented markets. There are two asset markets, such as equity and corporate bond markets, with cross-sections of excess returns $\boldsymbol{r}_t^1$ and $\boldsymbol{r}_t^2$. We assume the approximate factor structure for two asset markets, so $\boldsymbol{r}_t^1$ and $\boldsymbol{r}_t^2$ are driven by possibly different sets of latent factors $\tilde{\boldsymbol{v}}_t^1$ and $\tilde{\boldsymbol{v}}_t^2$,

$$\boldsymbol{r}_{\boldsymbol{t}}^j = \boldsymbol{\mu}_{\boldsymbol{r}}^j + \boldsymbol{\beta}_{\tilde{\boldsymbol{v}}}^j \tilde{\boldsymbol{v}}_t^j + \boldsymbol{w}_{\boldsymbol{rt}}^j, \quad \tilde{\boldsymbol{v}}_t^j \overset{\text{iid}}{\sim} \mathcal{N}(\boldsymbol{0}_{K_j}, \boldsymbol{I}_{K_j}), \quad \boldsymbol{w}_{\boldsymbol{rt}}^j \overset{\text{iid}}{\sim} \mathcal{N}(\boldsymbol{0}_{N_j}, \boldsymbol{\Sigma}_{\boldsymbol{wr}}^j), \text{ and} \tag{24}$$

$$\tilde{\boldsymbol{v}}_t^1 \perp \boldsymbol{w}_{\boldsymbol{t}}^1 \perp \boldsymbol{w}_{\boldsymbol{t}}^2, \quad \tilde{\boldsymbol{v}}_t^2 \perp \boldsymbol{w}_{\boldsymbol{t}}^1 \perp \boldsymbol{w}_{\boldsymbol{t}}^2, \quad \text{cov}\big(\tilde{\boldsymbol{v}}_t^1, \tilde{\boldsymbol{v}}_t^2\big) = \boldsymbol{\Sigma}_{\tilde{\boldsymbol{v}}}^{1,2}. \tag{25}$$

The conditions in equation (25) presume that idiosyncratic risks, $\boldsymbol{w}_{\boldsymbol{t}}^1$ and $\boldsymbol{w}_{\boldsymbol{t}}^2$, are uncorrelated and orthogonal to all systematic shocks. However, we allow the systematic components to correlate, so the correlation between $\tilde{\boldsymbol{v}}_t^1$ and $\tilde{\boldsymbol{v}}_t^2$ captures the comovement among asset returns in the two markets.

[12] For the international stock and currency markets, see Brandt et al. (2006), Bakshi et al. (2018), and Sandulescu et al. (2021). For corporate bond markets, see Schaefer and Strebulaev (2008), Kapadia and Pu (2012), Choi and Kim (2018), and Kelly et al. (2023).

Similar to equation (6), g_t is driven by asset pricing shocks

$$g_t = \sum_{s=0}^{\bar{S}} \tilde{\rho}_s \tilde{\boldsymbol{\eta}}_g^\top \tilde{\boldsymbol{v}}_{t-s} + w_{gt}, \quad \tilde{\boldsymbol{v}}_t = (\tilde{\boldsymbol{v}}_t^{1,\top}, \tilde{\boldsymbol{v}}_t^{2,\top})^\top; \tag{26}$$

therefore, the white noise innovation of g_t, $\epsilon_{g,t}$, is partially spanned by the contemporaneous $\tilde{\boldsymbol{v}}_t^1$ and $\tilde{\boldsymbol{v}}_t^2$.

We further extend asset pricing models in equations (2)–(3) to the two-market scenario,

$$\boldsymbol{\mu}_{\boldsymbol{r}}^j = \boldsymbol{\alpha}^j + \boldsymbol{\beta}_{\tilde{\boldsymbol{v}}}^j \boldsymbol{\lambda}_{\tilde{\boldsymbol{v}}}^j, \quad m_t^j = 1 - (\tilde{\boldsymbol{v}}_t^j)^\top \boldsymbol{\lambda}_{\tilde{\boldsymbol{v}}}^j, \tag{27}$$

where we assume that each asset's pricing error, α_n^j, is cross-sectionally IID and independent of factor loadings $\boldsymbol{\beta}_{\tilde{\boldsymbol{v}},\boldsymbol{n}}^j$. Ex ante, we are agnostic as to whether the markets are integrated or equivalently, whether the two SDFs m_t^1 and m_t^2 are perfectly correlated. Instead, we let the data speak on the degree of segmentation.

The existing literature on market segmentation can be summarized into two strands. The first one studies the correlation between the two SDFs and whether the SDF in one market can explain asset returns in another market. (See, e.g., Chen and Knez (1995), Brandt et al. (2006), Kelly et al. (2023), Sandulescu et al. (2021), and Sandulescu (2022)). Another strand of the literature tests instead whether systematic factors carry the same risk premia across markets (e.g., Choi and Kim (2018), Chaieb et al. (2021), and Patton and Weller (2022)).

Similar to the second strand of literature, we test whether g_t's risk premium, λ_g^S, is homogeneous in two asset markets. The risk premium of g_t in market j is defined as

$$\lambda_g^{S,j} = -\frac{\text{cov}(m_{t-1\to t+S}^j, g_{t-1\to t+S})}{1+S} = \frac{\sum_{\tau=0}^{S} \sum_{s=0}^{\tau} \tilde{\rho}_s}{1+S} \cdot \tilde{\boldsymbol{\eta}}_g^\top \text{cov}(\tilde{\boldsymbol{v}}_t, \tilde{\boldsymbol{v}}_t^j) \boldsymbol{\lambda}_{\tilde{v}}^j. \tag{28}$$

Formally, we aim to test whether the term structures of risk premia, $\lambda_g^{S,1} - \lambda_g^{S,2}$ $(0 \le S \le \bar{S})$, are statistically different in two asset markets.

We again rely on the rotation invariance property with slightly different formulations to identify risk premia. Let $\boldsymbol{v}_t^j = \boldsymbol{H}_j \tilde{\boldsymbol{v}}_t^j$, where $j \in \{1, 2\}$ and $\boldsymbol{H}_j$ is a $K_j \times K_j$ nonsingular matrix. The covariance matrix of $\boldsymbol{v}_t^j$ is $\boldsymbol{\Sigma}_v^j = \boldsymbol{H}_j \boldsymbol{H}_j^\top$. While the rotation invariant representations of $\boldsymbol{r}_t^j$, g_t, and m_t^j are similar to those in equation (8), we need to modify the representation of risk premia as follows:

$$\lambda_g^{S,j} = \frac{\sum_{\tau=0}^{S}\sum_{s=0}^{\tau}\tilde{\rho}_s}{1+S} \cdot \underbrace{\tilde{\boldsymbol{\eta}}_g^\top \boldsymbol{H}^{-1}}_{\boldsymbol{\eta}_g^\top}\underbrace{\mathrm{cov}(\boldsymbol{H}\tilde{\boldsymbol{v}}_t, \boldsymbol{H}_j\tilde{\boldsymbol{v}}_t^j)}_{\mathrm{cov}(\boldsymbol{v}_t,\boldsymbol{v}_t^j)}\underbrace{(\boldsymbol{H}_j^\top)^{-1}\boldsymbol{H}_j^{-1}}_{(\boldsymbol{\Sigma}_v^j)^{-1}}\underbrace{\boldsymbol{H}_j\boldsymbol{\lambda}_{\tilde{v}}^j}_{\boldsymbol{\lambda}_v^j}, \tag{29}$$

where $\boldsymbol{H}$ is a block diagonal matrix that contains $\boldsymbol{H}_1$ and $\boldsymbol{H}_2$ in the main-diagonal blocks.

One may notice that the formula of $\lambda_g^{S,j}$ in equation (28) differs from that of equation (7), in which we use only the latent factors in asset market j in the estimation. Proposition 2 shows that, at the population level, estimating $\lambda_g^{S,j}$ separately for each market j still leads to identical risk premia estimates as in equation (29).

Proposition 2. *Suppose that assumptions described in equations (24)–(28) hold, and we estimate g_t's risk premium separately in each market j using the following model:*

$$g_t = \sum_{s=0}^{\bar{S}} \rho_s (\boldsymbol{\eta}_g^j)^\top \boldsymbol{v}_{t-s}^j + w_{gt}^j, \quad \boldsymbol{v}_t^j \perp w_{gt}^j, \quad j \in \{1,2\}.$$

Then $\boldsymbol{\eta}_g^j = (\boldsymbol{\Sigma}_v^j)^{-1}\mathrm{cov}(\boldsymbol{v}_t^j, \boldsymbol{v}_t)\boldsymbol{\eta}_g$ and

$$\lambda_g^{S,j} = \frac{\sum_{\tau=0}^{S}\sum_{s=0}^{\tau}\rho_s}{1+S} \cdot (\boldsymbol{\eta}_g^j)^\top \boldsymbol{\lambda}_v^j = \frac{\sum_{\tau=0}^{S}\sum_{s=0}^{\tau}\rho_s}{1+S} \cdot \boldsymbol{\eta}_g^\top \mathrm{cov}(\boldsymbol{v}_t, \boldsymbol{v}_t^j)(\boldsymbol{\Sigma}_v^j)^{-1}\boldsymbol{\lambda}_v^j.$$

Internet Appendix IA.1.3 provides a simple proof. One may be tempted to estimate the posterior distribution of $\lambda_g^{S,j}$ separately for each market j, following the algorithm in Proposition 1. This would lead to unbiased risk premia estimates of $\lambda_g^{S,j}$. However, to provide proper credible intervals for $\lambda_g^{S,1} - \lambda_g^{S,2}$, we need to *jointly* model $\mathrm{cov}(g_t, \boldsymbol{v}_t^1)$ and $\mathrm{cov}(g_t, \boldsymbol{v}_t^2)$, captured by $\boldsymbol{\eta}_g^\top \boldsymbol{\Sigma}_v$. This is due to the fact that, if systematic factors are correlated, posterior draws of $\mathrm{cov}(g_t, \boldsymbol{v}_t^1)$ and $\mathrm{cov}(g_t, \boldsymbol{v}_t^2)$ (and, hence, $\lambda_g^{S,1}$ and $\lambda_g^{S,2}$) will not be independent. We revise the Gibbs sampler in Proposition 1 and present the new estimator in Proposition 3.

Proposition 3 (Gibbs sampler of the two-market model)**.** *Under the assumptions described in equations (24)–(28), the posterior distribution of model parameters is given by the following conditional distributions:*

(1) Conditional on the data $\{g_t\}_{t=1+\bar{S}}^T$ and latent factors $\{\boldsymbol{v}_t\}_{t=1}^T$, we sample $(\sigma_{wg}^2, \boldsymbol{\rho}_g, \boldsymbol{\eta}_g)$ in g_t's equation using equations (IA.1)–(IA.3). To identify $\boldsymbol{\rho}_g$ and $\boldsymbol{\eta}_g$, we normalize $\boldsymbol{\eta}_g$ after each posterior draw such that $\boldsymbol{\eta}_g^\top \boldsymbol{\eta}_g = 1$.

(2) In each asset market j, conditional on asset returns and latent factors, we update $(\boldsymbol{\Sigma}_{wr}^j, \boldsymbol{\mu}_r^j, \boldsymbol{\beta}_v^j)$ in $\boldsymbol{r}_t^j$'s equation using equations (IA.4)–(IA.5).

(3) In each asset market j, conditional on asset returns $\boldsymbol{r}_t^j$ and $(\boldsymbol{\mu}_r^j, \boldsymbol{\beta}_v^j, \boldsymbol{\Sigma}_{wr}^j)$, we first update the latent factors $\boldsymbol{v}_t^j$ using equation (14). Let $\boldsymbol{v}_t = (\boldsymbol{v}_t^{1,\top}, \boldsymbol{v}_t^{2,\top})^\top$. We next use equations (15)–(16) to update $(\boldsymbol{\mu}_v, \boldsymbol{\Sigma}_v)$.

(4) Based on the posterior draws from the time series steps (1)–(3), the posterior distribution of $\boldsymbol{\lambda}_v^j$ is a Dirac distribution at $(\boldsymbol{\beta}_v^{j,\top}\boldsymbol{\beta}_v^j)^{-1}\boldsymbol{\beta}_v^{j,\top}\boldsymbol{\mu}_r^j$ in each asset market j. In addition, the posterior distribution of the term structure of g_t's risk premia in market j is also a Dirac distribution at $\lambda_g^{S,j} = \frac{\sum_{\tau=0}^{S}\sum_{s=0}^{\tau}\rho_s}{1+S} \cdot \boldsymbol{\eta}_g^\top \mathrm{cov}(\boldsymbol{v}_t, \boldsymbol{v}_t^j)(\boldsymbol{\Sigma}_v^j)^{-1}\boldsymbol{\lambda}_v^j$.

Using the Gibbs sampler in Proposition 3, we can obtain the joint posterior distribution of $\lambda_g^{S,1} - \lambda_g^{S,2}$. In addition, we can provide the posterior distributions of other economic quantities that have been studied in the market segmentation literature, such as the correlation between m_t^1 and m_t^2 and the generalized correlation between $\boldsymbol{v}_t^1$ and $\boldsymbol{v}_t^2$.

Note that even if we do not reject the null hypothesis of $\lambda_g^{S,1} = \lambda_g^{S,2}$, it does not necessarily imply that two asset markets are integrated. A counterexample is that g_t, m_t^1, and m_t^2 are uncorrelated, which implies that $\lambda_g^{S,1} = \lambda_g^{S,2} = 0$ for all S. Given the abundant empirical evidence on market segmentation, it is unsurprising that we observe a correlation between m_t^1 and m_t^2 much lower than one. Instead, it is insightful to learn whether different economic risk factors play similar or distinct roles in the SDFs of two different markets, as such an analysis can shed light on how to build economic models with market segmentation.

One limitation of the approach above is the lack of time-varying risk premia but, as explained in Section 2.2, our method is generalizable to time-varying term structures. In contrast, Chaieb et al. (2021) and Patton and Weller (2022) allow for time-variation in segmented risk premia and build on the classical Fama-MacBeth regression. In particular, Patton and Weller (2022) devise a novel approach that simultaneously groups portfolio returns in several clusters and estimates the time-varying risk premia based on the clustered market segments. Unlike their paper, our estimates come from prespecified market segments, such as equity versus corporate bonds.

Our framework is novel in two aspects compared to the literature on market segmentation.

First, we begin with latent factors extracted from asset returns and regress observable factors onto them to obtain risk premia, which accounts for the omitted variable bias. In contrast, previous papers mostly rely on notable factor models, such as the three-factor model in Fama and French (1993). In this sense, we extend Giglio and Xiu (2021) into testing segmented risk premia. Furthermore, we add one crucial ingredient into our framework: the term structure of risk premia. As we show in empirical studies, many macro factors have negligible contemporaneous correlations with the SDFs, so their risk premia are close to zeros across asset markets. However, their long-horizon loadings and risk premia tend to be much more sizable, thereby enabling a more meaningful and powerful statistical test against the null hypothesis of homogeneous risk premia.

3 Simulations

This section studies the finite-sample properties of our Bayesian estimates outlined in Propositions 1 and 3 via Monte Carlo simulations. Throughout the simulations, we consider two sample sizes, $T \in \{200, 600\}$, matching the quarterly and monthly frequencies, respectively.

First, we simulate asset returns from a five-factor model as in equation (1). In the benchmark one-market simulations, we use the cross-section of Fama-French 275 portfolios (FF275), described in Internet Appendix IA.3. Factor loadings are calibrated as the eigenvectors corresponding to the five largest eigenvalues of the sample covariance matrix of asset returns, denoted as $\hat{\boldsymbol{\beta}}_{\tilde{v}}$. Since we always normalize latent factors such that they have an identify covariance matrix, $\hat{\boldsymbol{\beta}}_{\tilde{v}}^{\top}\hat{\boldsymbol{\beta}}_{\tilde{v}}$ is a diagonal matrix containing the largest five eigenvalues.

Using the same eigen-decomposition, we estimate the variance of idiosyncratic shocks for each asset, denoted by $\hat{\sigma}_{irt}^2, i = 1, \dots, N$. We allow for a non-diagonal covariance matrix of idiosyncratic shocks. Following Bai and Ng (2002), we simulate w_{irt} as follows:

$$w_{irt} = \hat{\sigma}_{irt} \cdot \Big[e_{it} + \sum_{j \neq 0, j=-J}^{J} \beta e_{i-j,t}\Big], \quad e_{it} \overset{\text{iid}}{\sim} \mathcal{N}(0, \frac{1}{1+2J\beta^2}), \tag{30}$$

where $J = \max\{10, \text{int}\big(N/20\big)\}$ and $\beta = 0.1$.[13]

[13] β cannot be too large since we need to ensure that the largest eigenvalue of $\hat{\boldsymbol{\Sigma}}_{wr}$ is less than the smallest eigenvalue of $\hat{\boldsymbol{\beta}}_{\tilde{v}}^{\top}\hat{\boldsymbol{\beta}}_{\tilde{v}}$. Otherwise, some common factors cannot be identified.

Expected returns are simulated via equation (2). Risk premia are estimated using the observed data, with the estimates denoted by $\hat{\boldsymbol{\lambda}}_{\tilde{v}}$. To ensure that $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}_{\tilde{v}}$ are orthogonal in simulations, we regress the estimated $\boldsymbol{\alpha}$ on $\boldsymbol{\beta}_{\tilde{v}}$ and extract the residual term, denoted by $\hat{\boldsymbol{\alpha}}$. With all the calibrated parameters, asset returns are simulated as follows:

$$\boldsymbol{r}_t = \hat{\boldsymbol{\alpha}} + \hat{\boldsymbol{\beta}}_{\tilde{v}}\hat{\boldsymbol{\lambda}}_{\tilde{v}} + \hat{\boldsymbol{\beta}}_{\tilde{v}}\tilde{\boldsymbol{v}}_t + \boldsymbol{w}_{rt}, \quad \tilde{\boldsymbol{v}}_t \overset{\text{iid}}{\sim} \mathcal{N}(\boldsymbol{0}_K, \boldsymbol{I}_K).$$

We consider both strong and useless factors. We first describe how we simulate a strong factor. For $T = 200$, we use nondurable consumption growth to estimate impulse responses, denoted by $\{\hat{\rho}_s\}_{s=0}^{\bar{S}}$, assuming the true $\bar{S} = 8$ (quarters). For $T = 600$, we use monthly industrial production growth to obtain the monthly impulse responses, and the true $\bar{S}$ is 16 (months). With these parameters, we simulate the strong g_t as follows:

$$g_t = c \cdot \sum_{s=0}^{\bar{S}} \hat{\rho}_s f_{t-s} + w_{gt}, \quad f_t = \frac{1}{\sqrt{6}}(2\tilde{v}_{1t} + \tilde{v}_{3t} + \tilde{v}_{5t}), \quad w_{gt} \overset{\text{iid}}{\sim} \mathcal{N}(0, \sigma_{wg}^2). \tag{31}$$

Equation (31) implies that the common component f_t relates to both large and small principal components (PCs) of asset returns. The first PC, $\tilde{v}_{1t}$, captures more variation in f_t than smaller PCs, consistent with the patterns that we observe in the data. We vary c and σ_{wg}^2 to consider different signal-to-noise ratios summarized by the time series fit $R_g^2 = 1 - \sigma_{wg}^2/\text{var}(g_t)$. When c is larger, or σ_{wg}^2 is smaller, the factor g_t correlates more with asset returns, so identifying its risk premium is ceteris paribus less challenging.

For the weak factor, we simulate f_t independently from a normal distribution, that is, $f_t \overset{\text{iid}}{\sim} \mathcal{N}(0, 1)$, so g_t is orthogonal to asset returns in this case. Nevertheless, the simulated weak factor g_t is autocorrelated, so we can use it to explore whether the Newey and West (1987) type of sandwich covariance matrix can deliver proper Bayesian credible intervals for factors with an autocorrelated measurement error.

3.1 Simulations in One Asset Market

First, we study the finite-sample performance of our Bayesian estimates in Proposition 1. We consider three distinct signal-to-noise ratios, that is, $R_g^2 \in \{30\%, 20\%, 10\%\}$. The pseudo-true number of latent factors is five; however, we explore three different numbers of factors in estimation, $K \in \{4, 5, 7\}$, to study the performance of our approach when we erroneously omit some priced factors ($K = 4$) or include redundant ones ($K = 7$).

Tables A1 and IA.III of the Internet Appendix report the empirical size of our test for strong factors in 1,000 simulations. We estimate the term structure of g_t's risk premia using $\bar{S} = 12$ for $T = 200$ and $\bar{S} = 24$ for $T = 600$. The performance of our Bayesian estimator is reassuring as long as we include all priced latent factors in the estimation ($K \geq 5$). Specifically, our method provides appropriate credible intervals for g_t's risk premia, even in an environment with a low signal-to-noise ratio and a small sample size. However, if we omit some priced factors (e.g., the number of factors is four), our Bayesian estimates are biased because the simulated g_t loads on the fifth PC of asset returns. For example, we observe over-rejections in the panels of four-factor models; hence, omitting the priced factor leads to the risk premia estimates not centering around the pseudo-true values. Nevertheless, including more latent factors than in the pseudo-true model has no sizable detrimental effect – suggesting that such an approach is conservative.

Can we recover the priced information embedded in g_t if we consider only the contemporaneous correlation between g_t and asset returns? To answer this question, we estimate the models with different numbers of lags $\bar{S}$. Figure 1 plots the average correlation between the true f_t and its estimate, $\hat{f}_t = \hat{\boldsymbol{\eta}}_g^\top \hat{\boldsymbol{v}}_t$. When we project g_t only on the current asset return shocks ($\bar{S} = 0$ in equation (9)), $\text{corr}(f_t, \hat{f}_t)$ is small, ranging from 0.4 to 0.65. As we include more lagged asset pricing information in g_t, this correlation coefficient significantly increases; hence, including the lagged asset return information is essential in identifying the priced shock driving the nontradable factor. Notably, the detrimental effect of including more lags than in the pseudo-true specification is generally very small.

Figure 2 reports the power of rejecting zero risk premia of strong factors.[14] The model with $\bar{S} = 0$ has low test power, even at the monthly frequency ($T = 600$) with $R_g^2 = 30\%$. In contrast, as we include more lagged latent factors in g_t's estimation, we considerably increase the test power. Hence, our proposed MA representation of g_t is the key to detecting significant risk premia in persistent factors.

Including more factors (e.g., in the seven-factor models) tends to be a conservative strategy since it delivers proper yet wider credible intervals of risk premia estimates. Nevertheless, it

[14]We report the power for $R_g^2 \in \{10\%, 20\%\}$ in Figure IA.1 of the Internet Appendix.

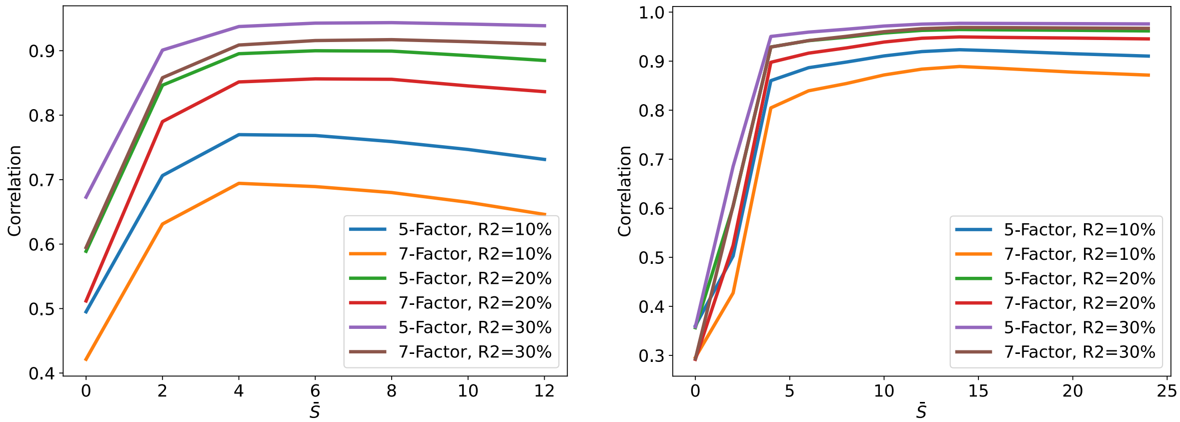

(a) $T = 200$ **(b)** $T = 600$

Figure 1: Posterior median of correlation coefficients between true and estimated f_t

The figure plots the average $\mathrm{corr}(\hat{f}_t, f_t)$ in 1,000 simulations, where $\mathrm{corr}(\hat{f}_t, f_t)$ quantifies the correlation between the true f_t and its estimate, $\hat{f}_t = \hat{\boldsymbol{\eta}}_g^\top \hat{\boldsymbol{v}}_t$. We consider strong factors, with $R_g^2 \in \{10\%, 20\%, 30\%\}$, and two sample sizes, $T \in \{200, 600\}$. In each simulated scenario, we estimate several model configurations with different numbers of factors and different $\bar{S}$.

comes at the cost of lowering the power of the test and the correlation between the true and estimated f_t. In the empirical application, we will explore whether our risk premia estimates are robust to adding more latent factors, acknowledging that more factors will increase the estimation uncertainty mechanically. Another critical parameter driving test power is the signal-to-noise ratio, R_g^2. As R_g^2 increases, our tests command larger power to reject the null of zero risk premia for priced factors.

In Tables IA.IV and IA.V of the Internet Appendix, we investigate useless factors that do not correlate with asset returns. The useless factors are assumed to be persistent, and a larger R_g^2 corresponds to a more persistent process. Past literature (e.g., Kan and Zhang (1999a,b)) points out the fragility of Fama-MacBeth and GMM estimates of risk premia in the presence of useless factors. It is worth noting that our Bayesian estimates do not suffer from this issue. The Bayesian credible intervals of useless factors' risk premia tend to be conservative, leading to a slight under-rejections of zero risk premia in small sample.

One potential concern is that including many lags of multiple latent factors might lead to severe overfitting of the data. To alleviate this concern, we report in the Internet Appendix (see Table IA.VI) the posterior means of R_g^2 in 1,000 simulations. Our simulation results suggest that the posterior means of R_g^2 are reasonably close to their pseudo-true values, even in a small

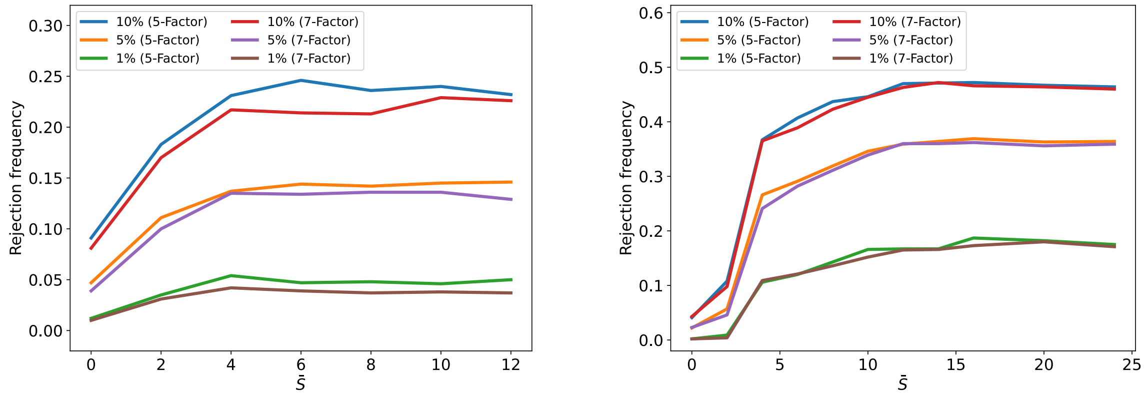

(a) $R_g^2 = 30\%$ & $T = 200$ **(b)** $R_g^2 = 30\%$ & $T = 600$

Figure 2: Power of identifying strong factors

The figure plots the frequency of rejecting the null hypothesis $H_0 : \lambda_g^{\bar{S}} = 0$, based on the 90%, 95%, and 99% credible intervals based on our Bayesian estimates in Proposition 1. $\lambda_g^{\bar{S}}$ is defined in equation (7). We consider strong factors, with $R_g^2 = 30\%$, and two sample sizes, $T \in \{200, 600\}$. In each simulated scenario, we estimate several model configurations with different numbers of factors and different $\bar{S}$. The number of Monte Carlo simulations is 1,000.

sample with 200 quarters. Hence, our approach does not lead to significantly inflated time series fits for g_t.

Moreover, we explore the performance of our Bayesian estimates for factors that correlate with only the contemporaneous asset return shocks (i.e., $\bar{S} = 0$ in the true data-generating process of g_t), which fits the model configuration studied in Giglio and Xiu (2021). In particular, we simulate priced factors using equation (31) by setting $\bar{S} = 0$ and tuning the value of c such that $R_g^2 \in \{10\%, 20\%, 30\%\}$. The simulation results are in Table IA.VII of the Internet Appendix. We study both the size and power of the test statistics based on (1) our Bayesian estimation in Proposition 1 (setting $\bar{S} = 0$) and (2) the frequentist test statistic in Theorem 1 of Giglio and Xiu (2021). Overall, our Bayesian test has almost identical size and power to the frequentist test in Giglio and Xiu (2021) in the special case of $\bar{S} = 0$.

Finally, in Appendix IA.2, we repeat our simulation study to examine the time-varying risk premia and their term structures as described in Section 2.2. Overall, size and power, as well as the correlation between filtered and calibrated latent processes (see Tables IA.X–IA.IX in the Appendix), are similar to those reported in this section. Despite the significant added generality, modeling the latent systematic risk drivers as following a VAR(1) process, we observe only a minimal degree of attenuation bias and increased posterior uncertainty for the estimated term structure of risk premia.

3.2 Simulations in Two Markets

This subsection explores the empirical performance of Proposition 3 in testing heterogeneous risk premia of g_t in two asset markets. In addition to the equity cross-section, we consider the cross-section of 40 corporate bond portfolios in Elkamhi, Jo, and Nozawa (2023) (EJN40).[15]

We assume a five-factor model for the corporate bond portfolios. We calibrate the EJN40 factor loadings, risk prices of latent factors, expected returns, and covariance matrix of returns' idiosyncratic errors, from the data in a method similar to the calibration of FF275.

We calibrate the data-generating process of g_t such that it is priced in FF275 but carries a zero risk premium in EJN40. In FF275, we assume that the correlation between its latent factors $\tilde{\boldsymbol{v}}_t^1$ and f_t is $\text{corr}(\tilde{\boldsymbol{v}}_t^1, f_t) = \hat{\boldsymbol{\eta}}_g^1 = \frac{1}{\sqrt{6}}(2,0,1,0,1)^\top$, where f_t is the common component driving asset returns and g_t. Under this calibration, g_t is priced in FF275. In contrast, we assume that $\text{corr}(\tilde{\boldsymbol{v}}_t^2, f_t) = \hat{\boldsymbol{\eta}}_g^2 = \frac{1}{\sqrt{(\lambda_{\tilde{v},1}^2)^2+(\lambda_{\tilde{v},3}^2)^2}}(\lambda_{\tilde{v},3}^2, 0, -\lambda_{\tilde{v},1}^2, 0, 0)^\top$, where $\lambda_{\tilde{v},1}^2$ and $\lambda_{\tilde{v},3}^2$ are, respectively, the risk prices of $\tilde{v}_{1t}^2$ and $\tilde{v}_{3t}^2$. This assumption ensures that g_t is not priced in EJN40. $\tilde{\boldsymbol{v}}_t^1$ and $\tilde{\boldsymbol{v}}_t^2$ are correlated, with correlation matrix ($\hat{\boldsymbol{\Sigma}}_{\tilde{v}}^{(1,2)}$) calibrated using the real data. Under these assumptions, we calibrate η_g and simulate f_t as follows:

$$\hat{\boldsymbol{\eta}}_g = \begin{bmatrix} \boldsymbol{I}_{K_1} & \hat{\boldsymbol{\Sigma}}_{\tilde{v}}^{(1,2)} \\ \hat{\boldsymbol{\Sigma}}_{\tilde{v}}^{(2,1)} & \boldsymbol{I}_{K_2} \end{bmatrix}^{-1} \cdot \begin{bmatrix} \hat{\boldsymbol{\eta}}_g^1 \\ \hat{\boldsymbol{\eta}}_g^2 \end{bmatrix}, \quad f_t = \left[\frac{\hat{\boldsymbol{\eta}}_g}{|\hat{\boldsymbol{\eta}}_g|}\right]^\top \tilde{\boldsymbol{v}}_t, \quad \text{and } g_t \text{ follows equation (31).}$$

We study in Tables IA.XI and IA.XII of the Internet Appendix the posterior coverage of our Bayesian credible intervals for $\lambda_g^{S,1} - \lambda_g^{S,2}$. Similar to the observations in the one-market simulations, omitting priced factors biases the risk premia estimates, leading to the over-rejection of the null hypothesis. However, in the five- and seven-factor cases, our Bayesian estimates provide valid credible intervals for the risk premium difference, even when the signal-to-noise ratio is low. Figures 3 and IA.2 in the Internet Appendix further explore the power of rejecting identical risk premia in two asset markets. Overall, the power of the test increases i) if we include more lags (larger $\bar{S}$) and do not use redundant latent factors and ii) if the signal-to-noise ratio (R_g^2) is larger.

[15] We thank the authors for generously providing us with the dataset. The detailed description of corporate bond data is in Internet Appendix IA.3.

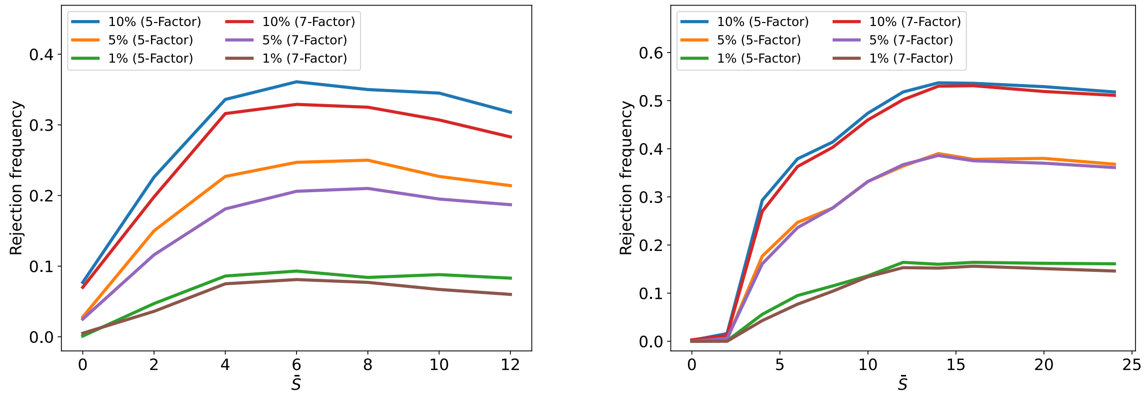

(a) $R_g^2 = 30\%$ & $T = 200$

(b) $R_g^2 = 30\%$ & $T = 600$

Figure 3: Power of identifying heterogeneous risk premia in two markets

The figure plots the frequency of rejecting the null hypothesis $H_0 : \lambda_g^{S,1} = \lambda_g^{S,2}$ based on the 90%, 95%, and 99% credible intervals based on our Bayesian estimates in Proposition 3. $\lambda_g^{\bar{S}}$ is defined in equation (7). We consider strong factors, with $R_g^2 = 30\%$, and two sample sizes, $T \in \{200, 600\}$. In each simulated scenario, we estimate several model configurations with different numbers of factors and different $\bar{S}$. The number of Monte Carlo simulations is 1,000.

4 Empirical Analysis

In this section, we apply our Bayesian framework to estimate the risk premia of both tradable and nontradable factors. We aim to investigate whether factors are priced, the term structure of the factors' risk premia, and whether the factors command heterogeneous risk premia in two potentially segmented asset markets.

4.1 Risk Premia Estimates in Equity Markets

We begin our empirical investigation with the risk premia in equity markets. Our analysis relies on a large cross-section of FF275, covering the period between Q3 1963 and Q4 2019. Throughout our paper, we standardize factors and returns to have unit variances per period. Definition, sample periods, and data sources of factors and test assets can be found in Internet Appendix IA.3.

To conduct our Bayesian estimation in Section 2, we need to determine the number of latent factors, K. We adopt the selection approach proposed by Giglio and Xiu (2021)[16] and estimate

[16] We follow the method in Internet Appendix I.1 of Giglio and Xiu (2021). That is, the selected number of factors is equal to $\hat{K} = \arg\min_{1\leq j\leq K_{\max}} \left[N^{-1}T^{-1}\gamma_j(\bar{\boldsymbol{R}}^\top \bar{\boldsymbol{R}}) + j \times \phi(n, T)\right] - 1$, where $\bar{\boldsymbol{R}}$ is a $T \times N$ matrix of

that the number of factors is five in FF275 at monthly or quarterly frequencies.

Moreover, we find that the first several latent factors explain most of the time series and cross-sectional variations. In the time series dimension, the first five PCs account for more than 92% of time series variations at monthly and quarterly frequencies. Adding the 6th and 7th PCs only marginally improves the time series fit. In the cross-sectional dimension, the five-, six, and seven-factor models explain 73.7%, 73.8%, and 74.6% (70.0%, 73.7%, and 73.9%) of cross-sectional variations in average returns at the monthly (quarterly) frequency. Therefore, the statistical test in Giglio and Xiu (2021), as well as time series and cross-sectional fit, indicate that the five-factor model is a reasonable benchmark; we thus adopt it in our baseline estimations (but also conduct robustness checks with $K = 6$ or 7).

We first explore Bayesian risk premia estimates of some canonical tradable factors and compare them with their time series average excess returns. Figure 4 plots the term structure of risk premia for Carhart (1997) four factors, whose risk premia are estimated using Proposition 1 ($\bar{S} = 24$ and $K = 5$). These tradable factors tend to have almost flat term structures of risk premia. The Bayesian point estimates (solid blue lines) have similar magnitudes as the time series Sharpe ratios (grey dotted lines), which are covered by the 68% Bayesian credible intervals (purple dotted lines). Therefore, our approach provides estimates very close to the time series averages of tradable factors in both economic and statistical sense.

Next, we study other economic variables and report their risk premia estimates in Table 1. For quarterly (monthly) variables, we conduct the Bayesian estimation as in Proposition 1, using a lag of 12 quarters (24 months) in g_t's equations. Several empirical findings are noteworthy.

First, many macro factors carry significant risk premia, including IP growth, GDP growth, durable and nondurable consumption growth, dividend growth, and macro PCs 1, 2, and 4 in the FRED-QD dataset of McCracken and Ng (2020).[17] More interestingly, most of them have increasing term structures of risk premia, as shown in Figure 5. At quarterly frequency ($S = 0$),

demeaned asset returns, $\gamma_j(\bar{\boldsymbol{R}}^\top \bar{\boldsymbol{R}})$ is the j-th eigenvalue of $\bar{\boldsymbol{R}}^\top \bar{\boldsymbol{R}}$, $\phi(n, T) = 0.5 \times \hat{\gamma} \times (\log(N) + \log(T))(N^{-\frac{1}{2}} + T^{-\frac{1}{2}})$, and $\hat{\gamma}$ is the median of the first $K_{\max}$ eigenvalues of $\bar{\boldsymbol{R}}^\top \bar{\boldsymbol{R}}$. We set $K_{\max}$ to 20.

[17] Dividend growth is the quarterly growth of the smoothed aggregate dividend payments made in the previous 12 months. We consider the smoothed annual dividends of the S&P 500 index in order to remove the mechanical seasonality in the dividend payments.

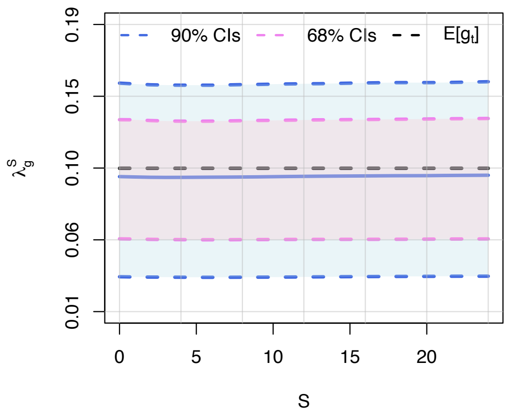

(a) MKT

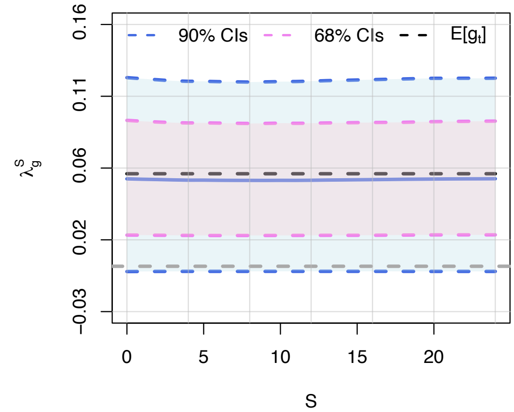

(b) SMB

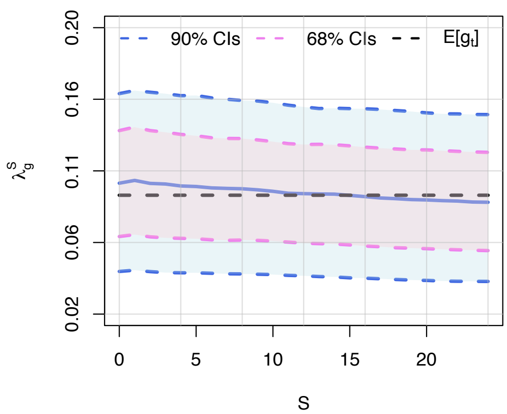

(c) HML

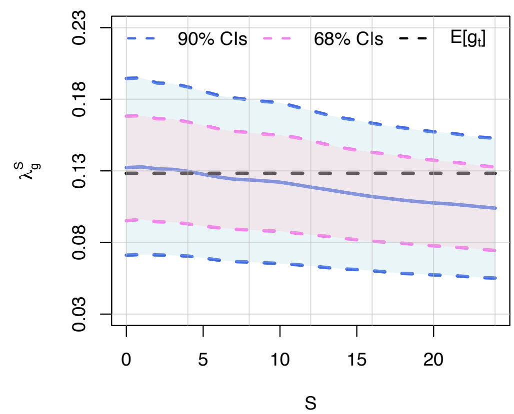

(d) MOM

Figure 4: Term structure of risk premia: Carhart four factors

Term structure of risk premia estimates (in Sharpe ratio units) using Proposition 1. The risk premium at horizon S (λ_g^S) is defined in equation (7). The cross-section of test assets consists of 275 Fama-French characteristic-sorted portfolios. We consider five-factor models for asset returns. We study monthly Carhart (1997) four factors, whose risk premia are estimated using a lag of 24 months in g_t's equations. We include their in-sample monthly Sharpe ratios (grey dotted lines). In addition to the point estimates, we report the 68% and 90% Bayesian credible intervals, highlighted in pink and blue, respectively. Definition and data sources of factors and test assets can be found in Appendix IA.3. Sample: July 1963 to December 2019.

most macroeconomic factors are weakly identified at best. However, risk premia carried by these macro factors are significant and as large as that of the market at business cycle frequencies (two–three years). Therefore, these macro factors are riskier from the perspective of long-term than short-term investors. The only exception among the priced macro factors is macro PC2, where we detect an almost flat term structure.

Second, the observations in Table 1 have direct implications for leading macro-finance models. Figure A1 of the Appendix plots the term structure of risk premia in the habit (Campbell

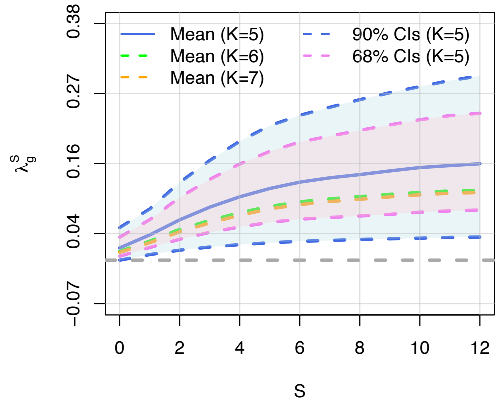

(a) GDP growth

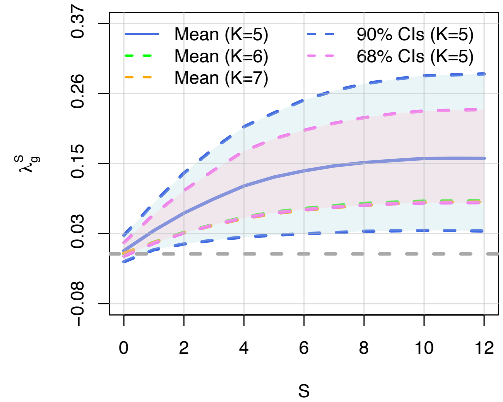

(b) IP growth

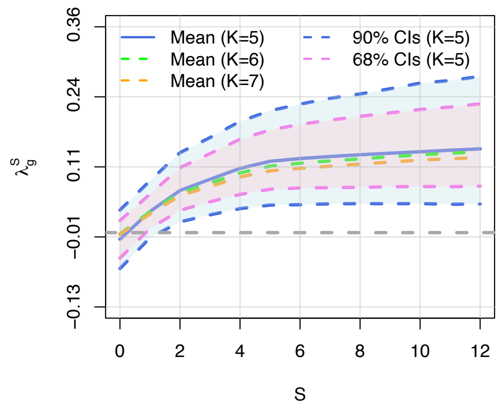

(c) Durable consumption growth

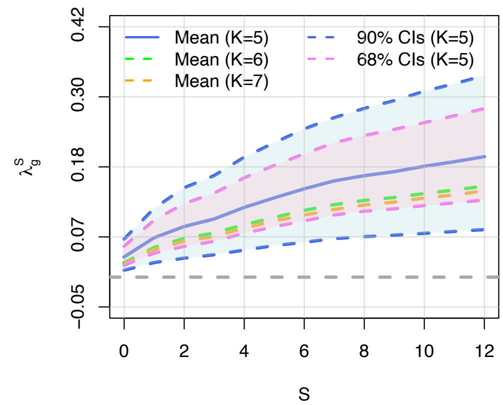

(d) Nondurable consumption growth

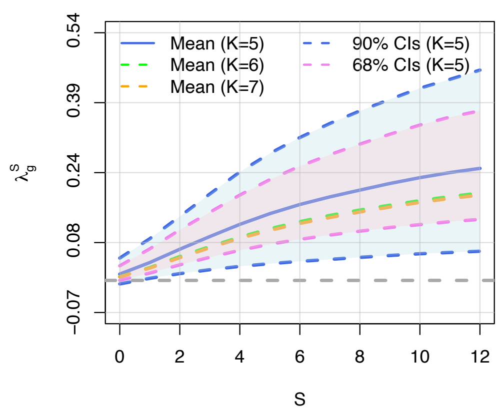

(e) Macro PC1 (FRED-QD)

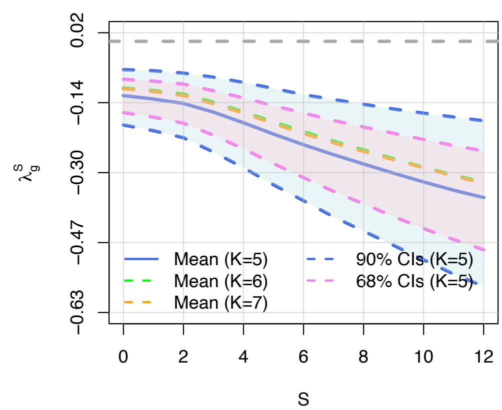

(f) Macro PC4 (FRED-QD)

Figure 5: Term structure of factor's risk premia: Some priced macro factors

This figure plots the term structure of risk premia estimates using Proposition 1, where the risk premium over S horizons (λ_g^S) is defined in equation (7). The cross-section of test assets consists of 275 Fama-French characteristic-sorted portfolios. We consider five-, six- and seven-factor models for asset returns. In addition to the point estimates, we show the 68% and 90% Bayesian credible intervals based on five-factor models, highlighted in pink and blue, respectively. Definition and data sources of factors and test assets can be found in Internet Appendix IA.3.

Table 1: Factors' risk premia: Five-factor models

	Panel A. Quarterly variables, $\bar{S} = 12$ quarters							
$S =$	0	2	4	6	8	10	12	R^2_g
AEM intermediary	0.083***	0.078**	0.078**	0.061	0.044	0.024	0.015	14.4%
Capital share growth	0.008	0.010	0.006	0.001	-0.004	-0.007	-0.012	7.6%
GDP growth	0.020	0.065**	0.102**	0.125**	0.137**	0.149**	0.155**	23.4%
IP growth	0.005	0.066**	0.109**	0.133**	0.147**	0.153**	0.154**	37.5%
Durable consumption growth	-0.011	0.074**	0.112***	0.130***	0.136***	0.141***	0.146***	18.3%
Nondurable consumption growth	0.034***	0.085***	0.117***	0.148***	0.170***	0.186***	0.202***	22.4%
Service consumption growth	0.004	0.009	0.015	0.021	0.026	0.030	0.032	9.9%
Nondurable + service	0.022	0.054	0.078	0.099	0.117	0.130	0.141	18.6%
Labor income growth	0.000	0.002	0.002	0.004	0.004	0.005	0.008	5.2%
Dividend growth of S&P500	0.005	0.017	0.046	0.096**	0.157***	0.218***	0.273***	51.1%
Macro PC1 (FRED-QD)	0.014	0.068**	0.122**	0.166**	0.197**	0.224**	0.244**	47.9%
Macro PC2 (FRED-QD)	0.076**	0.115**	0.115**	0.099*	0.082	0.066	0.051	37.1%
Macro PC3 (FRED-QD)	-0.002	-0.003	-0.003	-0.002	-0.001	-0.001	0.000	10.9%
Macro PC4 (FRED-QD)	-0.126***	-0.144***	-0.189***	-0.240***	-0.285***	-0.327***	-0.363***	47.3%
Macro PC5 (FRED-QD)	0.033	0.034	0.024	0.013	0.005	0.001	-0.002	29.3%
	Panel B. Monthly variables, $\bar{S} = 24$ months							
$S =$	0	4	8	12	16	20	24	R^2_g
Oil price change	-0.006	-0.029	-0.043	-0.048	-0.051	-0.051	-0.049	7.1%
TED spread change	0.000	0.000	0.002	0.002	0.002	0.003	0.003	8.9%
Nontraded HKM intermediary	0.075**	0.079**	0.076**	0.073**	0.072**	0.070**	0.069**	60.9%
Traded HKM intermediary	0.090***	0.091***	0.087***	0.082***	0.080***	0.077***	0.076***	71.0%
PS liquidity	0.039**	0.059**	0.069**	0.077**	0.086**	0.094**	0.100**	15.0%
$\Delta\log(\text{VIX})$	-0.112***	-0.068***	-0.053***	-0.041***	-0.036***	-0.031***	-0.027**	51.6%

The table reports Bayesian estimates of factors' risk premia using Proposition 1, where the risk premia over S horizons (λ_g^S) are defined in equation (7). The cross-section of test assets consists of 275 Fama-French characteristic-sorted portfolios. We consider a five-factor model for asset returns. Panel A tabulates the estimates of quarterly factors, using a lag of 12 quarters in g_t's equations. Panel B tabulates the estimates of monthly factors, using a lag of 24 months in estimation. We use Bayesian credible intervals to conduct hypothesis testing: If the 90% (95%, 99%) credible interval of g_t's risk premium does not contain zero, the risk premium estimate will be highlighted by * (**, ***). Definition and data sources of factors and test assets can be found in Internet Appendix IA.3.

and Cochrane (1999)) and long-run risk frameworks (Bansal and Yaron (2004)).[18] Specifically, the habit model implies a flat term structure of consumption risk premia, whereas it is upward-sloping in the long-run risk model. With respect to dividend growth, we consider the quarterly growth of the smoothed dividend payment (defined as the aggregate dividend payments made in the previous 12 months) to be consistent with our empirical analysis. Even though both models predict upward-sloping term structures of risk premia for smoothed dividend growth, the magnitudes and slopes are much more sizable in the long-run risk model than in the habit model. Overall, the long-run risk model tends to be more consistent with our estimates for nondurable consumption and dividend growth.

Third, researchers often fail to identify priced macro risks when studying only the contem-

[18] We discuss the calibrations in detail in Internet Appendix IA.4.

porary correlations between asset returns and macro factors, even when extracting their AR(1) innovations. The first column of Table 1, and Figure 5, indicate that the risk premia of GDP growth, IP growth, durable consumption growth, dividend growth, and macro PC1 are tiny and insignificant at $S = 0$. One concern of the analysis in Table 1 and Figure 5 is that we include many lags in the estimation, leading to noisier risk premia estimates. To alleviate this concern we repeat the analysis using $\bar{S} = 0$. Panel A of Table 2 shows that among the eight priced macro factors mentioned above, only macro PC2 and PC4 carry significant risk premia in this case. Panel B further extracts the AR(1) innovations in macro factors and estimates their risk premia by setting $\bar{S} = 0$. Similar to Panel A, we observe only macro PC2, PC4, and dividend growth (significant only in the seven-factor model) being priced, while all other macro factors have negligible and insignificant risk premia. Overall, these findings confirm our simulation based evidence that the long MA representation is crucial for correctly recovering the risk premia of persistent factors.

But why does including lagged asset return shocks in g_t's equation enable us to identify the priced risk? The time series fit, R_g^2, sheds light on this issue. For most traditional macro factors, R_g^2 values in Table 1 are considerably larger than those in Table 2. For instance, the contemporaneous five factors of asset returns explain only 2.8% of time series variations in macro PC1, but its R_g^2 increases to 47.9% in the estimation with $\bar{S} = 12$ quarters, hence greatly enhancing the signal-to-noise ratio and our ability to identify the risk premia.

Many macro variables are, in nature, persistent. The shocks to these variables are small per period, but the cumulative impulse responses can be sizable. Although the AR(1) model is often used in both empirical and theoretical works, extracting the AR(1) innovations is insufficient to recover the risk premia of many macro variables, either because the AR(1) shocks are inconsequential or the AR(1) assumption is questionable. Differently, the MA representation does not take a stance on their exact data-generating processes. We use both the current and lagged asset return innovations to extract the common priced component driving macro states and asset returns, hence recovering the risk premia.

Fourth, intermediary factors are priced in the cross-section. Specifically, Adrian, Etula, and Muir (2014) (AEM) intermediary factor commands a significantly positive risk premium at short horizons ($S \leq 4$ quarters). The nontraded He, Kelly, and Manela (2017) (HKM) intermediary

Table 2: Factors' risk premia: $\bar{S} = 0$

	$\mathbb{E}[\lambda_g \mid \mathcal{D}]$			$\mathbb{E}[R_g^2 \mid \mathcal{D}]$		
Number of factors:	5	6	7	5	6	7
			Panel A. Original factors			
AEM intermediary	0.138***	0.176***	0.176***	10.4%	12.2%	12.5%
Capital share growth	0.030	0.013	0.010	1.8%	2.7%	2.8%
GDP growth	-0.004	0.007	0.009	4.2%	4.3%	4.3%
IP growth	-0.031	0.005	0.005	2.9%	4.3%	4.3%
Durable consumption growth	-0.024	-0.013	-0.014	7.7%	7.9%	8.2%
Nondurable consumption growth	0.030	0.047	0.047	3.7%	4.0%	4.1%
Service consumption growth	0.007	0.051	0.052	4.0%	6.4%	6.5%
Nondurable + service	0.021	0.061	0.061	4.0%	5.9%	6.0%
Labor income growth	-0.010	0.033	0.026	1.5%	3.8%	8.7%
Dividend growth of S&P500	-0.004	-0.018	-0.013	5.0%	5.1%	6.8%
Macro PC1 (FRED-QD)	-0.013	0.020	0.024	2.8%	3.9%	4.6%
Macro PC2 (FRED-QD)	0.114***	0.087*	0.082*	21.3%	22.1%	23.0%
Macro PC3 (FRED-QD)	-0.053	-0.068*	-0.068*	4.3%	4.7%	4.8%
Macro PC4 (FRED-QD)	-0.138***	-0.150***	-0.152***	25.1%	25.4%	26.6%
Macro PC5 (FRED-QD)	0.040	0.078*	0.073	24.8%	25.5%	27.9%
Oil price change	-0.021	-0.022	-0.021	2.6%	4.5%	4.5%
TED spread change	-0.025	-0.029	-0.032	6.8%	10.9%	17.4%
Nontraded HKM intermediary	0.078**	0.081**	0.080**	60.4%	61.0%	61.2%
Traded HKM intermediary	0.088**	0.092***	0.091***	70.3%	71.0%	71.2%
PS liquidity	0.051***	0.049***	0.052***	11.9%	12.2%	12.9%
$\Delta\log(\text{VIX})$	-0.100***	-0.100***	-0.099***	42.8%	43.0%	43.1%
			Panel B. AR(1) shocks of macro factors			
GDP growth	-0.004	0.007	0.009	4.2%	4.3%	4.3%
IP growth	-0.025	-0.001	-0.003	3.6%	4.3%	4.9%
Durable consumption growth	-0.022	-0.010	-0.012	7.4%	7.6%	7.8%
Nondurable consumption growth	0.030	0.048	0.048	3.7%	4.1%	4.1%
Service consumption growth	0.008	0.054	0.055	3.7%	6.3%	6.4%
Nondurable + service	0.021	0.064	0.065	3.7%	6.0%	6.0%
Labor income growth	-0.011	0.031	0.024	1.4%	3.7%	8.2%
Dividend growth of SP500	0.039	0.063	0.067*	2.4%	3.6%	5.2%
Macro PC1 (FRED-QD)	0.002	0.006	0.003	6.1%	6.1%	7.1%
Macro PC2 (FRED-QD)	0.080*	0.051	0.046	23.9%	24.8%	27.6%
Macro PC3 (FRED-QD)	-0.043	-0.044	-0.047	3.3%	3.3%	4.1%
Macro PC4 (FRED-QD)	-0.130***	-0.142***	-0.147***	26.8%	26.9%	28.7%
Macro PC5 (FRED-QD)	0.026	0.038	0.032	33.8%	33.1%	37.2%
Oil price change	-0.030	-0.031	-0.030	3.2%	4.8%	4.8%

The table reports Bayesian estimates of (1) factors' risk premia and (2) time series fit R_g^2. Panel A considers the original variables that are identical to those in Tables 1 and IA.XIV, whereas Panel B studies the AR(1) shocks of some macro factors. We estimate model parameters using Proposition 1 by setting $\bar{S} = 0$. The cross-section of test assets consists of 275 Fama-French characteristic-sorted portfolios. We consider five-, six-, and seven-factor models for asset returns. For risk premia estimates, we use Bayesian credible intervals to conduct hypothesis testing: If the 90% (95%, 99%) credible interval of g_t's risk premium does not contain zero, the risk premium estimate will be highlighted by * (**, ***). Definition and data sources of factors and test assets can be found in Internet Appendix IA.3.

factor is strongly identified, with $R_g^2 = 62\%$, and its risk premia estimates are significantly positive at any horizon. Interestingly, the nontraded HKM factor commands almost the same risk premium as its tradable version, whose risk premium is almost identical to its time series average. Comparing the R_g^2 of AEM and HKM factors in Table 1 with those in Table 2, we find that lagged asset return innovations are not essential in driving intermediary factors.

Fourth, the term structure of VIX risk premia (more precisely, their absolute values) is downward-sloping. The mimicking portfolio hedging against monthly VIX changes earns a sizable risk premium of -0.11, but the two-year risk premium declines to only -0.03, although still significant. This observation is fully consistent with the previous literature (Eraker and Wu (2014), Dew-Becker et al. (2017), and Johnson (2017)), which estimates VIX risk premia using derivative contracts with different expiration dates.

Perhaps the most surprising empirical finding is that macro variables carry much more sizeable risk premia at long horizons ($S = 8$ to 12 quarters) than at quarterly frequency ($S = 0$). What is the economic intuition behind this phenomenon? To help answer this question, we plot in Figure 6 the MA component spanned by six priced macro variables and asset return factors, that is, $\sum_{s=0}^{\bar{S}} \rho_s \boldsymbol{\eta}_g^\top \boldsymbol{v}_{t-s}$. Strikingly, the MA components of all these six macro variables present clear business cycle patterns. The variables in Panels (a)–(e) are countercyclical, consistent with their positive risk premia: Long-horizon investors hedge against lower realizations of these macro factors. In contrast, macro PC4 is procyclical. As Panel (f) indicates, investors are averse to the spikes in this macro factor, so its risk premium is significantly negative.

Are the MA components of the priced macro factors similar? Table 3 shows that macro PC1, GDP growth, and IP growth have highly correlated MA components, often with correlation coefficients larger than 0.85, and their correlation with nondurable consumption is 70% or more. Nevertheless, the MA components of other macro variables, although correlated, seem to contain considerably independent information. In short, we detect some string commonality in the priced component of these macro variables, but they are not all alike.[19]

Which principal components of asset returns drive g_t? Table IA.XIII in the Internet Appendix reports the posterior means of the squared correlation[20] between the common component

[19]Table IA.XV and Figure IA.5 in the Internet Appendix repeat these analyses in six- and seven-factor models, showing very similar empirical patterns.

[20]We do not report the correlation since we cannot identify the sign of $\hat{\boldsymbol{\eta}}_g^\top \hat{\boldsymbol{v}}_t$.

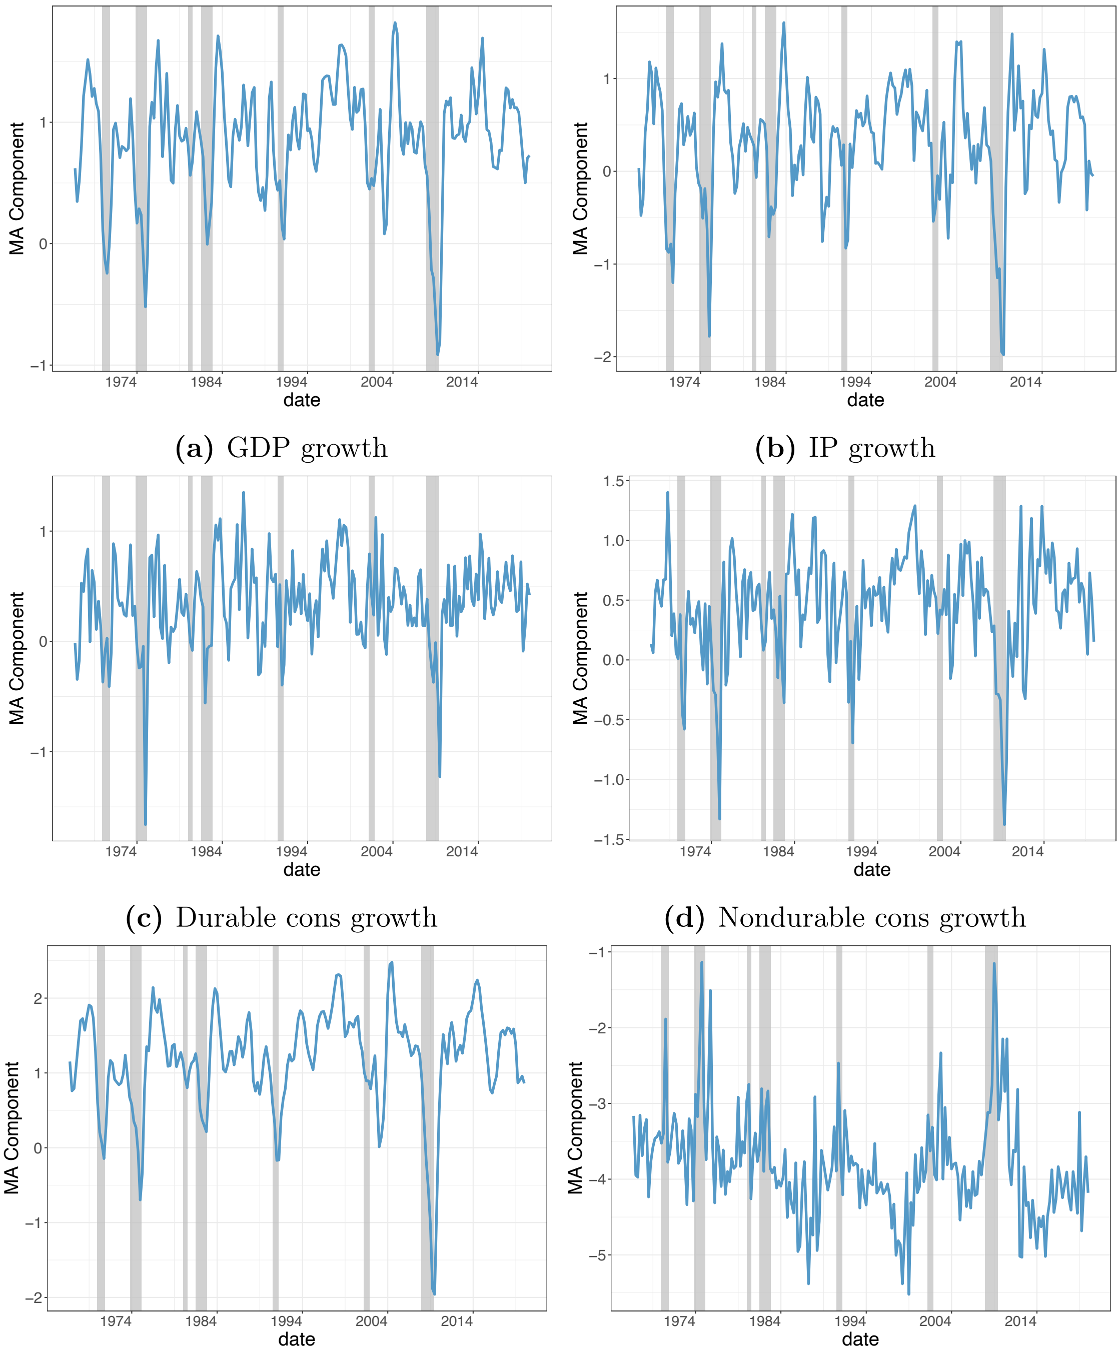

(a) GDP growth

(b) IP growth

(c) Durable cons growth

(d) Nondurable cons growth

(e) Macro PC1 (FRED-QD)

(f) Macro PC4 (FRED-QD)

Figure 6: Moving average components of some macro factors

This figure plots the time series of (posterior means of) moving average components spanned by asset returns' latent factors: $\sum_{s=0}^{\bar{S}} \rho_s \boldsymbol{\eta}_g^\top \boldsymbol{v}_{t-s}$, with $\bar{S} = 12$ quarters. The cross-section of test assets consists of 275 Fama-French characteristic-sorted portfolios. We consider five-factor models for asset returns. Additional plots (considering six- and seven-factor models, other macro factors) are in Table IA.5. Definition and data sources of factors and test assets can be found in Internet Appendix IA.3. Sample: Q3 1963 to Q4 2019.

Table 3: Are MA components of macro factors similar in five-factor models?

	GDP growth	IP growth	Durable	Nondurable	Service	Dividend	Macro PC1	Macro PC2	Macro PC4
GDP growth	1.00	0.90	0.69	0.70	0.62	0.16	0.90	0.43	-0.45
IP growth	0.90	1.00	0.73	0.70	0.58	0.04	0.85	0.41	-0.25
Durable	0.69	0.73	1.00	0.65	0.37	0.14	0.62	0.33	-0.19
Nondurable	0.70	0.70	0.65	1.00	0.60	0.32	0.73	0.34	-0.55
Service	0.62	0.58	0.37	0.60	1.00	0.27	0.73	0.14	-0.44
Dividend	0.16	0.04	0.14	0.32	0.27	1.00	0.38	-0.39	-0.59
Macro PC1	0.90	0.85	0.62	0.73	0.73	0.38	1.00	0.15	-0.51
Macro PC2	0.43	0.41	0.33	0.34	0.14	-0.39	0.15	1.00	-0.19
Macro PC4	-0.45	-0.25	-0.19	-0.55	-0.44	-0.59	-0.51	-0.19	1.00

The table reports the correlation among the moving average components spanned by asset returns' latent factors, $\sum_{s=0}^{\bar{S}} \rho_s \boldsymbol{\eta}_g^\top \boldsymbol{v}_{t-s}$, with $\bar{S} = 12$ quarters. The cross-section of test assets consists of 275 Fama-French characteristic-sorted portfolios. We consider five-factor models for asset returns. Definition and data sources of factors and test assets can be found in Internet Appendix IA.3.

estimates, $\hat{\boldsymbol{\eta}}_g^\top \hat{\boldsymbol{v}}_t$, and the first seven PCs of asset returns, where the posterior distributions of $\hat{\boldsymbol{\eta}}_g$ and $\hat{\boldsymbol{v}}_t$ are estimated using a seven-factor model. The first PC of asset returns is the most important, particularly for the priced factors. Specifically, PC1 of asset returns accounts for 65–90% of the time series variations in the common components of GDP growth, IP growth, nondurable consumption growth, dividend growth, macro PCs 1, 2, 4, HKM intermediary factors, the liquidity factor, and the VIX changes. Overall, the common component is spanned mainly by the first five PCs of asset returns.

However, several variables are closely related to PC6 and PC7 of equity portfolio returns. For example, these two small PCs explain 44% of the common component in labor income growth. Furthermore, PC6 of asset returns accounts for 24%, 38%, and 10% of common components in capital share growth, macro PC3, and oil price change. While labor income growth, capital share growth, and macro PC3 are not priced in six- and seven-factor models, the risk premia estimates of oil price change become significantly negative after we include PC6 of asset returns. Therefore, it is important to conduct robustness checks by considering different numbers of latent factors. We report the term structure of risk premia estimates based on six- and seven-factor models in the Internet Appendix. (See Table IA.XIV and Figures IA.3 and IA.4) The point estimates of most factors are nearly unchanged, but Bayesian credible intervals often become wider, consistent with the observations in simulation studies.

4.2 Time-Varying Term Structure of Macroeconomic Risk Premia

We now turn to the analysis of the time variation in the term structure of macroeconomic factors' risk premia, applying the method in Section 2.2. Since the dynamics of latent factors, $\boldsymbol{v}_t$, determines the time variation in factor risk premia, we first investigate whether the five largest PCs of asset returns can be predicted by their one-period lags and other external economic variables. Following past literature (e.g., Campbell and Vuolteenaho (2004), Campbell et al. (2013), and Gagliardini et al. (2016)), we include as external predictors the price-earning ratio as well as term, default, and value spreads.

Table IA.XVI in the Internet Appendix shows that external predictors have limited predictive power. In Panel A, we consider only the four external predictors. Although value and term spread can predict PC1 and PC5 to a certain extent, the adjusted R^2s are very smalll or even negative in these specifications. We further include the lagged return PCs in Panel B and observe economically sizable predictability. For example, the adjusted R^2 is above 6% for PC4 at both monthly and quarterly frequencies. In contrast, all external predictors are almost inessential in these regressions. Therefore, using them to model time-varying risk premia will introduce huge estimation noise, which can lead to attenuation bias in risk premia estimates. In the main text, we focus on the VAR(1) assumption for the latent factors without those four external economic variables. We report the corresponding empirical results with external predictors in the Internet Appendix.

Using the VAR(1) formulation for the latent systematic factors, we estimate the term structure of unconditional risk premia for the same set of variables as in Table 1. Figure IA.6 shows the empirical results, in which the blue lines and shaded areas present the estimates based on the conditional models. For comparison, we also include the previous estimates (the purple lines and areas) in Table 1 based on the unconditional models. The point estimates are almost identical in both conditional and unconditional models, although we occasionally detect some minor attenuations and wider confidence intervals due to the additional parameters in the VAR system. Overall, the risk premia estimates based on the unconditional models are able to deliver consistent estimates even if the true model is time-varying.[21]

[21] In the Internet Appendix, we further include the four external variables mentioned above to model the dynamics of latent factors, $\boldsymbol{v}_t$. Figures IA.7–IA.11 confirm that the term structures of unconditional risk premia

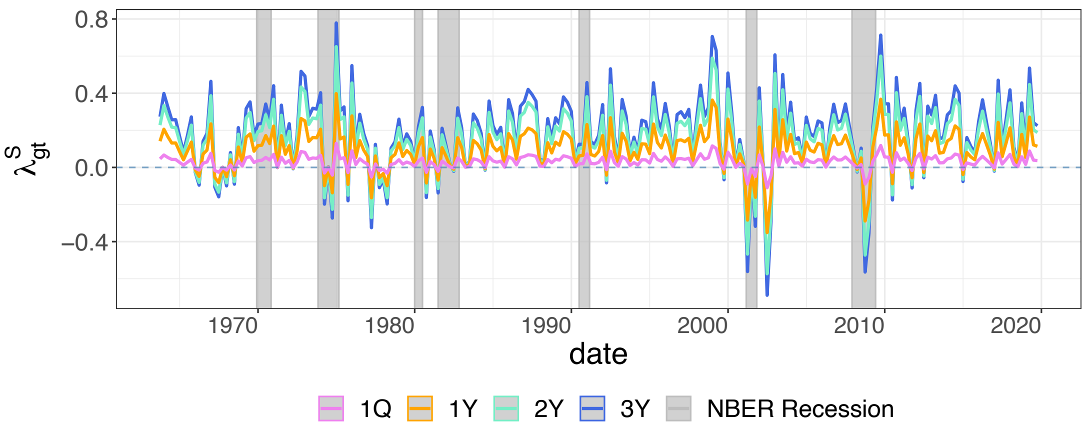

(a) Nondurable Consumption Growth

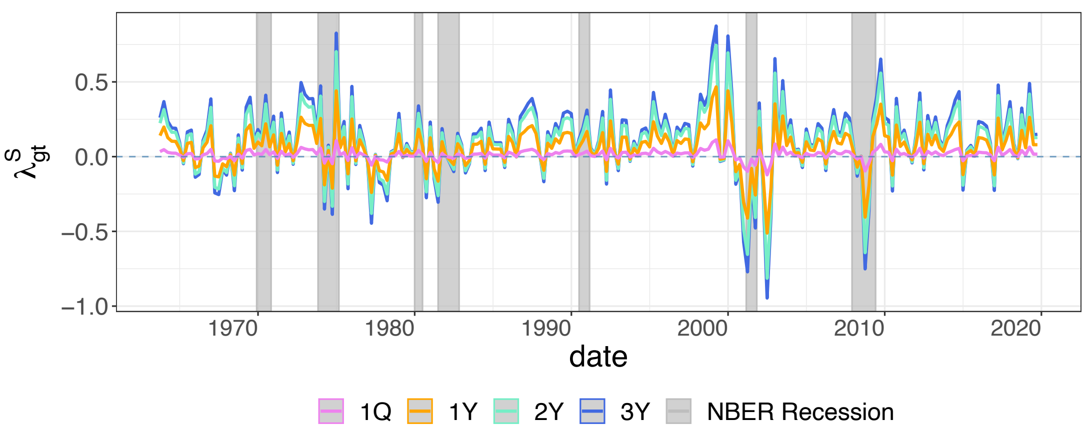

(b) GDP Growth

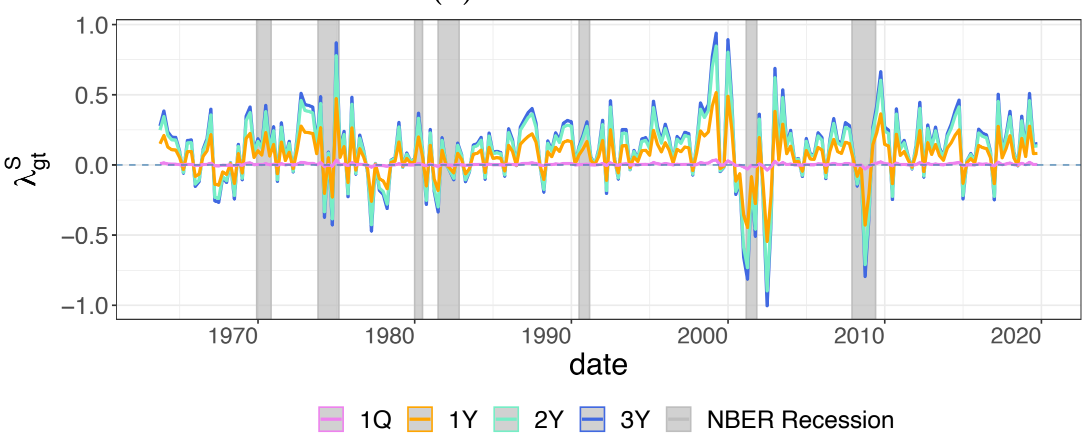

(c) Industrial Production Growth

Figure 7: Time-varying term structure of macroeconomic factor's risk premia

This figure plots the time-varying term structure of risk premia following the method in Section 2.2 and a VAR(1) for the latent systematic risk factors. Estimates are based on the composite cross-section of 275 Fama-French characteristic-sorted portfolios. Definition and data sources of factors and test assets can be found in Internet Appendix IA.3.

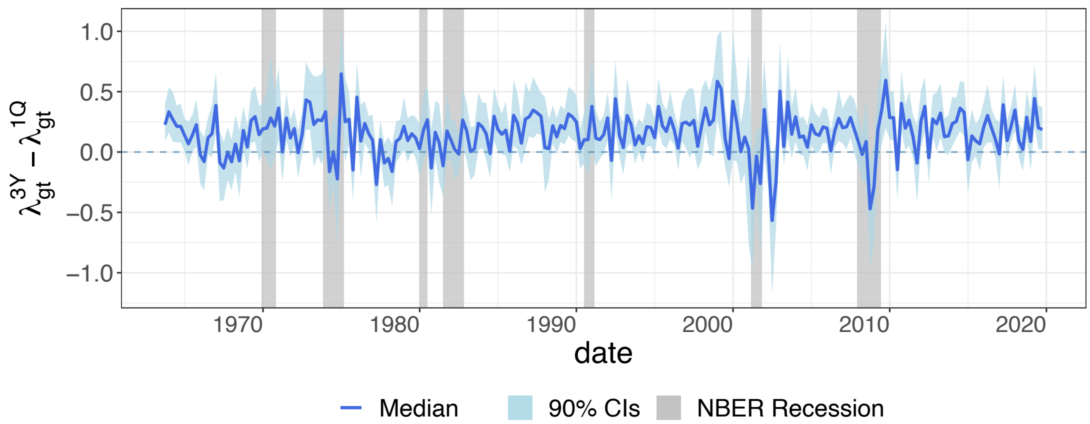

(a) Nondurable Consumption Growth

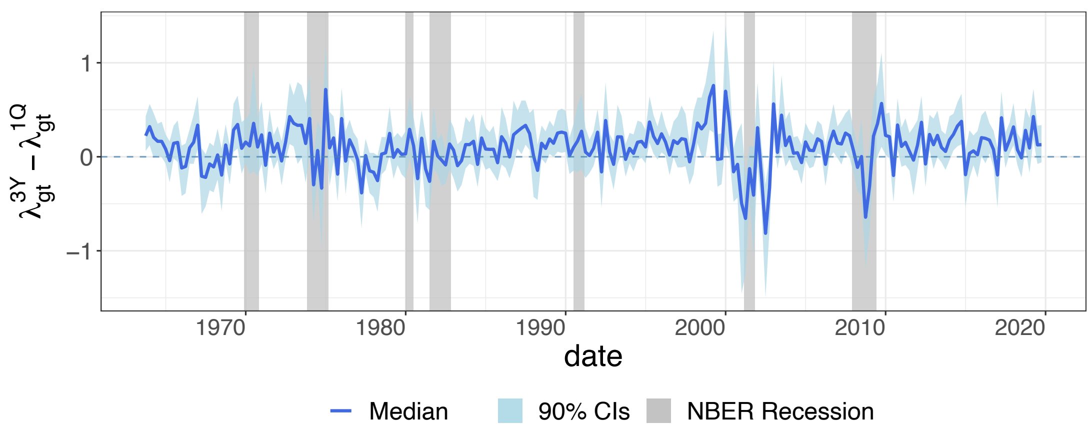

(b) GDP growth

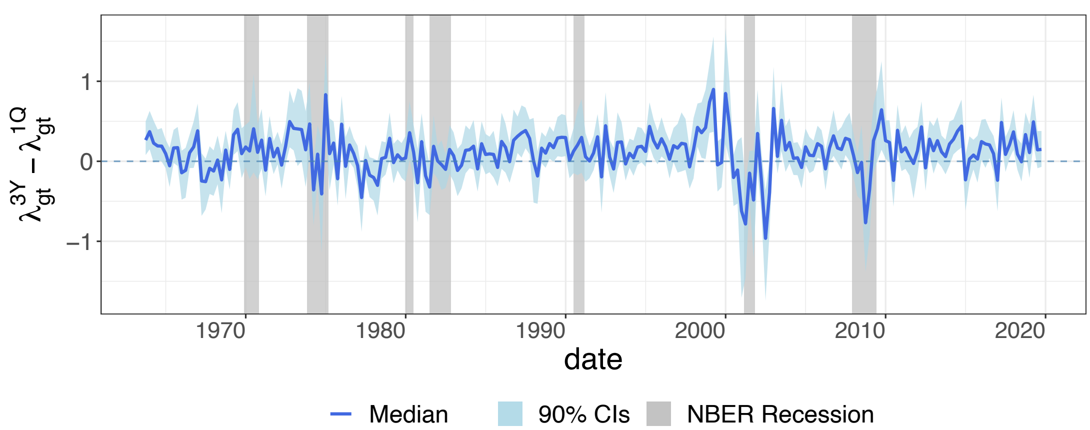

(c) Industrial Production Growth

Figure 8: Time-varying slope of the term structure of macroeconomic factor's risk premia

This figure plots the time-varying slope of the term structure of risk premia following the method in Section 2.2 and a VAR(1) for the latent systematic risk factors. The slope is defined as the difference between three-year and one-quarter conditional risk premia estimates. Blue solid lines show the posterior medians of the slopes, with the 90% posterior credible intervals denoted by the light blue shaded areas. Estimates are based on the composite cross-section of 275 Fama-French characteristic-sorted portfolios. Definition and data sources of factors and test assets can be found in the Internet Appendix IA.3.

Having established the robustness of the unconditional risk premia estimates, we proceed to explore the time-varying term structure of macro risk premia. Figure 7 reports the posterior means of the risk premia at one-quarter to three-year horizons for nondurable consumption, GDP, and industrial production growths (in Panels (a)–(c), respectively). The figure highlights a clear commonality in the business cycle behavior of the term structures of macroeconomic risk premia. Three observations are noteworthy.

First, the average level is strongly procyclical, with larger risk premia during expansion and significantly reduced, or even negative, during recession episodes. Second, as Figure 8 demonstrates, the term structure's slopes is also strongly procyclical, with the longer maturities characterized by substantially larger time variation over the business cycle and extremely high conditional Sharpe ratios immediately before recessions and at times of market crash episodes.

Interestingly, the pattern of inversion of the term structure of macro risks over the business cycle and their unconditional positive slopes mirror the findings of Bansal et al. (2021) for dividend strips with a two-state regime switching model estimated over a much smaller sample (2004:12–2017:01), as well as those obtained in Giglio et al. (2023) with an affine model for equity prices, dividends, and returns. These constitute two important external validations for our method and findings. Although our paper does not directly study the pricing and term structure of dividend strips, we can examine the risk premia of real dividend growth of the S&P500 index. As Panel (b) of Figures IA.12–IA.13 show, the slope of the term structure of dividend growth turns from positive to negative during the 2008 global financial crisis, consistent with the regime shift documented in the previous literature.

Third, short maturity (e.g., one-quarter) macro risk premia exhibit very small time variation, confirming that macroeconomic variables are weak factors at best at short horizons, even conditionally.

4.3 Segmented Risk Premia in Equity and Corporate Bond Markets

Equity and corporate bond securities are contingent claims on the same firms; hence, they should be jointly priced in the cross-section. (See Merton (1974)) If these two markets are fully integrated, we should observe that the same sources of risk carry identical compensations.

estimates are robust to different external variables.

However, financial frictions (e.g., slow-moving capital and rigid investment mandates) can lead to segmented risk premia of the same fundamental risk. In this subsection, we aim to answer the following question: Does the same factor carry significantly distinct risk premia in equity and corporate bond markets?

To address this question, we study the cross-section of 40 corporate bond portfolios in EJN40. The data sample of asset returns in this subsection ranges from February 1977 (the start date of corporate bond portfolios sorted by reversal) to December 2019. For quarterly variables, the sample starts from Q2 1977. More details about the corporate bond portfolios and their sources can be found in Internet Appendix IA.3. One caveat is that the point estimates of risk premia estimates in FF275 in this section are slightly different from those in Section 4.1 due to the smaller data sample.

Using the selection procedure of Giglio and Xiu (2021), we estimate that the number of latent factors is three in EJN40. We take a more conservative approach by studying five-, six-, and seven-factor models for both FF275 and EJN40. How (dis)integrated are equity and corporate bond markets? We estimate the generalized correlations among the largest five PCs of FF275 and EJN40 at the monthly frequency, finding that they are around 0.66, 0.30, 0.22, 0.09, and 0.06.[22] The low correlation coefficients (considerably smaller than ones) indicate that equity and corporate bond markets are mostly driven by different sources of systematic factors and, hence, disintegrated to a large extent.

Low correlations between equity and corporate bond systematic factors are not equivalent to segmented risk premia for every factor. Therefore, we formally test whether the same factor carries the same risk premium in these two markets and report in Table 4 and Figure 9[23] the estimates of risk premia using Proposition 3. The estimation is based on five-factor models for each of these two markets. For each factor, we report three quantities: the risk premia in FF275 and EJN40, as well as their difference.

We detect significantly segmented risk premia for several factors. For example, macro PC2 has positive yet insignificant risk premia in FF275, but it is priced in corporate bond portfolios.

[22]Suppose that $\hat{\boldsymbol{v}}_t^1$ and $\hat{\boldsymbol{v}}_t^2$ are the top five PCs of FF275 and EJN40. The generalized correlations between $\hat{\boldsymbol{v}}_t^1$ and $\hat{\boldsymbol{v}}_t^2$ are defined as the squared root of the eigenvalues of $\mathrm{cov}(\hat{\boldsymbol{v}}_t^1, \hat{\boldsymbol{v}}_t^2)^\top \mathrm{cov}(\hat{\boldsymbol{v}}_t^1)^{-1} \mathrm{cov}(\hat{\boldsymbol{v}}_t^1, \hat{\boldsymbol{v}}_t^2) \mathrm{cov}(\hat{\boldsymbol{v}}_t^2)^{-1}$. At the quarterly frequency, the generalized correlation coefficients are 0.78, 0.45, 0.29, 0.17, and 0.04.

[23]Plots for factors omitted from this figure can be found in Figures IA.14–IA.15 of the Internet Appendix.

As Table 4 demonstrates, its risk premia are significantly positive and upward-sloping in EJN40. The TED spread change, a proxy for funding liquidity shocks, has a flat term structure of significantly negative risk premia in EJN40, but it is not priced in FF275, which implies that funding liquidity is more salient in corporate bond markets.

The risk premia estimates of some factors even have opposite signs. The oil price change is one such example. While equity markets hedge against positive oil price changes, corporate bonds are not effective hedging tools since they tend to perform poorly as oil prices increase. Another example is macro PC3, which commands positive (negative) risk premia in FF275 (EJN40). Although the risk premia estimates are insignificant in both markets, the hypothesis of homogeneous risk premium is strongly rejected.

Interestingly, many factors carry similar term structures of risk premia in these two markets, both economically and statistically. Examples include GDP growth, IP growth, durable and nondurable consumption growth, dividend growth, macro PC1 and PC4, the HKM intermediary factors, and the VIX changes.[24] Remarkably, even though VIX is constructed using equity information (the risk-neutral variance of the equity index), it commands comparable risk premia in corporate bond markets.[25]

Table IA.XVII of the Internet Appendix repeats the same analysis using $\bar{S} = 0$. We emphasize three key findings. First, most macro variables (e.g., durable, nondurable, service consumption, and IP and GDP growth) do not carry significant risk premia in this table. This observation is consistent with our previous finding in the single market. Second, risk premia estimates based on contemporaneous correlations are sometimes counterintuitive. For instance, macro PC1, estimated using five-factor models, commands a significantly negative risk premium in EJN40. However, macro PC1 is countercyclical; hence, we would expect a positive risk premium, as seen in Table 4. Similarly, risk premia estimates of other countercyclical macro variables are mostly negative (yet insignificant) in corporate bond markets. Finally, fac-

[24]Service consumption growth has an upward-sloping term structure of significantly positive risk premia in FF275. This finding contradicts the previous observation based on the full sample (Q3 1963–Q4 2019), where the risk premium of service consumption growth is negligible. The inconsistency is likely reconciled by the fact that the relative importance of service industries has been increasing in the later subsample.

[25]We perform robustness checks by studying six- and seven-factor models and report the related results in Tables IA.XIX–IA.XX of the Internet Appendix. Overall, the point estimates of risk premia are similar to the results based on five-factor models, although with slightly wider credible intervals.

Table 4: Segmented markets: Difference of risk premia in equity and corporate bonds

		Panel A. Quarterly variables, $\bar{S} = 12$ quarters							
$S =$		0	2	4	6	8	10	12	R_g^2
AEM intermediary	FF275	0.046	0.053	0.054	0.046	0.034	0.023	0.020	16.9%
	EJN40	0.000	0.001	0.001	0.001	0.001	0.000	0.001	
	Diff	0.044	0.052	0.053	0.044	0.032	0.021	0.018	
Capital share growth	FF275	0.003	0.001	0.001	0.001	0.001	0.001	0.001	20.9%
	EJN40	0.003	0.001	0.001	0.001	0.000	0.000	0.000	
	Diff	0.000	0.000	0.000	0.000	0.000	0.001	0.001	
GDP growth	FF275	-0.022	0.003	0.036	0.061	0.079	0.094	0.105	40.8%
	EJN40	-0.033	0.006	0.052	0.086*	0.108**	0.126**	0.141**	
	Diff	0.005	0.000	-0.006	-0.013	-0.018	-0.023	-0.027	
IP growth	FF275	0.000	0.047	0.080	0.104	0.119	0.129	0.137	46.9%
	EJN40	0.000	0.068**	0.117**	0.152**	0.173**	0.186**	0.197**	
	Diff	0.000	-0.020	-0.037	-0.048	-0.055	-0.060	-0.063	
Durable consumption growth	FF275	-0.059**	0.030	0.070**	0.095***	0.118***	0.135***	0.147***	36.3%
	EJN40	-0.041*	0.021	0.048**	0.066**	0.082**	0.094**	0.103**	
	Diff	-0.015	0.007	0.019	0.026	0.033	0.039	0.042	
Nondurable consumption growth	FF275	0.013	0.056	0.098*	0.143**	0.175**	0.203**	0.231**	29.4%
	EJN40	0.007	0.034*	0.059**	0.086**	0.105**	0.121**	0.139**	
	Diff	0.003	0.019	0.037	0.055	0.070	0.082	0.093	
Service consumption growth	FF275	0.028	0.103*	0.149*	0.182*	0.211**	0.248**	0.284**	29.1%
	EJN40	0.007	0.035	0.052	0.064	0.075	0.088	0.101	
	Diff	0.016	0.061	0.090	0.110	0.127	0.149	0.172	
Nondurable + service	FF275	0.030	0.106**	0.166***	0.218***	0.260***	0.301***	0.344***	36.6%
	EJN40	0.013	0.050*	0.080**	0.106**	0.125**	0.146**	0.165**	
	Diff	0.013	0.050	0.080	0.108	0.130	0.152	0.174	
Labor income growth	FF275	-0.023	-0.006	0.000	0.007	0.015	0.021	0.028	20.2%
	EJN40	-0.003	-0.001	0.000	0.001	0.002	0.003	0.004	
	Diff	-0.016	-0.004	0.000	0.005	0.011	0.016	0.024	
Dividend growth of S&P500	FF275	0.000	0.001	0.008	0.027	0.052	0.078	0.104	61.5%
	EJN40	0.000	0.001	0.010	0.032	0.062	0.094	0.124	
	Diff	0.000	0.001	0.000	-0.001	-0.007	-0.012	-0.018	
Macro PC1 (FRED-QD)	FF275	0.023	0.093*	0.159*	0.217*	0.263*	0.300*	0.332*	64.6%
	EJN40	0.021	0.085**	0.148**	0.202**	0.245**	0.280**	0.310**	
	Diff	0.001	0.005	0.010	0.013	0.016	0.018	0.020	
Macro PC2 (FRED-QD)	FF275	0.009	0.026	0.032	0.036	0.037	0.038	0.039	70.1%
	EJN40	0.075***	0.204***	0.256***	0.278***	0.289***	0.298***	0.301***	
	Diff	-0.065**	-0.175**	-0.221**	-0.242**	-0.252**	-0.260**	-0.263**	
Macro PC3 (FRED-QD)	FF275	0.000	-0.030	-0.064	-0.101	-0.134	-0.164	-0.191	43.2%
	EJN40	0.001	0.021	0.042	0.061	0.079	0.094	0.107	
	Diff	-0.002	-0.059	-0.115	-0.165*	-0.211*	-0.253*	-0.288*	
Macro PC4 (FRED-QD)	FF275	-0.091**	-0.104**	-0.135**	-0.175**	-0.211**	-0.246**	-0.280**	50.2%
	EJN40	-0.080**	-0.093**	-0.122**	-0.156**	-0.190**	-0.224**	-0.253**	
	Diff	-0.008	-0.009	-0.012	-0.016	-0.021	-0.024	-0.028	
Macro PC5 (FRED-QD)	FF275	0.103**	0.132**	0.128*	0.116*	0.105	0.097	0.087	49.6%
	EJN40	0.079*	0.100*	0.095*	0.087	0.077	0.070	0.063	
	Diff	0.023	0.029	0.027	0.024	0.019	0.015	0.012	
		Panel B. Monthly variables, $\bar{S} = 24$ months							
$S =$		0	4	8	12	16	20	24	R_g^2
Oil price change	FF275	-0.020	-0.044*	-0.048*	-0.049*	-0.052*	-0.056*	-0.058*	17.8%
	EJN40	0.021	0.046	0.049	0.049	0.053	0.056	0.059	
	Diff	-0.042**	-0.090***	-0.096***	-0.099***	-0.105***	-0.111***	-0.117***	
TED spread change	FF275	0.013	0.014	0.014	0.014	0.015	0.015	0.015	20.7%
	EJN40	-0.099**	-0.103**	-0.108**	-0.108**	-0.111**	-0.113**	-0.117**	
	Diff	0.114**	0.117**	0.123**	0.124**	0.126**	0.128**	0.133**	
Nontraded HKM intermediary	FF275	0.095***	0.098***	0.096***	0.093***	0.090***	0.087***	0.086***	62.6%
	EJN40	0.081***	0.084***	0.081***	0.079***	0.077***	0.075***	0.074***	
	Diff	0.013	0.013	0.013	0.013	0.013	0.012	0.012	
Traded HKM intermediary	FF275	0.111***	0.115***	0.111***	0.106***	0.102***	0.099***	0.096***	72.0%
	EJN40	0.087***	0.091***	0.087***	0.083***	0.080***	0.078***	0.076***	
	Diff	0.024	0.025	0.023	0.022	0.022	0.021	0.021	
PS liquidity	FF275	0.022	0.027	0.034	0.038	0.043	0.047	0.049	10.7%
	EJN40	0.015	0.018	0.023	0.025	0.029	0.032	0.033	
	Diff	0.006	0.007	0.009	0.011	0.013	0.014	0.015	
$\Delta\log(\text{VIX})$	FF275	-0.110***	-0.065***	-0.051***	-0.039***	-0.033***	-0.028***	-0.025**	52.4%
	EJN40	-0.057*	-0.035*	-0.027*	-0.021*	-0.017*	-0.015*	-0.013*	
	Diff	-0.051	-0.031	-0.024	-0.018	-0.015	-0.013	-0.011	

The table reports Bayesian estimates of segmented risk premia using Proposition 3, where the risk premia over S horizons (λ_g^S) is defined in equation (28). The first cross-section of test assets consists of 275 Fama-French characteristic-sorted portfolios (FF275), whereas the second cross-section contains 40 corporate bond portfolios (EJN40) in Elkamhi et al. (2023). We also report the risk premia differences in these two asset markets. We consider a five-factor model for both equity and corporate bond returns.

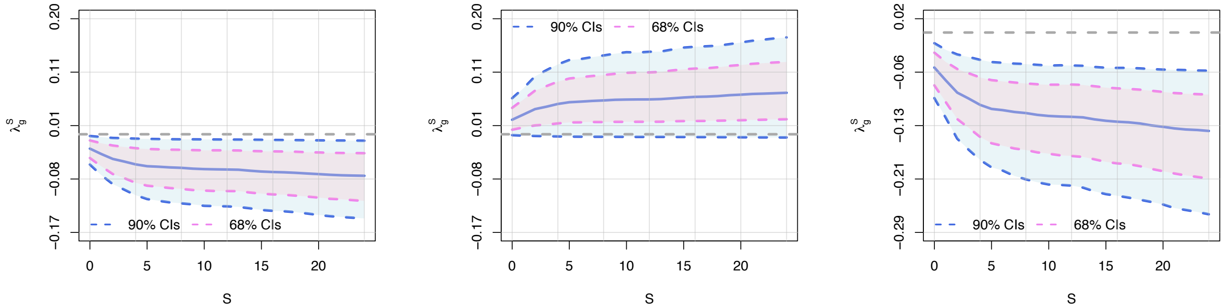

(a) Oil price: FF275 **(b)** Oil price: EJN40 **(c)** Oil price: Difference

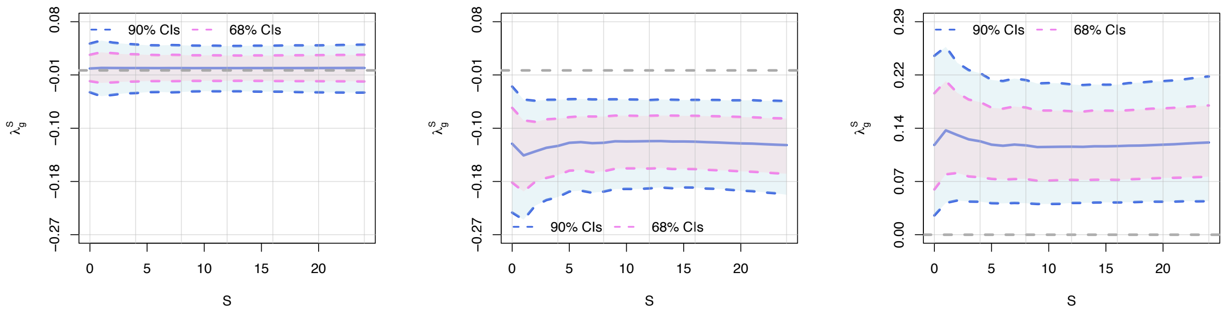

(d) TED_{ar1}: FF275 **(e)** TED_{ar1}: EJN40 **(f)** TED_{ar1}: Difference

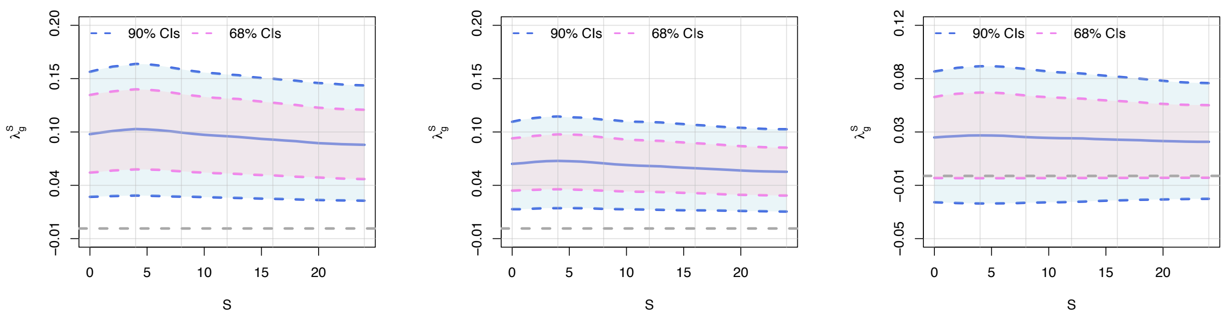

(g) HKM_{ntr}: FF275 **(h)** HKM_{ntr}: EJN40 **(i)** HKM_{ntr}: Difference

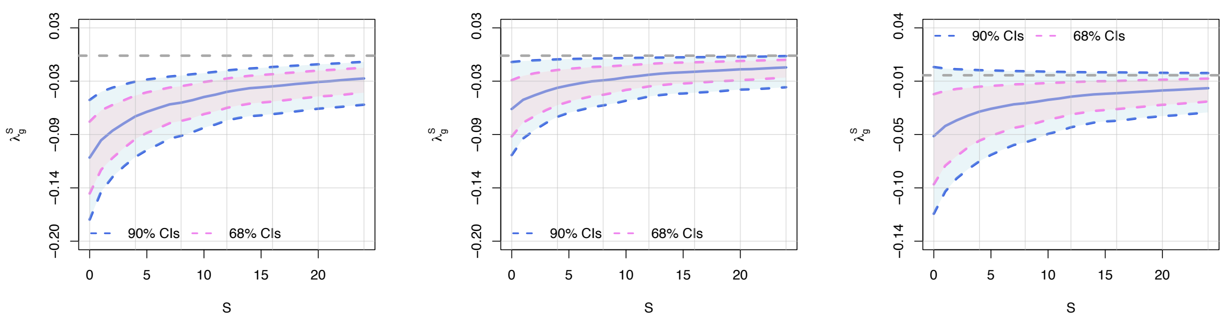

(j) $\Delta\log(\text{VIX})$: FF275 **(k)** $\Delta\log(\text{VIX})$: EJN40 **(l)** $\Delta\log(\text{VIX})$: Difference

Figure 9: Term structure of quarterly factor's risk premia in two markets: Five factors

This figure plots the term structure of risk premia estimates using Proposition 3, where the risk premia over S horizons (λ_g^S) is defined in equation (28). The first cross-section of test assets consists of 275 Fama-French characteristic-sorted portfolios (FF275), whereas the second cross-section contains 40 corporate bond portfolios (EJN40) in Elkamhi et al. (2023). We also report the risk premia differences in these two asset markets. We consider five-factor models for asset returns. In addition to the point estimates, we show the 68% and 90% Bayesian credible intervals, highlighted in pink and blue, respectively. Definition and data sources of factors and test assets can be found in Internet Appendix IA.3.

tors close to being serially uncorrelated seem to be estimated more precisely using $\bar{S}=0$, such as AEM/HKM intermediary factors, TED spread changes, and the nontraded liquidity factor. In particular, AEM and nontraded liquidity factors are priced in FF275 using $\bar{S}=0$ but not in Table 4. We also detect significantly segmented risk premia for the AEM intermediary factor in Table IA.XVII. It is apparent that the time series fits of AEM and TED spread changes, quantified by R_g^2, do not significantly improve after including lagged return shocks. A rule of thumb is that one can check whether the lagged asset pricing shocks $\{\boldsymbol{v}_{t-s}\}_{s=1}^{\bar{S}}$ can significantly enhance the time series fit R_g^2.

In summary, we detect some market segmentation between equity and corporate bond markets: These two markets are driven by considerably different systematic latent factors according to the low generalized correlations, and several factors indeed carry significantly different risk premia across asset classes. However, many factors have remarkably similar risk compensations, implying that they play similar roles in the marginal utility functions of investors in these two asset classes.

5 Conclusion

We propose a novel estimator of factors' risk premia, their term structure, and their time variation in a large cross-section of asset returns. The asset returns follow an approximate factor structure, whereas the tested factor can slowly adjust to the asset return systematic shocks, motivated by the Wold decomposition. The latter assumption allows the tested factors and asset returns to have rich dynamics but poses a challenge for the frequentist estimation. We tackle this challenge by taking a Bayesian perspective. Specifically, we derive a Gibbs sampler in which all conditional distributions of model parameters have standard closed forms, so our Bayesian estimator is straightforward to implement.

Our Bayesian framework has the frequentist three-pass procedure in Giglio and Xiu (2021) as a particular, unconditional and single-period, case. More precisely, we adopt their rotation invariance property but also show that both the conditional and unconditional term structures of risk premia of observable variables are invariant to arbitrary rotation of the latent factors. We show that risk premia of an economic state variable over multiple periods can be interpreted as the per-period Sharpe ratios of the mimicking portfolios that hedge against its multi-horizon innovations. In addition to the term structure and its time variation, we extend our framework to test for segmented risk premia across asset classes. Since the canonical paradigms (e.g., the two-step Fama-MacBeth estimator) can suffer from extreme bias due to weak factors, measurement errors, and omitted factors, an appropriate estimator and test should adequately account for these issues so that it does not over-reject the null hypothesis of homogeneous risk premia – and that is exactly what we deliver.

We first apply our method to a large equity cross-section. Our results suggest that, unconditionally, most macro variables have significantly upward-sloping term structures of risk premia. Although they are almost unpriced at quarterly horizons, their risk premia, measured over two- to three-year holding horizons, are comparable to many tradable anomalies in equity markets. In other words, macro risk strikes back at business cycle frequencies.

Meanwhile, we observe flat or downward-sloping unconditional term structures for other factors, such as VIX and intermediary factors. We go on to investigate the heterogeneous risk premia in equity and corporate bond markets. While we detect significant risk premia heterogeneity for, for example, oil price inflation, TED spread shocks, and two macro PCs in FRED-QD, most macro observables and intermediary factors command very similar risk premia in these two markets.

Furthermore, conditionally, the term structure of macro risk premia has a clear business cycle pattern: It is upward sloping during expansions and inverted during most recessions and at times of market crashes. In addition, average risk premia are strongly procyclical and particularly high at the onset of economic contractions.

Theoretical asset pricing models predict which economic state variables should be priced in the cross-section of asset returns. Given the rich set of new empirical facts that we uncover, we argue that when researchers evaluate their models, they should consider the heterogeneous factor risk premia across horizons, asset classes, and states of the business cycle.

References

Adrian, T., E. Etula, and T. Muir (2014). Financial intermediaries and the cross-section of asset returns. *Journal of Finance 69*(6), 2557–2596.

Ai, H., M. M. Croce, A. M. Diercks, and K. Li (2018). News shocks and the production-based term structure of equity returns. *Review of Financial Studies 31*(7), 2423–2467.

Ang, A., M. Piazzesi, and M. Wei (2006). What does the yield curve tell us about GDP growth? *Journal of Econometrics 131*(1-2), 359–403.

Asness, C. S., T. J. Moskowitz, and L. H. Pedersen (2013). Value and momentum everywhere. *Journal of Finance 68*(3), 929–985.

Avramov, D., S. Cheng, L. Metzker, and S. Voigt (2023). Integrating factor models. *Journal of Finance 78*(3), 1593–1646.

Bai, J. and S. Ng (2002). Determining the number of factors in approximate factor models. *Econometrica 70*(1), 191–221.

Bakshi, G., M. Cerrato, and J. Crosby (2018). Implications of incomplete markets for international economies. *Review of Financial Studies 31*(10), 4017–4062.

Bandi, F. M. and A. Tamoni (2023). Business-cycle consumption risk and asset prices. *Journal of Econometrics*.

Bansal, R., D. Kiku, and A. Yaron (2012). An empirical evaluation of the long-run risks model for asset prices. *Critical Finance Review 1*(1), 183–221.

Bansal, R., S. Miller, D. Song, and A. Yaron (2021). The term structure of equity risk premia. *Journal of Financial Economics 142*(3), 1209–1228.

Bansal, R. and A. Yaron (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance 59*(4), 1481–1509.

Barillas, F. and J. Shanken (2018). Comparing asset pricing models. *Journal of Finance 73*(2), 715–754.

Belo, F., P. Collin-Dufresne, and R. S. Goldstein (2015). Dividend dynamics and the term structure of dividend strips. *Journal of Finance 70*(3), 1115–1160.

Brandt, M. W., J. H. Cochrane, and P. Santa-Clara (2006). International risk sharing is better than you think, or exchange rates are too smooth. *Journal of Monetary Economics 53*(4), 671–698.

Brunnermeier, M. K. and L. H. Pedersen (2009). Market liquidity and funding liquidity. *Review of Financial Studies 22*(6), 2201–2238.

Bryzgalova, S. (2015). Spurious factors in linear asset pricing models. *LSE manuscript 1*(3).

Bryzgalova, S., J. Huang, and C. Julliard (2023a). Bayesian solutions for the factor zoo: We just ran two quadrillion models. *Journal of Finance 78*(1), 487–557.

Bryzgalova, S., J. Huang, and C. Julliard (2023b). Consumption in asset returns. *Available at SSRN 3783070*.

Campbell, J. Y. and J. H. Cochrane (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy 107*(2), 205–251.

Campbell, J. Y., S. Giglio, and C. Polk (2013). Hard times. *Review of Asset Pricing Studies 3*(1), 95–132.

Campbell, J. Y. and R. J. Shiller (1988). The dividend-price ratio and expectations of future dividends and discount factors. *Review of Financial Studies 1*(3), 195–228.

Campbell, J. Y. and T. Vuolteenaho (2004). Bad beta, good beta. *American Economic Review 94*(5), 1249–1275.

Carhart, M. M. (1997). On persistence in mutual fund performance. *Journal of Finance 52*(1), 57–82.

Chaieb, I., H. Langlois, and O. Scaillet (2021). Factors and risk premia in individual international stock returns. *Journal of Financial Economics 141*(2), 669–692.

Chamberlain, G. and M. Rothschild (1983). Arbitrage, factor structure, and mean-variance analysis on large asset markets. *Econometrica: Journal of the Econometric Society 51*(5), 1281–1304.

Chen, Z. and P. J. Knez (1995). Measurement of market integration and arbitrage. *Review of Financial Studies 8*(2), 287–325.

Chernov, M., L. A. Lochstoer, and S. R. H. Lundeby (2021). Conditional dynamics and the multihorizon risk-return trade-off. *The Review of Financial Studies 35*(3), 1310–1347.

Chib, S., X. Zeng, and L. Zhao (2020). On comparing asset pricing models. *Journal of Finance 75*(1), 551–577.

Choi, J. and Y. Kim (2018). Anomalies and market (dis)integration. *Journal of Monetary Economics 100*, 16–34.

Cochrane, J. (2009). *Asset Pricing: Revised Edition*. Princeton University Press.

Cohen, R. B., C. Polk, and T. Vuolteenaho (2009). The price is (almost) right. *Journal of Finance 64*(6), 2739–2782.

Connor, G. and R. A. Korajczyk (1986). Performance measurement with the arbitrage pricing theory: A new framework for analysis. *Journal of Financial Economics 15*(3), 373–394.

Connor, G. and R. A. Korajczyk (1988). Risk and return in an equilibrium APT: Application of a new test methodology. *Journal of Financial Economics 21*(2), 255–289.

Daniel, K. and D. Marshall (1997). Equity-premium and risk-free-rate puzzles at long horizons. *Macroeconomic Dynamics 1*(2), 452–484.

Dew-Becker, I. and S. Giglio (2016). Asset pricing in the frequency domain: Theory and empirics. *Review of Financial Studies 29*(8), 2029–2068.

Dew-Becker, I., S. Giglio, A. Le, and M. Rodriguez (2017). The price of variance risk. *Journal of Financial Economics 123*(2), 225–250.

Duffie, D. (2010). Presidential address: Asset price dynamics with slow-moving capital. *Journal of Finance 65*(4), 1237–1267.

Elkamhi, R., C. Jo, and Y. Nozawa (2023). A one-factor model of corporate bond premia. *Management Science*.

Eraker, B. and Y. Wu (2014). Explaining the negative returns to VIX futures and ETNs: An equilibrium approach. *Available at SSRN 2340070*.

Fama, E. F. and K. R. French (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics 33*(1), 3–56.

Gagliardini, P., E. Ossola, and O. Scaillet (2016). Time-varying risk premium in large cross-sectional equity data sets. *Econometrica 84*(3), 985–1046.

Garleanu, N. and L. H. Pedersen (2011). Margin-based asset pricing and deviations from the law of one price. *Review of Financial Studies 24*(6), 1980–2022.

Giglio, S., B. T. Kelly, and S. Kozak (2023). Equity term structures without dividend strips data. *Journal of Finance (forthcoming)*.

Giglio, S. and D. Xiu (2021). Asset pricing with omitted factors. *Journal of Political Economy 129*(7), 1947–1990.

Giglio, S., D. Xiu, and D. Zhang (2023). Test assets and weak factors. NBER Working Paper 29002.

Gospodinov, N., R. Kan, and C. Robotti (2014). Misspecification-robust inference in linear asset-pricing models with irrelevant risk factors. *Review of Financial Studies 27*(7), 2139–2170.

Gospodinov, N., R. Kan, and C. Robotti (2019). Too good to be true? Fallacies in evaluating risk factor models. *Journal of Financial Economics 132*(2), 451–471.

Greenwood, R., S. G. Hanson, and G. Y. Liao (2018). Asset price dynamics in partially segmented markets. *Review of Financial Studies 31*(9), 3307–3343.

Gromb, D. and D. Vayanos (2002). Equilibrium and welfare in markets with financially constrained arbitrageurs. *Journal of Financial Economics 66*(2-3), 361–407.

Gromb, D. and D. Vayanos (2018). The dynamics of financially constrained arbitrage. *Journal of Finance 73*(4), 1713–1750.

Hansen, L. and R. Jagannathan (1991). Implications of security market data for models of dynamic economies. *Journal of Political Economy 99*(2), 225–262.

Hansen, L. P., J. C. Heaton, and N. Li (2008). Consumption strikes back? Measuring long-run risk. *Journal of Political Economy 116*(2), 260–302.

Hasler, M. and R. Marfe (2016). Disaster recovery and the term structure of dividend strips. *Journal of Financial Economics 122*(1), 116–134.

He, Z., B. Kelly, and A. Manela (2017). Intermediary asset pricing: New evidence from many asset classes. *Journal of Financial Economics 126*(1), 1–35.

Jagannathan, R. and Y. Wang (2007). Lazy investors, discretionary consumption, and the cross-section of stock returns. *Journal of Finance 62*(4), 1623–1661.

Johnson, T. L. (2017). Risk premia and the VIX term structure. *Journal of Financial and Quantitative Analysis 52*(6), 2461–2490.

Kan, R., C. Robotti, and J. Shanken (2013). Pricing model performance and the two-pass cross-sectional regression methodology. *Journal of Finance 68*(6), 2617–2649.

Kan, R. and C. Zhang (1999a). GMM tests of stochastic discount factor models with useless factors. *Journal of Financial Economics 54*(1), 103–127.

Kan, R. and C. Zhang (1999b). Two-pass tests of asset pricing models with useless factors. *Journal of Finance 54*(1), 203–235.

Kapadia, N. and X. Pu (2012). Limited arbitrage between equity and credit markets. *Journal of Financial Economics 105*(3), 542–564.

Kelly, B., D. Palhares, and S. Pruitt (2023). Modeling corporate bond returns. *Journal of Finance 78*(4), 1967–2008.

Kelly, B. T., S. Pruitt, and Y. Su (2019). Characteristics are covariances: A unified model of risk and return. *Journal of Financial Economics 134*(3), 501–524.

Kleibergen, F. (2009). Tests of risk premia in linear factor models. *Journal of Econometrics 149*(2), 149–173.

Kleibergen, F. and Z. Zhan (2015). Unexplained factors and their effects on second pass R-squared's. *Journal of Econometrics 189*(1), 101–116.

Kleibergen, F. and Z. Zhan (2020). Robust inference for consumption-based asset pricing. *Journal of Finance 75*(1), 507–550.

Kozak, S., S. Nagel, and S. Santosh (2018). Interpreting factor models. *Journal of Finance 73*(3), 1183–1223.

Kozak, S., S. Nagel, and S. Santosh (2020). Shrinking the cross-section. *Journal of Financial Economics 135*(2), 271–292.

Kragt, J., F. De Jong, and J. Driessen (2020). The dividend term structure. *Journal of Financial and Quantitative Analysis 55*(3), 829–867.

Lettau, M. and S. Ludvigson (2001). Consumption, aggregate wealth, and expected stock returns. *Journal of Finance 56*(3), 815–849.

Lettau, M., S. C. Ludvigson, and S. Ma (2019). Capital share risk in us asset pricing. *Journal of Finance 74*(4), 1753–1792.

Liew, J. and M. Vassalou (2000). Can book-to-market, size and momentum be risk factors that predict economic growth? *Journal of Financial Economics 57*(2), 221–245.

Malloy, C. J., T. J. Moskowitz, and A. Vissing-Jørgensen (2009). Long-run stockholder consumption risk and asset returns. *The Journal of Finance 64*(6), 2427–2479.

McCracken, M. and S. Ng (2020). FRED-QD: A quarterly database for macroeconomic research. Technical report, National Bureau of Economic Research.

Merton, R. C. (1974). On the pricing of corporate debt: The risk structure of interest rates. *Journal of Finance 29*(2), 449–470.

Müller, U. K. (2013). Risk of Bayesian inference in misspecified models, and the sandwich covariance matrix. *Econometrica 81*(5), 1805–1849.

Newey, W. K. and K. D. West (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica 55*(3), 703–708.

Ortu, F., A. Tamoni, and C. Tebaldi (2013). Long-run risk and the persistence of consumption shocks. *Review of Financial Studies 26*(11), 2876–2915.

Parker, J. A. and C. Julliard (2005). Consumption risk and the cross section of expected returns. *Journal of Political Economy 113*(1), 185–222.

Pástor, L. and R. F. Stambaugh (2003). Liquidity risk and expected stock returns. *Journal of Political Economy 111*(3), 642–685.

Patton, A. J. and B. M. Weller (2022). Risk price variation: The missing half of empirical asset pricing. *Review of Financial Studies 35*(11), 5127–5184.

Sandulescu, M. (2022). How integrated are corporate bond and stock markets? *Available at SSRN 3528252*.

Sandulescu, M., F. Trojani, and A. Vedolin (2021). Model-free international stochastic discount factors. *Journal of Finance 76*(2), 935–976.

Schaefer, S. M. and I. A. Strebulaev (2008). Structural models of credit risk are useful: Evidence from hedge ratios on corporate bonds. *Journal of Financial Economics 90*(1), 1–19.

Shleifer, A. and R. W. Vishny (1997). The limits of arbitrage. *Journal of Finance 52*(1), 35–55.

van Binsbergen, J. H., M. Brandt, and R. Koijen (2012). On the timing and pricing of dividends. *American Economic Review 102*(4), 1596–1618.

van Binsbergen, J. H. and R. S. Koijen (2017). The term structure of returns: Facts and theory. *Journal of Financial Economics 124*(1), 1–21.

Appendices

Appendix A Additional Propositions and Proofs

Proposition A1. *As $N \to \infty$, $\mathbb{E}[\boldsymbol{r}_t]^\top \text{cov}(\boldsymbol{r}_t)^{-1}\boldsymbol{\beta}_{\tilde{v}} \to \boldsymbol{\lambda}_{\tilde{v}}^\top$ under the following assumptions:*

i. The eigenvalues of $\boldsymbol{\beta}_{\tilde{v}}^\top\boldsymbol{\beta}_{\tilde{v}}$ explode as $N \to \infty$, whereas $\boldsymbol{\Sigma}_{wr}$ has bounded eigenvalues: $\gamma_{min}(\boldsymbol{\beta}_{\tilde{v}}^\top\boldsymbol{\beta}_{\tilde{v}}) = O_p(N)$ and $\gamma_{max}(\boldsymbol{\Sigma}_{wr}) = O_p(1)$;

ii. $\boldsymbol{\beta}_{\tilde{v}}$ and $\boldsymbol{\Sigma}_{wr}$ are of full rank;

iii. Asset returns and their expectations follow equations (1) and (2).

A.1 Proof of Proposition A1

Assumptions in equation (2) imply that $\mathbb{E}[\boldsymbol{r}_t]^\top \text{cov}(\boldsymbol{r}_t)^{-1}\boldsymbol{\beta}_{\tilde{v}} = \underbrace{\boldsymbol{\alpha}^\top \text{cov}(\boldsymbol{r}_t)^{-1}\boldsymbol{\beta}_{\tilde{v}}}_{(I)} + \underbrace{\boldsymbol{\lambda}_{\tilde{v}}^\top\boldsymbol{\beta}_{\tilde{v}}^\top \text{cov}(\boldsymbol{r}_t)^{-1}\boldsymbol{\beta}_{\tilde{v}}}_{(II)}$.
Since we assume that α_i is cross-sectionally independent of $\boldsymbol{\beta}_{\tilde{v},i}$, (I) will disappear as $N \to \infty$, as follows:

$$(I) = \frac{1}{N}\boldsymbol{\alpha}^\top\Big[\frac{1}{N}\text{cov}(\boldsymbol{r}_t)\Big]^{-1}\boldsymbol{\beta}_{\tilde{v}} \to \boldsymbol{0}_K^\top.$$

Assumptions in equation (1) imply that $\text{cov}(\boldsymbol{r}_t) = \boldsymbol{\beta}_{\tilde{v}}\boldsymbol{\beta}_{\tilde{v}}^\top + \boldsymbol{\Sigma}_{wr}$. Using the Woodbury matrix identity, we can rewrite the inverse of $\text{cov}(\boldsymbol{r}_t)$ as follows:

$$\text{cov}(\boldsymbol{r}_t)^{-1} = \boldsymbol{\Sigma}_{wr}^{-1} - \boldsymbol{\Sigma}_{wr}^{-1}\boldsymbol{\beta}_{\tilde{v}}(\boldsymbol{I}_K + \boldsymbol{\beta}_{\tilde{v}}^\top\boldsymbol{\Sigma}_{wr}^{-1}\boldsymbol{\beta}_{\tilde{v}})^{-1}\boldsymbol{\beta}_{\tilde{v}}^\top\boldsymbol{\Sigma}_{wr}^{-1}.$$

$$\begin{aligned}
\implies (II) &= \boldsymbol{\lambda}_{\tilde{v}}^\top\boldsymbol{\beta}_{\tilde{v}}^\top\big[\boldsymbol{\Sigma}_{wr}^{-1} - \boldsymbol{\Sigma}_{wr}^{-1}\boldsymbol{\beta}_{\tilde{v}}(\boldsymbol{I}_K + \boldsymbol{\beta}_{\tilde{v}}^\top\boldsymbol{\Sigma}_{wr}^{-1}\boldsymbol{\beta}_{\tilde{v}})^{-1}\boldsymbol{\beta}_{\tilde{v}}^\top\boldsymbol{\Sigma}_{wr}^{-1}\big]\boldsymbol{\beta}_{\tilde{v}} \\
&= \boldsymbol{\lambda}_{\tilde{v}}^\top\big[\boldsymbol{\beta}_{\tilde{v}}^\top\boldsymbol{\Sigma}_{wr}^{-1}\boldsymbol{\beta}_{\tilde{v}} - \boldsymbol{\beta}_{\tilde{v}}^\top\boldsymbol{\Sigma}_{wr}^{-1}\boldsymbol{\beta}_{\tilde{v}}(\boldsymbol{I}_K + \boldsymbol{\beta}_{\tilde{v}}^\top\boldsymbol{\Sigma}_{wr}^{-1}\boldsymbol{\beta}_{\tilde{v}})^{-1}\boldsymbol{\beta}_{\tilde{v}}^\top\boldsymbol{\Sigma}_{wr}^{-1}\boldsymbol{\beta}_{\tilde{v}}\big] \quad (\text{let } \boldsymbol{A} = (\boldsymbol{\beta}_{\tilde{v}}^\top\boldsymbol{\Sigma}_{wr}^{-1}\boldsymbol{\beta}_{\tilde{v}})^{-1}) \\
&= \boldsymbol{\lambda}_{\tilde{v}}^\top\big[\boldsymbol{A}^{-1} - \boldsymbol{A}^{-1}(\boldsymbol{I}_K + \boldsymbol{A}^{-1})^{-1}\boldsymbol{A}^{-1}\big] = \boldsymbol{\lambda}_{\tilde{v}}^\top(\boldsymbol{A} + \boldsymbol{I}_K)^{-1}.
\end{aligned}$$

Since we assume that the eigenvalues of $\boldsymbol{\beta}_{\tilde{v}}^\top\boldsymbol{\beta}_{\tilde{v}}$ will explode as $N \to \infty$, whereas $\boldsymbol{\Sigma}_{wr}$ has bounded eigenvalues, $\boldsymbol{A} \to \boldsymbol{0}$ as $N \to \infty$. This further implies that $(II) \to \boldsymbol{\lambda}_{\tilde{v}}^\top$.

A.2 Estimating Time-Varying Risk Premia in Section 2.2

In estimation, we identify a linear rotation of $\tilde{\boldsymbol{v}}_t$: $\boldsymbol{v}_t = \boldsymbol{H}\tilde{\boldsymbol{v}}_t = \boldsymbol{H}\boldsymbol{\mu}_{\tilde{v},t-1} + \boldsymbol{H}\boldsymbol{\epsilon}_{\tilde{\boldsymbol{v}}\boldsymbol{t}} = \boldsymbol{\mu}_{v,t-1} + \boldsymbol{\epsilon}_{vt}$, which implies that $\boldsymbol{\Sigma}_{\epsilon v} = \text{cov}(\boldsymbol{\epsilon}_{vt}) = \boldsymbol{H}\boldsymbol{H}^\top$. We generalize the rotation invariance to identify the time-varying risk premia as follows:

$$\boldsymbol{r}_t = \boldsymbol{\alpha} + \underbrace{\boldsymbol{\beta}_{\tilde{v}}\boldsymbol{H}^{-1}}_{\boldsymbol{\beta}_v}\underbrace{\boldsymbol{H}\boldsymbol{\lambda}_{\tilde{v}}}_{\boldsymbol{\lambda}_v} + \underbrace{\boldsymbol{\beta}_{\tilde{v}}\boldsymbol{H}^{-1}}_{\boldsymbol{\beta}_v}\underbrace{\boldsymbol{H}\tilde{\boldsymbol{v}}_t}_{\boldsymbol{v}_t} + \boldsymbol{w}_{rt}, \quad g_t = \mu_g + \sum_{s=0}^{\bar{S}} \tilde{\rho}_s \underbrace{\tilde{\boldsymbol{\eta}}_g^\top \boldsymbol{H}^{-1}}_{\boldsymbol{\eta}_g^\top}\underbrace{\boldsymbol{H}\boldsymbol{\epsilon}_{\tilde{v},t-s}}_{\boldsymbol{\epsilon}_{v,t-s}} + w_{gt},$$

$$m_t = 1 - \boldsymbol{\lambda}_v^\top (\boldsymbol{H}^{-1})^\top \boldsymbol{H}^{-1}\boldsymbol{\epsilon}_{vt} - \boldsymbol{\mu}_{v,t-1}^\top (\boldsymbol{H}^{-1})^\top \boldsymbol{H}^{-1}\boldsymbol{\epsilon}_{vt} = 1 - \boldsymbol{\lambda}_v^\top \boldsymbol{\Sigma}_{\epsilon v}^{-1}\boldsymbol{\epsilon}_{vt} - \boldsymbol{\mu}_{v,t-1}^\top \boldsymbol{\Sigma}_{\epsilon v}^{-1}\boldsymbol{\epsilon}_{vt}, \text{ and} \tag{A1}$$

$$\lambda_{g,t-1}^S = \frac{\sum_{\tau=0}^S \sum_{s=0}^\tau \tilde{\rho}_s}{1+S} \cdot \underbrace{\tilde{\boldsymbol{\eta}}_g^\top \boldsymbol{H}^{-1}}_{\boldsymbol{\eta}_g^\top}\underbrace{\boldsymbol{H}(\boldsymbol{\lambda}_{\tilde{v}} + \boldsymbol{\mu}_{\tilde{v},t-1})}_{\boldsymbol{\lambda}_v + \boldsymbol{\mu}_{v,t-1}};$$

therefore, the time-varying risk premia, $\lambda_{g,t-1}^S$, are well-defined.

Proposition A2 (Gibbs sampler of the time-varying model)**.** *Under the assumptions in equations (19)–(23), the posterior distribution of the model parameters can be sampled from the following conditional distributions:*

(1) Conditional on the data, $\{g_t\}_{t=1+\bar{S}}^T$, and shocks to latent factors, $\{\boldsymbol{\epsilon}_{vt}\}_{t=1}^T$, the parameters of the g_t process (σ_{wg}^2, $\boldsymbol{\rho}_g$, and $\boldsymbol{\eta}_g$) follow the normal-inverse-gamma distribution in equations (IA.1)–(IA.3) of Internet Appendix IA.1.1. The only difference is that we replace $\boldsymbol{v}_t$ with $\boldsymbol{\epsilon}_{vt}$ in equations (IA.1)–(IA.3). For point identification purposes, draws of $\boldsymbol{\rho}_g$ and $\boldsymbol{\eta}_g$ are normalized such that $\boldsymbol{\eta}_g^\top \boldsymbol{\eta}_g = 1$.

(2) Conditional on asset returns, $\{\boldsymbol{r}_t\}_{t=1}^T$, and latent factors, $\{\boldsymbol{v}_t\}_{t=1}^T$, the parameters of the $\boldsymbol{r}_t$ process ($\boldsymbol{\Sigma}_{wr}$ and $\boldsymbol{B}_r^\top = (\boldsymbol{\mu}_r, \boldsymbol{\beta}_v)$) follow the normal-inverse-Wishart distribution in equations (IA.4)–(IA.5) of Internet Appendix IA.1.1.

(3) Conditional on asset returns and $(\boldsymbol{\mu}_r, \boldsymbol{\mu}_v, \boldsymbol{\beta}_v, \boldsymbol{\Sigma}_{wr})$, the latent factors, $\boldsymbol{v}_t$, can be sampled from the normal-inverse-Wishart distribution in equation (IA.6).

(4) Conditional on latent factors, $\{\boldsymbol{v}_t\}_{t=1}^T$, the model parameters in the VAR(q) system of $\boldsymbol{v}_t$ can be obtained from equations (IA.9)–(IA.10). The conditional mean of $\boldsymbol{v}_t$ equals the first K elements of $\boldsymbol{\phi}_0 + \boldsymbol{\phi}_1\boldsymbol{x}_{t-1} + \cdots + \boldsymbol{\phi}_q\boldsymbol{x}_{t-q}$, and the first K variables in $\boldsymbol{\epsilon}_{xt}$ are shocks

to priced systematic factors, $\boldsymbol{\epsilon}_{vt}$. We can also obtain the unconditional mean of $\boldsymbol{v}_t$ as the first K elements in $(\boldsymbol{I} - \boldsymbol{\phi}_1 - \cdots - \boldsymbol{\phi}_q)^{-1}\boldsymbol{\phi}_0$.

(5) Conditional on the posterior draws from the time series steps (1)–(4), the posterior distribution of $\boldsymbol{\lambda}_v$ is a Dirac distribution at $(\boldsymbol{\beta}_v^\top \boldsymbol{\beta}_v)^{-1}\boldsymbol{\beta}_v^\top \boldsymbol{\mu}_r$, yielding a Dirac conditional posterior for the term structure of g_t's risk premia at $\lambda_{g,t-1}^S = \sum_{\tau=0}^S \sum_{s=0}^{\tau} \frac{\rho_s \boldsymbol{\eta}_g^\top (\boldsymbol{\lambda}_v + \boldsymbol{\mu}_{v,t-1})}{1+S}$, where $0 \le S \le \bar{S}$.

Appendix B Additional Figures and Tables

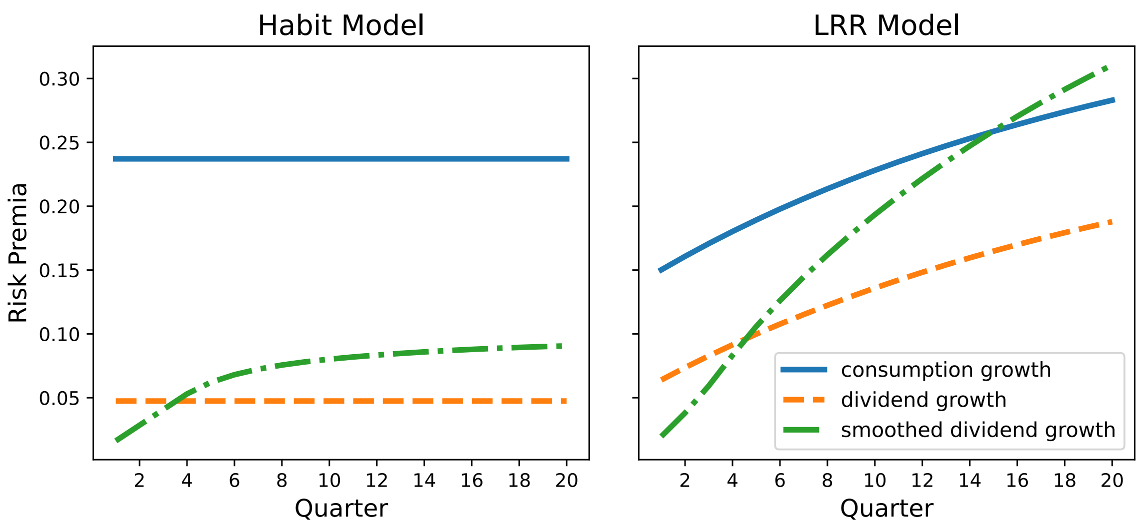

Figure A1: Term Structure of Risk Premia in Habit and Long-Run Risk Models

The figure plots the term structure of risk premia implied by the habit model of Campbell and Cochrane (1999) (left panel) and the long-run risk model (right panel) of Bansal and Yaron (2004). We consider three macro variables: (1) quarterly consumption growth, (2) quarterly dividend growth, and (3) quarterly growth in the smooth dividend payment (the aggregate dividend payments made in the previous 12 months). Risk premia are normalized by the quarterly volatility of the macro variables. Calibration and derivation details can be found in Internet Appendix IA.4.

Table A1: Testing risk premia of strong factors at quarterly frequencies ($T = 200$)

	$S=0$	1	2	3	4	5	6	7	8	9	10	11	12
							Panel A: $R_g^2 = 30\%$						
							Number of Factors = 5						
10%	0.114	0.101	0.101	0.100	0.104	0.102	0.104	0.101	0.102	0.098	0.104	0.102	0.104
5%	0.059	0.049	0.050	0.054	0.049	0.051	0.051	0.050	0.053	0.053	0.052	0.054	0.052
1%	0.011	0.011	0.009	0.010	0.010	0.011	0.012	0.010	0.010	0.008	0.009	0.009	0.009
							Number of Factors = 4						
10%	0.297	0.278	0.292	0.298	0.288	0.281	0.292	0.288	0.286	0.289	0.288	0.290	0.295
5%	0.190	0.197	0.192	0.185	0.187	0.191	0.187	0.190	0.187	0.191	0.191	0.195	0.191
1%	0.062	0.078	0.075	0.080	0.078	0.078	0.073	0.072	0.074	0.068	0.066	0.074	0.074
							Number of Factors = 7						
10%	0.110	0.092	0.089	0.089	0.092	0.098	0.096	0.097	0.094	0.091	0.092	0.098	0.099
5%	0.054	0.045	0.046	0.047	0.045	0.045	0.042	0.042	0.043	0.043	0.042	0.045	0.046
1%	0.009	0.011	0.010	0.010	0.009	0.008	0.008	0.007	0.008	0.007	0.008	0.009	0.008
							Panel B: $R_g^2 = 20\%$						
							Number of Factors = 5						
10%	0.128	0.136	0.141	0.143	0.143	0.141	0.140	0.142	0.143	0.147	0.144	0.138	0.139
5%	0.062	0.074	0.077	0.073	0.075	0.076	0.073	0.072	0.070	0.075	0.073	0.071	0.071
1%	0.007	0.016	0.015	0.017	0.017	0.017	0.014	0.017	0.020	0.016	0.017	0.014	0.014
							Number of Factors = 4						
10%	0.278	0.296	0.296	0.305	0.296	0.298	0.290	0.298	0.299	0.304	0.305	0.300	0.296
5%	0.163	0.206	0.211	0.210	0.209	0.203	0.205	0.211	0.201	0.208	0.207	0.207	0.202
1%	0.028	0.079	0.077	0.074	0.079	0.075	0.078	0.079	0.078	0.079	0.078	0.081	0.077
							Number of Factors = 7						
10%	0.117	0.132	0.135	0.130	0.130	0.130	0.133	0.131	0.135	0.138	0.140	0.136	0.136
5%	0.054	0.075	0.071	0.069	0.066	0.068	0.071	0.065	0.061	0.064	0.075	0.069	0.067
1%	0.007	0.017	0.014	0.012	0.015	0.014	0.014	0.018	0.017	0.019	0.018	0.018	0.019
							Panel C: $R_g^2 = 10\%$						
							Number of Factors = 5						
10%	0.073	0.138	0.136	0.127	0.134	0.137	0.133	0.134	0.139	0.140	0.136	0.134	0.139
5%	0.033	0.072	0.070	0.064	0.070	0.073	0.076	0.081	0.075	0.080	0.076	0.075	0.071
1%	0.003	0.006	0.014	0.012	0.013	0.012	0.012	0.010	0.013	0.013	0.012	0.009	0.013
							Number of Factors = 4						
10%	0.138	0.258	0.252	0.252	0.260	0.258	0.253	0.251	0.258	0.269	0.263	0.259	0.253
5%	0.059	0.148	0.154	0.138	0.157	0.158	0.156	0.158	0.156	0.155	0.145	0.150	0.147
1%	0.009	0.039	0.039	0.037	0.046	0.050	0.045	0.044	0.043	0.039	0.043	0.040	0.035
							Number of Factors = 7						
10%	0.071	0.145	0.137	0.125	0.134	0.138	0.140	0.139	0.136	0.141	0.129	0.140	0.144
5%	0.018	0.068	0.057	0.054	0.056	0.070	0.062	0.066	0.071	0.068	0.071	0.071	0.064
1%	0.002	0.009	0.012	0.007	0.017	0.013	0.013	0.010	0.011	0.009	0.009	0.008	0.008

The table reports the frequency of rejecting the null hypothesis $H_0 : \lambda_g^S = \lambda_g^{S,\star}$ based on the 90%, 95%, and 99% credible intervals of our Bayesian estimates in Proposition 1. λ_g^S is defined in equation (7), and $\lambda_g^{S,\star}$ is λ_g^S's pseudo-true value. We consider strong factors, with $R_g^2 \in \{10\%, 20\%, 30\%\}$. We simulate quarterly observations of g_t and $\boldsymbol{r}_t$ by assuming that i) the true number of latent factors is 5, ii) the time series sample size is 200 quarters, and iii) the true $\bar{S} = 8$. We estimate several model configurations with different numbers of factors (4, 5, and 7) and $\bar{S} = 12$. The number of Monte Carlo simulations is 1,000.