Currency Return Dynamics: What Is the Role of U.S. Macroeconomic Regimes?*

Guanhao Feng Jingyu He Junye Li Lucio Sarno Qianshu Zhang

First Draft: May 2023; This Draft: August 24, 2025

Abstract

This paper investigates the influence of U.S. macroeconomic fundamentals on currency pricing factor models and factor risk premia. We develop a regime-switching model that endogenously detects macroeconomic regimes and estimates regime-specific models. Our empirical analysis identifies significant regime shifts, revealing three recurring regimes driven by U.S. inflation and interest rate trends. First, the Carry factor demonstrates robustness across regimes, while the importance of Value and Momentum factors varies. Second, high-inflation regimes are linked to currency mispricing and elevated idiosyncratic volatility. Finally, we identify a novel anomaly tied to macroeconomic regime shifts, unexplained by standard factor models.

Keywords: business cycles; conditional models; currency returns; regime shifts; risk premia.

JEL Classification: F31, G12, G15.

*We thank Victor DeMiguel, Stefano Giglio, Yongmiao Hong, Raymond Kan, Bryan Kelly, Ingomar Krohn (discussant), Sining Liu (discussant), Tao Liu (discussant), Zhenya Liu (discussant), Semyon Malamud, Andreas Neuhierl (discussant), Takeshi Osada (discussant), Shrihari Santosh, Yushan Tang (discussant), Qi Xu and seminar and conference participants at the 2024 SoFiE Annual Conference, China International Conference in Finance, 8th Annual World Symposium on Investment Research, AsianFA Annual Conference, 5th Fintech Forum (Renmin University), Sydney Banking and Financial Stability Conference, Australasian Finance and Banking Conference, Big Data and Econometric Conference (Xiamen University), 2023 Econometric Workshop for Time-varying Coefficient Models (Hunan University), CityU HK, and Zhejiang University for their insightful comments. For the purpose of open access, the authors have applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising from this submission. Feng (E-mail: gavin.feng@cityu.edu.hk), He (E-mail: jingyuhe@cityu.edu.hk) and Zhang (E-mail: qszhang7-c@my.cityu.edu.hk) are at City University of Hong Kong; Li (E-mail: li_junye@fudan.edu.cn) is at Fudan University; Sarno (E-mail: l.sarno@jbs.cam.ac.uk) is at University of Cambridge and Centre for Economic Policy Research (CEPR).

1 Introduction

The risk-return trade-off in currency markets remains a key focus in asset pricing research. Studies show that common risk factors like Carry (e.g., Lustig, Roussanov, and Verdelhan, 2011), Value (e.g., Menkhoff et al., 2017), and Momentum (e.g., Menkhoff et al., 2012b) help explain currency portfolio return dynamics and cross-sectional variation.[1] Despite these advancements, there is limited understanding of whether these factors maintain stable explanatory power and consistent risk premia across different macroeconomic regimes.

The gap is especially important because theories of exchange rate determination emphasize the critical role of macroeconomic fundamentals. Exchange rate dynamics are shaped both by domestic fundamentals and global financial cycles, which are significantly influenced by U.S. monetary policy (e.g., Miranda-Agrippino and Rey, 2020; Akinci and Queralto, 2024). Consequently, U.S. macroeconomic states play a pivotal role in global economic stability, with monetary policy decisions by the Federal Reserve frequently inducing international spillovers.[2] Notably, Lustig, Roussanov, and Verdelhan (2014) document that US economic downturns and heightened uncertainty elevate currency risk premia, driving countercyclical currency returns through shifts in US macroeconomic states. Dahlquist and Hasseltoft (2020) find that macroeconomic trends forecast currency returns, while Colacito et al. (2020) highlight output gaps as superior predictors compared to interest rate differentials. These findings emphasize the role of macroeconomic states in developing currency pricing models.

Empirical evidence underscores the time-varying nature of currency pricing. For example, Christiansen et al. (2011) show that regime-dependent models better capture the Carry factor by accounting for systematic risk, while Lettau et al. (2014) document that currencies with higher downstate market betas earn higher returns. Similarly,

[1] Recent studies employing latent factor models (Nucera, Sarno, and Zinna, 2024; Li, Sarno, and Zinna, 2024) underscore the importance of these factors in explaining variations in currency pricing.

[2] The U.S. economy exerts substantial spillover effects on global markets, as reflected in the pronounced co-movement of equity and bond prices across international economies (e.g., Rapach, Strauss, and Zhou, 2013; Albagli, Ceballos, Claro, and Romero, 2019).

Liu et al. (2024) find that conditional models outperform static ones in explaining currency returns. However, one challenge in asset pricing models is effectively integrating macroeconomic variables into conditional frameworks.

We aim to address key limitations in the literature by linking macroeconomic regimes to currency risk-return dynamics. To advance this understanding, we explore three interconnected research questions: (i) How does the relationship between currency risk and return change over time? (ii) To what extent are these changes associated with shifts in the macroeconomic states? (iii) Do currency factor risk premia vary across distinct macroeconomic regimes? To address these questions, we introduce a Bayesian conditional asset pricing model featuring regime-switching. The model identifies macroeconomic regimes and estimates regime-specific factor premia, providing a unified framework for analyzing the panel of asset returns. By incorporating macroeconomic states into asset pricing, our model offers new insights into the dynamic relationship between currency risk and return.

A key modeling challenge is efficiently detecting regimes and associating them with macroeconomic states. Research using advanced econometrics methods has investigated currency returns, factor risk premia, and idiosyncratic volatility (e.g., Bekaert, 1995; Kho, 1996; Johnson, 2002; Lustig et al., 2019; Dahlquist and Pénasse, 2022). These studies frequently apply time-varying parameter VAR or GARCH models for returns but provide limited insights into macroeconomic drivers. Markov regime-switching models are commonly used in exchange rate research (e.g., Engel and Hamilton, 1990; Clarida et al., 2003; Ang and Timmermann, 2012), but they produce latent regimes that are hard to interpret or tie to macroeconomic fundamentals. Another strand studies structural breaks in asset prices (e.g., Smith and Timmermann, 2021), but such breaks are often defined by calendar time and lack macroeconomic interpretation.

We address this by adapting and modifying the panel tree framework (Cong, Feng, He, and Li, 2023), which enables endogenous regime detection through a data-driven methodology. This approach involves selecting a large number of macroeco-

nomic states, such as high- and low-inflation regimes. It contrasts with methods that rely on a limited number of predefined regimes or detect structural breaks based on calendar time. Our framework integrates regime detection, factor selection, and risk premia estimation within a unified Bayesian approach. By employing a characteristics-based panel regression to model asset returns, we estimate regime-specific mispricings, idiosyncratic volatility, and factor risk premia.

Empirical highlights. We examine monthly exchange rate data for 47 currencies against the U.S. dollar from 1986 to 2023, investigating how U.S. macroeconomic states shape currency returns for U.S.-based investors. For these investors, idiosyncratic country risks are likely to diversify, leaving U.S.-specific risks and global factors as primary drivers. We focus on eight U.S. macroeconomic variables: interest rates, inflation, industrial production, nonfarm payrolls, unemployment, M2 money supply, oil prices, and the VIX index. Additionally, we examine eleven currency characteristics: forward premia (interest rate differentials for carry trades), momentum (short-, medium-, and long-term), valuation metrics, net foreign assets, domestic currency liabilities, long-term yields, term spreads, currency volatility, and currency market beta.

First, our empirical analysis provides robust evidence of regime-switching patterns in the dynamics of currency returns. We observe a marked improvement in the statistical model fitness when incorporating regime-switching. The detected regimes are explicitly defined by U.S. macroeconomic states: high inflation, low inflation with low interest rates, and low inflation with high interest rates. We demonstrate that neglecting regime switching significantly inflates mispricing estimates in low-inflation periods while understating them during high-inflation periods. This underscores the critical role of accounting for macroeconomic regimes in currency pricing.

Second, we document significant variation in factor selection and risk premia across regimes, emphasizing temporal shifts in characteristics-based models. The Carry factor consistently exhibits high risk premia across all regimes, reaffirming its central role in currency pricing (e.g., Lustig and Verdelhan, 2007; Lustig et al., 2011). How-

ever, the importance of other factors varies with macroeconomic states. For instance, LDC and market beta factors show elevated risk premia during high-inflation periods. Conversely, the Value factor gains prominence in low-inflation, low-interest-rate regimes, while the Volatility factor becomes critical in low-inflation, high-interest-rate environments. These results underscore the sensitivity of currency returns to business cycle phases and the pronounced regime-dependence of currency risk premia.

Third, we find that high-inflation regimes are often associated with currency mispricing and elevated idiosyncratic volatility, highlighting the uncertainty and divergence in currency valuations during periods of economic instability. This phenomenon is linked to the increased difficulty investors face in accurately forecasting exchange rates under such conditions. These observations contribute to the understanding of how currency mispricing (Bartram et al., 2025) and volatility (Della Corte et al., 2021) evolve across varying macroeconomic states.

Building on these insights, we develop an investment strategy that capitalizes on cross-sectional variation in currencies' sensitivity to *regime-switching risk*. This risk is particularly relevant for long-term investors with infrequent portfolio rebalancing, as it undermines the stability of expected returns. Short-term investors face challenges in detecting regime shifts and swiftly adjusting risk premia estimates. We define sensitivity as the difference between regime-specific Sharpe ratios from regime-switching models and unconditional Sharpe ratios from non-regime-switching models. Currencies with larger model-implied Sharpe ratio differentials show greater sensitivity to macroeconomic regime shifts, reflecting higher exposure to regime-switching risk and demanding higher expected returns as compensation. Empirical evidence reveals a monotonic relationship between sensitivity-sorted portfolios and returns. The proposed long-short strategy delivers significant abnormal returns, unexplained by currency common factors. Our currency regime-switching risk aligns with the equity break risk highlighted by Smith and Timmermann (2022), which introduces a factor capturing stocks' sensitivity to structural breaks in equity common factors.

Regime-switching model. We introduce a tree-based Bayesian regime-switching model designed to segment time series into recurrent regimes based on macroeconomic indicator thresholds. By employing a spike-and-slab prior, the model identifies sparse factor structures specific to each regime, facilitating efficient selection of key characteristics and estimation of risk premia.[3] A low selection probability for a factor suggests its risk premium is negligible due to the dominance of other factors within the regime. A key innovation of our approach is the adoption of a *global* split criterion, which optimizes the marginal likelihood of regime-specific models. This method mitigates overfitting and improves parameter estimation in regime detection by leveraging model uncertainty within a Bayesian framework.

The proposed model addresses key limitations of existing regime-switching methods. Traditional models, such as Markov-switching processes (e.g., Hamilton, 1989), often yield latent regimes that are challenging to interpret and only weakly linked to macroeconomic states *ex post*. Additionally, they fail to capture relationships between non-adjacent regimes. Another branch of threshold models identifies regimes when a variable crosses thresholds, but is usually limited to deterministic regime detection (e.g., Massacci, 2017; Liu and Chen, 2020). In contrast, our data-driven framework identifies regimes by estimating regime-specific factor models for the cross section. This approach characterizes detected regimes explicitly via macroeconomic states, providing a nuanced understanding of regime shifts. For example, regimes can be defined by disjoint threshold values of interest rates and inflation, enabling improved forecasting and interpretability of regime dynamics for currency returns.

Other related literature. The literature has documented various profitable currency investment strategies, including currency Carry (e.g., Lustig et al., 2011; Menkhoff et al., 2012a; Bakshi and Panayotov, 2013; Lettau et al., 2014; Koijen et al., 2018; Bekaert and Panayotov, 2020), currency value (e.g., Asness et al., 2013; Menkhoff et al., 2017),

[3] Giannone, Lenza, and Primiceri (2021) utilize the spike-and-slab prior to compare sparse and dense representations in empirical studies of economics and finance. Similarly, Bryzgalova, Huang, and Julliard (2023) apply this methodology within a two-pass cross-sectional regression.

and currency momentum (e.g., Menkhoff et al., 2012b; Asness et al., 2013). Nucera, Sarno, and Zinna (2024) identify three to four latent factors in a broad cross section of currency portfolios. Beyond the Dollar factor, they provide robust evidence for Carry and Momentum factors but report a weak Value factor (see also Li et al., 2024). In contrast, Chernov et al. (2023) examine G10 currencies and develop an efficient portfolio to price currency risks. They find Carry and Value factors dominate priced risks, with Momentum contributing less. These differences may stem from ignoring regime shifts in currency markets. In addition to the novel regime-switching model, our test assets also challenge conventional approaches by incorporating emerging markets and weaker factors evident in individual currency returns.

Our paper also contributes to the conditional asset pricing literature by emphasizing structural breaks and regime shifts, as highlighted in prior studies (e.g., Ang and Bekaert, 2002; Smith and Timmermann, 2021). While advanced methods have been developed to model time-varying coefficients in conditional asset pricing (e.g., Kelly et al., 2019; Cui et al., 2024), comparatively less attention has been given to the development of time-varying or regime-based factor models.

Bessembinder, Burt, and Hrdlicka (2024) highlights the importance of conditional models in capturing cross-sectional return shifts driven by evolving economic and market structures, especially during firm entry and exit. Similarly, Li et al. (2023) and Cui et al. (2023) document substantial time variations in characteristics-based return predictability and beta-pricing models. Kim, Korajczyk, and Neuhierl (2021) estimate conditional factor loadings, estimate conditional alphas using firm-level characteristics, and construct arbitrage portfolios. He, Su, and Yu (2024) examine how conditional anomaly returns change in response to revisions in expectations of macroeconomic states. Building on this foundation, our paper addresses a key gap in the literature by examining regime-dependent factor dynamics in currency markets. We provide novel insights into how regime-dependent factor models and risk premia change under varying macroeconomic states.

Finally, our approach leverages advancements in tree-based methods in finance, including those by Cong et al. (2023) and Cong et al. (2025).[4] These methods utilize a goal-oriented clustering framework to partition cross sections or time periods (regimes) based on economic objectives. The approaches vary significantly in model choices and targeted outcomes. Cong et al. (2023) introduce novel factor models by clustering stock return cross sections, offering fresh insights into return dynamics. Bie et al. (2024) employ the dynamic Nelson-Siegel yield curve model to effectively capture time variations in bond yields. To enhance asset pricing model tests, Cong et al. (2025) propose clustering individual stocks to construct basis portfolios, serving as characteristics-based test assets for evaluating model performance.

The paper is structured as follows. Section 2 introduces the regime-switching model. Section 3 describes the data. Section 4 presents the main empirical findings. Section 5 discusses an anomaly capturing regime risk. Finally, Section 6 concludes.

2 Methodology

Let $r_{i,t}$ denote the return of currency i at time t. Notably, the individual currency returns are unbalanced across time periods. Let $\mathbf{x}_{t-1} = (x_{1,t-1}, \cdots, x_{M,t-1})$ be a vector of M lagged macroeconomic variables and $\mathbf{Z}_{i,t-1} = (z_{i,1,t-1}, \cdots, z_{i,K,t-1})$ be a vector of K currency characteristics for individual currency i, which may influence the variation in currency returns. A regime is characterized as:

$$\mathcal{R}_j = \{t \mid \cup_{\text{some } m} \{x_{m,t-1} \leq c_{m,l}\} \cup_{\text{some other } m} \{x_{m,t-1} > c_{m,l}\}\},$$

which is a set of time periods during which the values of various macroeconomic variables $x_{m,t-1}$ (e.g., inflation) are below or above specified thresholds $c_{m,l}$. In the subsequent analysis, we refer to $x_{m,t-1} \leq c_{m,l}$ as a split rule, $x_{m,t-1}$ as a split variable, and $c_{m,l}$

[4]We introduce a tree-based Bayesian regime-switching model that integrates decision trees with Bayesian methods. Building on XBART (He et al., 2019; He and Hahn, 2023a), our approach incorporates Bayesian variable selection to enhance the detection of macroeconomic regimes. In clustering-based asset pricing, Creal and Kim (2021) analyze the risk-return tradeoff in currency markets using Bayesian Additive Regression Trees (BART; Chipman et al., 2010), with a focus on cross-sectional aspects while largely overlooking macroeconomic states and regime shifts.

as a split value. Notably, while this formulation resembles traditional threshold models, a key advantage of our approach is its ability to accommodate regimes defined by interactions among multiple macroeconomic variables, which standard threshold models are not equipped to capture.

Macroeconomic time series often exhibit trends over time, requiring normalization before analysis, as detailed in Section 3.3. Using a tree clustering algorithm, the number of macroeconomic variables defining a regime can be determined. These regimes may consist of non-contiguous periods characterized by either low or high levels of specific macroeconomic variables. The characteristics-based model, incorporating all identified regimes, is presented in Section 2.1. Procedures for regime detection and parameter estimation are discussed in subsequent subsections.

2.1 Full Factor Model with Regimes Given

We defer the discussion of regime identification to the next section. For now, we assume there are $j = 1, \cdots, J$ regimes over the time horizon, and the assignment of each period to a regime is given. The full model with regimes is

$$
\begin{aligned}
& r_{i,t} = r_{z,t} + \boldsymbol{a}(\mathbf{x}_{t-1}) + \boldsymbol{\lambda}(\mathbf{x}_{t-1})^{\mathsf{T}} \mathbf{Z}_{i,t-1} + \epsilon_{i,t}, \\
\text{where} \quad & \boldsymbol{a}(\mathbf{x}_{t-1}) = \sum_{j=1}^{J} \mathbb{1}_{\{t \in \mathcal{R}_j\}} \alpha_{j,i}, \quad \boldsymbol{\lambda}(\mathbf{x}_{t-1}) = \sum_{j=1}^{J} \mathbb{1}_{\{t \in \mathcal{R}_j\}} \boldsymbol{\lambda}_j, \\
& \epsilon_{i,t} \overset{\text{independent}}{\sim} N(0, \sigma_{i,t}^2), \quad \sigma_{i,t}^2 = \sum_{j=1}^{J} \mathbb{1}_{\{t \in \mathcal{R}_j\}} \sigma_{j,i}^2.
\end{aligned}
\tag{1}
$$

where the common factor $r_{z,t}$ captures the variation of the level factor in the cross section, following Smith and Timmermann (2022).[5] The indicator function, $\mathbb{1}_{\{t \in \mathcal{R}_j\}}$, equals to 1 if period t belongs to regime $\mathcal{R}_j$ and 0 otherwise. Each period t is assigned to one and only one regime. The vector $\mathbf{Z}_{i,t-1}$ is lagged characteristics of currency i at $t-1$. Following Fama and French (2020) and Smith and Timmermann (2022), the coefficient vector $\boldsymbol{\lambda}_j$ (with K elements) represents cross-sectional factor premia with

[5] Consistent with their methodology, we demean currency excess returns cross-sectionally.

given time-varying factor loadings (currency characteristics). The terms $\alpha_{j,i}$ and $\sigma_{j,i}$ capture the mispricing and idiosyncratic volatility of currency i under regime j.[6] The full model specification with all characteristics for the analysis is shown in Eq. (6).

In particular, the model parameters $\boldsymbol{a}$, $\boldsymbol{\sigma}^2$, and $\boldsymbol{\lambda}$ are regime-specific, enabling the model to capture the time-varying dynamics of currency returns over time. These parameters are dynamically linked to macroeconomic states, with regimes defined by aggregating the corresponding indicator functions. Essentially, we estimate a conditional characteristics-based model where the risk premia are driven by macroeconomic variables in a non-linear fashion. For comparison, we refer to the baseline model without regime-switching as the unconditional model in the following discussion.

Efficiently searching for and characterizing regimes using macro variables poses a modeling challenge. To address this, we leverage the panel tree framework that generalizes the traditional threshold-type model. Specifically, our algorithm learns a decision tree, represented as a non-linear function $\mathcal{T}(\mathbf{x}_{t-1})$. This function takes the value of a vector comprising all macro variables at period $t-1$ and assigns a corresponding regime label. Importantly, this function returns only one label, ensuring no regime overlap. A multi-way split at each node is unnecessary, as sequential binary splits can represent the same structure equivalently.

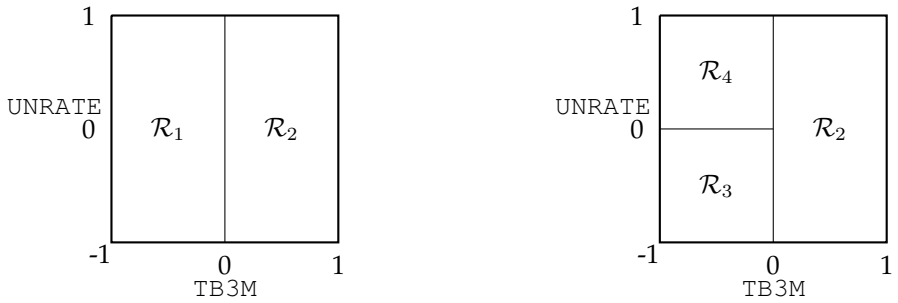
Figure 1 provides a simple example illustrating the definition of two and three regimes based on macro variables, such as the interest rate and unemployment rate. Using binary split rules, the entire space can be partitioned into distinct regimes. For instance, regime $\mathcal{R}_4$ can be described as $\mathcal{R}_4 = \{t \mid \texttt{TB3M}_{t-1} \leq 0 \text{ and } \texttt{UNRATE}_{t-1} > 0\}$, corresponding to periods when the short-term interest rate is below its 10-year historical average, and the unemployment rate is above its 10-year average. The interpretation of the other regimes follows a similar logic.

Thus far, the number of regimes and their corresponding split rules have been assumed to be predetermined. Section 2.2 presents the prior specification and marginal

[6]For robustness check, we also consider a version of model with extra latent factor $\epsilon_{i,t} = \delta_i F_t + e_{i,t}$ to capture the residual correlation as in Smith and Timmermann (2022). The resulting regime detection in the empirical study remains identical.

Figure 1: **Partition Plot for Macroeconomic-Instrumented Regimes**

The two partition plots illustrate how we determine regimes based on macro variables. Left: The partition plot with the first split at TB3M. Right: Based on the first split, we further split at UNRATE, resulting in three non-overlapping regimes labeled as $\mathcal{R}_2$, $\mathcal{R}_3$, $\mathcal{R}_4$. Both TB3M and UNRATE are demeaned based on the past ten years.



likelihood for the Bayesian variable selection model. Building on this, Section 2.3 introduces how the Bayesian regime-switching model detects regimes. Additionally, we propose a Gibbs sampler designed for regime detection, parameter estimation, and risk premia selection within each regime in Appendix A.III.

2.2 Bayesian Variable Selection and Marginal Likelihood

For any specific regime j and currency i, the full model in Eq. (1) becomes

$$r_{i,t} = r_{z,t} + \alpha_{j,i} + \boldsymbol{\lambda}_j^{\mathsf{T}} \mathbf{Z}_{i,t-1} + \epsilon_{i,t}, \tag{2}$$

since all indicators for other regimes take the value zero. We aim to estimate the risk premia $\boldsymbol{\lambda}_j$ of all factors within the j-th regime. Given the potentially large number of characteristics, we employ a Bayesian spike-and-slab prior (George and McCulloch, 1993) for variable selection. This prior allows for the possibility that the risk premium of some factors may be zero within the regime, specifically

$$\pi(\lambda_{j,k} \mid \boldsymbol{\gamma}_j) = (1 - \gamma_{j,k}) N(0, \xi_0^2) + \gamma_{j,k} N(0, \xi_1^2),$$

which is a mixture of normal priors with different variances. The "spike" component has a small prior variance ξ_0^2, concentrated around zero, indicating that the corresponding factor's risk premium is heavily shrunk toward zero. In contrast, the "slab"

component has a large prior variance ξ_1^2, allowing for a wider distribution around zero and suggesting minimal shrinkage for non-zero risk premiums. The latent dummy indicator variables $\boldsymbol{\gamma}_j = (\gamma_{j,1}, \cdots, \gamma_{j,K})$ determine whether the prior for each factor risk premium $\lambda_{j,k}$ belongs to the slab or spike category. Specifically, when $\gamma_{j,k} = 1$, the factor risk premium $\lambda_{j,k}$ follows the slab prior, indicating minimal shrinkage. Conversely, when $\gamma_{j,k} = 0$, $\lambda_{j,k}$ follows the spike prior, shrinking the coefficient toward zero.

This mechanism facilitates model selection by identifying relevant factors. Averaging posterior samples yields the posterior mean of $\gamma_{j,k}$ within $[0, 1]$, which can be interpreted as a selection probability. A low selection probability implies a high posterior likelihood that the factor's risk premium is zero, conditional on the presence of other factors. This approach provides a systematic way to discard irrelevant factors while retaining significant ones in the model.

It is important to note that the risk premium $\boldsymbol{\lambda}_j$ in Eq. (2) is regime-specific and independent of currency i. Therefore, estimating $\boldsymbol{\lambda}_j$ requires considering its joint selection across all currencies. The traditional spike-and-slab prior only applies to variable selection on a single regression. We modify it to estimate the panel regression model, where the risk premia estimate (the slope coefficient) is the same in the cross section, while alphas $\alpha_{j,i}$ and residual variance $\sigma_{j,i}^2$ are allowed to be currency-specific. The priors for these parameters are specified as follows:

$$\pi(\alpha_{j,i} \mid \sigma_{j,i}^2) = N(0, \sigma_\alpha^2 \sigma_{j,i}^2),$$
$$\pi(\sigma_{j,i}^2) = \text{inverse-Gamma}(v_0, S_0).$$

For the prior distribution of $\gamma_{j,k}$, we introduce an additional parameter w_k, which has an economically meaningful interpretation, representing the investor's belief about the probability that the risk premium of the k-th factor is non-zero. Three types of investors are taken into consideration: a confident investor with $w_k = 0.9$, who believes that most factors are relevant; an agnostic investor with $w_k = 0.5$, who is indifferent toward all factors; and a skeptic investor with $w_k = 0.1$, who believes that each fac-

tor has only a 10% chance of being included in the model. The prior for $\boldsymbol{\gamma}$ follows a binomial distribution, with w_k as its parameter:

$$\pi(\boldsymbol{\gamma}_j) = \prod_{k=1}^{K} w_k^{\gamma_{j,k}} (1 - w_k)^{(1-\gamma_{j,k})}.$$

The regime-switching model partitions the entire time series into multiple regimes, fitting a regime-specific model as defined in Eq.(2). Accurate regime detection enhances the overall fit of the data by leveraging these tailored models. The frequentist approach typically evaluates model fitness using the likelihood function, which simultaneously assesses data partitions and estimates model parameters. However, noisy plug-in parameter estimates can compromise the reliability of the likelihood function in evaluating regime effectiveness.

To address this issue, we adopt a Bayesian perspective by computing the marginal likelihood of the model. This approach integrates out the parameters for their prior distributions. Specifically, for any regime j, we stack all the data within the j-th regime in matrix form, $\mathbf{R}_{j,i}$ and $\mathbf{Z}_{j,i}$, respectively. The marginal likelihood of the model for any regime $\mathcal{R}_j$ is given by the product of individual marginal likelihood as follows:

$$\begin{aligned} p(\mathcal{R}_j) := \prod_i p(\mathbf{R}_{j,i} \mid \mathbf{Z}_{j,i}) = \int \prod_i p(\mathbf{R}_{j,i} \mid \mathbf{Z}_{j,i}, \boldsymbol{\gamma}_j, \alpha_{j,i}, \boldsymbol{\lambda}_j, \sigma_{j,i}^2) \\ \times \pi(\alpha_{j,i} \mid \sigma_{j,i}^2) \pi(\boldsymbol{\lambda}_j \mid \boldsymbol{\gamma}_j) \pi(\sigma_{j,i}^2 \mid \boldsymbol{\gamma}_j) \pi(\boldsymbol{\gamma}_j) d\alpha_{j,i} d\boldsymbol{\lambda}_j d\sigma_{j,i}^2 d\boldsymbol{\gamma}_j. \end{aligned} \tag{3}$$

Notably, the marginal likelihood integrates out unknown parameters in the likelihood function with respect to prior distributions. In addition, with K factors, there are 2^K possible model specifications, which are also averaged out using marginal likelihood. This Bayesian strategy accounts for uncertainties in both parameter estimation and model selection, relying solely on the data and prior distributions. By leveraging the marginal likelihood, we can effectively separate the steps of regime detection, parameter estimation, and model selection.

The marginal likelihood in Eq. (3) has no closed-form solution. Thus we follow

Chib (1995) to calculate it numerically using posterior samples obtained from the corresponding Gibbs sampler. Detailed derivations are provided in Appendix A.I. For simplicity, in the following sections, we denote $p(\mathcal{R}_j)$ as the marginal likelihood evaluation of Eq. (3), using the data within regime $\mathcal{R}_j$.

2.3 Macro-Instrumented Regime Detection

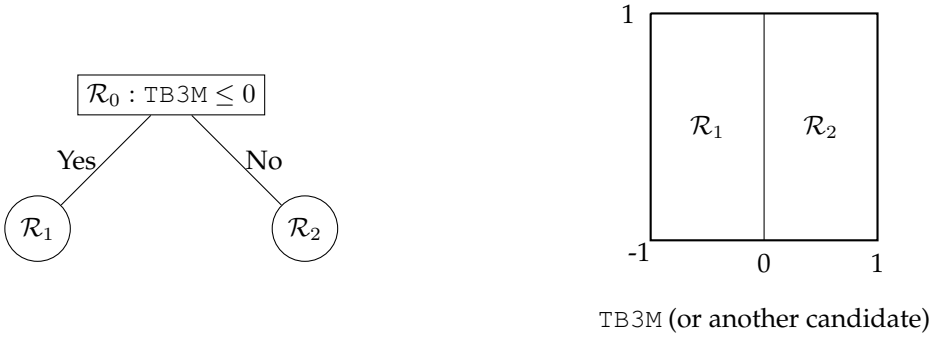
This section introduces our Bayesian regime-switching approach, which leverages a tree-based algorithm to endogenously identify macroeconomic regimes in currency return dynamics. The core idea is to iteratively partition the time series into regimes that maximize model fit, as measured by the marginal likelihood of the panel return model in Eq. (3).

To illustrate the procedure, Figure 2 presents a candidate for the first split: a single macroeconomic variable and cut value (e.g., the previous 10-year demeaned 3-month Treasury bill rate, `TB3M`, splitting at 0) that divide the sample into two regimes, denoted $\mathcal{R}_1$ and $\mathcal{R}_2$. They assign each calendar month to one of each regime according to the macroeconomic values. Thus the regimes are recurrent in the time horizon. For each regime, we estimate the factor model in Eq. (3) separately. Each split candidate yield different regimes, and the goal is to identify splits that capture structural changes in the relationship between characteristics and currency returns across macroeconomic states, specifically, improve the overall fitness to the data. The marginal likelihood in Eq. (3) thus serves as a natural measure of the quality of a split candidate. By iterating through all potential split candidates, we select the one that maximizes the overall marginal likelihood.

In particular, let $\mathcal{C} = \{\tilde{c}_j\}$ denote the set of all split rule candidates. To simplify the computation, we evaluate only a single cut value, 0, for each macroeconomic variable. Since all variables are demeaned using a 10-year rolling window, a cut point of 0 reflects whether the macro variable has been at a high or low level over the past decade. Each candidate splits the root node into two child leaf nodes, $\mathcal{R}_1$ and $\mathcal{R}_2$ (see Figure 2). The joint marginal likelihood of the complete dataset of currency returns is

Figure 2: **Illustration of one candidate for the first split**

To calculate the split criterion and find the optimal cutpoint, let's examine one potential split candidate: TB3M≤ 0. TB3M is demeaned based on the past ten years and can be replaced by any other variable.



the product of the marginal likelihoods for the two child nodes:

$$l(c_j) = p(\mathcal{R}_1) \times p(\mathcal{R}_2) = p(\mathbf{R}^L \mid \mathbf{Z}^L) \times p(\mathbf{R}^R \mid \mathbf{Z}^R), \tag{4}$$

where the superscripts L and R indicate data observations from the left and right child nodes (regimes). Each split rule creates a unique data partition, with the evaluation of Eq. (4) adjusting accordingly. Higher values indicate better model fit post-partition, offering a natural metric for assessing split candidates, termed the split criterion. For further details on additional splits and the stopping rule, see Appendix A.II.

The tree splits sequentially, and the following splits are similarly searched. At each j-th split, global split criteria and stopping rules are evaluated for all leaf nodes, selecting the candidate that optimizes the criterion. The process stops when a stopping criterion is met or conditions such as a minimum leaf size or maximum tree depth are satisfied. After $J-1$ iterations, the algorithm produces J regimes. After constructing the tree and identifying regimes, we estimate model parameters for each regime. We fit the Bayesian panel regression model in Eq. (2) within each regime and generate posterior samples using the Gibbs sampler outlined in Appendix A.III.

The Internet Appendix includes a simulation study to validate our methodology. This study examines three types of factors: strong, weak, and useless, which are characterized by varying levels of risk premia in the data generating process. Our approach can effectively distinguishes among these factor types. Strong factors con-

sistently exhibit selection probabilities near 1, indicating high confidence in their importance. Weak factors show more variability, with selection probabilities ranging between 0 and 1, reflecting their marginal significance. Useless factors are reliably excluded, with selection probabilities close to 0. Additional details are provided in Section IA.I of the Internet Appendix.

Lastly, our approach aligns with the goal-oriented clustering framework of Cong et al. (2023) but diverges in model selection. While Cong et al. (2023) employ a time-series model of equity returns on traded factors, we estimate cross-sectional slope factor models for currencies (Eq. (1)), capturing both common risk premia and idiosyncratic mispricing and volatility. This distinction in modeling has significant implications for clustering, yielding distinct marginal likelihoods. Unlike their framework, our marginal likelihood lacks a closed-form solution due to the complexity of cross-sectional return dynamics, requiring advanced computational methods.

3 Data

This section describes the data used in our empirical analysis, including spot and forward exchange rates, currency characteristics, and key U.S. macroeconomic variables employed for regime detection. The sample spans the period from 1986 to 2023.

3.1 Exchange Rates and Currency Returns

We gather spot and one-month forward exchange rates against the U.S. dollar (USD) for 47 economies.[7] The data are sourced from Barclays Bank International (BBI) and WM/Reuters via *Datastream*. The euro series begins in January 1999, with individual euro area countries excluded thereafter, retaining only the euro series. Following standard practice in the literature, we exclude observations during periods of signifi-

[7]The list includes Australia, Austria, Belgium, Brazil, Bulgaria, Canada, Croatia, Cyprus, Czech Republic, Denmark, Euro Area, Finland, France, Germany, Greece, Hong Kong, Hungary, Iceland, India, Indonesia, Ireland, Israel, Italy, Japan, Kuwait, Malaysia, Mexico, Netherlands, New Zealand, Norway, Philippines, Poland, Portugal, Russia, Saudi Arabia, Singapore, Slovakia, Slovenia, South Africa, South Korea, Spain, Sweden, Switzerland, Taiwan, Thailand, the United Kingdom, and Turkey.

cant covered interest rate parity (CIP) failures for a currency.[8] From a U.S. investor's perspective, exchange rates are quoted as USD per unit of foreign currency (FCU), i.e., FCU/USD. A rise in spot or forward rates reflects foreign currency appreciation relative to the USD or, equivalently, USD depreciation.

Currency excess returns are calculated using end-of-month spot and forward exchange rates. The excess return from purchasing a foreign currency in the forward market at month t and selling it in the spot market at month $t+1$ is defined as

$$rx_{i,t+1} = \frac{S_{i,t+1} - F_{i,t}}{S_{i,t}},$$

where $S_{i,t}$ and $F_{i,t}$ denote the spot and forward exchange rates, respectively, for currency i at month t. The excess return can be decomposed into its spot-return and forward-premium components as follows:

$$rx_{i,t+1} = \frac{S_{i,t+1} - S_{i,t}}{S_{i,t}} + \frac{S_{i,t} - F_{i,t}}{S_{i,t}} \approx r_{i,t+1} + \big(i_{i,t} - i_t\big), \tag{5}$$

where $r_{i,t+1}$ represents the spot exchange rate return, while $i_{i,t}$ and i_t denote the foreign and U.S. dollar interest rates, respectively. The approximation in Eq. (5) follows the standard practice of using forward rates instead of interest rate differentials to calculate excess returns, i.e., $fp_{i,t} \equiv \frac{S_{i,t}-F_{i,t}}{S_{i,t}} \approx i_{i,t} - i_t$, under covered interest parity.

3.2 Currency Characteristics

We follow the literature to construct various currency characteristics, including Carry (interest rate differential) (e.g., Lustig et al., 2011; Menkhoff et al., 2012a; Lettau et al., 2014; Koijen et al., 2018), short (1-month)-, medium (6-month)-, and long (12-month)-term currency momentum (e.g., Menkhoff et al., 2012b; Asness et al., 2013), currency value (e.g., Asness et al., 2013; Menkhoff et al., 2017), global imbalances based

[8] We follow Nucera et al. (2024) and their references to exclude below observations: Indonesia (01/12/1997 to 31/07/1998; 01/02/2001 to 31/05/2005); Malaysia (01/05/1998 to 30/06/2005); Russia (01/12/2008 to 30/01/2009; 03/11/2014 to 27/02/2015); South Africa (01/08/1985 to 30/08/1985; 01/01/2002 to 31/05/2005); Turkey (01/11/2000 to 30/11/2001).

on net foreign assets and liabilities in domestic currencies (Della Corte et al., 2016), long-term yields (Della Corte et al., 2016), term spread (Lustig et al., 2019), currency volatility (Menkhoff et al., 2012a), and currency market beta (Lettau et al., 2014; Verdelhan, 2018). See Table IA.2 in the Internet Appendix for detailed variable definitions. In what follows, we refer to these characteristics as Carry, MOM1, MOM6, MOM12, Value, NFA, LDC, LT Yield, Term Spread, Volatility, and MKT_Beta, respectively. To ensure comparability across the panel, we rank all characteristics cross-sectionally each month and normalize these ranks to $[-1, 1]$, with the cross-sectional median set to 0. Missing values are imputed as 0. Using cross-sectional ranks mitigates the influence of data errors and outliers in individual characteristics.

Table 1: **Summary Statistics**

The table shows summary statistics for currency excess returns (%) and 11 characteristics, including the number of observations, mean, standard deviation, median, 10th percentile, and 90th percentile. The dataset includes 11 characteristics, defined in Table IA.2. The monthly sample is from 1986 to 2023.

	Num	Mean	Std	Med	p10	p90
Excess Return (%)	12356	0.14	3.04	0.06	-3.24	3.55
Carry (%)	12356	0.18	0.55	0.06	-0.18	0.58
MOM1 (%)	12318	0.14	3.04	0.06	-3.25	3.54
MOM6 (%)	12126	0.89	7.94	0.36	-7.97	10.00
MOM12 (%)	11893	1.93	11.73	0.54	-11.35	16.10
Value	12344	0.02	0.23	0.00	-0.26	0.34
NFA	12346	-0.16	1.15	0.11	-1.39	0.59
LDC (million)	12346	152.03	894.01	2.83	0.17	97.90
LT Yield	10443	5.28	3.46	4.87	1.27	9.69
Term Spread	9110	0.82	1.68	0.85	-0.97	2.64
Volatility	12356	0.01	0.03	0.01	0.00	0.02
MKT_Beta	11630	1.01	0.63	1.07	0.09	1.74

Table 1 presents the summary statistics of currency excess returns and 11 characteristics. The pooled excess return averages 0.14% with a standard deviation of 3.04%, and its 10th and 90th percentiles are -3.24% and 3.55%, respectively.

3.3 U.S. Macroeconomic Variables

We consider eight leading U.S. macroeconomic indicators to define major economic regimes. These indicators include the 3-month interest rate, inflation, industrial

production, nonfarm payrolls, unemployment rate, M2, crude oil price index, and VIX.[9] Figure IA.1 in the Internet Appendix presents the time series of these macroeconomic variables, with NBER recessions overlaid. Clear periodical trends are evident in many macroeconomic time series during both boom and recession periods.

All macroeconomic indicators are standardized by demeaning, subtracting the 10-year rolling window average. To further smooth, we apply a 12-month smoothing to the macro variables. We use a simple weighted moving average smoothing, which puts more weight on recent data and less on past data, proportional to the number of months elapsed since 12 months ago. A positive value of a macro indicator indicates that its current level exceeds its 10-year average.[10] This standardized approach enhances the comparability of macroeconomic indicators over time, enabling the tree clustering method to effectively identify macroeconomic regimes from recent data.

4 Empirical Results

This section presents our key empirical findings. Section 4.1 presents the baseline model estimation and selection without considering regime shifts. Section 4.2, we examine the marginal likelihood gains and the resulting macroeconomic-instrumented regimes. Section 4.3 investigates the selection of regime-specific risk premia, and Section 4.4 analyzes cross-sectional variation of mispricing and idiosyncratic volatility.

4.1 Full Sample Risk Premia Estimation

We evaluate the overall performance of currency factors while excluding regime shifts, utilizing the full sample. We re-express model in Eq. (2) for individual currency

[9] For crude oil price, we calculate the annual percent change. For VIX before 1990, we impute with the S&P 500 Index Realized Volatility. Inflation is the annual percent change of CPI. Industrial production, M2, and nonfarm payrolls are their annual percent change. All data are from FRED.

[10] For instance, if the interest rate is greater than 0, it suggests that the current interest rate is above its average over the past decade, signaling a high interest rate regime in the economy.

returns with full list of characteristics:

$$
\begin{aligned}
r_{i,t} = \alpha_{i,t} + r_{z,t} + \lambda_{\text{Carry}}\text{Carry}_{i,t-1} + \lambda_{\text{MOM1}}\text{MOM1}_{i,t-1} + \lambda_{\text{MOM6}}\text{MOM6}_{i,t-1} \\
+ \lambda_{\text{MOM12}}\text{MOM12}_{i,t-1} + \lambda_{\text{Value}}\text{Value}_{i,t-1} + \lambda_{\text{NFA}}\text{NFA}_{i,t-1} \\
+ \lambda_{\text{LDC}}\text{LDC}_{i,t-1} + \lambda_{\text{LT_Yield}}\text{LT_Yield}_{i,t-1} + \lambda_{\text{Term_Spread}}\text{Term_Spread}_{i,t-1} \\
+ \lambda_{\text{Volatility}}\text{Volatility}_{i,t-1} + \lambda_{\text{MKT_Beta}}\text{MKT_Beta}_{i,t-1} + \epsilon_{i,t},
\end{aligned}
$$

where $r_{z,t}$ can be interpreted as a Dollar (DOL) factor, representing the common variation in currency excess returns.

Posterior samples of the model are drawn using the modified Gibbs sampler outlined in Appendix A.III. Table 2 summarizes the results of factor premia estimation. It presents the posterior factor selection probability (probability of non-zero risk premia) and the posterior means of factor premia, with 5% and 95% credible intervals provided in parentheses. Posterior evaluations are conducted across a wide range of the parameter w (as defined in Subsection 2.2), which reflects the prior belief in the usefulness of a factor. The results are provided for three types of priors: a confident prior ($w = 0.9$), an agnostic prior ($w = 0.5$), and a skeptical prior ($w = 0.1$). From the perspective of a panel predictive regression, there is strong return predictability for currencies based on those significant characteristic loadings (with a selection probability above 90%), which represent our risk premia estimates.

Notably, the Carry factor exhibits statistically significant and economically substantial risk premia, with nearly 100% selection probabilities regardless of investors' prior beliefs. This highlights its role as a strong factor in the cross section of individual currency returns. The factors MOM1 and Value follow, both show statistically significant positive risk premia and high selection probabilities for the confident prior. However, as the investor's prior beliefs become more skeptical, the Value premium diminishes more rapidly than that of MOM1. This finding aligns with Nucera, Sarno, and Zinna (2024), which indicates that DOL, Carry, and MOM1 are strong factors in the currency pricing kernel, while Value is relatively weak. In contrast, using a small cross

Table 2: **Factor Premia and Selection Probabilities without Regime Shifts**

The table displays risk premia estimation and selection results with no regime shifts. Risk premia estimates and selection probabilities are presented based on prior selection probability (with $w = 0.9$, $w = 0.5$, and $w = 0.1$). The left panel shows the results of the risk premia estimate (with 5% and 95% credible intervals shown in parentheses), while the right panel presents the results of the corresponding selection probability. The number represents the mean (5% and 95% quantiles in the parentheses) of 2000 MCMC post-convergence samples of λ and γ. The monthly sample is from 1986 to 2023.

	Risk Premium λ (%)			Selection Probability (%)		
Factors	$w = 0.9$	$w = 0.5$	$w = 0.1$	$w = 0.9$	$w = 0.5$	$w = 0.1$
Carry	0.433 [0.35, 0.52]	0.416 [0.34, 0.50]	0.412 [0.34, 0.49]	100	100	100
MOM1	0.129 [0.06, 0.19]	0.124 [0.04, 0.20]	0.094 [0.03, 0.18]	94	78	37
MOM6	0.104 [0.02, 0.19]	0.056 [-0.00, 0.15]	0.047 [0.00, 0.10]	82	26	6
MOM12	-0.038 [-0.12, 0.03]	-0.015 [-0.06, 0.03]	-0.009 [-0.05, 0.03]	47	8	0
Value	0.099 [0.03, 0.16]	0.060 [0.01, 0.13]	0.046 [0.00, 0.09]	85	28	4
NFA	-0.045 [-0.09, -0.00]	-0.038 [-0.08, -0.00]	-0.034 [-0.07, -0.00]	51	12	1
LDC	0.106 [0.00, 0.22]	0.046 [-0.01, 0.16]	0.028 [-0.02, 0.08]	79	22	3
LT Yield	-0.061 [-0.17, 0.02]	-0.028 [-0.11, 0.03]	-0.021 [-0.07, 0.03]	61	12	1
Term Spread	0.007 [-0.04, 0.06]	0.008 [-0.04, 0.05]	0.008 [-0.03, 0.05]	30	5	0
Volatility	0.043 [-0.02, 0.13]	0.013 [-0.03, 0.06]	0.010 [-0.03, 0.05]	48	5	0
MKT_Beta	0.095 [0.01, 0.20]	0.046 [-0.01, 0.13]	0.033 [-0.01, 0.08]	76	20	3

section of ten individual currencies from developed countries, Chernov, Dahlquist, and Lochstoer (2023) find that the Carry and Value factors account for large portions of the priced risk, followed by Momentum.

When the investor holds a strong prior belief in the model, the MOM6 and LDC factors demonstrate significant positive risk premia, with selection probabilities approaching 80%. However, as the investor adopts increasingly skeptical prior beliefs, both the associated risk premia and selection probabilities decline substantially. The factor risk premium of currency volatility is not statistically significant after control-

ling for other factors and decreases substantially, demonstrating a markedly low selection probability as the investor's prior beliefs become more skeptical.

Table 3: **Cross-sectional Alpha and Sigma Estimation without Regime Shifts**

The table displays the cross-sectional distribution of posterior estimates for α (mispricing, in %) and σ (idiosyncratic volatility, in %) without considering any regime shifts. The top panel shows the results of the α distribution, while the bottom panel presents the σ distribution. The mean, standard deviation, and quantiles of the cross-sectional distribution are provided. The posterior means of α and σ estimates are calculated from 2,000 MCMC post-convergence samples. The monthly sample is from 1986 to 2023.

	Prior	Mean	Std	5%	10%	25%	Median	75%	90%	95%
α (mispricing)	w = 0.9	0.127	0.141	-0.050	-0.040	0.025	0.100	0.215	0.300	0.375
	w = 0.5	0.127	0.134	-0.061	-0.040	0.035	0.110	0.215	0.310	0.367
	w = 0.1	0.128	0.134	-0.064	-0.034	0.040	0.120	0.215	0.314	0.357
σ (idiosyncratic volatility)	w = 0.9	2.193	0.812	1.389	1.520	1.610	2.000	2.575	3.314	3.812
	w = 0.5	2.196	0.813	1.392	1.516	1.610	2.000	2.575	3.314	3.815
	w = 0.1	2.196	0.813	1.392	1.526	1.610	2.000	2.570	3.320	3.815

To gain insight into mispricing and idiosyncratic volatility in the cross section, Table 3 provides a detailed summary of the cross-sectional mispricing (α) and idiosyncratic volatility (σ) distributions for the full sample period from 1986 to 2023, without considering any regime shifts. For various priors, the means of α are around 0.13%, with a standard deviation in the cross section of approximately 0.14%. The α percentiles show a negative lower end (around -0.06%), a slightly positive median (0.10%), and a positive upper end (0.36%). The cross-sectional means and standard deviations of idiosyncratic volatility are nearly the same across different priors, around 2.2%. The percentiles indicate a wide spread from a lower bound of about 1.38% to an upper bound of about 3.82%, with a median of about 2%. In summary, α and σ exhibit similar estimations regardless of prior beliefs about factor importance.

4.2 Macroeconomic Regime Detection

Next, we examine the pricing performance of currency factors while accounting for regime shifts. Importantly, our model can detect regimes by iteratively splitting the entire history of observations using U.S. macroeconomic states, based on enhance-

ments in Bayesian marginal likelihood, which remains unaffected by model and parameter uncertainties. This approach enables us to capture currency return dynamics in response to changing macroeconomic states.

In the regime-switching literature, the number of regimes is typically limited to two or three for ease of interpretability. Therefore, we restrict the maximum number of splits to two, allowing for three distinct macro-instrumented regimes in our implementation. Notably, our regimes are defined using macro variables, meaning that each regime may correspond to several disjoint periods over the time horizon.[11] Figure 3 presents the partition plot and ultimately detected regimes for the entire sample from 1986 to 2023. The number of months in each regime is also displayed: 160 months in Regime 1, 281 months in Regime 2, and only 15 months in Regime 3. It turns out that inflation (CPI) and the 3-month interest rate (TB3M) are the two most important business cycle indicators that define the regimes. The tree algorithm first splits the entire sample based on the inflation rate (CPI) into high- and low-inflation periods, and then further splits the low-inflation periods based on interest rates into low- and high-rate periods, resulting in three regimes as follows:

- High inflation regime (Regime 1): CPI > 0;
- Low inflation and low interest rate regime (Regime 2): CPI < 0 and TB3M < 0;
- Low inflation and high interest rate regime (Regime 3): CPI < 0 and TB3M > 0.

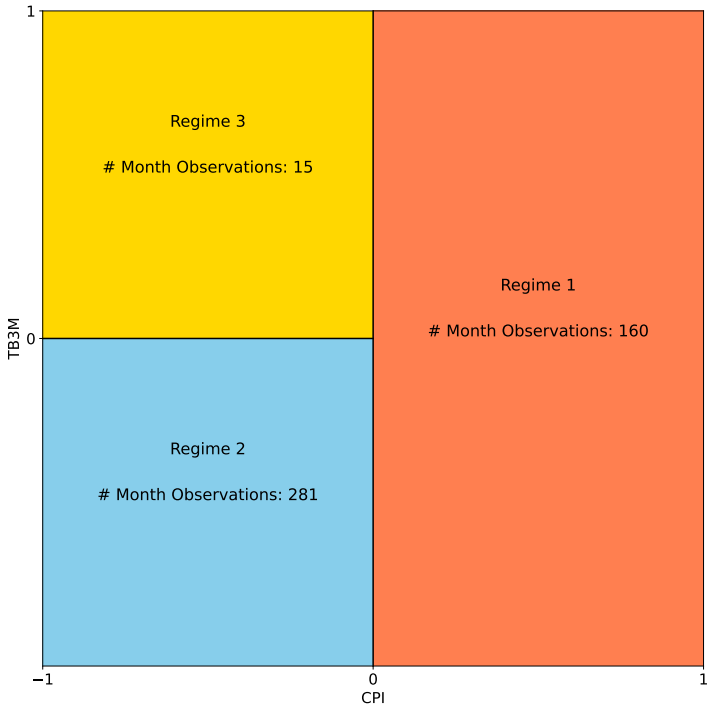
The samples from individual regimes are subsequently employed to identify factors specific to each regime and estimate their associated risk premia. Figure 4 illustrates the regime-switching patterns observed in the data.[12] Notably, Regime 1, characterized by high inflation and marked in red, includes significant events such as the

[11] Our model demonstrates a substantial improvement in Bayesian marginal likelihood compared to models without regime shifts, with the log marginal likelihood increasing by 172.5 after the first split and 173.09 after the second.

[12] Regime 1 predominantly spans the periods from 1990 to 1992, 2000 to 2002, 2005 to 2007, 2017 to 2020, and 2021 to 2023; Regime 2 includes the periods from 1986 to 1990, 1992 to 2000, 2002 to 2004, 2009 to 2017, and 2020 to 2021; and Regime 3 encompasses all remaining periods.

Figure 3: **Partitioning Macroeconomic Time Series**

The figure illustrates the detected macroeconomic regimes over the monthly sample period from 1986 to 2023, presented as a partition plot. The data is first partitioned based on standardized inflation (CPI), followed by a subdivision of low-inflation periods into regimes characterized by high and low interest rates. This process yields three distinct regimes. Macroeconomic indicators are standardized by subtracting the 10-year rolling average and applying smoothing, with a partition threshold of zero.

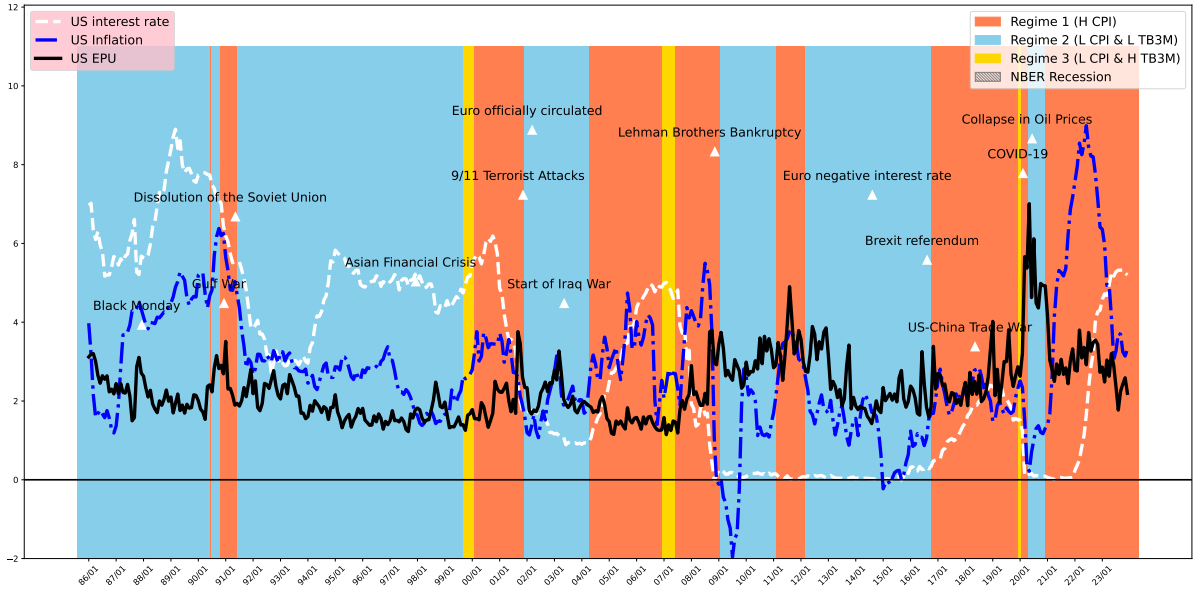


Gulf War (1990-1991) and the early phase of the COVID-19 pandemic (2020). These periods are indicative of economic instability and elevated price levels, consistent with a high inflation environment. This regime typically coincides with peaks in the business cycle, reflecting an overheating economy that may lead to asset bubbles.

Regime 2 (low inflation and low interest rates), depicted in blue, is marked by persistently subdued inflation, historically low nominal rates, and episodic macro-financial disruptions. Notable events include Black Monday (1987), the Soviet Union's dissolution (1991), the Asian Financial Crisis (1997–1998), the 9/11 attacks (2001), Lehman Brothers' collapse (2008), and the Brexit referendum (2016). These episodes highlight financial markets' susceptibility to external shocks despite stable inflation and accommodative monetary policy. This regime aligns with the recovery phase of the business cycle, transitioning into expansion, where stimulative monetary conditions foster borrowing, spending, and sustainable growth.

Figure 4: **Macroeconomic-Instrumented Regimes for Currency Returns**

The figure illustrates detected economic regimes from 1986 to 2023 using monthly data. Regime 1 (red) represents a high-inflation environment, Regime 2 (blue) corresponds to periods of low inflation and low interest rates, and Regime 3 (yellow) denotes low inflation coupled with high interest rates. Shaded areas mark NBER-defined recession periods. The figure also presents time-series plots of U.S. interest rates, inflation, and the scaled EPU index (Baker et al., 2016), alongside notable global economic events.



Regime 3, characterized by low inflation and high interest rates, occurs infrequently and typically for short durations. Notably, during the late 1990s, early 2007, and prior to the COVID-19 pandemic, the U.S. Federal Reserve raised interest rates to temper stock market booms despite the presence of moderate inflation. This regime represents the slowdown and contraction phase of the business cycle, characterized by diminished economic growth. The figure illustrates U.S. interest rates (dashed white line), U.S. inflation (dashed-dotted blue line), and U.S. Economic Policy Uncertainty (EPU) (solid black line) (Baker et al., 2016). Spikes in EPU are closely associated with major global events and transitions between economic regimes, highlighting elevated uncertainty during these periods.

4.3 U.S. Macroeconomic Regimes and Factor Premia

Macroeconomic regimes, defined by varying business cycle conditions shaped by inflation and interest rates, exhibit distinct dynamics in currency returns. We further estimate currency factor premia under each regime using three types of investor

prior beliefs — confident, agnostic, and skeptical — reflecting the varying degrees of belief in the explanatory power of currency return dynamics. We present the regime-dependent time-varying factor premia estimates in Table 4.[13]

In Regime 1, marked by high inflation, the Carry factor consistently delivers statistically significant positive risk premia across all priors: 0.33% (confident prior), 0.32% (agnostic prior), and 0.32% (skeptical prior), underscoring its robustness to varying investor beliefs. The MKT Beta factor also exhibits significant positive premia, ranging from 0.47% (skeptical prior) to 0.52% (confident prior), exceeding Carry's performance and highlighting the profitability of beta strategies in high-inflation environments. Additionally, the risk premium of the LDC factor is substantial across all three types of priors, ranging from 0.09% (for the skeptical prior) to 0.22% (for the confident prior). While MOM1 is statistically significant for the confident prior, it decreases significantly for the other priors. In contrast, the risk premia of other factors, including the Value factor, are comparatively small and statistically insignificant.

Regime 2 corresponds to periods of low inflation and low interest rates. Again, the Carry factor stands out with a high positive risk premium of around 0.43% for all types of priors. These factor premia are larger than those in Regime 1. However, we find that the factor risk premium of the LDC is much smaller than that in Regime 1 and is not statistically significant for any prior. Additionally, the factor risk premium of MKT Beta is also not statistically significant for all three types of priors. The MOM1 and MOM6 factors exhibit significant premia, although these decline for the skeptical prior. The Value factor emerges in this more stable regime, delivering a statistically significant positive risk premium under the confident and agnostic priors, ranging from 0.06% (for the agnostic prior) to 0.10% (for the confident prior), though it remains lower than the premium for MOM6.

Regime 3, with low inflation and high interest rates, shows more complex results. Carry continues to generate a positive and statistically significant risk premium, but

[13]We report posterior estimates derived from 2,000 Markov Chain Monte Carlo (MCMC) draws, following the exclusion of burn-in samples.

Table 4: **Factor Premia under All Regimes**

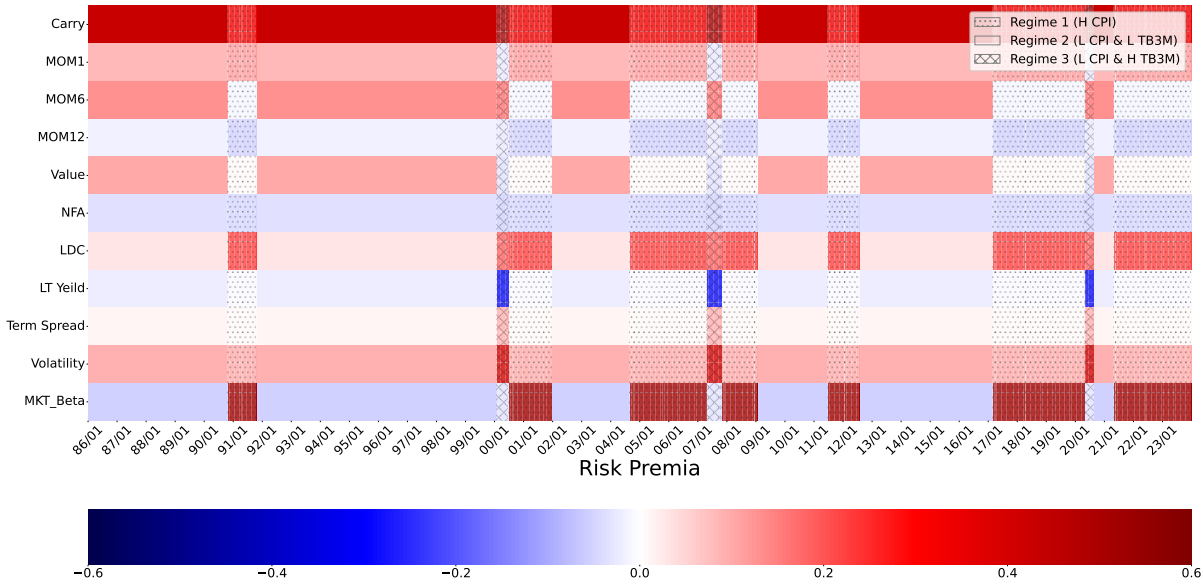
The table reports estimated factor premia (%) across three regimes detected via inflation and interest rate data (see Figure 3). Regime 1 reflects high inflation, Regime 2 corresponds to low inflation and low interest rates, and Regime 3 represents low inflation and high interest rates. Risk premia estimates are presented with 5% and 95% credible intervals, based on prior selection probabilities: confident ($w = 0.9$), agnostic ($w = 0.5$), and skeptical ($w = 0.1$). Values represent the mean (with 5% and 95% quantiles in parentheses) from 2,000 post-convergence MCMC samples of λ.

	Regime 1 (H CPI)			Regime 2 (L CPI & L TB3M)			Regime 3 (L CPI & H TB3M)		
Factors	w = 0.9	w = 0.5	w = 0.1	w = 0.9	w = 0.5	w = 0.1	w = 0.9	w = 0.5	w = 0.1
Carry	0.333	0.321	0.318	0.430	0.426	0.427	0.501	0.269	0.048
	[0.20, 0.46]	[0.20, 0.44]	[0.20, 0.44]	[0.33, 0.54]	[0.32, 0.52]	[0.33, 0.52]	[0.06, 0.90]	[-0.03, 0.75]	[-0.04, 0.35]
MOM1	0.115	0.063	0.034	0.082	0.061	0.046	-0.025	0.004	0.004
	[0.01, 0.22]	[-0.01, 0.19]	[-0.02, 0.09]	[0.01, 0.17]	[-0.00, 0.16]	[0.00, 0.11]	[-0.27, 0.19]	[-0.07, 0.08]	[-0.05, 0.06]
MOM6	-0.010	-0.005	-0.004	0.129	0.088	0.050	0.168	0.057	0.016
	[-0.11, 0.07]	[-0.06, 0.05]	[-0.05, 0.04]	[0.02, 0.23]	[0.00, 0.21]	[-0.00, 0.14]	[-0.16, 0.63]	[-0.04, 0.38]	[-0.04, 0.08]
MOM12	-0.053	-0.020	-0.013	-0.014	-0.002	0.005	-0.027	0.015	0.009
	[-0.18, 0.04]	[-0.09, 0.03]	[-0.06, 0.03]	[-0.10, 0.05]	[-0.05, 0.05]	[-0.04, 0.05]	[-0.45, 0.34]	[-0.09, 0.23]	[-0.05, 0.06]
Value	0.005	0.003	0.003	0.103	0.066	0.040	-0.034	-0.023	-0.009
	[-0.06, 0.08]	[-0.04, 0.05]	[-0.04, 0.05]	[0.02, 0.18]	[0.00, 0.16]	[-0.01, 0.09]	[-0.34, 0.24]	[-0.20, 0.05]	[-0.06, 0.05]
NFA	-0.048	-0.026	-0.024	-0.036	-0.024	-0.020	-0.055	-0.018	0.001
	[-0.13, 0.01]	[-0.08, 0.02]	[-0.07, 0.02]	[-0.10, 0.01]	[-0.07, 0.02]	[-0.06, 0.02]	[-0.25, 0.08]	[-0.14, 0.05]	[-0.05, 0.06]
LDC	0.224	0.186	0.097	0.029	0.008	0.005	0.185	0.085	0.015
	[0.05, 0.36]	[0.01, 0.35]	[-0.01, 0.31]	[-0.04, 0.13]	[-0.04, 0.06]	[-0.04, 0.05]	[-0.05, 0.52]	[-0.04, 0.42]	[-0.05, 0.07]
LT Yield	-0.008	-0.004	-0.001	-0.020	-0.012	-0.007	-0.284	-0.085	0.001
	[-0.15, 0.11]	[-0.06, 0.05]	[-0.05, 0.05]	[-0.14, 0.06]	[-0.07, 0.04]	[-0.06, 0.04]	[-0.72, 0.06]	[-0.55, 0.05]	[-0.05, 0.06]
Term Spread	0.005	0.003	0.003	0.010	0.007	0.007	0.095	0.071	0.029
	[-0.06, 0.08]	[-0.05, 0.05]	[-0.04, 0.05]	[-0.05, 0.08]	[-0.04, 0.06]	[-0.04, 0.05]	[-0.17, 0.45]	[-0.05, 0.42]	[-0.04, 0.22]
Volatility	0.102	0.036	0.019	0.092	0.063	0.039	0.397	0.314	0.105
	[-0.02, 0.25]	[-0.03, 0.18]	[-0.03, 0.07]	[-0.01, 0.20]	[-0.01, 0.18]	[-0.01, 0.09]	[0.05, 0.68]	[-0.01, 0.61]	[-0.03, 0.48]
MKT_Beta	0.521	0.482	0.473	-0.055	-0.038	-0.033	-0.033	-0.015	-0.009
	[0.37, 0.68]	[0.35, 0.62]	[0.34, 0.59]	[-0.17, 0.03]	[-0.14, 0.03]	[-0.09, 0.02]	[-0.33, 0.23]	[-0.21, 0.08]	[-0.06, 0.05]

it is only significant for the confident prior. For the confident prior, its risk premium is 0.50%, while for the agnostic and skeptical priors, it is 0.27% and 0.05%, respectively. It suggests the robustness of the carry trade, which remains unaffected by this reduced growth regime. Although it shrinks under a more skeptical prior, this is due to the relatively short time period of this regime, which limits the likelihood's ability to counteract the prior. In this regime, the Volatility factor becomes more important, generating a high statistically significant risk premium of 0.40% for the confident prior. None of the other factors is significant, and some are even negative.

Figure 5: **Time-varying Risk Premia under Regime Shifts**

The figure presents the estimated risk premia (%) under a confident prior ($w = 0.9$) across three distinct regimes detected using U.S. inflation and interest rate data. These regimes are illustrated in Figure 3.



From the perspective of regime-dependent panel predictive regression, there is significant return predictability for currencies driven by risk premia estimates, as evidenced by the 90% credible intervals. Figure 5 illustrates how factor premia vary across macroeconomic regimes, highlighting both common and unique factors under different conditions. These findings emphasize the importance of regime-specific analysis for investors with confident priors.

Factor selection. Table 5 summarizes the selection probabilities of factors across regimes under three prior beliefs, highlighting their importance across macroeconomic states.

Consistent with the factor risk premium estimates, in Regime 1, the Carry factor and MKT Beta are selected with nearly 100% probability across all three priors, highlighting their robustness in the high-inflation environment. Among other significant factors, LDC shows a 96% selection probability for the confident prior, while MOM1 exhibits an 81% probability. However, the selection probabilities for both factors diminish considerably as the investor's prior beliefs become more skeptical.

In Regime 2, only Carry is selected with a probability of 100% for all three priors. The MOM6 and Value factors are selected with a probability larger than 80% for the confident priors, and for the agnostic and skeptical priors, their selection probabilities decrease. Other factors, such as Volatility and MOM1, exhibit selection probabilities exceeding 70% under a confident prior. However, these probabilities decline significantly when investors adopt a less confident prior.

In Regime 3, no factors exhibit a selection probability of 100%. Under the confident prior, it is observed that the Carry and Volatility factors are selected with high probabilities, exceeding 90%. For the agnostic prior, the Volatility factor has a selection probability of 80%, while the Carry factor is selected with a probability of 61%; all other factors exhibit probabilities below 50%. In contrast, under the skeptical prior, all factors are selected with negligible probabilities.

To sum up, factor premia vary significantly across regimes driven by inflation and interest rate dynamics, highlighting the influence of business cycle conditions on currency returns. The Carry factor demonstrates consistently high selection probabilities. This finding suggests that the interest rate differential is a significant source of risk. Empirical evidence supports that the carry trade compensates investors for bearing this risk. In Regime 1 (high inflation), Market Beta, MOM1 (short-term Momentum), and Liquidity (LDC) factors dominate due to their sensitivity to inflationary pressures. In Regime 2 (low inflation, low interest rates), MOM6 (medium-term Momentum), Value, MOM1, and Volatility factors play a central role, reflecting their importance in stable economic conditions. In Regime 3 (low inflation, high interest rates), the Volatil-

Table 5: **Factor Selection Probability under All Regimes**

The table reports factor selection probabilities (%) across three regimes detected using U.S. inflation and interest rate data (see Figure 3). Results are based on varying priors, with probabilities computed as the mean of 2,000 MCMC post-convergence samples of γ. The monthly sample is from 1986 to 2023.

	Regime 1 (H CPI)			Regime 2 (L CPI & L TB3M)			Regime 3 (L CPI & H TB3M)		
Factors	w = 0.9	w = 0.5	w = 0.1	w = 0.9	w = 0.5	w = 0.1	w = 0.9	w = 0.5	w = 0.1
Carry	100	99	99	100	100	100	96	61	11
MOM1	81	31	4	72	29	7	59	13	2
MOM6	43	6	0	87	47	9	75	22	4
MOM12	57	10	1	40	6	0	68	20	2
Value	36	5	0	82	35	4	66	15	2
NFA	54	9	1	45	7	1	60	14	3
LDC	96	76	32	47	7	0	77	31	3
LT Yield	49	7	0	44	7	1	83	28	2
Term Spread	37	6	1	35	5	0	68	27	7
Volatility	71	17	1	72	32	4	96	80	25
MKT_Beta	100	100	100	56	18	4	62	19	3

ity factor becomes prominent, particularly under an agnostic prior, highlighting risks associated with slower economic growth.

4.4 Mispricing and Idiosyncratic Volatility

The analysis shows that currency risk premia are sensitive to macroeconomic regime shifts. Using model-implied parameters α_i and σ_i, we can analyze their cross-sectional distributions to investigate potential mispricing and idiosyncratic volatility during regime transitions. Della Corte et al. (2021) uncover a currency volatility factor by sorting currencies based on the slope of their implied volatility term structure, while Bartram et al. (2025) show that mispricing explains currency return predictability. Given these empirical facts, we further evaluate whether the explanatory power of currency risk factors shifts across macroeconomic regimes.

Table 6 Panel A presents the estimated cross-sectional alphas across the three economic regimes, and Figure 6 Panel A plots the kernel density estimation. Regime 1, characterized by high inflation, shows significant positive mean alphas, approximately 0.27% across all three prior beliefs, greater in absolute value than the overall mispricing in Table 3. This suggests substantial underpricing on average, potentially

due to restrictive monetary policy during high inflation, which increases the likelihood of mispricing. The 95% quantile is notably positive (0.55%), while the 5% quantile is slightly negative (-0.012% to 0.000%). Inflationary environments are often characterized by greater volatility and uncertainty about future inflation trends. This creates challenges for market participants in forming accurate expectations about currency risk premia, contributing to inefficiencies (see, Campbell and Vuolteenaho, 2004).

Table 6: **Cross-Sectional Alpha and Sigma Estimation under All Regimes**

The table presents the cross-sectional distributions of posterior estimates for α (mispricing, Panel A, in %) and σ (idiosyncratic volatility, Panel B, in %) across three regimes detected based on U.S. inflation and interest rate data (see Figure 3). For each regime, the table reports the cross-sectional mean, standard deviation, and quantiles of the distribution.

	Prior	Mean	Std	5%	Median	95%	Mean	Std	5%	Median	95%
		Panel A: α (mispricing)					Panel B: σ (idiosyncratic volatility)				
Regime 1 (H CPI)	w = 0.9	0.269	0.173	-0.012	0.255	0.570	2.363	0.964	1.252	2.135	4.025
	w = 0.5	0.274	0.165	-0.008	0.275	0.552	2.362	0.972	1.242	2.125	4.025
	w = 0.1	0.280	0.157	0.000	0.280	0.538	2.364	0.970	1.250	2.125	4.045
Regime 2 (L CPI & L TB3M)	w = 0.9	-0.001	0.105	-0.164	-0.010	0.138	2.148	0.776	1.432	1.970	3.312
	w = 0.5	-0.000	0.104	-0.158	-0.010	0.150	2.147	0.782	1.424	1.960	3.320
	w = 0.1	0.002	0.104	-0.150	-0.010	0.158	2.146	0.782	1.422	1.970	3.320
Regime 3 (L CPI & H TB3M)	w = 0.9	0.029	0.057	-0.047	0.025	0.127	1.888	0.613	1.206	1.725	3.063
	w = 0.5	0.030	0.060	-0.040	0.030	0.118	1.871	0.606	1.174	1.695	3.030
	w = 0.1	0.033	0.063	-0.040	0.020	0.135	1.880	0.603	1.160	1.700	2.984

Conversely, Regime 2, defined by low inflation and interest rates, shows a smaller cross-sectional mean alpha, ranging from 0.002% (for the skeptical prior) to -0.001% (for the confident prior), indicating substantially reduced mispricing. Low inflation and low interest rates create a stable economic environment, which reduces uncertainty and helps align currency prices with their fundamental values, minimizing mispricing. Regime 3, characterized by low inflation and high interest rates, exhibits slightly higher mean alpha values than Regime 2, with a notably smaller cross-sectional standard deviation (0.06%) compared to Regime 2 (0.10%). Across all priors, the 5% quantile alpha is about -0.04% and the 95% quantile alpha is about 0.12%. A low-

inflation environment coupled with high interest rates often signals central banks prioritizing long-term risks, such as asset bubbles or financial instability, over short-term inflation concerns.[14] This enhances the risk-return dynamics in currency markets, aligning more closely with risk premia and subtle underpricing. Mispricing in Regimes 2 and 3 is lower than the full-period estimate (0.127) in Table 3, highlighting that ignoring regime shifts can cause significant errors, including overestimation during these periods and underestimation in high-inflation regimes.

Figure 6: **Cross-Sectional Alpha and Sigma Distribution in Each Regime**

This figure presents the cross-sectional distributions of posterior estimates for α (mispricing, Panel A) and σ (idiosyncratic volatility, Panel B) across three regimes detected based on inflation and interest rate data (see Figure 3) utilizing a confident prior ($w = 0.9$). The distributions are estimated employing kernel density estimation. The monthly sample is from 1986 to 2023.

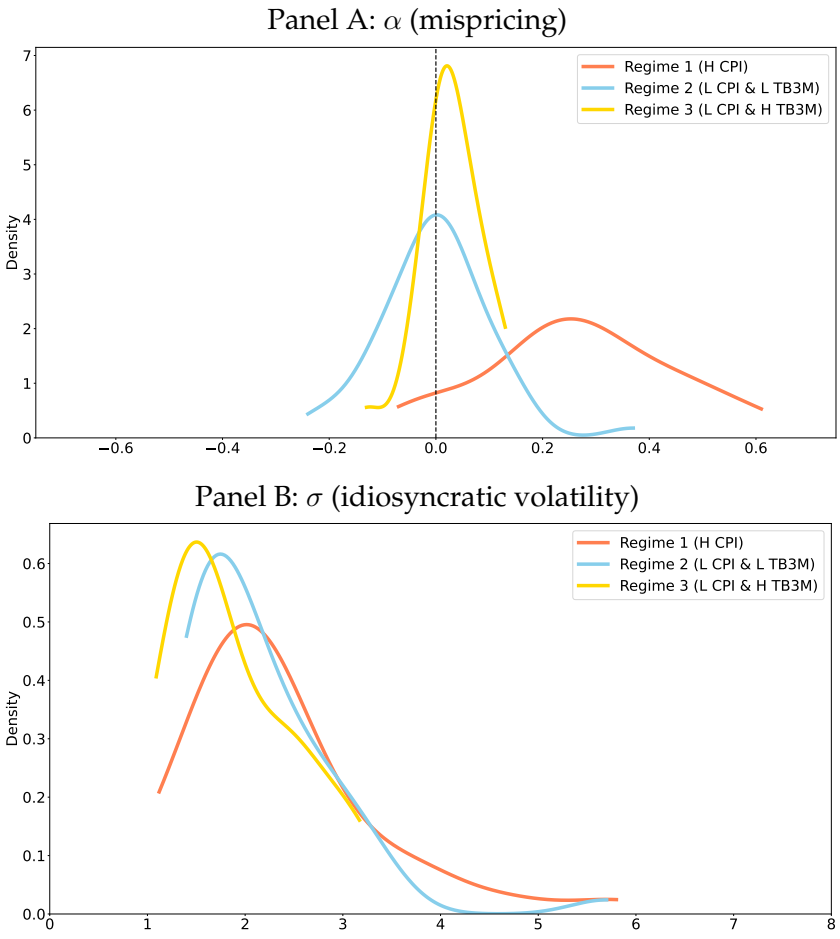


Table 6 Panel B shows the averages of cross-sectional idiosyncratic volatility across

[14]For instance, Miao, Shen, and Wang (2019) demonstrate that central banks effectively mitigate asset price bubble volatility by raising interest rates in response to growing bubbles.

economic regimes, and Figure 6 Panel B plots their kernel density estimation. In Regime 1, average idiosyncratic volatility is about 2.36%, with 5% and 95% quantiles at 1.2% and 4.0%. This suggests that currency-specific risk reaches its peak during the high inflation regime, aligning with Vavra (2014), who finds that greater volatility is associated with increasing inflation. The average idiosyncratic volatility decreases to less than 2% across all three priors in Regime 3, marked by low inflation and high interest rates. This decline suggests reduced currency-specific risks in this slowdown regime. The 5% and 95% quantiles are 1.2% and 3.0%. Regime 2, with low inflation and interest rates, shows moderate idiosyncratic volatility. In summary, inflation is crucial for currency idiosyncratic volatility. High inflation is closely linked to heightened volatility, driven by greater economic uncertainty and market instability, whereas low inflation is generally associated with reduced volatility. Overall, our findings reveal significant variations in mispricing and idiosyncratic volatility over time and across the cross section of currency returns.

5 Regime Risk Anomaly

We document empirical evidence of regime-switching dynamics in factor premia and idiosyncratic volatility. These regime shifts introduce an additional layer of risk for investors, as changes in risk premia and volatility can significantly impact portfolio outcomes. If an investor constructs a model that disregards regime shifts and forms portfolios based solely on estimated risk premia, a future macroeconomic regime change could lead to altered risk premia, potentially resulting in unexpected losses. Our concept of currency regime-switching risk is closely related to the equity break risk identified by Smith and Timmermann (2022). Their break risk factor arises from stocks' sensitivity to structural breaks in market conditions and the time-varying nature of risk premia associated with equity common factors.

Sensitivity measure. To measure individual currencies' sensitivity to regime shifts, we define the regime-specific implied Sharpe ratio (*RISR*) as the regime-dependent

model-implied return divided by the regime-dependent model-implied idiosyncratic volatility. Specifically, for currency i at time t when it is in regime j given macro condition $\mathbf{x}_{t-1}$ with characteristics $\boldsymbol{Z}_{i,t-1}$ and factor premia $\boldsymbol{\lambda}_j$, we have

$$RISR_{i,t,j} = \frac{\boldsymbol{Z}_{i,t-1}^{\mathsf{T}} \boldsymbol{\lambda}_j}{\sigma_{i,j}}.$$

Notice that the value of $RISR_{i,t,j}$ changes over time and depends on the currency characteristics. Similarly, under the unconditional model (without regime shifts) estimated with the full sample, there is only one constant risk premium $\boldsymbol{\lambda}$ and idiosyncratic volatility σ_i for the entire sample period. We define such unconditional implied Sharpe ratio (*ISR*) as follows:

$$ISR_{i,t} = \frac{\boldsymbol{Z}_{i,t-1}^{\mathsf{T}} \boldsymbol{\lambda}}{\sigma_i}.$$

We further define the difference in implied Sharpe Ratio (*DISR*) as,

$$DISR_{i,t} = RISR_{i,t,j} - ISR_{i,t},$$

which measures the difference in model-implied Sharpe ratios between the regime-switching model and the unconditional model. A higher *DISR* for a currency signifies a greater divergence between the two models, underscoring heightened exposure to regime-switching risk. As a result, such a currency commands a higher risk-adjusted return, thereby introducing a novel dimension to the risk-return trade-off.

Out-of-Sample evaluation. In Section 4, we identify three regimes based on inflation and interest rates over the full sample period. These regimes, however, may shift with changes in the data sample period. To evaluate the robustness of our proposed method, we perform an out-of-sample evaluation while addressing potential look-ahead bias by computing risk anomalies in real-time. The regime-switching model is initially estimated using a 14-year training period, after which it is updated monthly with historical data. The out-of-sample evaluation commences in the year 2000.

At the end of each month, regimes are detected using lagged macroeconomic variables, and a regime-dependent model is estimated for currency return dynamics, incorporating factor premia and idiosyncratic volatility. The regime for the subsequent month is determined based on lagged macroeconomic states. *RISR* and *ISR* for each currency are dynamically updated throughout the sample, based on lagged characteristics. We focus on investors with strong priors that currency return dynamics are primarily driven by these characteristics. Conversely, for investors with weaker priors, factor premia shrink toward zero, complicating the identification of regime-specific risks relative to unconditional premia. Investors trusting the characteristics-based model are more attuned to regime-switching risk exposure.

Table 7: **Summary Statistics: *DISR*-sorted Portfolios**

The table summarizes statistics for *DISR*-sorted quintile currency portfolios over 2000–2023. Portfolios P1–P5 represent quintiles, with P5 comprising the top 20% currencies by *DISR*. The high-minus-low (HML) portfolio is formed by taking a long position in P5 and shorting P1. Reported metrics include monthly returns (in percentages), autocorrelation (AC1), annualized Sharpe ratios, and maximum drawdowns. Newey-West t-ratios are shown in parentheses.

	P1	P2	P3	P4	P5	HML
Average Return	-0.04	0.03	0.07	0.10	0.27	0.31
t stat	(-0.31)	(0.20)	(0.45)	(0.71)	(1.80)	(2.78)
Median Return	0.09	0.19	0.08	0.02	0.22	0.32
SD	2.44	2.31	2.16	2.22	2.34	2.13
Skewness	-0.63	-0.47	-0.35	0.21	-0.04	0.01
Kurtosis	3.79	2.00	0.73	2.28	2.80	2.19
AC1	-0.01	0.07	0.12	0.01	0.10	-0.04
Sharpe Ratio	-0.06	0.04	0.10	0.15	0.39	0.51
Max Drawdown	0.28	0.25	0.21	0.17	0.19	0.12

We implement a monthly rebalanced portfolio sorting methodology based on *DISR* values. At the end of each month following the initial training period, currencies are sorted into quintile portfolios according to their *DISR* rankings. Table 7 presents summary statistics for *DISR*-sorted quintile currency portfolios over the out-of-sample period from 2000 to 2023. Portfolio P5 represents the top 20% of currencies ranked by *DISR*, while P1 corresponds to the bottom 20%. The results indicate that average portfolio returns exhibit a monotonic increase across quintiles, rising from P1 to

P5. Furthermore, the annualized Sharpe ratio improves significantly, increasing from -0.06 for P1 to 0.39 for P5. These findings underscore the potential utility of *DISR*-based portfolio sorting in enhancing risk-adjusted returns.

The high-minus-low (*DISR-HML*) strategy involves taking a long position in P5 and shorting P1. This strategy yields a statistically significant monthly average return of 0.31%, with a t-statistic of 2.78, exceeding the typical threshold for significance. Additionally, the strategy produces a high Sharpe ratio of 0.51. Moreover, the annualized maximum drawdown of the *DISR-HML* strategy is relatively modest at approximately 12%, which helps mitigate regime risk. The investment performance metrics highlight the robustness of abnormal returns produced by the *DISR-HML* strategy.

It is important to evaluate the marginal contribution of regime risk relative to other common currency risk factors. Table 8 presents the results from factor-spanning regressions of the *DISR-HML* strategy on commonly used currency factors. We regress the regime risk anomaly returns on three models: (i) Dollar and Carry, (ii) a four-factor model including Dollar, Carry, Momentum, and Value (Nucera et al., 2024), and (iii) a comprehensive model incorporating all factors considered in our study. Across all models, the alpha estimates are statistically significant and economically meaningful, ranging from 0.29% to 0.33% per month. Notably, the four-factor model achieves an R^2 of only 0.6%, indicating limited explanatory power for the time variation in the regime risk anomaly. When all factors are included, almost all factor loadings are statistically insignificant (except Term Spread), and the adjusted R^2 remains low at 5.7%. These findings underscore the distinct nature of the regime risk anomaly, which cannot be explained by established currency factors.

In summary, these findings show that changes in macro-instrumented regimes related to risk premia and volatility introduce a new type of undiversified risk — regime risk — into the currency market's risk-return dynamics. This risk arises from shifts in macroeconomic states, as identified by our model. Currencies more sensitive to macroeconomic fluctuations are expected to achieve superior risk-adjusted returns as

Table 8: **Testing the Regime Risk Anomaly**

The table reports alphas (%) and factor loadings from regressions of regime risk anomaly returns on standard currency factors over 2000–2023. We examine three models: (1) Dollar and Carry, (2) Dollar, Carry, Momentum, and Value, and (3) a comprehensive model including all factors studied. All factors are constructed using high-minus-low (5-1) portfolios based on established characteristics, except for DOL, which represents the cross-sectional mean return (see Table IA.3 in the Internet Appendix). Monthly returns are in percentages, with Newey-West t-ratios in parentheses.

	DOL + Carry	DOL+Carry +MOM1+Value	All factors
Alpha	0.308	0.287	0.334
	(2.26)	(2.30)	(2.75)
DOL	-0.088	-0.071	-0.331
	(-0.65)	(-0.54)	(-1.57)
Carry	0.061	0.069	0.150
	(0.69)	(0.80)	(0.92)
MOM1		0.032	-0.002
		(0.29)	(-0.02)
Value		-0.110	-0.129
		(-0.96)	(-1.16)
MOM6			-0.039
			(-0.28)
MOM12			0.158
			(1.20)
NFA			0.064
			(0.37)
LDC			-0.043
			(-0.36)
LT Yield			-0.194
			(-1.13)
Term Spread			-0.327
			(-2.43)
Volatility			0.124
			(0.76)
MKT_Beta			-0.134
			(-0.90)
Adj R^2 (%)	0.0	0.6	5.7

compensation for their exposure to regime-switching risk. For instance, Mueller et al. (2017) propose a strategy of shorting the U.S. dollar while holding long positions in other currencies. Their approach generates higher excess returns, particularly around monetary policy announcements, reflecting compensation for event risk.

Our findings provide empirical evidence that macroeconomic regime-switching risk is often overlooked in currency pricing research, despite its significant implications for understanding currency returns. To address this gap, we introduce a regime-switching model that quantifies this uncertainty and offers a fresh perspective on the role of macroeconomic states. However, identifying the key macroeconomic states driving currency returns remains challenging due to the dynamic and time-varying nature of regime shifts. These monthly regime updates, tested out-of-sample in our model, highlight the inherent variability in macroeconomic states.

6 Conclusion

Macroeconomic fundamentals are widely recognized as key determinants of exchange rates (e.g., Lustig et al., 2014). However, their influence on currency risk-return factors remains underexplored. This study introduces a novel regime-switching framework to investigate how U.S. macroeconomic fundamentals impact currency pricing factor models. The framework endogenously detects macroeconomic regimes and estimates regime-specific models, offering new insights into the interaction between macroeconomic states and currency risk premia.

Our findings offer novel insights into the interplay between macroeconomic fundamentals and currency pricing: (i) Incorporating regime shifts significantly improves the statistical fit of currency pricing factor models, highlighting its importance in capturing return dynamics. (ii) The detected regimes and their transitions align closely with U.S. macroeconomic states, particularly reflecting the interaction between inflation and interest rates. (iii) Regime-dependent shifts in factor selection and risk premia exhibit pronounced time variation in characteristics-based models. (iv) The Carry factor demonstrates robustness across regimes, while the importance of Value and Momentum factors varies. (v) High-inflation regimes are linked to currency mispricing and elevated idiosyncratic volatility.

Finally, our analysis uncovers a novel regime risk anomaly driven by cross-sectional differences in currencies' sensitivities to macroeconomic regime shifts. We construct

a long-short portfolio that takes long positions in currencies with high exposure to regime-switching risk and short positions in those with low exposure. This portfolio generates significant excess returns unexplained by standard currency factor models, highlighting regime-switching risk as a distinct factor in currency markets.

References

Akinci, Ö. and A. Queralto (2024). Exchange rate dynamics and monetary spillovers with imperfect financial markets. *Review of Financial Studies 37*(2), 309–355.

Albagli, E., L. Ceballos, S. Claro, and D. Romero (2019). Channels of US monetary policy spillovers to international bond markets. *Journal of Financial Economics 134*(2), 447–473.

Ang, A. and G. Bekaert (2002). Regime switches in interest rates. *Journal of Business & Economic Statistics 20*(2), 163–182.

Ang, A. and A. Timmermann (2012). Regime changes and financial markets. *Annu. Rev. Financ. Econ. 4*(1), 313–337.

Asness, C., T. Moskowitz, and L. Pedersen (2013). Value and Momentum Everywhere. *Journal of Finance 68*, 929–985.

Baker, S. R., N. Bloom, and S. J. Davis (2016). Measuring economic policy uncertainty. *Quarterly Journal of Economics 131*(4), 1593–1636.

Bakshi, G. and G. Panayotov (2013). Predictability of currency carry trades and asset pricing implications. *Journal of Financial Economics 110*(1), 139–163.

Bartram, S. M., L. Djuranovik, A. Garratt, and Y. Xu (2025). Mispricing and risk premia in currency markets. *Journal of Financial and Quantitative Analysis 60*(2), 695–733.

Bekaert, G. (1995). The time variation of expected returns and volatility in foreign-exchange markets. *Journal of Business & Economic Statistics 13*(4), 397–408.

Bekaert, G. and G. Panayotov (2020). Good Carry, Bad Carry. *Journal of Financial and Quantitative Analysis 55*(4), 1063–1094.

Bessembinder, H., A. P. Burt, and C. M. Hrdlicka (2024). More factors matter and factors matter more than you might think: The role of time variation in factor premia. Technical report, Arizona State University.

Bie, S., F. X. Diebold, J. He, and J. Li (2024). Machine Learning and the Yield Curve: Tree-Based Macroeconomic Regime Switching. Technical report, City University of Hong Kong.

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

Campbell, J. Y. and T. Vuolteenaho (2004). Inflation illusion and stock prices. *American Economic Review 94*(2), 19–23.

Chernov, M., M. Dahlquist, and L. A. Lochstoer (2023). Pricing Currency Risks. *Journal of Finance 78*, 693–730.

Chib, S. (1995). Marginal likelihood from the Gibbs output. *Journal of the American Statistical Association 90*(432), 1313–1321.

Chipman, H. A., E. I. George, and R. E. McCulloch (2010). BART: Bayesian additive regression trees. *Annals of Applied Statistics 4*, 266–298.

Christiansen, C., A. Ranaldo, and P. Söderlind (2011). The time-varying systematic risk of carry trade strategies. *Journal of Financial and Quantitative Analysis 46*(4), 1107–1125.

Clarida, R. H., L. Sarno, M. P. Taylor, and G. Valente (2003). The out-of-sample success of term structure models as exchange rate predictors: a step beyond. *Journal of International Economics 60*(1), 61–83.

Colacito, R., S. Riddiough, and L. Sarno (2020). Business Cycles and Currency Returns. *Journal of Financial Economics 137*(3), 659–678.

Cong, L., G. Feng, J. He, and X. He (2025). Growing the Efficient Frontier on Panel Trees. *Journal of Financial Economics 167*, 104024.

Cong, L. W., G. Feng, J. He, and J. Li (2023). Sparse Modeling Under Grouped Heterogeneity with an Application to Asset Pricing. Technical report, National Bureau of Economic Research.

Creal, D. and J. Kim (2021). Empirical Asset Pricing with Bayesian Regression Trees. Technical report, University of Notre Dame.

Cui, L., G. Feng, and Y. Hong (2024). Regularized gmm for time-varying models with applications to asset pricing. *International Economic Review 65*(2), 851–883.

Cui, L., G. Feng, Y. Hong, and J. Yang (2023). Time-Varying Factor Selection: A Sparse Fused GMM Approach. Technical report, City University of Hong Kong.

Dahlquist, M. and H. Hasseltoft (2020). Economic Momentum and Currency Returns. *Journal of Financial Economics 136*(1), 152–167.

Dahlquist, M. and J. Pénasse (2022). The missing risk premium in exchange rates. *Journal of Financial Economics 143*(2), 697–715.

Della Corte, P., R. Kozhan, and A. Neuberger (2021). The cross-section of currency volatility premia. *Journal of Financial Economics 139*(3), 950–970.

Della Corte, P., T. Ramadorai, and L. Sarno (2016). Volatility Risk Premia and Exchange Rate Predictability. *Journal of Financial Economics 120*, 21–40.

Della Corte, P., S. J. Riddiough, and L. Sarno (2016). Currency Premia and Global Imbalances. *Review of Financial Studies 29*, 2161–2193.

Engel, C. and J. D. Hamilton (1990). Long Swings in the Dollar: Are They in the Data and Do Markets Know It? *American Economic Review 80*(4), 689–713.

Fama, E. F. and K. R. French (2020). Comparing cross-section and time-series factor models. *Review of Financial Studies 33*(5), 1891–1926.

George, E. I. and R. E. McCulloch (1993). Variable selection via Gibbs sampling. *Journal of the American Statistical Association 88*(423), 881–889.

Giannone, D., M. Lenza, and G. E. Primiceri (2021). Economic predictions with big data: The illusion of sparsity. *Econometrica 89*(5), 2409–2437.

Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 357–384.

He, J. and P. R. Hahn (2023a). Stochastic tree ensembles for regularized nonlinear regression. *Journal of the American Statistical Association 118*(541), 551–570.

He, J. and P. R. Hahn (2023b). Stochastic tree ensembles for regularized nonlinear regression. *Journal of the American Statistical Association 118*(541), 551–570.

He, J., S. Yalov, and P. R. Hahn (2019). XBART: Accelerated Bayesian additive regression trees. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 1130–1138.

He, W., Z. Su, and J. Yu (2024). Macroeconomic perceptions, financial constraints, and anomalies. *Journal of Financial Economics 162*, 103952.

Johnson, T. C. (2002). Volatility, momentum, and time-varying skewness in foreign exchange returns. *Journal of Business & Economic Statistics 20*(3), 390–411.

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.

Kho, B.-C. (1996). Time-varying risk premia, volatility, and technical trading rule profits: Evidence from foreign currency futures markets. *Journal of Financial Economics 41*(2), 249–290.

Kim, S., R. A. Korajczyk, and A. Neuhierl (2021). Arbitrage portfolios. *Review of Financial Studies 34*(6), 2813–2856.

Koijen, R. S., T. J. Moskowitz, L. H. Pedersen, and E. B. Vrugt (2018). Carry. *Journal of Financial Economics 127*(2), 197–225.

Lettau, M., M. Maggiori, and M. Weber (2014). Conditional Risk Premia in Currency Markets and Other Asset Classes. *Journal of Financial Economics 114*, 197–225.

Li, J., L. Sarno, and G. Zinna (2024). Skewness Risk Premia and the Cross-Section of Currency Returns. Technical report, Fudan University.

Li, S. Z., P. Yuan, and G. Zhou (2023). Pockets of Factor Pricing. Technical report, Washington University in St. Louis.

Liu, S., T. Maurer, A. Vedolin, and Y. Zhang (2024). Dollar and Carry Redux. Technical report, University of Hong Kong.

Liu, X. and R. Chen (2020). Threshold factor models for high-dimensional time series. *Journal of Econometrics 216*(1), 53–70.

Lustig, H., N. Roussanov, and A. Verdelhan (2011). Common risk factors in currency markets. *Review of Financial Studies 11*, 527–553.

Lustig, H., N. Roussanov, and A. Verdelhan (2014). Countercyclical Currency Risk Premia. *Journal of Financial Economics 11*, 527–553.

Lustig, H., A. Stathopoulos, and A. Verdelhan (2019). The term structure of currency carry trade risk premia. *American Economic Review 109*(12), 4142–4177.

Lustig, H. and A. Verdelhan (2007). The cross section of foreign currency risk premia and consumption growth risk. *American Economic Review 97*(1), 89–117.

Massacci, D. (2017). Least squares estimation of large dimensional threshold factor models. *Journal of Econometrics 197*(1), 101–129.

Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf (2012a). Carry trades and global FX volatility. *Journal of Finance 67*, 681–718.

Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf (2012b). Currency Momentum Strategies. *Journal of Financial Economics 106*, 620–684.

Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf (2017). Currency Value. *Review of Financial Studies 30*, 416–441.

Miao, J., Z. Shen, and P. Wang (2019). Monetary policy and rational asset price bubbles: Comment. *American Economic Review 109*(5), 1969–1990.

Miranda-Agrippino, S. and H. Rey (2020). US monetary policy and the global financial cycle. *Review of Economic Studies 87*(6), 2754–2776.

Mueller, P., A. Tahbaz-Salehi, and A. Vedolin (2017). Exchange rates and monetary policy uncertainty. *Journal of Finance 72*(3), 1213–1252.

Nucera, F., L. Sarno, and G. Zinna (2024). Currency risk premiums redux. *Review of Financial Studies 37*(2), 356–408.

Rapach, D. E., J. K. Strauss, and G. Zhou (2013). International stock return predictability: What is the role of the United States? *Journal of Finance 68*(4), 1633–1662.

Smith, S. C. and A. Timmermann (2021). Break risk. *Review of Financial Studies 34*(4), 2045–2100.

Smith, S. C. and A. Timmermann (2022). Have risk premia vanished? *Journal of Financial Economics 145*(2), 553–576.

Vavra, J. (2014). Inflation dynamics and time-varying volatility: New evidence and an ss interpretation. *Quarterly Journal of Economics 129*(1), 215–258.

Verdelhan, A. (2018). The share of systematic variation in bilateral exchange rates. *Journal of Finance 73*(1), 375–418.

Appendix

A.I Derivation of Marginal Likelihood

For any model M, data y, and parameter vector θ, from Bayes rule:

$$p(y, \theta \mid M) = p(y \mid \theta, M)p(\theta \mid M) = p(\theta \mid y, M)p(y \mid M)$$

and thus

$$p(y \mid M) = \frac{p(y \mid \theta, M)p(\theta \mid M)}{p(\theta \mid y, M)}$$

Taking logs of both sides, we obtain

$$\ln p(y \mid M) = \ln p(y \mid \theta, M) + \ln p(\theta \mid M) - \ln p(\theta \mid y, M)$$

The log marginal likelihood on the left-hand side does not involve θ, and so the equation must hold for any value of θ. Thus, for purposes of computational reliability, let us choose a $\widehat{\theta}$ of high posterior density, say, posterior mean:

$$\ln p(y \mid M) = \ln p(y \mid \widehat{\theta}, M) + \ln p(\widehat{\theta} \mid M) - \ln p(\widehat{\theta} \mid y, M)$$

The first two terms on the right-hand side of the equality are readily available. That is, $\ln p(y \mid \widehat{\theta}, M)$ is the log likelihood evaluated at $\widehat{\theta}$ and $\ln p(\widehat{\theta} \mid M)$ is the log prior ordinate at $\widehat{\theta}$. The last term $\ln p(\widehat{\theta} \mid y, M)$ requires more care, since the normalizing constant of the posterior is typically unknown. We can evaluate it from the Gibbs sampling output from Chib (1995).

Our parameter set $\boldsymbol{\theta} = [\boldsymbol{a}, \boldsymbol{\lambda}, \boldsymbol{\sigma}, \boldsymbol{\gamma}]$. First, write

$$p(\widehat{\boldsymbol{\theta}} \mid y) = p(\widehat{\boldsymbol{\sigma}} \mid y)p(\widehat{\boldsymbol{a}} \mid \widehat{\boldsymbol{\sigma}}, y)p(\widehat{\boldsymbol{\gamma}} \mid \widehat{\boldsymbol{a}}, \widehat{\boldsymbol{\sigma}}, y)p(\widehat{\boldsymbol{\lambda}} \mid \widehat{\boldsymbol{\gamma}}, \widehat{\boldsymbol{a}}, \widehat{\boldsymbol{\sigma}}, y)$$

$p(\widehat{\boldsymbol{\lambda}} \mid \widehat{\boldsymbol{\gamma}}, \widehat{\boldsymbol{a}}, \widehat{\boldsymbol{\sigma}}, y)$ is known, as it is simply an ordinate of the complete conditional for $\boldsymbol{\lambda}$,

$$p(\widehat{\boldsymbol{\lambda}} \mid \widehat{\boldsymbol{\gamma}}, \widehat{\boldsymbol{a}}, \widehat{\boldsymbol{\sigma}}, y) = \mathcal{MVN}\left(\widehat{\boldsymbol{V}}_{\lambda,j}\left(\sum_{i=1}^{N} \widehat{\sigma}_{j,i}^{-2}\mathbf{Z}_{j,i}^{T}\left(\mathbf{R}_{j,i} - \widehat{\alpha}_{j,i}\mathbf{1}\right)\right), \left(\sum_{i=1}^{N} \widehat{\sigma}_{j,i}^{-2}\mathbf{Z}_{j,i}^{T}\mathbf{Z}_{j,i} + V_0^{-1}\right)^{-1}\right)$$

for rest parameters, approximate the density from R Gibbs samples:

$$\begin{aligned}
p(\widehat{\boldsymbol{\sigma}} \mid y) &\approx \frac{1}{R}\sum_{r=1}^{R} p(\widehat{\boldsymbol{\sigma}} \mid \boldsymbol{a}^{(r)}, \boldsymbol{\lambda}^{(r)}, \boldsymbol{\gamma}^{(r)}, y) \\
&= \frac{1}{R}\sum_{r=1}^{R}\prod_{i=1}^{N} IG\left(v_i, \frac{1}{2}(||\mathbf{R}_{j,i} - \alpha_{j,i}^{(r)}\mathbf{1} - \mathbf{Z}_{j,i}\boldsymbol{\lambda}_{j,i}^{(r)}||_2^2 + \frac{\alpha_{j,i}^{2(r)}}{\sigma_\alpha^2} + S_0)\right) \\
p(\widehat{\boldsymbol{a}} \mid \widehat{\boldsymbol{\sigma}}, y) &\approx \frac{1}{R}\sum_{r=1}^{R} p(\widehat{\boldsymbol{a}} \mid \boldsymbol{\gamma}^{(r)}, \boldsymbol{\lambda}^{(r)}, \widehat{\boldsymbol{\sigma}}, y) \\
&= \frac{1}{R}\sum_{r=1}^{R}\prod_{i=1}^{N} \mathcal{N}\left(\mu_{\alpha,j,i}^{-1}\mathbf{1}^{\mathsf{T}}(\mathbf{R}_{j,i} - \mathbf{Z}_{j,i}\boldsymbol{\lambda}_j^{(r)}), \frac{\widehat{\sigma}_{j,i}^2}{\mu_{\alpha,j,i}}\right) \\
p(\widehat{\boldsymbol{\gamma}} \mid \widehat{\boldsymbol{a}}, \widehat{\boldsymbol{\sigma}}, y) &\approx \frac{1}{R}\sum_{r=1}^{R} p(\widehat{\boldsymbol{\gamma}} \mid \boldsymbol{\lambda}^{(r)}, \widehat{\boldsymbol{a}}, \widehat{\boldsymbol{\sigma}}, y) \\
&= \frac{1}{R}\sum_{r=1}^{R} \textsc{Ber}\left(\frac{\xi_1^{-1}\exp\left(-\frac{1}{2\xi_1^2}\lambda_{j,k}^{2(r)}\right)w_k}{\xi_1^{-1}\exp\left(-\frac{1}{2\xi_1^2}\lambda_{j,k}^{2(r)}\right)w_k + \xi_0^{-1}\exp\left(-\frac{1}{2\xi_0^2}\lambda_{j,k}^{2(r)}\right)(1-w_k)}\right),
\end{aligned}$$

where $\widehat{\boldsymbol{a}}, \widehat{\boldsymbol{\lambda}}, \widehat{\boldsymbol{\gamma}}, \widehat{\boldsymbol{\sigma}}$ are the corresponding posterior mean from Gibbs output, and $\boldsymbol{a}^{(r)}$, $\boldsymbol{\lambda}^{(r)}$, $\boldsymbol{\gamma}^{(r)}$, $\boldsymbol{\sigma}^{(r)}$ are the r-th post-convergence Gibbs sample from additional reduced runs. Conditional posterior distributions are given in Appendix A.III.

A.II Early Stop Criterion and Further Split

To prevent overfitting, we consider the possibility of definitively stopping the splitting process at the current node and terminating the algorithm, which we call the *null split rule* option. If the current node does not split, all the data will stay in the same leaf node and fit one model. Therefore, the split criterion for the null split rule is

$$l(\emptyset) = |\mathcal{C}|\left(\frac{(1+d)^{\tilde{a}_2}}{\tilde{a}_1} - 1\right)p(\mathbf{R} \mid \mathbf{Z}) = |\mathcal{C}|\left(\frac{(1+d)^{\tilde{a}_2}}{\tilde{a}_1} - 1\right)p(\mathcal{R}_0), \tag{A.1}$$

where $|\mathcal{C}|$ is the total number of all candidates, and d is the depth of the current node. Note that the root node has a depth of 1, and nodes that are one extra level beneath it have one more depth. The extra term before the marginal likelihood controls the strength of the regularization of the tree size. Specifically, $\tilde{a}_1$ and $\tilde{a}_2$ are two prior hyperparameters for the tree shrinkage. Following the XBART framework of He et al. (2019) and He and Hahn (2023b) and the BART of Chipman et al. (2010), we choose $\tilde{a}_1 = 0.95$ and $\tilde{a}_2 = 2$. Intuitively, as $\tilde{a}_2 = 2$, the penalty of growing the tree (adding one extra depth) increases exponentially with the current depth. Deeper nodes tend to stop, and thus, regularization plays a role.

After calculating the split criterion for all candidates and the stop option, we assume that the prior probability of each option is proportional to 1. Therefore, the posterior of selecting one of them follows the Bayes rule as

$$L(c_j) = \frac{l(c_j)}{\sum_{j'} l(c_{j'}) + l(\emptyset)}, \quad L(\emptyset) = \frac{l(\emptyset)}{\sum_{j'} l(c_{j'}) + l(\emptyset)}. \tag{A.2}$$

Eq. (A.2) guides the selection of the best option (either a split candidate or a stop option) at each iteration. The option that maximizes the criterion is chosen. The algorithm proceeds to the next split if the selected option is a split rule candidate. If it is the stop option, the current node ceases to split further, and the algorithm concludes.

The second split follows a similar process and can occur at either the left or right child nodes resulting from the first split. The tree grows iteratively, considering split candidates for both nodes together. The set of candidates expands to $\mathcal{C} = \{c_k^{\mathcal{R}_1}\} \cup \{c_k^{\mathcal{R}_2}\}$, where the superscript indicates the node to be split.

Each candidate, whether at the root's left or right child nodes, generates two new leaf nodes, resulting in three. The joint marginal likelihood or split criterion, corresponding to Eq. (4), is:

$$l(c_j^{\mathcal{R}_1}) = p(\mathcal{R}_3) \times p(\mathcal{R}_4) \times p(\mathcal{R}_2), \quad l(c_j^{\mathcal{R}_2}) \quad = p(\mathcal{R}_1) \times p(\mathcal{R}_5) \times p(\mathcal{R}_6). \tag{A.3}$$

The split criterion of the stop option is[15]

$$l(\emptyset^{\mathcal{R}_1}) = |\mathcal{C}| \left(\frac{(1+d)^{\tilde{a}_2}}{\tilde{a}_1} - 1 \right) p(\mathcal{R}_1) \times p(\mathcal{R}_2), \quad l(\emptyset^{\mathcal{R}_2}) = |\mathcal{C}| \left(\frac{(1+d)^{\tilde{a}_2}}{\tilde{a}_1} - 1 \right) p(\mathcal{R}_2) \times p(\mathcal{R}_1),$$

where $\emptyset^{\mathcal{R}_1}$ and $\emptyset^{\mathcal{R}_2}$ denote the option of stopping split $\mathcal{R}_1$ or $\mathcal{R}_2$, respectively.

The calculation for all candidates in $\mathcal{C} = \{c_j^{\mathcal{R}_1}\} \cup \{c_j^{\mathcal{R}_2}\}$ and two stop options proceed similarly as in Eq. (A.2),

$$\begin{aligned} W &= \sum_{j'} l(c_{j'}^{\mathcal{A}_1}) + \sum_{j'} l(c_{j'}^{\mathcal{A}_2}) + l(\emptyset^{\mathcal{A}_1}) + l(\emptyset^{\mathcal{A}_2}) \\ L(c_j^{\mathcal{A}_1}) &= \frac{l(c_j^{\mathcal{A}_1})}{W}, L(\emptyset^{\mathcal{A}_1}) = \frac{l(\emptyset^{\mathcal{A}_1})}{W}, L(c_j^{\mathcal{A}_2}) = \frac{l(c_j^{\mathcal{A}_2})}{W}, L(\emptyset^{\mathcal{A}_2}) = \frac{l(\emptyset^{\mathcal{A}_2})}{W}. \end{aligned} \tag{A.4}$$

We select the option that maximizes Eq. (A.4) as the second split. Unlike typical tree algorithms in machine learning that recursively define the split criterion based on data within the node, our criterion is defined *globally*. Eq. (A.3) encompasses all resulting leaf nodes and the entire dataset. This approach prevents myopic local optimization and mitigates overfitting.

A.III Gibbs sampler

The posterior inference of parameters for the characteristics-based model within each regime is conducted using Markov Chain Monte Carlo (MCMC) methods with a Gibbs sampler. Representing the stacked data within regime j of individual currency i ($1 \leq i \leq N$) in matrix form $\mathbf{R}_{j,i}$ (cross-sectionally demeaned return) and $\mathbf{Z}_{j,i}$, we have

$$\mathbf{R}_{j,i} = \alpha_{j,i}\mathbf{1} + \mathbf{Z}_{j,i}\boldsymbol{\lambda}_j + \boldsymbol{\epsilon}_{j,i},$$

the full conditionals are given as follows,

1. Update $\alpha_{j,i}$. Different from the slope coefficient λ_j, the intercept differs across individual currencies, indicating different levels of individual mispricing. For

[15] $d = 2$ for the second split, since nodes $\mathcal{R}_1$ or $\mathcal{R}_2$ has depth 2.

each j and $i = 1, ..., N$,

$$\alpha_{j,i} \mid \boldsymbol{\lambda_j}, \sigma^2_{j,i}, \boldsymbol{\gamma}_j, \mathbf{R}_{j,i}, \mathbf{Z}_{j,i} \sim \mathcal{N}\left(\tilde{\alpha}_{j,i}, \frac{\sigma^2_{j,i}}{\mu_{\alpha,j,i}}\right)$$

where $\mu_{\alpha,j,i} = \mathbf{1}^T\mathbf{1} + \frac{1}{\sigma^2_\alpha}$ and $\tilde{\alpha}_{j,i} = \mu^{-1}_{\alpha,j,i}\mathbf{1}^\mathsf{T}(\mathbf{R}_{j,i} - \mathbf{Z}_{j,i}\boldsymbol{\lambda}_j)$.

2. Update $\boldsymbol{\lambda}_j = \{\lambda_{j,k}\}_{k=1,\cdots,K}$. Note that this step updates the regression coefficients of all variables. For $i = 1, ..., N$, $k = 1, 2, \cdots, K$,

$$\boldsymbol{\lambda}_j \mid \alpha_{j,i}, \sigma^2_{j,i}, \boldsymbol{\gamma}_j, \mathbf{R}_{j,i}, \mathbf{Z}_{j,i} \sim \mathcal{MVN}\left(\tilde{\boldsymbol{\lambda}}_j, \boldsymbol{V}_{\lambda,j}\right),$$

where $\boldsymbol{V}_{\lambda,j} = \left(\sum_{i=1}^N \sigma^{-2}_{j,i}\mathbf{Z}^T_{j,i}\mathbf{Z}_{j,i} + V_0^{-1}\right)^{-1}$, V_0^{-1} is a diagonal matrix where the corresponding k-th diagonal element is $\xi^2_{\gamma,k}$, and $\tilde{\boldsymbol{\lambda}}_j = \boldsymbol{V}_{\lambda,j}\left(\sum_{i=1}^N \sigma^{-2}_{j,i}\mathbf{Z}^T_{j,i}\left(\mathbf{R}_{j,i} - \alpha_{j,i}\mathbf{1}\right)\right)$.

3. Update $\boldsymbol{\gamma}_{j,k}$, for $k = 1, 2, \cdots, K$.

$$\gamma_{j,k} \mid \boldsymbol{\lambda}_j, \mathbf{R}_j, \mathbf{Z}_j \sim \textsc{Ber}\left(\frac{\xi_1^{-1}\exp\left(-\frac{1}{2\xi_1^2}\lambda^2_{j,k}\right)w_k}{\xi_1^{-1}\exp\left(-\frac{1}{2\xi_1^2}\lambda^2_{j,k}\right)w_k + \xi_0^{-1}\exp\left(-\frac{1}{2\xi_0^2}\lambda^2_{j,k}\right)(1-w_k)}\right).$$

4. Update $\sigma^2_{j,i}$, for $i = 1, ..., N$

$$\sigma^2_{j,i} \mid \alpha_{j,i}, \boldsymbol{\lambda}_j, \boldsymbol{\gamma}_j, \mathbf{R}_{j,i}, \mathbf{Z}_{j,i} \sim IG(v_i, S_i),$$

where $v_i = \frac{1}{2}(T_i + 1 + v_0)$, $S_i = \frac{1}{2}\left(||\mathbf{R}_{j,i} - \alpha_{j,i}\mathbf{1} - \mathbf{Z}_{j,i}\boldsymbol{\lambda}_{j,i}||_2^2 + \frac{\alpha^2_{j,i}}{\sigma^2_\alpha} + S_0\right)$, and T_i is number of data observations for i-th currency in the leaf node.

Internet Appendices for "Currency Return Dynamics: What Is the Role of U.S. Macroeconomic Regimes?"

IA.I Simulation

Our Gibbs sampler, described in Appendix A.III, employs a modified spike-and-slab prior to estimate alpha and idiosyncratic volatility for individual assets while simultaneously identifying consistent risk premia across the cross section. This approach improves upon conventional methods, such as those by George and McCulloch (1993), which typically sample along a single dimension and overlook cross-sectional dependencies. To assess the effectiveness of our method, we conduct a comprehensive simulation study.

The cross section of asset returns is designed to replicate empirical properties observed in 47 currencies. Specifically, we consider $K = 10$ characteristics, $N = 50$ assets, and $T = 600$ time periods. Our model incorporates three types of factors: strong, weak, and irrelevant. Strong factors exhibit large risk premia, weak factors have smaller risk premia, and irrelevant factors contribute no risk premia.

The data-generating process is as follows. We fix the risk premia of 10 factors at $\lambda = [0.5, 0.4, 0.2, 0.1, 0.07, 0.05, 0.02, 0.01, 0, 0]$. The first three factors are classified as strong contributors, the next five as weak contributors, and the last two as irrelevant. For each currency i, we simulate α_i, σ_i, and ten characteristic levels denoted by $\bar{z}_i$.

$$\begin{aligned} \alpha_i &\sim N(0, 0.01^2) \\ \sigma_i &\sim |N(0, 0.5^2)| \\ \bar{z}_i &\sim N(0, 0.5^2) \end{aligned} \tag{IA.1}$$

At each time t, we generate the cross-sectional characteristics $z_{i,t}$, rank them to ensure they are evenly distributed within the range $[-1, 1]$, and subsequently generate the residual term $\epsilon_{i,t}$.

$$\mathbf{z}_{i,t-1} \sim N(\bar{z}_i, 0.1^2), \quad \epsilon_{i,t} \sim N(0, \sigma_i^2) \tag{IA.2}$$

By multiplying the factor loadings (characteristics) with the risk premia at time t, and subsequently adding the intercept term and the innovation component, one obtains the simulated returns denoted by $R_{i,t}$ for asset i at time t.

$$R_{i,t} = \alpha_i + \sum_{i=1}^{K} \lambda_i \times z_{i,t-1} + \epsilon_{i,t} \tag{IA.3}$$

In total, we analyze 600×50 currency return observations. Each MCMC run consists of 2,000 samples, and posterior inference for selection probabilities is based on the sample mean. The simulation is repeated 100 times, and the results are reported as averages.

Table IA.1 presents the posterior factor premia estimates and selection probabilities from 100 simulation runs. The results indicate that the modified spike-and-slab Gibbs sampler effectively estimates factor premia with high accuracy, demonstrating its robustness in identifying relevant factors. Strong factors exhibit selection probabilities consistently close to 1, while weak factors show probabilities varying between 0 and 1. Useless factors are consistently unselected, with probabilities remaining near 0. These findings highlight the methodology's ability to differentiate between strong, weak, and irrelevant factors, providing a reliable approach for factor selection in financial models.

References

Della Corte, P., S. J. Riddiough, and L. Sarno (2016). Currency Premia and Global Imbalances. *Review of Financial Studies 29*, 2161–2193.

Frazzini, A. and L. H. Pedersen (2014). Betting against beta. *Journal of Financial Economics 111*(1), 1–25.

George, E. I. and R. E. McCulloch (1993). Variable selection via Gibbs sampling. *Journal of the American Statistical Association 88*(423), 881–889.

Lettau, M., M. Maggiori, and M. Weber (2014). Conditional Risk Premia in Currency Markets and Other Asset Classes. *Journal of Financial Economics 114*, 197–225.

Table IA.1: **Simulation Results**

The table summarizes results from a simulation study estimating risk premia and selection probabilities across 100 runs. The factor structure includes three strong, five weak, and several extraneous factors. Values represent sample averages of $\boldsymbol{\lambda}_k$ and $\boldsymbol{\gamma}_k$ (for $k = 1, \ldots, 16$), based on 2000 post-convergence MCMC samples.

	True Premia (%)	Estimated Premia (%)			Selection Probability (%)		
Factors		w = 0.9	w = 0.5	w = 0.1	w = 0.9	w = 0.5	w = 0.1
			Strong Factors				
Z1	0.5	0.436	0.436	0.436	100	100	100
Z2	0.4	0.339	0.339	0.339	100	100	100
Z3	0.2	0.174	0.174	0.174	98.5	98.6	98.5
			Weak Factors				
Z4	0.1	0.079	0.079	0.079	12.1	12.1	11.9
Z5	0.07	0.056	0.056	0.056	2.2	2.1	2.2
Z6	0.05	0.039	0.039	0.038	1	1	1
Z7	0.02	0.017	0.017	0.017	0.5	0.5	0.5
Z8	0.01	0.008	0.008	0.008	0.2	0.2	0.2
			Useless Factors				
Z9	0	0.002	0.002	0.003	0.4	0.4	0.4
Z10	0	0.001	0.001	0.001	0.5	0.5	0.5

IA.II Additional Tables and Figures

Figure IA.1: **Time-Series of Macro Variables**

The figure shows eight U.S. macroeconomic indicators. The crude oil price index is displayed on a logarithmic scale. Inflation is the annual percent change of CPI. The shaded areas represent NBER recession periods. VIX is imputed by S&P 500 Index Realized Volatility in the early stage. The data source is FRED. The monthly data sample spans from 1986 to 2023.

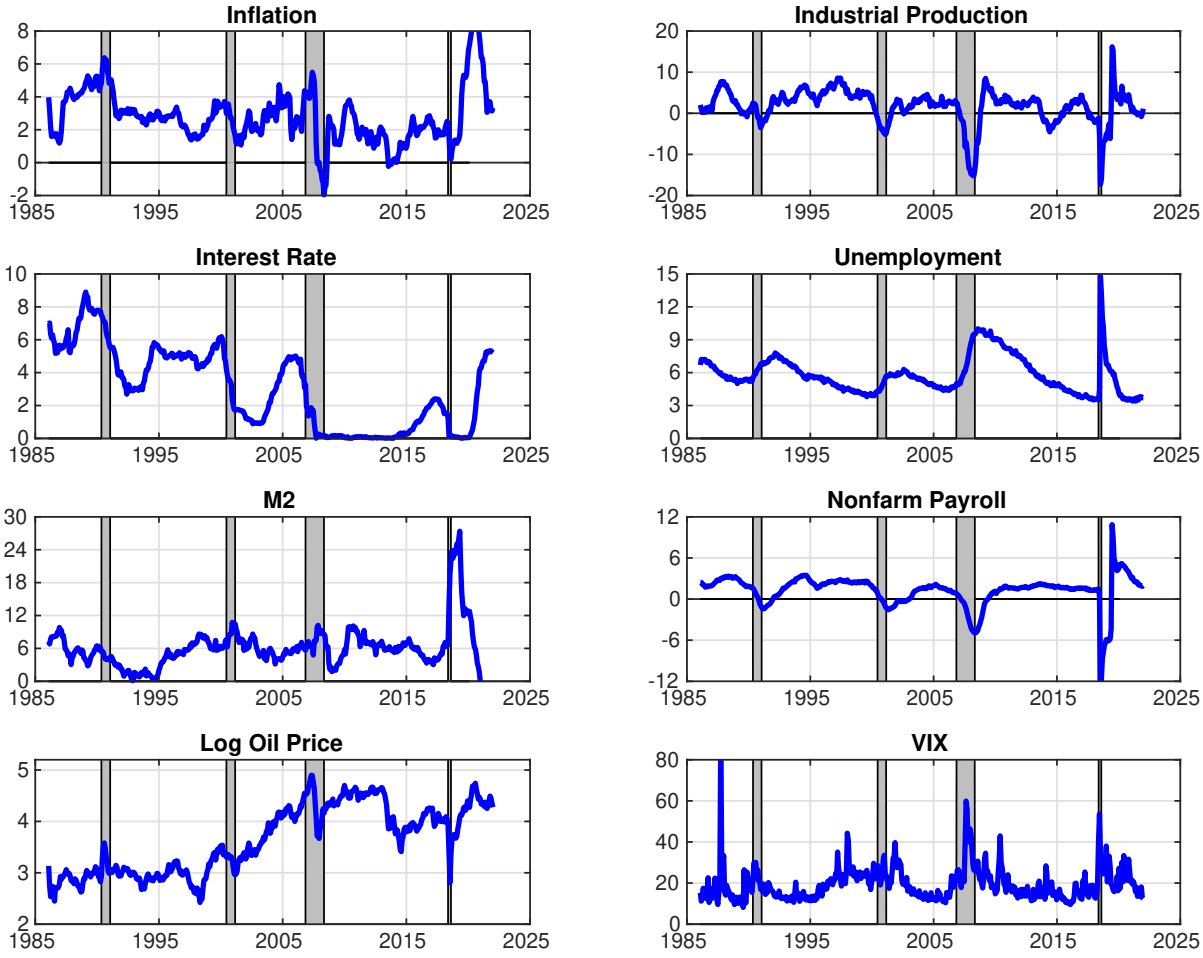


Table IA.2: **Description of Data Source**

The table displays the data source of currency characteristics. There are 11 characteristics in total (Carry, 1-, 6-, 12-month momentum, value, net foreign asset, liabilities in domestic currency, long-term yield, term spread, volatility, and market beta). The monthly data sample spans from 1986 to 2023.

Characteristics	Calculation	Source
Carry	$(S_t - F_t)/S_t$	Authors' calculation based on spot and forward exchange rate quotes (midquote) from Barclays Bank International (BBI), Reuters and WM/Reuters accessed via Datastream.
MOM1	X_t	Authors' calculation based on spot and forward exchange rate quotes (midquote) from Barclays Bank International (BBI), Reuters and WM/Reuters accessed via Datastream.
MOM6	$X_{t-6:t}$	Authors' calculation based on spot and forward exchange rate quotes (midquote) from Barclays Bank International (BBI), Reuters and WM/Reuters accessed via Datastream.
MOM12	$X_{t-12:t}$	Authors' calculation based on spot and forward exchange rate quotes (midquote) from Barclays Bank International (BBI), Reuters and WM/Reuters accessed via Datastream.
Value	-($log(S_t * CPI_t/CPI_{US,t}) - log(S_{t-5yr} * CPI_{t-5yr}/CPI_{US,t-5yr})$)	CPI data is from IMF International Financial Statistics. Only for Taiwan does CPI data come from national statistics.
NFA	-(Total assets excl. gold-Total liabilities)/GDP (in US$)	The data are kindly provided by https://www.brookings.edu/articles/the-external-wealth-of-nations-database/. For NFA are correlated with Carry, in the empirical part, we assign the top 50% NFA to -1 and the bottom 50% to 1 following the original paper Della Corte et al. (2016)
LDC	Total liabilities * Exchange rate (end of period)	The data are kindly provided by https://www.brookings.edu/articles/the-external-wealth-of-nations-database/.
LT Yield	i_{10yr}	For the 10-year rates, we use the Long-term interest rates available on OECD Monthly Monetary and Financial Statistics.
Term Spread	$-(i_{10yr} - i_{3m})$	For the 10-year (3-month) rates, we use the Long (Short)-term interest rates available on OECD Monthly Monetary and Financial Statistics.
Volatility	$vol(r_{t-3m:t})$	The currency volatility is measured by the standard deviation of daily exchange rate return in the previous 3 months
MKT_Beta	β	The currency market beta is calculated by rolling a 3-year regression of currency return on the cross-sectional mean of currency return. We use negative beta in the empirical part following Frazzini and Pedersen (2014).

Table IA.3: **Summary Statistics: All Factor Portfolios**

The table presents summary statistics for currency factors based on all characteristics outlined in Section 5. Currencies are sorted by corresponding traits, with the first portfolio (P1) comprising the bottom 20% of currencies with low characteristics and the last portfolio (P5) containing the top 20% with high characteristics. The factor represents a long-short strategy, buying P5 and selling P1. The DOL factor is the cross-sectional mean return of the currencies. Returns are reported in percent, with Newey-West t-ratios in parentheses. Additional metrics include the first-order autocorrelation coefficient (ac1), annualized Sharpe ratio, and annualized maximum drawdown. The LDC factor exhibits low mean return and Sharpe ratio due to differences in sorting methodology; Della Corte et al. (2016) sorts LDC jointly with NFA, whereas we treat them as distinct characteristics, using only the original LDC. The market beta factor shows negative mean returns over the full sample period, reflecting high volatility and conditional performance (Lettau et al., 2014). Monthly returns are expressed in percent, and the sample spans 1986 to 2023.

	DOL	Carry	MOM1	MOM6	MOM12	Value	NFA	LDC	LT Yield	Term Spread	Volatility	MKT_Beta
Mean Return	0.18	0.62	0.38	0.33	0.19	0.14	0.25	0.01	0.29	0.26	0.24	-0.06
t stat	(1.80)	(4.85)	(3.53)	(3.21)	(1.70)	(1.26)	(2.52)	(0.17)	(2.61)	(2.64)	(1.73)	(-0.39)
Median Return	0.28	0.69	0.22	0.41	0.29	0.04	0.32	0.03	0.31	0.34	0.30	-0.08
SD	1.96	2.47	2.48	2.55	2.47	2.39	2.02	1.59	2.19	1.92	2.81	3.29
Skewness	-0.31	-0.69	0.43	-0.22	-0.24	0.31	-0.28	-0.32	-0.49	-0.24	-0.09	0.16
Kurtosis	1.01	1.89	2.17	0.78	0.62	2.06	2.08	2.46	2.18	2.75	1.25	1.41
AC1	0.08	0.15	-0.04	-0.03	-0.02	0.03	0.06	-0.01	0.08	0.07	0.05	0.08
Sharpe Ratio	0.31	0.87	0.53	0.45	0.27	0.20	0.44	0.03	0.46	0.46	0.29	-0.07
Max Drawdown	0.19	0.22	0.16	0.17	0.19	0.19	0.19	0.19	0.20	0.23	0.24	0.32