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: January 14, 2025

Abstract

This paper investigates the time-varying impact of U.S. macroeconomic regimes on the risk factors driving currency return dynamics. We introduce a novel regime-switching model that endogenously identifies regimes based on thresholds of macroeconomic variables. The empirical analysis provides robust evidence of regime shifts in the currency risk-return relationship, with U.S. inflation and interest rates as crucial variables defining three economically interpretable regimes. The Carry factor is pervasive across all regimes, whereas the importance of other factors (e.g., Value and Momentum) varies between regimes. Finally, we identify a regime risk anomaly that appears related to macroeconomic regime uncertainty risk.

Keywords: Business Cycles; Currency Returns; Panel Tree; Regime Shifts; Risk Premia.

JEL Classification: F31, G12, G15.

*We thank Raymond Kan, Yongmiao Hong, Zhenya Liu (discussant), Qi Xu, Sining Liu (discussant), Yushan Tang (discussant), Tao Liu (discussant), Takeshi Osada (discussant), Ingomar Krohn (discussant) and seminar and conference participants at the City University of Hong Kong, Zhejiang University, 2024 SoFiE Annual Conference, 2024 China International Conference in Finance, 2024 AsianFA Annual Conference,5th Fintech Forum (RUC), 2024 Sydney Banking and Financial Stability Conference, 2024 Australasian Finance and Banking Conference, 2024 Big Data and Econometric Conference (XMU), 2023 Econometric Workshop for Time-varying Coefficient Models (HNU) for their valuable comments. 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 relationship in currency markets is a central topic of interest in the finance literature. Numerous studies have highlighted that several common factors, such as Carry (Lustig et al., 2011), Value (Menkhoff et al., 2017), and Momentum (Menkhoff et al., 2012b), effectively capture the dynamics of currency portfolio returns and explain their cross-sectional variation. Moreover, recent work employing latent factor approaches (see, e.g., Nucera et al., 2024; Li et al., 2024) also identifies these factors as key drivers of currency pricing. However, little, if anything, is known about whether these factors display stable explanatory power and risk premia across different macroeconomic regimes.

This gap is important, given that theories of exchange rate determination highlight the crucial role of macroeconomic fundamentals — especially those of the U.S. economy — in driving exchange rate dynamics. The U.S. economy generates substantial spillovers to global markets and induces strong co-movements in asset prices across countries (e.g., Rapach et al., 2013; Akinci and Queralto, 2024). U.S. macroeconomic conditions are particularly relevant due to their significant role in driving global economic fluctuations and the necessity for policymakers worldwide to respond to U.S.-originated shocks, such as those stemming from Federal Reserve actions. Notably, Lustig, Roussanov, and Verdelhan (2014) show that U.S. macroeconomic variables strongly predict expected currency returns, emphasizing their role in currency market predictability.

Empirical evidence highlights the significance of time variation in currency pricing. For instance, Christiansen et al. (2011) demonstrate that regime-dependent models better explain the Carry factor with time-varying systematic risk, and Accominotti et al. (2019) find that Carry returns occur exclusively within floating exchange rate regimes while being essentially zero in fixed regimes. Liu et al. (2024) show that conditional factor models outperform time-invariant models, and previous studies also document the time-varying nature of risk premia for currency factors (e.g., Kho, 1996; Lettau et al., 2014; Dahlquist and Pénasse, 2022).

Ang and Timmermann (2012) examine the impact of regime changes on equilibrium asset prices and demonstrate the existence of regimes across various asset return time series, including fixed income, equities, and currency markets. Business cycle fluctuations and macroeconomic-driven regime shifts can lead to variations in the currency return dynamics, factor risk premia, mispricing, and idiosyncratic volatility (e.g., Bekaert, 1995; Johnson, 2002; Colacito et al., 2020). Although regime-switching models for exchange rates have been proposed (e.g., Engel and Hamilton, 1990; Clarida et al., 2003), these models often rely on latent regimes that are challenging to interpret and difficult to connect to macroeconomic conditions explicitly.

Amidst this backdrop, we investigate three related questions: (i) Do regime shifts in the currency risk-return relationship exist? (ii) How are these shifts associated with the macroeconomic environment? (iii) Do the risk premia associated with currency factors vary across different regimes? To address these questions, we develop a Bayesian regime-switching model that simultaneously identifies regimes and estimates regime-dependent factor risk premia. A key challenge in designing this model is the parsimonious detection of regimes and their association with macroeconomic variables. We tackle this challenge by adapting the panel tree framework (Cong, Feng, He, and Li, 2023), which allows for the endogenous determination of macroeconomic regimes. Rather than pre-specifying regimes, as is common in the current empirical literature, we employ a data-driven approach to regime selection. Within this framework, we jointly model the dynamics of currency returns and their characteristics, while permitting mispricing, volatility, and factor risk premia to vary across regimes defined by recurrent shifts based on the threshold values of macroeconomic variables.

Empirical Highlights. We analyze monthly exchange rate data for 47 currencies relative to the U.S. dollar from 1986 to 2023. From a U.S. investor's perspective, we focus on how U.S. macroeconomic conditions impact currency returns. For such an investor trading widely across currencies, country-specific risks average out, leaving U.S.-specific risk and exposure to common risk factors as key concerns. We consider eight U.S. macroeconomic variables: interest rates, inflation, industrial production, nonfarm payrolls, unemployment rate, M2, oil prices, and the VIX index; and eleven currency characteristics, including carry trade (forward premia), momentum (short-, medium-, and long-term), value, net foreign assets, liabilities in domestic currencies (LDC), long-term yields, term spreads, currency volatility, and currency market beta.

First, our empirical analysis reveals strong evidence of regime-switching patterns in the dynamics of currency returns. The regime-switching model demonstrates a significant improvement in the Bayesian marginal likelihood compared to models that do not incorporate regime-switching, highlighting the statistical significance of accounting for distinct regimes in currency pricing. Moreover, the identified regimes are explicitly defined by U.S. macroeconomic conditions. Specifically, we identify three non-overlapping regimes: a high-inflation regime, a low-inflation and low-interest-rate regime, and a low-inflation and high-interest-rate regime.

Second, we find significant regime-dependent changes in factor selection and risk premia, suggesting that the characteristics-based models undergo substantial variations over time. While the Carry factor consistently generates high risk premia across all regimes, reaffirming its central role in currency pricing (e.g., Lustig and Verdelhan, 2007; Lustig et al., 2011), the importance of other factors varies with macroeconomic conditions. For example, the LDC and market beta factors exhibit high risk premia in the high-inflation regime, the Value factor becomes more important in the low-inflation and low-interest-rate regime, and the Volatility factor becomes critical in the low-inflation and high-interest-rate regime. These findings indicate that currency return dynamics respond to shifts in business cycle conditions.

Third, our findings highlight substantial variations in mispricing and idiosyncratic volatility over time and in the cross section of individual currency excess returns. In particular, the high-inflation regime is characterized by elevated levels of both mispricing and idiosyncratic volatility, reflecting increased uncertainty and divergence in currency valuations during periods of macroeconomic instability. These findings are consistent with Bartram et al. (2023) and Della Corte et al. (2021), which find time variation in currency mispricing and volatility.

Building on those observations, we investigate a long-short strategy that exploits cross-sectional differences in currencies' sensitivity to regime-switching risk. This sensitivity is measured as the difference between regime-specific and unconditional model-implied Sharpe ratios, calculated from regime-switching and non-regime-switching models, respectively. Currencies that display larger differences in model-implied Sharpe ratios are more sensitive to regime changes, and thus they are more exposed to macroeconomic regime-switching uncertainty risk, earning higher expected returns. Portfolios sorted based on this sensitivity exhibit monotonically increasing returns, with the long-short strategy generating significant average excess returns. These excess returns cannot be explained by existing common currency factor models and, in this sense, they appear to be an anomaly related to regime risk.

Methodology Innovations. Our tree-based Bayesian regime-switching model partitions the time horizon into recurrent regimes, allowing for simultaneous estimation of regime-dependent characteristics-based panel regression models. These regimes are characterized by threshold values of macroeconomic variables, such as interest rates and inflation. This approach integrates factor selection, risk premium estimation, and regime detection into a unified framework. To facilitate the factor selection of the regime-dependent models and factor risk premia estimation, we adopt the Bayesian spike-and-slab prior.[1] A low selection probability for a factor implies that its risk premium is more likely to be zero when controlling for other factors in the regime. Another key innovation of our approach is using a *global* split criterion based on the marginal likelihood, which helps mitigate the concern of over-fitting and the potentially adverse impact of noisy parameter estimation on regime detection while accounting for both parameter and model uncertainties.

The proposed model addresses several limitations of existing regime-switching methods. Traditional models, such as those based on Markov-switching processes (e.g., Hamilton, 1989), often fail to capture relationships between non-adjacent regimes. Those models tend to generate latent regimes that are challenging to interpret and can only be loosely linked to macroeconomic conditions *ex-post*.[2] In contrast, our data-driven framework identifies regimes across a large cross section of assets by finding regime-specific factor-selection models. This approach explicitly characterizes the detected regimes using macroeconomic variables, offering a more nuanced understanding of regime changes.[3] 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.

Related Literature. Few papers in the currency literature account for regime changes when examining the currency risk-return relationship.[4] Early studies, such as En-

[1] Giannone, Lenza, and Primiceri (2021) also employ the spike-and-slab prior to compare sparse and dense representations in economics and finance. Bryzgalova, Huang, and Julliard (2023) apply this spike-and-slab prior in a two-pass stochastic discount factor setting.

[2] Threshold models identify regimes when a variable crosses thresholds but are usually limited to single time-series analysis and deterministic regime detection (e.g., Massacci, 2017; Liu and Chen, 2020).

[3] Our model acts as a structural break framework when regimes are defined by calendar months.

[4] Regime-switching models are widely applied in fixed income (e.g., Ang et al., 2008) and equities (e.g., Ang and Bekaert, 2002; Smith and Timmermann, 2021).

gel and Hamilton (1990), Johnson (2002), and Clarida et al. (2003), employ Markov-switching models to model and forecast exchange rate returns. However, these models define regimes as latent states, which are often difficult to interpret and connect to economic fundamentals. Bekaert (1995) studies time variation in conditional means and variance in currency markets, while Kho (1996) evaluates time-varying risk premia and volatility in currency futures markets. Lettau et al. (2014) find that the conditional downside risk CAPM can jointly explain the cross section of currencies, equity, equity index options, commodities, and sovereign bond returns. Nevertheless, these studies lack an interpretation of macroeconomic conditions for regimes.

The literature has documented various profitable currency investment strategies, including currency Carry (e.g., Sarno et al., 2012; Menkhoff et al., 2012a; Bakshi and Panayotov, 2013; Lettau et al., 2014; Koijen et al., 2018; Bekaert and Panayotov, 2020), currency momentum (e.g., Menkhoff et al., 2012b; Asness et al., 2013), and currency value (e.g., Asness et al., 2013; Menkhoff et al., 2017). A recent study by Nucera, Sarno, and Zinna (2024) identifies three or four latent factors using a large cross section of currency portfolios. Aside from the Dollar factor, they find strong evidence of Carry and Momentum factors but a relatively weak Value factor (also see Li et al., 2024). In contrast, Chernov et al. (2023), focusing on ten major currencies (G10), report that the Carry and Value factors account for significant portions of priced risk, followed by Momentum. These divergent findings may arise from ignoring regime changes in currency markets. Our paper uses a novel methodology on individual currency excess returns to examine regime shifts of the characteristics-based model in pricing currency risks. Moreover, our test assets present greater challenges than those based solely on major currencies or portfolios, as we include emerging markets and weaker factors observed in individual currency returns.

Tree-based machine learning methods have gained attention in pricing equity and bond risks due to their nonlinear modeling capabilities using high-dimensional characteristics (see, e.g., Gu et al., 2020; Bianchi et al., 2021; Cong et al., 2024; Bryzgalova et al., 2024; Li and Tang, 2024). Patton and Simsek (2023) combine the generalized autoregressive score models with decision trees and random forests to improve forecasts. Our paper contributes to this growing literature with a focus on regime-switching of factor structure in the currency markets.[5] For parallel development, Cong et al. (2023) examine asset heterogeneity and select uncommon factor models in equity returns, and Bie et al. (2024) apply to the dynamic Nelson-Siegel model for the yield curve. Extending Cong et al. (2023), the model further allows for heterogeneous idiosyncratic volatility, a key component of our regime risk anomaly. Finally, while Creal and Kim (2021) explore the risk-return tradeoff in currency markets using Bayesian Additive Regression Trees (BART, Chipman et al., 2010) and Patton and Weller (2022) extend the K-means to group assets and investigate risk price heterogeneity in the equity market, both studies primarily focus on the cross-section dimension and largely neglect macroeconomic information and regime changes. By addressing these gaps, our paper provides new insights into regime-dependent factor dynamics for currency.

The rest of the paper proceeds as follows. Section 2 presents our regime-switching model. Section 3 describes the data. Section 4 discusses the main empirical findings of factor risk premia on currency regime shifts. In Section 5, we construct an anomaly that captures regime risk. In Section 6, we connect our model to structural break models. Finally, Section 7 concludes the paper.

[5]We propose a tree-based Bayesian regime-switching model that combines decision trees and Bayesian methods, building upon and extending XBART (He et al., 2019; He and Hahn, 2023) by incorporating Bayesian variable selection and emphasizing the detection of macroeconomic regimes.

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:

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

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. In the subsequent analysis, we refer to $x_{m,t-1} \leq c_m$ as a split rule, $x_{m,t-1}$ as a split variable, and c_m as a split value.

Since macroeconomic time series often exhibit trends over time, they need to be normalized prior to use. Details on the data and normalization procedures are provided in Section 3.3. As discussed in subsequent sections, the number of macro variables that define a regime can be determined using a tree clustering algorithm. These regimes may consist of non-contiguous periods characterized by either low or high levels of a specific macroeconomic variable.

Next, we present the full characteristics-based model that incorporates all regimes as given in Section 2.1, and then discuss the procedures for regime detection and parameter estimation in subsequent subsections.

2.1 Full Factor Model with Regimes Given

Suppose there are $j = 1, \cdots, J$ regimes over the time horizon, and the assignment of each period to a regime is predetermined. The full model with regimes is

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

The indicator function, $\mathbb{1}_{\{t\in\mathcal{R}_j\}}$, denotes whether the t-th period belongs to regime $\mathcal{R}_j$. Each period t is assigned to one and only one regime. The vector $\mathbf{Z}_{i,t-1}$ represents the lagged characteristics for the currency i at time $t-1$. Following Fama and French (2020) and Smith and Timmermann (2022), the coefficient vector $\boldsymbol{\lambda}_j$, with K elements, is interpreted as cross-sectional factor risk premia with predetermined time-varying factor loadings (currency characteristics). The terms $\alpha_{j,i}$ and $\sigma_{j,i}$ denote the mispricing and idiosyncratic volatility of currency i within regime j.

In particular, the model parameters $\boldsymbol{\alpha}$, σ^2, and $\boldsymbol{\lambda}$ are regime-specific, enabling the model to capture the evolving dynamics of currency returns over time. These parameters are dynamically linked to macroeconomic variables, 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.

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.

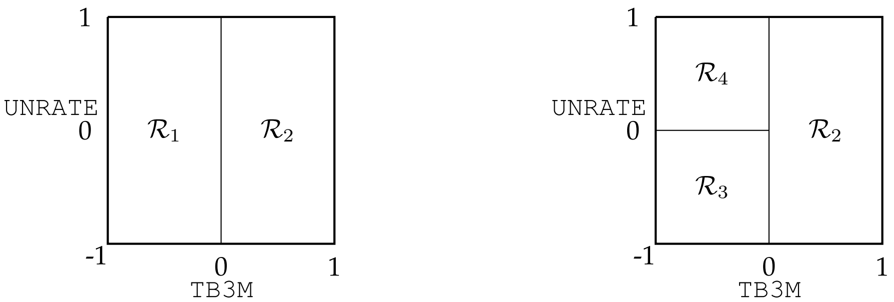

Thus far, the number of regimes and their corresponding split rules have been assumed to be predetermined. Section 2.2 presents the prior specification and marginal likelihood for the Bayesian variable selection model. Building on this, Section 2.3 introduces how the Bayesian regime-switching model detects 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 Equation (1) becomes

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

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 risk premium of the corresponding factor is shrunk towards zero. In contrast, the "slab" component has a large prior variance ξ_1^2, allowing for a wider distribution around zero and suggesting that the risk premium is non-zero with minimal shrinkage. The latent dummy indicator variables $\boldsymbol{\gamma}_j = (\gamma_{j,1}, \cdots, \gamma_{j,K})$ play a central role. Each $\gamma_{j,k}$ takes a value of 1 or 0, indicating whether the prior for the corresponding risk premium $\lambda_{j,k}$ belongs to the slab or the spike. Specifically, when $\gamma_{j,k} = 1$, the factor risk premium $\lambda_{j,k}$ follows the slab prior, resulting in minimal shrinkage. Conversely, when $\gamma_{j,k} = 0$, $\lambda_{j,k}$ follows the spike prior, shrinking the coefficient toward zero. This mechanism effectively helps select or discard factors in the model. By averaging the posterior samples, the posterior mean of $\gamma_{j,k}$ will fall within $[0, 1]$, which can naturally be interpreted as a selection probability. A low selection probability for a factor indicates a high posterior probability that its risk premium is zero in this regime, conditional on the presence of other factors in the regime.

It is important to note that the risk premium $\boldsymbol{\lambda}_j$ in Equation (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 $\sigma^2_{j,i}$ are allowed to be currency-specific. The priors for these parameters are specified as follows:

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

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 factor has only a 10% chance of being included in the model. The prior for γ 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 requires partitioning the entire time series into multiple regimes, and fitting a regime-specific model as per Equation (2). Effective regime detection should improve the overall fit of the data by utilizing these regime-specific models. The frequentists' approach typically relies on the likelihood function to evaluate the model's fitness for different partitions of the data. However, the likelihood evaluates both the data partitions and model parameters estimation, and noisy plug-in parameter estimates can distort its ability to assess 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:

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

Notably, the marginal likelihood integrates out unknown parameters in the likelihood function for 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 Equation (3) has no closed-form solution. To address this, we follow Chib (1995) and perform a numerical calculation 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 $p(\mathcal{R}_0)$ as the marginal likelihood evaluation of Equation (3), using the data within regime $\mathcal{R}_0$.

2.3 Macro-Instrumented Regime Detection

This section introduces our Bayesian regime-switching approach, which employs a tailored tree growth algorithm to partition currency return data based on marginal likelihood. Figure 1 illustrates an example of the clustering results. This decision tree divides the time horizon into distinct regimes according to values of macroeconomic variables, specifically, the 3-month Treasury bill rate (`TB3M`). The algorithm begins by applying an initial split to the entire dataset, partitioning it into subspaces corresponding to specific time series regimes. This process is iterative, with each child node further divided until predefined stopping conditions are satisfied. At each step, the algorithm evaluates all potential split rules using the marginal likelihood of the model in Equation (3) and selects the rule that maximizes the marginal likelihood gain.

To determine the first split, we evaluate all potential split rule candidates. These candidates consist of pairs of macroeconomic variables and their corresponding cut values, which divide the currency return time series into distinct regimes. Figure 2 illustrates an example of such a candidate split. Each pair of a macro variable and the corresponding cut value generates distinct regimes, denoted as $\mathcal{R}_1$ and $\mathcal{R}_2$. The key question is: which candidate should be chosen?

We aim to identify regime-switching patterns in the characteristics-based model for currency returns, i.e., partitioning the entire time series into multiple regimes, fitting a regime-specific model, and achieving an overall improved fit of currency returns. Statistically, this involves optimizing the marginal likelihood of the model, as defined in Equation (3), for each regime. The marginal likelihood thus serves as a natural measure of the quality of a split candidate.

When a candidate partitions the entire time series into two regimes, it implies that these regimes should be modeled separately, as described in Equation (1). To evaluate this partition, we calculate the marginal likelihood for each regime separately and then take their product to obtain the overall marginal likelihood for the entire dataset. By iterating through all potential split candidates, we select the one that maximizes the overall marginal likelihood.

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.

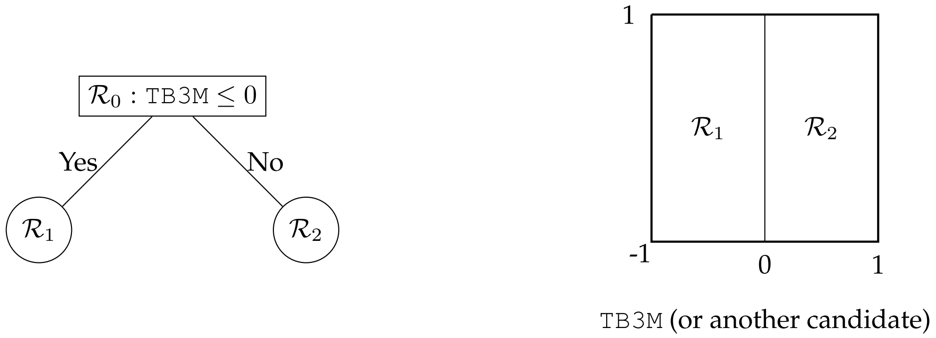

In particular, let $\mathcal{C} = \{\tilde{c}_j\}$ denote the set of all split rule candidates. To simplify the computation, we evaluate only a single cut value, 0, for each macroeconomic variable. Since all variables are demeaned using a 10-year rolling window, a cut point of 0 reflects whether the 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 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 results in a different data partition, and the evaluation of Equation (4) varies accordingly. A larger value indicates a better overall model fit after the partition, making it a natural metric for assessing split candidates, which we refer to as the split criterion. For details on further splits and the stopping criterion, please see Appendix A.II.

The tree grows similarly for all subsequent splits. At each split (the j-th split), we evaluate the global split criterion and stopping options for all existing leaf nodes. The candidate that optimizes the criterion is then selected. The algorithm continues until a stopping option is chosen or certain predetermined stop conditions are satisfied, such as reaching a minimum number of data points in a leaf node or the maximum depth of the tree. After $J-1$ iterations, the algorithm ends, resulting in J regimes.

After growing the tree and identifying the regimes, we estimate the model parameters for each regime. Specifically, we fit the Bayesian panel regression model in Equation (2) for each regime and draw posterior samples using the Gibbs sampler detailed 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 — characterized by varying levels of risk premia in the data-generating process. The results demonstrate that this approach effectively distinguishes among these factor types. Strong factors consistently 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.

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 collect spot and one-month forward exchange rates with respect to the U.S. dollar (USD) for the following 47 economies.[6] The data are from Barclays Bank International (BBI) and WM/Reuters via *Datastream*. The euro series is included starting from January 1999, and individual euro area countries are excluded thereafter, retaining only the euro series.

Following standard practice in the literature, we exclude observations during periods of significant covered interest rate parity (CIP) failures for a currency.[7] From the perspective of a U.S. investor, exchange rates are expressed as units of USD per unit of foreign currency (FCU), i.e., FCU/USD. An increase in spot or forward rates indicates an appreciation of the foreign currency relative to the U.S. dollar or, equivalently, a depreciation of the U.S. dollar.

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

[6] 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.

[7] We follow Nucera, Sarno, and Zinna (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).

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 Equation (5) holds because we follow the common practice of using forward rates rather than interest rate differentials to compute 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 rate parity.

3.2 Currency Characteristics

Following the literature, we construct various currency characteristics. We consider Carry (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 on net foreign assets and liabilities in domestic currencies (Della Corte et al., 2016), long-term yields (Ang and Chen, 2010; 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 make the values of characteristics comparable across the panel, we rank all characteristics cross-sectionally for each month and normalize the ranks to the range of $[-1, 1]$, with their cross-sectional median set to 0. Any missing values are then imputed as 0. One advantage of using the cross-sectional ranks of characteristics is that they largely mitigate the impact of potential data errors and outliers in individual characteristics (see, e.g., Feng, He, Polson, and Xu, 2024).

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. There are 11 characteristics, whose definitions are provided in Table IA.2. The 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.[8] 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.

[8]For crude oil price, we calculate the annual percent change. For VIX before 1990, we impute with 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.

All macroeconomic indicators are standardized by demeaning, subtracting the 10-year rolling window average. To further smoothing, 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.[9] This standardized approach ensures the comparability of macroeconomic indicators over time, which is essential for the tree clustering method to detect macroeconomic regimes based on recent data effectively.

4 Empirical Results

This section presents our key empirical findings. First, Section 4.1 presents the baseline model estimation and selection without considering regime shifts. Next, in 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. Finally, Section 4.4 analyzes cross-sectional mispricing and heteroskedasticity.

4.1 Full Sample Risk Premia Estimation

We evaluate the overall performance of currency factors while excluding regime shifts, utilizing the full sample. We estimate the characteristics-based model in Equa-

[9]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.

tion (2) for individual currency returns as follows:

$$
\begin{aligned}
r_{i,t} = \alpha_{i,t} + r_{z,t} &+ \lambda_{\text{Carry}}\text{Carry}_{i,t-1} + \lambda_{\text{MOM1}}\text{MOM1}_{i,t-1} + \lambda_{\text{MOM6}}\text{MOM6}_{i,t-1} \\
&+ \lambda_{\text{MOM12}}\text{MOM12}_{i,t-1} + \lambda_{\text{Value}}\text{Value}_{i,t-1} + \lambda_{\text{NFA}}\text{NFA}_{i,t-1} \\
&+ \lambda_{\text{LDC}}\text{LDC}_{i,t-1} + \lambda_{\text{LT\_Yield}}\text{LT\_Yield}_{i,t-1} + \lambda_{\text{Term\_Spread}}\text{Term\_Spread}_{i,t-1} \\
&+ \lambda_{\text{Volatility}}\text{Volatility}_{i,t-1} + \lambda_{\text{MKT\_Beta}}\text{MKT\_Beta}_{i,t-1} + \epsilon_{i,t},
\end{aligned}
$$

where we include a common factor $r_{z,t}$ to capture the variation of the level factor in the cross section, following the approach of Pesaran (2006) and Smith and Timmermann (2022).[10] In this context, $r_{z,t}$ 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 risk premia estimation. It presents the posterior factor selection probability (probability of non-zero risk premia) and the posterior means of factor risk 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 ($w = 0.1$).

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,

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

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

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.

	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

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 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.

For the investor's confident prior, the MOM6 and LDC factors exhibit statistically significant positive risk premia, with selection probabilities around 80%. However, as the investor's prior beliefs become more skeptical, both their risk premia and selection probabilities decline sharply. The factor risk premium of currency volatility is not statistically significant after controlling 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 with No Regime**

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.

	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 switching. Importantly, our model can detect regimes by iteratively splitting the entire history of observations using U.S. macroeconomic variables, based on enhancements 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 conditions.

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;

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

- 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.

Figure 3: **Partitioning Macroeconomic Time Series**

The figure illustrates the detected regimes over the monthly sample period from 1986 to 2023, presented as a partition plot. The algorithm first partitions the data based on inflation (CPI) and subsequently divides the low-inflation periods into two regimes based on high and low interest rates, resulting in three non-overlapping regimes. The macroeconomic indicators are standardized by subtracting the previous 10-year rolling window average and applying smoothing, with the split point set at zero.

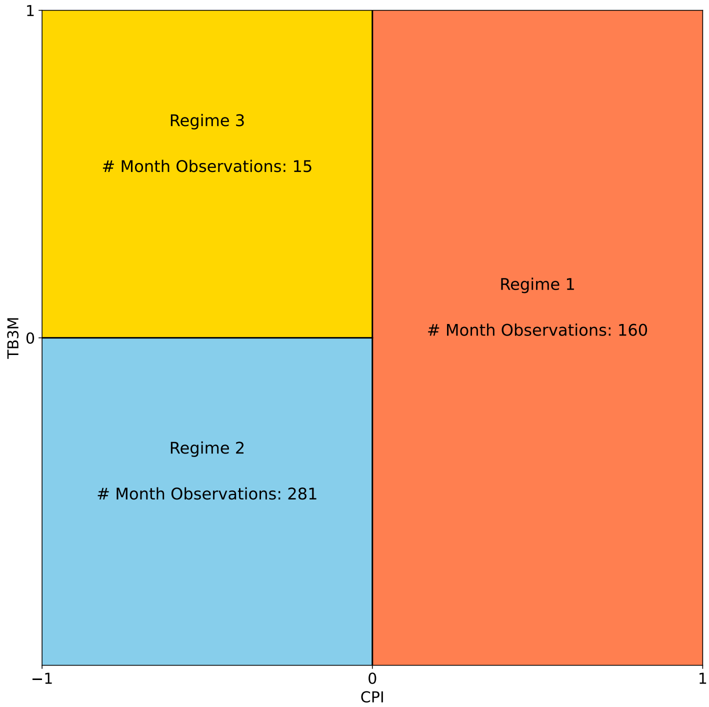

The samples from individual regimes are then used to select factors specific to each regime and estimate their risk premia. Figure 4 presents the regime-switching pattern.[12] We note that Regime 1 (high inflation), marked in red, encompasses events such as the Gulf War (1990-1991) and the early phase of the COVID-19 pandemic (2020); these periods are characterized by economic instability and rising prices, reflecting a high inflation environment. This regime is usually at the peak of the business

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

cycle, with an overheating economy, which can lead to asset bubbles.

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

The figure illustrates the detected regimes over time using the monthly sample from 1986 to 2023. Regime 1 (red) represents the high-inflation regime, Regime 2 (blue) corresponds to the low-inflation and low-interest-rate regime, and Regime 3 (yellow) represents the low-inflation and high-interest-rate regime. Shaded areas indicate NBER recession periods. The figure includes time-series plots of U.S. interest rates, inflation, and the EPU index (scaled by dividing by 50 for alignment with the y-axis) (Baker et al., 2016), and key global events affecting the world economy.

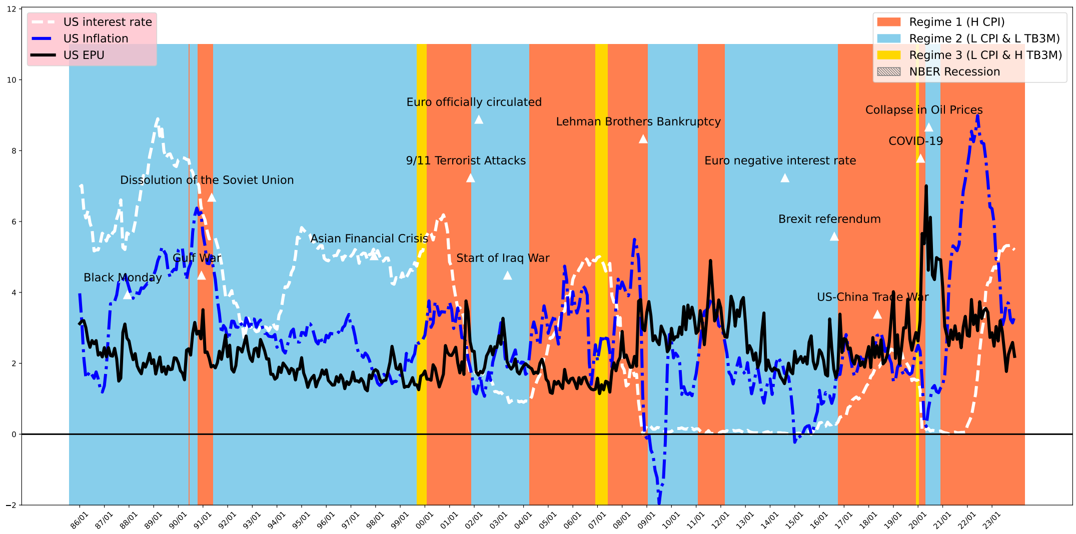

Regime 2 (low inflation and low interest rates), depicted in blue, includes key events such as Black Monday (1987), the dissolution of the Soviet Union (1991), the Asian Financial Crisis (1997–1998), the 9/11 terrorist attacks (2001), the Lehman Brothers bankruptcy (2008), and the Brexit referendum (2016). It is marked by low inflation, low interest rates, and intermittent financial disruptions. This regime is in the phase of recovery from a recession and expansion within the business cycle with stimulated borrowing, spending, and moderate growth.

Regime 3 (low inflation and high interest rates), highlighted in yellow, covers the introduction of the Euro (2002) and the start of the Iraq War (2003). These events occurred during periods of low inflation but rising interest rates, often in response to geopolitical tensions and policy shifts. This regime is in the slowdown and contraction phase of the business cycle, with reduced growth. The figure also tracks 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), illustrating the correlation between these indicators and the economic regimes. Notably, spikes in EPU correspond with major events and shifts between regimes, indicating heightened economic uncertainty during these periods.

4.3 U.S. Macroeconomic Regimes and Factor Risk Premia

Each macroeconomic regime is characterized by distinct business cycle conditions, influenced by inflation rates and interest rate levels, suggesting that the dynamics of currency returns may vary across these regimes. We further estimate currency factor risk 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 summarize the posterior results from 2,000 Markov Chain Monte Carlo (MCMC) sample draws after burn-in samples, and present the factor risk premia estimates for each regime in Table 4.

In Regime 1, characterized by high inflation, the Carry factor consistently exhibits highly statistically significant positive risk premia for all priors: 0.33% for the confident prior, 0.32% for the agnostic prior, and 0.32% for the skeptical prior, indicating its robustness to investor's prior belief. The MKT Beta factor also shows a significant positive risk premium, ranging from 0.47% (for the skeptical prior) to 0.52% (for the confident prior), surpassing that of Carry, reflecting the potential high profitability of the beta strategy in high inflation conditions. 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. Overall, the MKT Beta, Carry, and LDC factors stand out, providing statistically significant

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

The table presents the estimation results of factor risk premia (in %) for three detected regimes based on inflation and interest rate data (see Figure 3). Regime 1 represents high inflation, Regime 2 corresponds to low inflation and low interest rates, and Regime 3 represents low inflation and high interest rates. Each regime group displays risk premia estimates (with 5% and 95% credible intervals in parentheses) according to prior selection probabilities: confident prior with $w = 0.9$, agnostic prior with $w = 0.5$, and skeptical prior with $w = 0.1$. The numbers indicate the mean (with 5% and 95% quantiles in parentheses) of 2,000 MCMC post-convergence 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]

and economically substantial risk premia in a high-inflation regime. In contrast, the risk premia of other factors, including the Value factors, 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 risk 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 generates 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 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 only 0.05%, respectively. 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 are significant, and some are even negative.

Figure 5 illustrates the variation of factor risk premia across macroeconomic regimes for investors with a confident prior, highlighting both common and uncommon factors in different regimes. These variations indicate that currency return dynamics change in response to shifting business cycle conditions, primarily influenced by U.S. inflation rates and interest rate levels.

Table 5 summarizes the selection probabilities of factors across all regimes under

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

The figure shows regimes and heatmap of risk premia estimate (%) under confident prior ($w = 0.9$) for three detected regimes identified using U.S. inflation and interest rate data (see Figure 3).

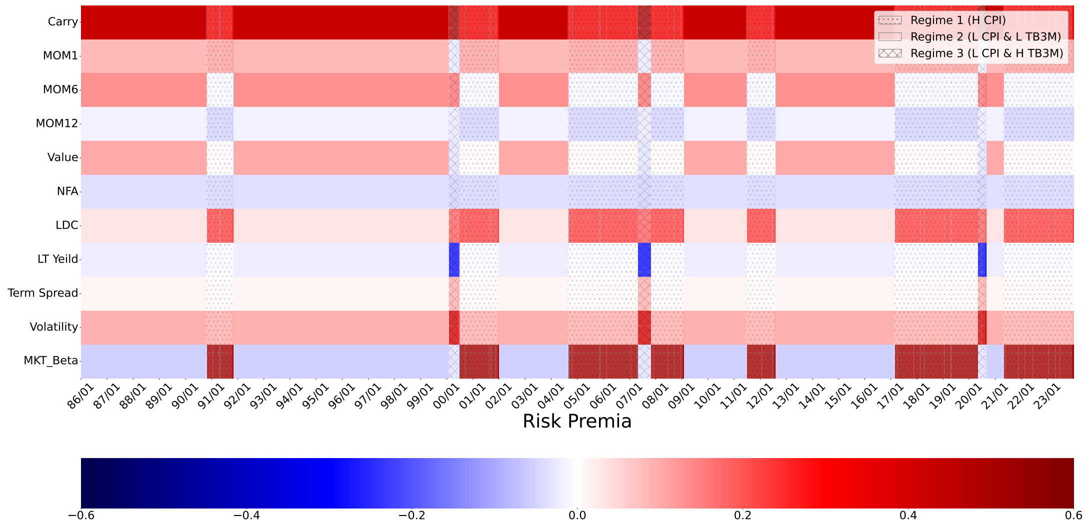

the three types of prior beliefs, further illustrating the strength of factors in different macroeconomic conditions. Consistent with the factor risk premium estimates, in Regime 1, the Carry factor and MKT Beta are selected with nearly 100% probability across all three priors, highlighting their robustness in the high-inflation environment. Among other significant factors, LDC exhibits a high selection probability of 96% for the confident prior, while MOM1 has a selection probability of 81%. 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 prior, their selection probabilities decrease. Other factors, such as Volatility and MOM1, also have selection probabilities larger than 70% for the confident prior; however, their selection probabilities decrease drastically when investors are less confident a priori.

In Regime 3, none of the factors are selected with a 100% probability. However, we find that the Carry and Volatility factors still show a probability of more than 90% for the confident prior. For the agnostic prior, the Volatility factor is selected with a probability of 80%, the Carry factor is 61%, and all other factors show a selection probability of less than 50%. For the skeptical prior, all factors are selected with very low probabilities.

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

The table shows factor selection probabilities (%) for three detected regimes identified using U.S. inflation and interest rate data (see Figure 3). We present results based on different priors. The numbers represent the mean of 2,000 MCMC post-convergence samples of γ.

	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

In summary, factor risk premia vary significantly across the identified regimes based on inflation and interest rate indicators. The Carry factor consistently exhibits high selection probabilities in all regimes. Beyond that, the MKT Beta, MOM1, and LDC factors become pivotal in Regime 1, reflecting their sensitivity to high inflation. In Regime 2 (low inflation and low interest rates), the MOM6 and Value factors emerge as significant, alongside the Volatility and MOM1 factors. In Regime 3 (low inflation and high interest rates), the Volatility factor gains prominence, particularly for the agnostic prior, reflecting its heightened sensitivity to the regime with reduced growth.

4.4 Mispricing and Idiosyncratic Volatility

The previous section shows that currency risk premia vary across regimes. With our model incorporating currency-specific α_i and σ_i, we analyze the cross-sectional distributions of α and σ estimates to study potential mispricing and idiosyncratic volatility during regime shifts. We find that the explanatory power of currency risk factors also varies 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. The 95% quantile is notably positive (0.55%), while the 5% quantile is slightly negative (-0.012% to 0.000%).

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. Regime 3, characterized by low inflation and high interest rates, displays slightly larger positive mean alpha values than Regime 2, while the cross-sectional standard deviation (approximately 0.06%) is notably smaller than in Regime 2 (approximately 0.10%). This suggests mispricing is lower during low inflation, and the risk-return relationship aligns more closely with risk premia. Across all priors, the 5% quantile alpha is about -0.04% and the 95% quantile alpha is about 0.12%. Mispricing in Regimes 2 and 3 is lower than that for the full period in Table 3.

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

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

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

	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

at 1.2% and 4.0%. This suggests that currency-specific risk reaches its peak during the high inflation regime. 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.

Overall, our findings reveal significant variations in mispricing and idiosyncratic volatility over time and in the cross section of currency returns. These results align with Bartram et al. (2023) and Della Corte et al. (2021), who document that currency mispricing and volatility exhibit time variation.

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 identified 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.

Panel A: α (mispricing)

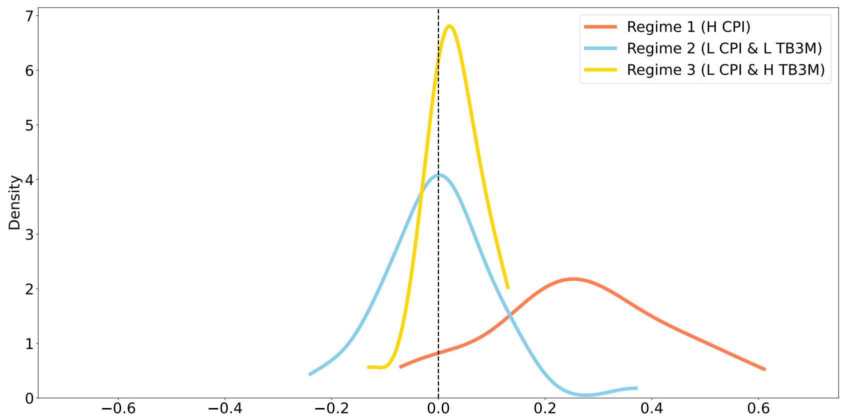

Panel B: σ (idiosyncratic volatility)

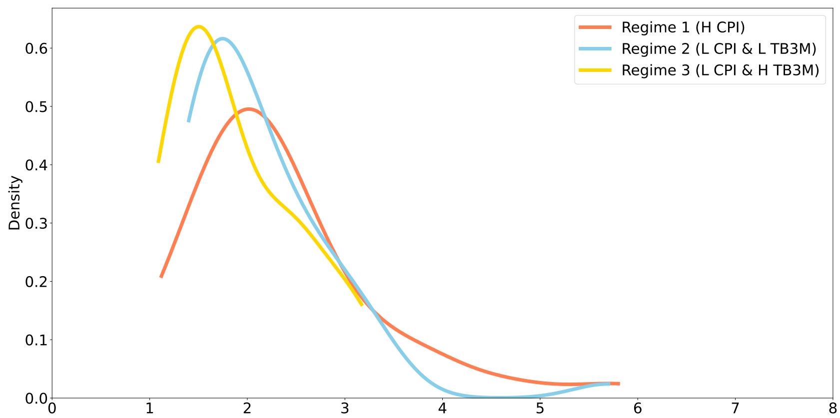

5 Regime Risk Anomaly

We have provided empirical evidence of regime-switching patterns in factor risk premia and idiosyncratic volatility. We highlight that exposure to potential regime-switching behavior in risk premia and volatility induces an additional source of risk for investors. For example, if an investor estimates a model without accounting for regime changes and constructs portfolios based on the estimated risk premia, a future shift in the macroeconomic regime could result in significantly different risk premia, potentially leading to unexpected loss to the investor. This regime-switching risk parallels the break risk documented for equity by Smith and Timmermann (2022), which captures the structural break risk for individual stocks.

These findings suggest that macro-instrumented regime shifts in risk premia and volatility introduce new undiversified risk to the risk-return trade-offs in the currency market—regime risk. This is not merely a modeling uncertainty but is also economically meaningful, since the regimes identified by our model align with shifts in macroeconomic conditions. Currencies more sensitive to significant macroeconomic changes are likely to deliver higher risk-adjusted performance as compensation for their exposure to regime uncertainty. For example, Mueller, Tahbaz-Salehi, and Vedolin (2017) document a trading strategy of shorting the U.S. dollar and taking long positions in other currencies, generating higher excess returns on days with scheduled Federal Open Market Committee announcements, which can be interpreted as compensation for monetary policy uncertainty.

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 x_{t-1}, we have

$$\text{RISR}_{i,t,j} = \frac{z_{i,t-1}\lambda_j}{\sigma_{i,j}}.$$

Similarly, under the unconditional (non-regime) model estimated with the full sample, there is only one constant risk premium λ and idiosyncratic volatility σ_i for the entire sample period. We define such unconditional implied Sharpe ratio (ISR) as follows:

$$\text{ISR}_{i,t} = \frac{z_{i,t-1}\lambda}{\sigma_i}.$$

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

$$\text{DISR}_{i,t} = \text{RISR}_{i,t,j} - \text{ISR}_{i,t},$$

which measures the difference in model-implied Sharpe ratios between the regime-switching model and the unconditional model. A higher DISR for a currency indicates a larger disagreement between the two models, reflecting greater exposure to regime-switching risk. Consequently, such a currency demands higher risk-adjusted compensation, highlighting a novel source of risk in the risk-return trade-offs.

In Section 4, we identified three regimes based on inflation and interest rates using the full sample period. However, these regimes are not necessarily fixed and may vary if the sample period changes. To assess the robustness of our proposed approach, we conduct an out-of-sample study and calculate risk anomalies in real-time to avoid look-ahead bias. Specifically, we estimate the regime-switching model using an initial 14-year training period and subsequently update the model every month with historical data. The out-of-sample analysis begins in 2000.

At the end of each month, we identify regimes according to lagged macroeconomic variables and estimate the regime-dependent model for cross-sectional currency returns, including factor risk premia and idiosyncratic volatility. The regime for the following out-of-sample month is determined based on the lagged macroeconomic conditions. Subsequently, the RISR and ISR for each currency are calculated using the lagged characteristics, and updated as we move through the sample.

We focus on investors with a confident prior belief that characteristics play a significant role in explaining the cross-sectional variation in currency returns. In contrast, for investors with less confident priors, the risk premia of most factors tend to shrink toward zero, making it more challenging to capture regime-specific risks compared to unconditional risk premia. Intuitively, investors who trust the characteristics-based model are more likely to be concerned about their exposure to regime-switching risk.

To examine the implications of regime-switching risk, we implement a monthly rebalanced portfolio sort 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 values. The regime risk anomaly is then constructed as the return difference between the top DISR portfolio (high regime risk) and the bottom DISR portfolio (low regime risk).

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

The table presents summary statistics for the DISR-sorted quintile currency portfolios. Portfolios P1-P5 represent the five sorted portfolios, where P5 includes the top 20% currencies with the highest DISR values. The high minus low (HML) portfolio is constructed by going long on P5 and shorting P1. Furthermore, the table reports the autocorrelation (AC1), the annualized Sharpe ratio, and the annualized maximum drawdown. Monthly returns are expressed as percentages, and the Newey-West t-ratios (Newey and West, 1987) are presented in parentheses. The sample period spans from 2000 to 2023.

	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

Table 7 provides summary statistics for the DISR-sorted quintile currency portfolios, where P5 is the portfolio of the top 20% currencies with the highest DISR values, and P1 includes those with the lowest DISR values. The high minus low (DISR-HML) strategy takes a long position in P5 while shorting P1. First, over the out-of-sample period from 2000 to 2023, average portfolio returns monotonically increase from P1 to P5. The annualized Sharpe ratio significantly improves from P1 (-0.06) to P5 (0.39). Second, the DISR-HML strategy yields a monthly average return of 0.31%, which is statistically significant, with a t-statistic of 2.78. The DISR-HML strategy produces a high Sharpe ratio of 0.51. Third, the annualized maximum drawdown of the DISR-HML strategy is relatively modest, at approximately 12%, hedging the regime risk. Overall, these investment performance statistics underscore the significance of the abnormal returns generated by the DISR-HML strategy based on the regime risk.

Beyond the investment gains generated by the DISR-HML strategy, it is essential to evaluate the marginal contribution of regime risk relative to other common currency risk factors. Table 8 presents the results from the factor-spanning regressions of the DISR-HML strategy on commonly used currency factors. We regress the regime risk anomaly returns on three-factor models: (i) Dollar and Carry, (ii) a four-factor model with Dollar, Carry, Momentum, and Value (Nucera, Sarno, and Zinna, 2024), and (iii) a model that includes all factors considered in our study. Across all models, the alpha estimates are both statistically significant and economically meaningful, ranging from 0.29% to 0.33% per month. The four-factor model achieves an R^2 of 0.6%, indicating that it explains little of the time-series variation in the regime risk anomaly. When all factors are included, we find that almost all factor loadings are statistically insignificant (except Term Spread), and the adjusted R^2 remains notably low, around 5.7%. These results suggest that the regime risk anomaly is distinct and cannot be explained by other established currency factors.

These results provide empirical evidence that macroeconomic regime-switching risk remains underexplored in the existing currency pricing literature. Although our regime-switching model offers a method to quantify this uncertainty risk, identifying the most relevant macroeconomic conditions affecting currency returns remains challenging because of the changing regime that conditions macroeconomic variables in our monthly updated out-of-sample model training.

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

The table presents estimates of alphas (in %) and factor loadings derived from the factor-spanning regressions of the regime risk anomaly returns on commonly used currency factors. We consider three-factor model specifications: (1) Dollar and Carry, (2) a four-factor model comprising Dollar, Carry, Momentum, and Value, and (3) a model that includes all factors considered in our paper. All factors are constructed from high minus low (5-1) portfolios based on previously utilized characteristics, except for DOL, which denotes the cross-sectional mean return (see Table IA.3 in the Internet Appendix). Monthly returns are expressed as percentages, and the Newey-West t-ratios (Newey and West, 1987) are reported in parentheses. The sample is from 2000 to 2023.

	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

6 Connection to Structural Break Models

As noted by Ang and Timmermann (2012), regime changes can be broadly categorized as recurring (e.g., recessions versus expansions) or unique (e.g., structural breaks). Our regime-switching models effectively capture these abrupt behavioral shifts and the persistence of new dynamics in currency returns and macroeconomic variables for several periods following a regime change.

Moreover, our general model can be extended to a structural break framework. Specifically, using calendar months as split candidates instead of macroeconomic variables, the regime-switching model transitions into a structural break model, where the identified split months represent breakpoints. To implement this, we allow for up to nine structural breaks. When applied exclusively to currency returns and their characteristics, the structural break model offers a new perspective, potentially uncovering new insights beyond regimes.

Figure 7 presents the estimated break dates from 1986 to 2023. Distinct regimes are represented by different colors, segmented by the identified structural break dates. Shaded areas denote NBER recession periods. Figure 8 displays the corresponding heatmap of time-varying risk premia estimation for the confident prior specification. Detailed risk premia estimation results, along with selection probabilities (in parentheses), are provided in Table IA.4 in the appendix.

We also include time series plots of the U.S. EPU index, inflation and interest rate levels. The first break occurred during the stock market crash of October 1987, commonly referred to as Black Monday. Before this break, the Carry and MKT Beta factors exhibited positive premia, while mid- and long-term momentum factors showed negative premia, as depicted in Figure 8. This event triggered a significant increase in economic uncertainty, as evidenced by a sharp spike in the EPU index. In response,

Figure 7: **Time Series Plot of Structural Break**

The figure displays the estimated break dates based on the complete monthly sample from 1986 to 2023, with different colors representing distinct regimes segmented by the break date points. Shaded areas denote NBER recession periods. Additionally, the figure includes a time series of U.S. interest rates, U.S. inflation, the U.S. EPU index (scaled by dividing by 50 for alignment with the y-axis) (Baker et al., 2016), and global events that impact the world economy.

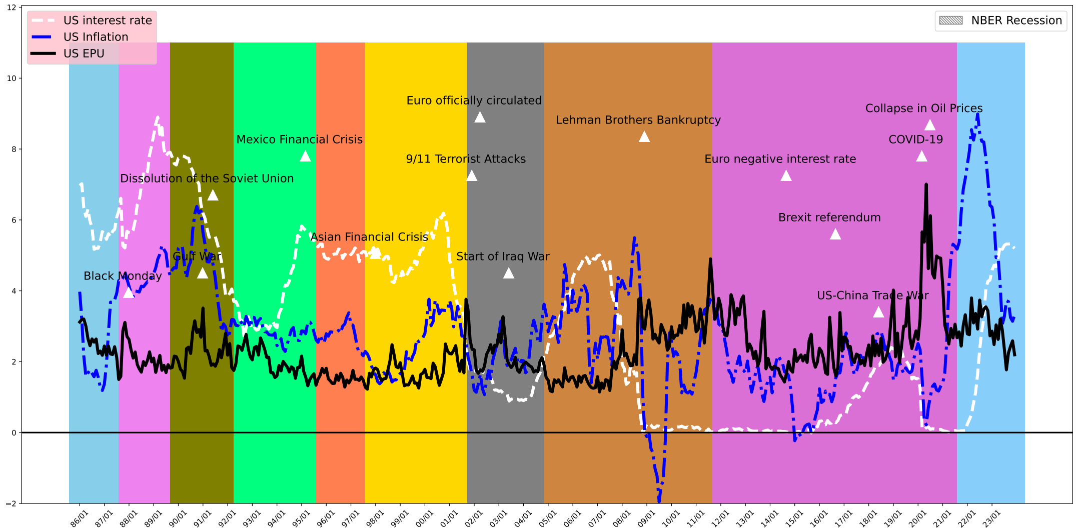

the Federal Reserve adjusted interest rates to stabilize the financial system, leading to a contraction in most risk premia, including Carry.

The second and third breaks occurred before and after the onset of the Gulf War in January 1991 and the dissolution of the Soviet Union in December 1991, as highlighted in Figure 7. These events raised oil prices and global economic uncertainty, prompting interest rate adjustments and influencing the EPU index. Throughout this period, Carry showed positive performance.

The fourth break followed the 1995 Mexican Financial Crisis, which caused severe financial instability in emerging markets and triggered global spillover effects. This period saw a significant rise in the EPU index and adjustments in U.S. interest rates. Nearly all factors experienced a contraction in premia, with the Value factor particularly affected.

The fifth break coincided with the Asian Financial Crisis. Following this, the

Figure 8: **Time-varying Risk Premia under Structural Break**

The figure illustrates the detected regimes (see Figure 7) over time, utilizing the complete monthly sample from 1986 to 2023, along with a heatmap of risk premia estimates (%) under a confident prior (w = 0.9). The regimes are identified based on calendar month breaks (see Figure 7).

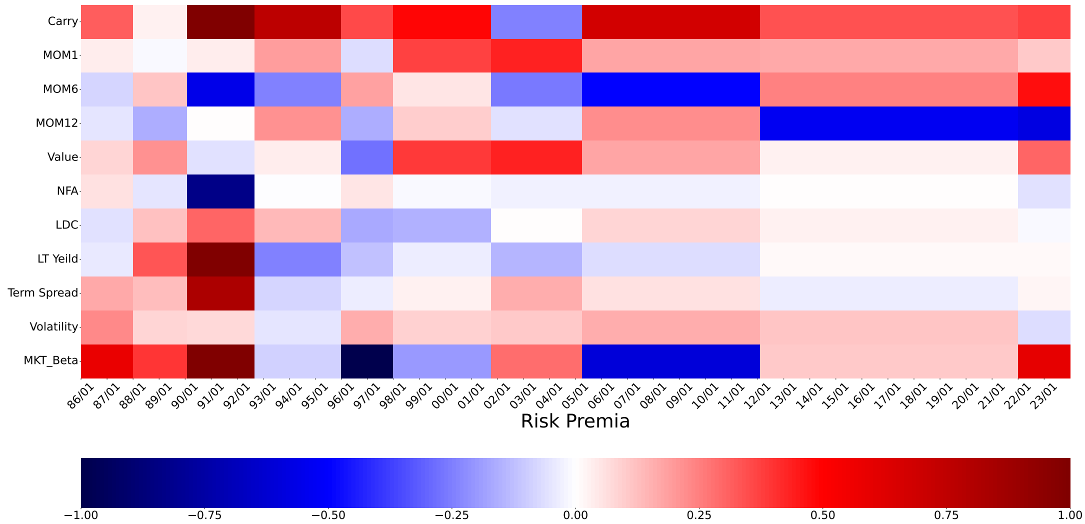

Carry, Momentum (MOM1), and Value factors recovered, demonstrating positive premia. The sixth break occurred in late 2001, coinciding with the 9/11 terrorist attacks. Risk premia contracted broadly during this period, with the Carry factor even turning negative. The seventh break, aligned with the lead-up to the Iraq War in 2003, saw MOM6 and MKT Beta factors exhibit negative premia.

The eighth break, following the 2008 Global Financial Crisis, marked a prolonged regime defined by negative interest rates in the Eurozone, in addition to Brexit, and the U.S.-China Trade War. During this period, most factor risk premia declined relative to the previous regime, with a particularly sharp reduction in MOM12.

The ninth break coincided with the COVID-19 pandemic and the collapse of oil prices, triggering global economic disruptions, a sharp increase in the EPU index, and aggressive monetary policy responses, including significant interest rate cuts. A rapid economic recovery followed, accompanied by rising interest rates. During this period, the Carry, MOM6, and Value factors exhibited strong risk premia.

In summary, the structural break model effectively captures major economic events and institutional changes. It identifies breakpoints that largely align with those detected by the previous regime-switching model while also providing new insights from multiple events. Although regime-switching and structural breaks are distinct concepts, they can coexist within our tree-based approach.

7 Conclusion

Theories of exchange rate determination highlight the importance of macroeconomic fundamentals in driving exchange rate dynamics. The empirical asset pricing literature in currency markets has identified the key role of interpretable common factors in understanding the pricing of currencies (e.g., Lustig et al., 2011; Menkhoff et al., 2017, 2012b). In this paper, we address whether these factors exhibit stable explanatory power and risk premia across varying macroeconomic regimes, using a novel tree-based Bayesian regime-switching model. This framework allows for regime changes in the importance of factors that drive currency returns, with the regimes being driven by interpretable macroeconomic conditions.

First, we document a significant improvement in Bayesian marginal likelihood compared to models that do not account for regime-switching, underscoring the importance of modeling distinct regimes in currency pricing. Second, the identified regimes and their shifts are closely tied to U.S. macroeconomic conditions, including a high-inflation regime, a low-inflation and low-interest-rate regime, and a low-inflation and high-interest-rate regime. Third, we find significant regime-dependent changes in factor selection and risk premia, indicating that characteristics-based models evolve substantially over time. Notably, the Carry factor remains dominant across all regimes, exhibiting a significant risk premium and high selection probability, while the importance of other factors, such as Value and Momentum, varies across regimes.

The empirical analysis also uncovers a regime risk anomaly that arises from cross-sectional differences in currencies' sensitivities to regime-switching risk, specifically regarding exposure to macroeconomic regime uncertainty risk. A long-short portfolio that buys the most sensitive currencies to regime-switching risk and sells the least sensitive currencies generates significant average excess returns, which cannot be explained by common currency factor models. Finally, we extend the tree-based approach to structural break models, capturing major economic events and institutional changes. These models identify breakpoints that align with the regime-switching model while offering new insights from multiple major events.

References

Accominotti, O., J. Cen, D. Chambers, and I. W. Marsh (2019). Currency regimes and the carry trade. *Journal of Financial and Quantitative Analysis 54*(5), 2233–2260.

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

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

Ang, A., G. Bekaert, and M. Wei (2008). The term structure of real rates and expected inflation. *Journal of Finance 63*(2), 797–849.

Ang, A. and J. Chen (2010). Yield curve predictors of foreign exchange returns. In *AFA 2011 Denver Meetings Paper*.

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 (2023). Mispricing and risk premia in currency markets. *Journal of Financial and Quantitative Analysis*, 1–93.

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.

Bianchi, D., M. Büchner, and A. Tamoni (2021). Bond risk premiums with machine learning. *Review of Financial Studies 34*(2), 1046–1089.

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.

Bryzgalova, S., M. Pelger, and J. Zhu (2024). Forest through the trees: Building cross-sections of stock returns. *Journal of Finance, Forthcoming*.

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. W., G. Feng, J. He, and X. He (2024). Growing the Efficient Frontier on Panel Trees. *Journal of Financial Economics, Forthcoming*.

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.

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.

Feng, G., J. He, N. G. Polson, and J. Xu (2024). Deep learning in characteristics-sorted factor models. *Journal of Financial and Quantitative Analysis 59*(7), 3001–3036.

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.

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

Gu, S., B. Kelly, and D. Xiu (2020). Empirical asset pricing via machine learning. *Review of Financial Studies 33*(5), 2223–2273.

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 (2023). 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.

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

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.

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. *Working Paper*.

Li, S. Z. and Y. Tang (2024). Automated volatility forecasting. *Management Science, Forthcoming*.

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.

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

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

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

Patton, A. J. and Y. Simsek (2023). Generalized autoregressive score trees and forests. *arXiv preprint arXiv:2305.18991*.

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

Pesaran, M. H. (2006). Estimation and inference in large heterogeneous panels with a multifactor error structure. *Econometrica 74*(4), 967–1012.

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.

Sarno, L., P. Schneider, and C. Wagner (2012). Properties of foreign exchange risk premiums. *Journal of Financial Economics 105*(2), 279–310.

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.

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 $\theta = [\alpha, \lambda, \sigma, \gamma]$. First, write

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

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

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

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

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

where $\widehat{\alpha}, \widehat{\lambda}, \widehat{\gamma}, \widehat{\sigma}$ are the corresponding posterior mean from Gibbs output, and $\alpha^{(r)}$, $\lambda^{(r)}$, $\gamma^{(r)}$, $\sigma^{(r)}$ 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 (2023) 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

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

Equation (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 Equation (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[13]

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

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 Equation (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 Equation (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*. Equation (A.3) encompasses all resulting leaf nodes and the entire dataset. This approach prevents myopic local

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

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 portfolio i ($1 \leq i \leq N$) in matrix form $\mathbf{R}_{j,i}$ and $\mathbf{Z}_{j,i}$, we have

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

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 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 $k = 1, 2, \cdots, K$.

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

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

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

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

Internet Appendix

(Not for Publication)

IA.I Simulation

Our Gibbs sampler in Appendix A.III uses a modified spike-and-slab prior, estimating alpha and idiosyncratic volatility for each individual, then estimating the same risk premia across the cross section. Hence, this method differs from conventional approaches (such as George and McCulloch (1993)) that solely rely on sampling along a single dimension, disregarding the cross-sectional information. To test the validity of our method, we conduct a simulation. The cross section of asset returns resembles the empirical properties of 47 currencies. We consider $K = 10$ characteristics, $N = 50$ assets, and $T = 600$ time periods. Specifically, we consider three kinds of factors: strong factors, weak factors, and useless factors, where the strong factors have large risk premia, the weak factors have smaller risk premia, and the useless factors have no risk premia when generating data.

The data-generating process is as follows. First, we fix the risk premia of 10 factors at $\lambda = [0.5, 0.4, 0.2, 0.1, 0.07, 0.05, 0.02, 0.01, 0, 0]$, then the first three factors are strong, the following five factors are weak, while the last two factors are useless. For each individual currency i, we first simulate α_i, σ_i, and ten characteristics level $\bar{z}_i$

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

then at each time t, we cross setionally generate characteristics $z_{i,t}$,rank them to evenly distributed in $[-1, 1]$ and generate $\epsilon_{i,t}$.

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

By multiplying the factor loadings (characteristics) with the risk premia at time t, and adding the intercept and innovation, we get the simulated returns $R_{i,t}$ for 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 have 600 $\times$ 50 currency returns observations. Each run contains 2000 MCMC samples, and the posterior inference of selection probability is based on the sample average. We repeat the simulation 100 times and take the average.

Table IA.1: **Simulation Results**

The table displays the simulation results for risk premia estimation and selection probability of 100 simulation runs. The factor consists of 3 strong factors,5 weak factors, and the remaining useless factors. The number is the sample average of 100 simulation runs' $\boldsymbol{\lambda}_k$ and $\boldsymbol{\gamma}_k (k = 1...16)$ from 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

Table IA.1 reports posterior factor risk premia estimate and selection probabilities of 100 simulation runs. We can see that factor risk premia are estimated accurately in this simulation study. These results validate our modified spike-and-slab Gibbs sampler. Regardless of the presence of strong, weak, useless, or their combination, our methodology can effectively detect them through sampling. The selection probabilities for strong factors consistently concentrate around 1, and those for weak factors may vary between 0 and 1. However, the useless factors are consistently unselected with very low selection probabilities close to 0.

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.

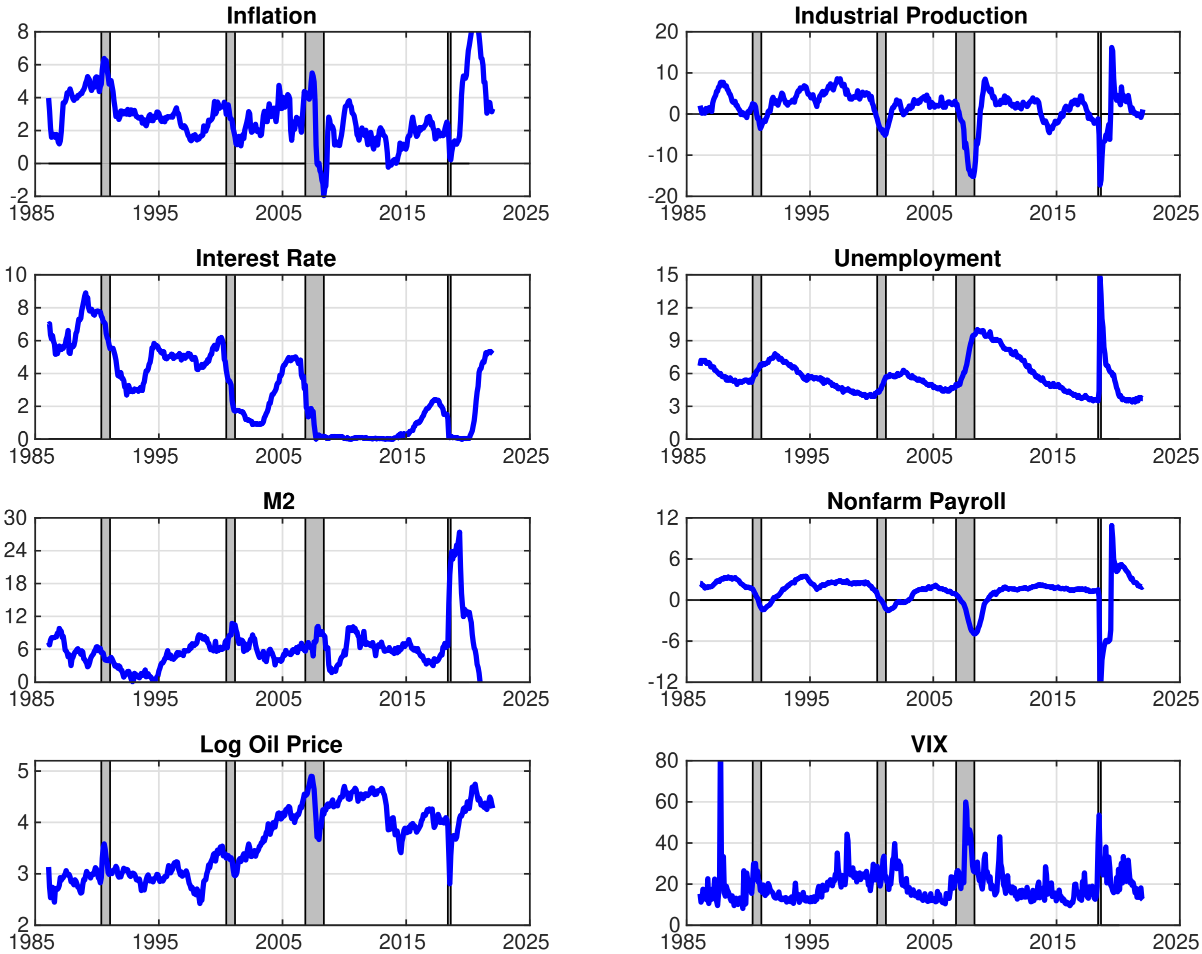

Table IA.2: **Description of Data Source**

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 shows summary statistics of currency factors based on all characteristics used, as discussed in Section 5. The currencies are sorted by corresponding characteristics. The first portfolio (P1) contains the top 20% of all currencies with low characteristics, whereas the last portfolio (P5) contains the top 20% of all currencies with high characteristics. The factor is a long-short strategy that buys P5 and sells P1. The DOL factor is the cross-sectional mean return of the currencies. The returns are expressed in percent, and the Neway-West t-ratios (Newey and West, 1987) are presented in parentheses. The table also reports the first-order autocorrelation coefficient (ac1), the annualized Sharpe ratio, and the annualized maximum drawdown. LDC shows low mean return and Sharpe ratio because the original paper (Della Corte et al., 2016) sorts LDC interact with NFA, while we treat them as separate characteristics. We only use the original LDC. The market beta factor shows even a negative mean in the full sample period since it's highly volatile and may show some power under certain conditions (Lettau et al., 2014). The monthly returns are expressed in percent. The sample spans the period from 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

Table IA.4: **Factor Risk Premia and Selection under Structural Breaks**

The table presents the estimation results for risk premia (%) and selection probabilities (%, in parentheses) for ten regimes identified by calendar month breaks (see Figure 7). Each regime group provides risk premia estimates based on prior selection probabilities: confident prior (w = 0.9), agnostic prior (w = 0.5), and skeptic prior (w = 0.1). The values represent the means of the MCMC post-convergence samples for λ and γ (shown in parentheses). The monthly sample is from 1986 to 2023.

Date	1986-01 to 1988-11			1988-12 to 1991-12			1992-01 to 1994-06			1994-07 to 1997-12			1998-01 to 2000-01		
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	w = 0.9	w = 0.5	w = 0.1	w = 0.9	w = 0.5	w = 0.1
Carry	0.352 (89)	0.146 (43)	0.023 (7)	0.315 (88)	0.225 (60)	0.127 (28)	0.487 (98)	0.517 (97)	0.480 (95)	-0.245 (87)	-0.172 (55)	-0.023 (5)	0.368 (88)	0.252 (59)	0.049 (11)
MOM1	-0.069 (66)	-0.004 (12)	0.003 (1)	0.037 (57)	0.012 (13)	0.007 (2)	0.370 (98)	0.391 (96)	0.290 (70)	0.433 (99)	0.388 (92)	0.212 (51)	0.109 (69)	0.058 (25)	0.030 (8)
MOM6	0.182 (79)	0.087 (34)	0.047 (12)	-0.080 (68)	-0.020 (18)	0.001 (1)	0.049 (64)	0.024 (17)	0.015 (4)	-0.262 (89)	-0.183 (61)	-0.044 (11)	0.469 (94)	0.354 (69)	0.082 (15)
MOM12	-0.158 (77)	-0.001 (24)	0.010 (4)	-0.054 (71)	-0.006 (21)	0.003 (2)	0.097 (70)	0.036 (22)	0.003 (2)	-0.061 (67)	-0.022 (19)	-0.010 (2)	-0.580 (97)	-0.377 (71)	-0.036 (7)
Value	-0.280 (81)	-0.164 (44)	-0.068 (13)	0.080 (67)	0.014 (16)	0.001 (2)	0.387 (97)	0.290 (81)	0.068 (16)	0.431 (98)	0.338 (84)	0.065 (15)	0.304 (92)	0.152 (49)	0.046 (11)
NFA	0.051 (59)	0.004 (12)	-0.003 (1)	0.061 (59)	0.012 (11)	0.006 (0)	-0.011 (47)	-0.007 (10)	-0.007 (1)	-0.028 (55)	-0.008 (11)	0.002 (1)	-0.058 (60)	-0.016 (12)	-0.009 (1)
LDC	-0.164 (77)	-0.102 (33)	-0.063 (12)	-0.055 (62)	-0.009 (15)	-0.001 (2)	-0.152 (73)	-0.085 (30)	-0.034 (7)	0.005 (62)	0.004 (13)	0.001 (1)	-0.015 (64)	-0.029 (19)	-0.009 (2)
LT Yield	-0.121 (79)	-0.030 (25)	-0.004 (2)	-0.046 (72)	0.010 (25)	0.018 (4)	-0.037 (72)	-0.083 (37)	-0.081 (16)	-0.146 (74)	-0.121 (41)	-0.038 (9)	0.015 (73)	0.050 (30)	0.036 (8)
Term Spread	-0.036 (66)	-0.006 (18)	-0.003 (2)	0.166 (77)	0.063 (26)	0.022 (6)	0.027 (58)	0.012 (14)	0.008 (1)	0.159 (75)	0.062 (27)	0.009 (1)	0.023 (71)	0.012 (23)	0.005 (2)
Volatility	0.158 (77)	0.079 (28)	0.049 (11)	0.231 (81)	0.047 (23)	0.007 (1)	0.089 (66)	0.059 (27)	0.022 (4)	0.102 (71)	0.016 (15)	0.004 (1)	-0.067 (72)	-0.007 (18)	0.001 (1)
MKT_Beta	-1.135 (98)	-1.108 (93)	-0.538 (42)	0.582 (97)	0.473 (89)	0.474 (83)	-0.201 (80)	-0.166 (49)	-0.125 (30)	0.285 (87)	0.114 (39)	0.017 (3)	0.597 (97)	0.570 (93)	0.603 (95)

Date	1998-01 to 2001-09			2001-10 to 2005-03			2005-04 to 2012-01			2012-02 to 2021-12			2022-01 to 2023-12		
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	w = 0.9	w = 0.5	w = 0.1	w = 0.9	w = 0.5	w = 0.1
Carry	0.759 (100)	0.702 (100)	0.658 (100)	0.676 (99)	0.649 (100)	0.630 (100)	0.341 (98)	0.340 (97)	0.344 (95)	0.023 (51)	0.013 (11)	0.015 (4)	1.150 (100)	1.141 (100)	1.220 (100)
MOM1	0.192 (83)	0.100 (40)	0.032 (7)	0.174 (85)	0.084 (37)	0.021 (3)	0.167 (90)	0.118 (54)	0.075 (23)	-0.014 (36)	-0.006 (5)	-0.006 (0)	0.038 (63)	0.009 (13)	-0.004 (3)
MOM6	-0.246 (86)	-0.049 (22)	-0.006 (1)	-0.494 (99)	-0.390 (95)	-0.222 (57)	0.249 (96)	0.239 (83)	0.158 (47)	0.116 (79)	0.031 (16)	0.013 (1)	-0.557 (98)	-0.532 (91)	-0.278 (49)
MOM12	0.213 (84)	0.055 (25)	0.015 (3)	0.220 (83)	0.080 (33)	0.009 (3)	-0.533 (100)	-0.517 (100)	-0.454 (100)	-0.159 (90)	-0.060 (29)	-0.033 (5)	0.006 (75)	-0.001 (26)	-0.005 (4)
Value	0.036 (57)	0.009 (12)	0.002 (1)	0.175 (81)	0.038 (18)	0.012 (1)	0.026 (46)	0.012 (8)	0.007 (1)	0.218 (99)	0.213 (94)	0.159 (64)	-0.059 (69)	-0.020 (21)	-0.001 (2)
NFA	-0.008 (43)	-0.003 (8)	-0.003 (1)	-0.024 (47)	-0.008 (8)	-0.005 (1)	0.004 (37)	0.002 (6)	0.003 (0)	-0.053 (56)	-0.034 (13)	-0.030 (1)	-0.831 (100)	-0.899 (100)	-0.972 (100)
LDC	0.139 (73)	0.049 (23)	0.019 (3)	0.081 (63)	0.028 (16)	0.009 (1)	0.024 (48)	0.008 (8)	0.002 (0)	0.122 (80)	0.055 (28)	0.026 (3)	0.304 (82)	0.149 (39)	0.097 (16)
LT Yield	-0.243 (84)	-0.175 (54)	-0.080 (20)	-0.063 (62)	-0.022 (17)	-0.005 (1)	0.011 (51)	0.009 (12)	0.003 (1)	0.330 (98)	0.316 (98)	0.283 (88)	1.080 (98)	1.021 (94)	0.511 (52)
Term Spread	-0.082 (67)	-0.036 (21)	-0.012 (3)	0.058 (58)	0.018 (15)	0.007 (1)	-0.036 (50)	-0.013 (10)	-0.009 (1)	0.126 (82)	0.080 (42)	0.031 (3)	0.826 (99)	0.797 (96)	0.500 (64)
Volatility	-0.052 (61)	-0.012 (11)	-0.006 (1)	0.164 (78)	0.033 (18)	0.008 (2)	0.113 (71)	0.026 (14)	0.014 (2)	0.083 (67)	0.032 (17)	0.012 (0)	0.072 (70)	0.011 (22)	-0.002 (1)
MKT_Beta	-0.086 (66)	-0.013 (12)	-0.003 (1)	-0.608 (99)	-0.726 (100)	-0.706 (100)	0.107 (67)	0.018 (12)	0.009 (1)	0.394 (99)	0.347 (98)	0.263 (80)	1.272 (100)	1.242 (100)	1.046 (89)