VIX Futures Term-Structure and Currency Returns

Mohammed Mehdi Kaebi[1]
Insper

Abstract

I use the VIX futures term-structure to build novel currency risk factors and show that they are priced in the cross-section of currency returns and provide information for return predictability. Currencies more exposed to the level and curvature risk-factors offer a lower risk premium, acting as an insurance, while those more exposed to the slope factor offer a higher risk-premium. A portfolio that buys the high risk factor exposure currencies and shorts the low risk factor exposure currencies captures risk-return dispersion, and the excess returns of these strategies can be understood as compensation for a globally traded shock. I rationalize these results with a simple model where experts have volatility hedging demands and arbitrageurs face funding and margin costs.

Keywords: Exchange rates, currency returns, VIX futures, international finance, asset pricing.

JEL Classification: G12, G15, F31.

---

[1]Quatá Street, 300 – Vila Olímpia, São Paulo – SP, Brazil. Email: mohammedmk@al.insper.edu.br. I thank Massimiliano Bondatti, Christopher Polk, Ruy Ribeiro, Marco Bonomo, participants of Insper's Macro & Finance Workshop and participants of the 32nd AEFIN Finance Forum for comments and feedback on this paper. I also thank the São Paulo Research Foundation (FAPESP) for their support under grants 2024/01046-9 and 2024/16089-5, and ANBIMA for their support under the 2024 XX ANBIMA Capital Markets Prize (2024 XX Prêmio ANBIMA de Mercado de Capitais). This version: December 2025.

1. Introduction

I propose novel tradable currency risk factors using information from the VIX futures term-structure, showing that they are priced in the cross-section of currency excess returns and provide incremental information for currency return prediction. Using the first three principal components of the VIX futures term-structure, namely the level, slope and curvature, I document that currencies have heterogeneous exposures (betas) to innovations in these components. Sorting currencies on these exposures yields high-minus-low beta sorted portfolios that capture systematic risk–return differentials. These portfolios proxy for the global component of the volatility term-structure, which is the priced source of volatility risk, and they provide information beyond the carry and dollar factors of Lustig et al., 2011.

In Fama and MacBeth, 1973 regressions, I find that the currency risk factors built on the level and curvature of the volatility term-structure act as an insurance, demanding payment of a risk premium of about -5% to -8% per year with a Sharpe ratio of -0.9. The risk factor built on the slope of the term-structure, on the other hand, pays a positive risk premium of about 6% per year with a Sharpe ratio of 0.8. These risk factors provide incremental information to the carry and dollar factors and are not tilted toward either emerging or advanced currencies, displaying stable and balanced portfolio compositions over time. In predictive regressions for next-month currency returns, I find that innovations to the volatility term-structure components provide statistically significant predictive information even when controlling for lagged returns, in both pooled and currency fixed-effects specifications.

These empirical findings are rationalized with a simple two-agent model in which horizon-specific volatility hedging demand interacts with intermediary supply frictions to generate prices of volatility term-structure risk. In the model, experts hold risky assets whose returns covary with shocks to expected volatility at different horizons, and they hedge these exposures using volatility-linked contracts that are traded only in the Home economy. Their inelastic demand to volatility hedging across horizons creates maturity-specific imbalances in the market for these claims. Arbitrageurs absorb these imbalances but face maturity-dependent capital costs, so the equilibrium prices of volatility term-structure risk arise from a scarcity mechanism analogous to preferred-habitat and intermediary-based asset pricing models. Countries differ in how their currency co-move with the global volatility term-structure component, leading to heterogeneous exposures across currencies and therefore to currency risk premia.

The volatility term-structure is empirically proxied by the VIX futures term-structure. VIX futures are forward contracts on 30-day equity-market implied volatility and constitute the primary tradable instruments for accessing volatility-risk exposures across maturities. A large literature shows that these contracts embed both forecasts of future realized volatility and, importantly, variance-risk premia. Zhou, 2018 provides a comprehensive overview of variance-risk premia across maturities and asset classes. Johnson, 2017 demonstrates that movements in the level, slope, and curvature of the VIX futures term-structure largely reflect compensation for bearing systematic volatility risk, while Aït-Sahalia et al., 2020 show that investors exhibit a strong and maturity-dependent willingness to insure against volatility risk, especially after market downturns. Consistent with this interpretation, Dew-Becker et al., 2017 show that VIX futures and variance swaps command time-varying variance-risk premia across maturities, and Egloff et al., 2010 show that movements in the term-structure of variance swaps are driven by multiple components of priced volatility risk. In another dimension, Bollerslev et al., 2009 document that variance-risk premia help explain cross-sectional differences in expected stock returns, while Londono and Zhou, 2017 show that variance-risk premia help explain currency return predictability and the forward-premium puzzle.

Institutional investors, such as market-makers, volatility arbitrage funds, structured product desks, and multi-asset managers use VIX futures because they provide linear and standardized access to forward implied volatility and function as effective hedges for volatility-risk exposures arising from their portfolios. These institutions routinely face maturity-specific volatility sensitivities that cannot be neutralized efficiently using only options. VIX futures allow them to offload these risks by taking positions along the volatility term-structure: buying near-dated contracts to insure against volatility spikes, using calendar spreads to hedge differences between short- and long-horizon expectations, and employing butterfly structures to manage curvature risk. In practice, VIX futures serve as one of the main vehicles through which volatility risks of different horizons are transferred across investors, generating persistent horizon-specific demand that governs the shape and dynamics of the volatility term-structure.

Intuitively, if currencies co-move with innovations in the volatility term-structure, they can effectively act as a hedge against volatility risk premia and expected future volatility. Therefore, I can empirically assess the pricing implications for currency excess returns. I show that if a country's volatility term-structure contains a global component, then a currency's heterogeneous exposures to term-structure factors can help extract the global component of volatility term-structure risk, exploring a risk-based explanation of exchange rate puzzles. Previous research has tried to explore a similar relationship, but using other measures of risk. Looking at tail risk and the cross-section of currency returns, Fan et al., 2022 build an equity tail risk factor using S&P500 options data and show that it is priced in the cross-section of currency returns. Others have also explored the predictive power of equity data to FX markets, where Londono and Zhou, 2017 show that the US equity variance risk premium helps predict currency returns in the time-series. Complementary evidence shows that equity downside risk helps explain the cross-section of currency returns, as in Lettau et al., 2014 and Dobrynskaya, 2014. Despite these findings, the way investors price volatility and variance risk premia across horizons, and the implications for currency market dynamics, remains underexplored.

I address this gap in the literature by exploring how the VIX futures term-structure influences currency excess returns. The key question of this paper is whether innovations to the volatility term-structure are priced in the cross-section of currency returns and if they provide information for currency return predictability. This paper aims to make several contributions to the existing literature. First, it proposes novel global currency risk factors from volatility derivatives that are priced in a broad cross-section of currency returns, contributing to the literature that tries to understand the cross-section of currency returns. Second, it demonstrates that country exposure to the volatility term-structure can identify global components in currency-level risk using a simple two-agent model. Third, it establishes a tradable currency portfolio built on exposures to equity-market derivatives factors, enhancing our understanding of the interconnection between equity and currency markets. By constructing novel currency risk factors based on equity derivatives, this paper provides a comprehensive analysis that not only enhances our understanding of currency market dynamics but also offers practical insights for academics and practitioners.

2. Model

I develop a simple stylized two-agent model linking horizon-specific volatility hedging demand to currency risk premia in a world with a large Home economy and smaller Foreign economies. Experts in all countries have institutional, inelastic hedging demand for volatility at multiple horizons, while arbitrageurs are risk-neutral intermediaries who absorb residual volatility positions subject to maturity-dependent balance-sheet and margin costs. Only the Home economy has a deep market for tradable volatility-linked contracts, and these contracts span the global component of volatility term-structure risk. Foreign experts also trade these Home claims to hedge their exposure to global volatility shocks. When expert hedging demand differs from outstanding supply, arbitrageurs must take net positions, generating scarcity in volatility claims and equilibrium prices of risk for level, slope, and curvature shocks. Differences in country exposures to global and local components of volatility term-structure risk then translate into currency risk premia.

The modeling framework builds on several strands of the asset-pricing literature. First, I adopt the preferred-habitat and demand-based equilibrium logic of Vayanos and Vila, 2021, in which inelastic investor demand interacts with intermediary balance-sheet constraints to generate risk premia. Second, the supply side of the model follows the intermediary-based asset-pricing tradition of Gârleanu and Pedersen, 2011 and He and Krishnamurthy, 2013, where arbitrageurs face maturity-dependent capital costs and absorb investor demand at a price that reflects their risk-bearing capacity. Investors' desire to hedge volatility shocks across horizons is consistent with the broader literature documenting that volatility is an important priced risk, such as Drechsler and Yaron, 2010. Finally, in linking volatility hedging to exchange rates, I follow the international asset-pricing logic of Lustig et al., 2011 and Verdelhan, 2018, where currency premia arise from differences in countries' stochastic discount factor (SDF) loadings on global risk factors. The model integrates these ideas to deliver a simple mechanism through which horizon-specific volatility hedging and intermediary frictions jointly generate prices of risk for volatility term-structure components and, ultimately, currency risk premia.

2.1. Intuition

Experts, such as institutional investors, insurers, and pension funds, face volatility-sensitive liabilities and portfolio exposures that give rise to hedging demand at specific horizons buckets. They desire insurance against unexpected movements in volatility across different horizons. In the model, these hedging needs are represented by three maturity buckets of volatility-linked claims, which correspond to shifts in the level, slope, and curvature of the volatility term-structure. Arbitrageurs intermediate these trades but face maturity-dependent funding and margin costs. When hedging demand is concentrated at specific horizons and arbitrageur supply is costly, those volatility buckets become scarce and command nonzero prices of risk. Because the Home economy hosts the only deep tradable volatility term-structure market, its volatility-linked claims span both domestic and global volatility components, and experts globally use these claims to hedge global volatility shocks. Local volatility shocks are not spanned by traded contracts and therefore do not earn risk premia. Exchange rates inherit these priced global volatility term-structure risks through differences in the SDF.

2.2. Assets

All assets are traded at $t = 1$ and deliver payoffs at $t = 2$. In each country $c \in \{H, F\}$ there is a domestic risky asset with excess payoff $A^{(c)}$ and a domestic riskless asset yielding $1 + r^c$. Only the Home economy has a liquid market for traded volatility-linked claims. I represent their maturity-bucket payoffs by $TS = (TS_L, TS_S, TS_C)'$ and their respective innovations as $\Delta TS = (\Delta TS_L, \Delta TS_S, \Delta TS_C)'$[2], where L, S, C denote buckets relating to the level, slope and curvature of volatility-linked claims.

In an interconnected world, I assume that the innovations to the volatility term-structure of each country contains a global and local component

$$\Delta TS_t^{(c)} = \varphi_c \Delta TS_t^g + \Delta TS_t^{\mathrm{loc},c}, \tag{1}$$

where $\varphi_c \geq 0$ measures how much country c loads on global volatility term-structure shocks ΔTS^g, and $\Delta TS^{\mathrm{loc},c}$ is a local volatility term-structure shock.

2.3. Agents

The model is composed by experts and arbitrageurs, where experts perform asset allocation in risky and riskless assets, and wish to hedge their volatility exposures. Arbitrageurs supply the volatility claims to experts, but face capital and margin costs.

Experts in the Home and Foreign economies begin by allocating wealth between risky assets and the domestic riskless asset using standard mean–variance principles. This first-stage portfolio choice, which I do not model explicitly, generates overall portfolio exposure to volatility term-structure shocks, because the returns of risky assets co-move with innovations in the volatility term-structure. Let q_A denote the vector of risky-asset holdings chosen in this step. The resulting volatility exposure of the risky portfolio is captured by the covariance term $\Sigma_{TS,A} q_A$, where $\Sigma_{TS,A} = \mathrm{Cov}(TS, A)$ measures how risky-asset payoffs co-move with the volatility buckets.

In the second stage, experts use the Home volatility-linked contracts TS to hedge this portfolio-level volatility exposure across horizons. Their liability structures and institutional constraints impose inelastic[3] horizon-specific hedging mandates $h = (h_L, h_S, h_C)'$ for level,

[2]Since TS spans ΔTS, there exists a nonsingular matrix B such that $TS = B\Delta TS$. Normalizing $B = I$ yields $TS = \Delta TS$ and $\mathrm{Var}(TS) = \Sigma_{\Delta TS}$.

[3]Hedging demand is assumed inelastic for simplicity. In Appendix A I provide a simple extension with a downward sloping hedging demand.

slope and curvature volatility buckets. Experts therefore choose q_{TS} to align their net volatility $\Sigma_{TS,A}q_A + \Sigma_{\Delta TS}q_{TS}$ with the target h by solving a tracking-error minimization problem

$$\min_{q_{TS}} \|\Sigma_{TS,A}q_A \ + \ \Sigma_{\Delta TS}q_{TS} \ - \ h\|^2 . \tag{2}$$

This two-step structure follows the delegated-risk-management logic of Gârleanu and Pedersen, 2011 and the preferred-habitat interpretation of Vayanos and Vila, 2021, where portfolio exposures induce inelastic hedging needs in specific volatility directions. Solving[4] this optimization problem, experts choose

$$q^*_{TS} = \Sigma^{-1}_{\Delta TS}(h - \Sigma_{TS,A}q_A). \tag{3}$$

Since q_A is fixed and $\Sigma_{TS,A}q_A$ is absorbed into the experts' effective hedging demand, I simplify notation by writing the reduced-form inelastic demand as $q^*_{TS} \equiv h$. Experts in all countries use the Home volatility claims TS to hedge global volatility shocks.

Arbitrageurs are risk-neutral intermediaries who take net positions in TS to absorb demand imbalances, but face maturity-dependent capital and margin costs:

$$C(x) = \frac{\kappa}{2}x'Wx. \tag{4}$$

x represents the vector of positions in the three volatility-linked claims. $\kappa > 0$ scales overall intermediary tightness and captures balance-sheet, margin, and capital constraints that make supplying volatility exposure costly. A high κ means arbitrageurs are capital constrained or liquidity-constrained, so supplying volatility-linked contracts is expensive, while a small κ means arbitrageurs are unconstrained and volatility hedging is cheaply absorbed. $W \succ 0$ is a matrix that encodes maturity-specific supply frictions, where its diagonal elements govern the cost of supplying each bucket individually, while the off-diagonal elements capture cross-horizon interactions[5].

[4] $\Sigma_{\Delta TS}$ is positive definite and full rank, therefore invertible. So the minimization problem yields a unique solution with zero tracking error.

[5] The matrix W captures maturity-specific and cross-horizon supply costs. Large diagonal entries reflect that some buckets (e.g., short maturity) are more capital- or margin-intensive to supply, while off-diagonal terms allow exposure in one bucket to affect the cost of supplying another. This parsimonious structure follows preferred-habitat and intermediary-based models.

This quadratic specification yields the standard linear supply function

$$x = \kappa^{-1} W^{-1} p_{TS}, \tag{5}$$

used in preferred-habitat and intermediary-based asset-pricing models such as Vayanos and Vila, 2021 and Gârleanu and Pedersen, 2011. Prices are normalized so that $p_{TS} = \lambda_{\Delta TS}$ is the price-of-risk vector for ΔTS.

2.4. Market Clearing and Stochastic Discount Factor (SDF)

Let Θ denote the exogenous net supply of the volatility-linked claims TS. h denotes experts' reduced-form target position (hedging demand), where $h > 0$ corresponds to experts wanting to be long volatility claims and $h < 0$ to experts wanting to be short, and x denotes the net position held by arbitrageurs.

Market clearing requires that the total quantity held by experts and arbitrageurs equals the outstanding supply. Equivalently, arbitrageurs must absorb the residual position implied by expert demand relative to supply:

$$x = h - \Theta \tag{6}$$

Arbitrageur positions are constrained by balance-sheet costs. Using the arbitrageur supply function in Equation 5, market clearing implies

$$\kappa^{-1} W^{-1} \lambda_{\Delta TS} = h - \Theta. \tag{7}$$

Solving yields the equilibrium price of volatility term-structure risk,

$$\lambda_{\Delta TS} = \kappa W (h - \Theta). \tag{8}$$

Equation 8 shows that each volatility term-structure bucket price of risk is proportional to the net scarcity faced by arbitrageurs, measured by the residual demand $h - \Theta$ that must be absorbed on their balance sheets, and scaled by the maturity-dependent supply frictions encoded in W. An increase in expert hedging demand h, holding supply fixed, raises the amount of risk intermediaries must bear and therefore increases the price of risk.

This equilibrium relation is structurally analogous to the preferred-habitat and intermediary-based pricing mechanisms in Vayanos and Vila, 2021, Gârleanu and Pedersen, 2011, and related literature, in which prices of risk are determined by investor demand imbalances relative to intermediary risk-bearing capacity. In the present setting, horizon-specific volatility hedging demand interacts with maturity-dependent supply constraints to generate the vector of volatility term-structure prices of risk $\lambda_{\Delta TS}$.

Any shock carrying a price of risk appears linearly in the log SDF. In addition to volatility term-structure risk, SDFs generally load on broader macroeconomic and financial risks. I therefore incorporate two standard reduced-form sources of discount-factor volatility: a global shock u_g, capturing worldwide fluctuations in economic activity, financial conditions, or global marginal utility; and a country-specific shock u_c, capturing domestic variation in consumption, productivity, or financial frictions. These shocks are not microfounded within the term-structure hedging mechanism above, but represent the usual non-volatility factors appearing in international SDF specifications, such as Lustig et al., 2011. They are assumed orthogonal to ΔTS and ensure that the SDF captures the full set of macroeconomic risks relevant for exchange-rate determination. Country c's log nominal SDF at $t = 2$ is thus:

$$-m_{c,2} = i_{c,1} + a_{c,1} + \delta_c u_{c,2} + \gamma_c u_{g,2} + \theta_c' \Delta TS_2^{(c)}, \tag{9}$$

where $i_{c,1}$ is the risk-free rate, $a_{c,1}$ is a normalization constant, (δ_c, γ_c) are the loadings on the country-specific (u_c) and global (u_g) shocks, and θ_c determines how the SDF loads on volatility term-structure innovations.

2.5. Currency Pricing, Conditional Betas, and Risk Factors

Following the Asset Market View of exchange rates, as in Backus et al., 2001, the log change in the nominal exchange rate (s_t) between the Foreign and Home currency is equal to the difference of the log nominal SDFs of the two countries, where the exchange rate is expressed in foreign currency units per one unit of Home currency:

$$\Delta s_{F,2} = m_{H,2} - m_{F,2}. \tag{10}$$

Thus the excess return to a Home investor who holds foreign currency is

$$\begin{aligned} rx_{F,2} &= -\Delta s_{F,2} + i_{F,1} - i_{H,1} \\ &= a_{H,1} - a_{F,1} + \delta_H u_{H,2} - \delta_F u_{F,2} + (\gamma_H - \gamma_F) u_{g,2} + \theta_H' \Delta TS_2^{(H)} - \theta_F' \Delta TS_2^{(F)} \end{aligned} \tag{11}$$

Decomposing the volatility term-structure shock into a global and local component, as in Equation 1, the excess return becomes:

$$
\begin{aligned}
rx_{F,2} = a_{H,1} - a_{F,1} - & \overbrace{\left(\delta_F u_{F,2} + \theta_F' \Delta TS_2^{loc,F}\right)}^{\text{Foreign country shocks}} + \overbrace{\left(\delta_H u_{H,2} + \theta_H' \Delta TS_2^{loc,H}\right)}^{\text{Home country shocks}} \\
+ & \underbrace{\left(\varphi_H \theta_H' - \varphi_F \theta_F'\right) \Delta TS_2^{g}}_{\text{Global volatility term-structure shock}} + \underbrace{\left(\gamma_H - \gamma_F\right) u_{g,2}}_{\text{Other global shocks}} .
\end{aligned} \tag{12}
$$

So, the currency excess return for a Home investor who invests in foreign currency F is influenced by several factors. These include shocks related to both Foreign and Home SDFs, as well as local components of the volatility term-structure. Additionally, it is impacted by shocks to the global component of the SDF and the global component of the volatility term-structure. For international volatility term-structure pricing, however, only the global component will earn a risk premium, since the local components can be diversified away. Therefore, the loading on the country c SDF will be

$$
\theta_c = \lambda_{\Delta TS}. \tag{13}
$$

So, the global volatility term-structure shock term can be written as

$$
\left(\varphi_H \theta_H' - \varphi_F \theta_F'\right) \Delta TS_2^g = \left(\varphi_H - \varphi_F\right) \lambda_{\Delta TS}' \, \Delta TS_2^g, \tag{14}
$$

such that the volatility-driven contribution to the excess return is proportional to the cross-country exposure wedge $(\varphi_H - \varphi_F)$ times the priced global shock $\lambda_{\Delta TS}' \Delta TS_2^g$.

However, the global component of volatility term-structure shock is not observable. Instead, I observe the Home volatility term-structure innovations $\Delta F^{(H)}$, which itself contains the global component. Since the Home volatility claims market is large and liquid, $\Delta TS^{(H)}$ spans ΔTS^g, while the local components can be diversified in the cross-section. Therefore, the conditional beta of currency c with respect to the observable Home volatility term-structure innovations,

$$
\boldsymbol{\beta}_{\Delta TS,c} = \mathrm{Cov}(rx_c, \Delta TS) \left[\mathrm{Var}(\Delta TS)\right]^{-1}, \tag{15}
$$

is an empirically valid proxy for the unobservable beta with respect to the global volatility term-structure component ΔTS^g. Moreover, since

$$\beta_{\Delta TS^g,c} \propto (\varphi_H - \varphi_c)\lambda_{\Delta TS}, \tag{16}$$

sorting currencies by $\boldsymbol{\beta}_{\Delta TS,c}$ produces the same ordering as sorting by $\beta_{\Delta TS^g,c}$. Hence long–short portfolios formed from $\boldsymbol{\beta}_{\Delta TS,c}$ capture exposure to the priced global volatility term-structure factor. This follows because the Home term-structure spans the global component of volatility risk, while country-specific local components are diversified away. The construction is analogous to the global factor extraction methods in Verdelhan, 2018 and Lustig et al., 2011.

Therefore, the long-short portfolio of buying high volatility term-structure beta currencies and selling low volatility term-structure beta currencies will be a proxy for the global component of volatility term-structure risk:

$$fTS_t = \frac{1}{N_T}\sum_{c\in T} rx_{c,t} - \frac{1}{N_B}\sum_{c\in B} rx_{c,t}, \tag{17}$$

where N_T and N_B represent the number of currencies in the high and low beta portfolios, respectively. With a high enough number of currencies, each portfolio leg is diversified enough to make the portfolio dominated by the global volatility term-structure component and not by the local components. Therefore, this simple model suggests that long-short beta-sorted currency portfolios can be used to proxy for the global volatility term-structure factor. In the remainder of this paper, I use the US as the Home country and take the perspective of the US investor, although this framework allows any reference currency to be used.

3. Data, risk factor formation and portfolio composition

3.1. Data

I obtain spot (s_t, in log) and 1-month forward exchange rates (f_t, in log) with respect to USD from Barclays Bank International (BBI), World Markets Company/Refinitiv (WMR) and Refinitiv, all via Datastream. I set the order of preference as WMR being the most preferable source and BBI as the least preferable, and I change the source of the data whenever it becomes available from a preferred source, as is commonly done in the literature. I also get data on real effective exchange rates from the IMF International Financial Statistics.

The exchange rate data is for 48 countries[6]: Australia, Austria, Belgium, Brazil, Bulgaria, Canada, Croatia, Cyprus, Czech, Denmark, Egypt, Euro, Finland, France, Germany, Greece, Hong Kong, Hungary, Iceland, India, Indonesia, Ireland, 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, Turkey, United Kingdom, Ukraine. Appendix B provides further details on the data, such as start and end date for each currency, development group (advanced or emerging), Datastream mnemonics and periods of covered interest parity violations.

I use daily VIX futures contracts with maturities from one to nine months to capture the volatility term-structure. VIX futures are contracts on S&P500 implied volatility and the data are sourced from Datastream. Trading in VIX futures begins in March 2004 with maturities up to four months, while contracts with longer maturities are introduced gradually through 2006. To ensure a balanced maturity structure, the baseline sample begins in October 2006 and ends in March 2024.

Given the shorter availability of exchange-traded VIX futures, I implement two robustness datasets. Appendix E uses option-replicated VIX futures from Johnson, 2017, constructed from S&P500 index options following the CBOE VIX methodology, providing a daily synthetic VIX futures series from January 1996 to July 2019. Appendix F uses monthly forward variance claims on S&P500 variance inferred from variance-swap quotes, spanning December 1995 to September 2013, obtained from Dew-Becker et al., 2017.

Figure 1 displays selected VIX futures term-structures, and Figure 2 plots the complete time-series across maturities. Under normal conditions, the term-structure is upward sloping, with longer maturities priced above shorter ones, reflecting the mean-reverting nature of volatility and the premium investors demand for holding long-term volatility exposure. During periods of market stress the entire curve shifts upward and flattens or inverts as short-term contracts surge relative to longer maturities. These episodes signal elevated near-term uncertainty and a compression of volatility risk premia, as short-term protection becomes disproportionately expensive and the usual compensation for bearing long-horizon volatility risk temporarily vanishes.

[6] The set of currencies effectively included depends on the rolling window length specified in Subsection 3.2, which determines the starting point at which the risk factor becomes available. The list shown here represents the full set of currencies that can potentially enter the construction. Additionally, I follow Kroencke et al., 2014 and Corte et al., 2016, among others, and leave out indicated countries due to large deviations of covered interest rate parity. See Table B.3 in Appendix B for details.

[Figure 1 about here.]

[Figure 2 about here.]

3.2. *Construction of VIX term-structure currency risk factors*

In order to extract the common components that drive variations in the VIX futures term-structure, I compute an expanding-window Principal Component Analysis (PCA)[7] on daily VIX futures term-structures, and use the first three principal components sampled at month-end. Signs of the factor loadings are chosen such that increases in PC1 and PC3 and decreases in PC2 correspond to the standard stress-period pattern of higher overall volatility (shift of the curve), front-end spikes (flattening or inversion of the curve), and greater curvature in the term-structure. In order to ensure enough data for the principal components, I start with a window of one year. Then at each following date I use all available past data, so that after the initial startup period the procedure has no look-ahead bias.

Figure 3 shows the factor loadings of each contract on the first three principal components using the full-sample (last-date) decomposition. Together, they explain about 99% of the variation in the curve. All expirations load positively and by similar amounts on PC1, identifying it as the level factor. PC2 displays increasing loadings across maturities, capturing the slope. PC3 loads positively on the short and long ends and negatively in the middle, identifying it as the curvature factor. Figure 4 shows that PC1 exhibits sharp upward spikes during stress periods (2008–2009, 2011, 2020), consistent with its interpretation as the level of the term-structure. Outside those episodes, PC1 remains near a lower, more stable baseline. PC2 steps down markedly in the same stress periods, reflecting short-end implied volatility rising relative to the long end, and then gradually reverts upward as conditions stabilize. PC3 shows brief, high-amplitude movements only around the major dislocations, indicating temporary increases in curvature when the short and long maturities move jointly against the middle.

[Figure 3 about here.]

[7] Following the term-structure literature (Joslin et al., 2011), let Y contain n maturities and let W be the $n \times 3$ matrix of the first three eigenvectors of $\mathrm{cov}(Y)$. Rows correspond to maturities ordered from shortest ($i = 1$) to longest ($i = n$). Signs are fixed by enforcing positive loadings on the longest maturity. The components are then scaled into canonical level, slope, and curvature factors using Level scale $= \sum_{i=1}^{n} W_{i1}$, Slope scale $= W_{n,2} - W_{1,2}$ and Curvature scale $= W_{n,3} - 2W_{m,3} + W_{1,3}$, where m indexes the central maturity. The principal components are divided by these scalars so that PC1 corresponds to a parallel shift, PC2 to a tilt, and PC3 to a second-difference movement.

[Figure 4 about here.]

To evaluate whether these common volatility term-structure factors are priced in the cross-section of currency returns, I sort currencies into five[8] portfolios based on their lagged exposure to innovations in the term-structure principal components, $\beta_{TS^j,i}$. For that, I compute each currency's monthly log excess return of buying a foreign currency in the forward market and then selling it in the spot market after one month, $rx_t = f_{t-1} - s_{t-1} - \Delta s_t$, and estimate the following rolling 48-month[9] regression:

$$rx_{i,t} = \alpha_i + \beta_{DOL,i} DOL_t + \beta_{TS^j,i} \Delta TS_t^j + \varepsilon_{i,t}, \qquad j \in TS \tag{18}$$

where DOL is the Dollar factor proposed by Lustig et al., 2011, used to control for the overall variation in currency excess returns, and $TS = \{\text{PC1}, \text{PC2}, \text{PC3}\}$ represents the principal components extracted from VIX futures term-structure and Δ denotes the monthly variation.

Then, at the end of month t, I sort currencies based on $\beta_{TS^j,i}$ into five portfolios and compute the equal-weighted currency excess return of each portfolio in the following month. $\beta_{TS^j,i}$-sorts provide tradable diversified currency portfolios that proxy for the global component of volatility term-structure risk. In terms of the model, these portfolios group currencies with similar exposures to the priced global volatility term-structure factor. The proposed currency risk factor for each VIX futures term-structure common component is thus the high-minus-low (HML) beta-sorted portfolio, denoted as $fPC1$ for the level factor, $fPC2$ for slope and $fPC3$ for curvature.

Figure 5 shows wide and persistent dispersion in exposures across the five currency portfolios, with a clear and rather stable ranking for each component. For level (PC1), the spread between high- and low-beta portfolios remains broadly steady over most of the sample, narrowing only temporarily around major volatility events. Slope (PC2) betas exhibit similar patterns, with a noticeable compression of betas during the 2014–2016 period and again around 2020, after which the dispersion gradually re-emerges. Curvature (PC3) shows the largest swings, but even there the overall pattern is similar.

[Figure 5 about here.]

Summary statistics for the beta-sorted currency portfolios are shown in Table 1. In all

[8] In Appendix D.1 I repeat the same analysis but with currencies sorted into three portfolios and in Appendix D.2 with currencies sorted into ten portfolios. Results are analogous.

[9] Results are robust to changes in the rolling window and are available upon request.

three dimensions, the portfolios display clear monotonic patterns in average excess returns and Sharpe ratios that align with the underlying beta rankings. For the level (PC1) and curvature (PC3) sorts, expected returns and Sharpe ratios decline steadily as beta increases, while the slope (PC2) sort shows the opposite pattern. These patterns indicate that currencies with stronger exposure to each volatility term-structure factor earn systematically different excess returns, consistent with compensation for exposure to priced shocks. In each case, the return differential is concentrated in the extremes: the highest-beta portfolios for level (PC1) and curvature (PC3), and the lowest-beta portfolio for slope (PC2), generate the bulk of the cross-sectional spread. The resulting HML portfolios deliver large average excess returns of around 7% annualized in absolute value, with statistically significant t-statistics and strong Sharpe ratios. Portfolio skewness is negative across most portfolios, consistent with downside tail risk typical of currency strategies, but the HML slope (PC2) portfolio exhibits positive skewness.

[Table 1 about here.]

The results in Table 2 show that the returns of the beta-sorted portfolios load on the Carry factor of Lustig et al., 2011, and several of these loadings are statistically significant. However, the estimated coefficients are uniformly well below one in magnitude. This means that even when the Carry factor helps explain part of the variation in these portfolios, the movements are far from one-to-one, and the factor does not account for the dominant share of their returns. The HML portfolios also load significantly on Carry, yet the coefficients remain modest. Taken together, these results show that the beta-sorted portfolios are not simply repackaging the traditional carry trade, they contain variation linked to exposures to volatility term-structure rather than being driven primarily by interest rate differentials.

[Table 2 about here.]

3.3. Composition of currency risk factors

It is important to understand the underlying portfolio compositions over time in order to see how the portfolios behave and what type of information they capture. Figure 6 shows the portfolio assignment for each beta-sorted portfolio over time. There is substantial time-variation in how individual currencies are ranked by their exposures to innovations in the three VIX term-structure components. Although some broad patterns appear, for example, traditional safe-haven currencies such as JPY and CHF often sit in the high-beta portfolios for level (PC1), while higher-yielding emerging-market currencies more frequently populate the low-beta groups, the assignments are far from static. ISK is a striking example: it alternates between low-, mid-, and high-beta portfolios for level (PC1) and slope (PC2), but appear more consistently in higher-beta portfolios for curvature (PC3). JPY, CHF, and DKK, by contrast, show stable rankings for level (PC1) but much more variation in their slope (PC2) and curvature (PC3) rankings, indicating distinct sensitivities to different parts of the VIX futures curve. These patterns demonstrate that the portfolio ranks trace evolving sensitivities to VIX term-structure innovations rather than fixed "safe" or "risky" labels, and that currencies load differently on level, slope, and curvature components rather than following a stable macro-financial grouping.

[Figure 6 about here.]

Figure 7 makes the underlying factor structure even clearer by showing whether each currency enters the long, neutral, or short leg of the proposed HML risk factors. The heatmaps show that the composition of the long–short factors is neither constant nor mechanically aligned with carry-trade positions: safe-haven currencies do not always appear in the long leg, high-yielders do not always appear in the short leg, and many currencies move across legs as volatility conditions change. GBP, ZAR, and ISK frequently switch sides in all three HML risk factors, entering the long leg when their slope betas spike and falling into the short or neutral leg when exposures compress as their underlying risk sensitivities shift.

[Figure 7 about here.]

To assess whether the factors embed any systematic country-type bias, it is useful to track how the share of advanced and emerging currencies evolves over time. Figure 8 and Figure 9 show that the advanced–emerging shares vary across the sample but exhibits no persistent dominance of one group. In the aggregate view, the share of advanced and emerging currencies fluctuates across all three VIX components, without settling into a stable pattern in which one country group systematically drives the factor. The long- and short-leg decomposition reinforces this point: while certain episodes feature a higher representation of emerging currencies or, at other times, a greater presence of advanced currencies, neither pattern persists. Both legs of the factors remain relatively balanced across the sample, indicating that the strategy is not simply financing positions in advanced low-yield currencies to invest in high-yield emerging currencies. Instead, the composition shifts with changes in volatility-risk exposures rather than a structural divide between advanced and emerging markets.

[Figure 8 about here.]

[Figure 9 about here.]

4. The price of global volatility term-structure risk in the cross section of currency returns

4.1. Explaining currency excess returns with VIX term-structure currency risk factors

As suggested by the stylized model in Section 2, the return of the long-short beta-sorted portfolio can capture the global component in the volatility term-structure factor. Therefore, with the volatility term-structure risk factors in hand, I can then evaluate if they are priced in the cross-section of currency returns by estimating their associated price of risk and their ability to explain the cross-section of currency excess returns. For that, I use the two-pass methodology put forward by Fama and MacBeth, 1973. Since the influential study of Lustig and Verdelhan, 2007, this method has been widely employed in currency asset pricing precisely because it links asset returns to risk factors through the SDF. Following Cochrane, 2005, the fundamental pricing equation states that the price of any asset equals the expected discounted payoff:

$$p_t \;=\; \mathbb{E}_t[\, m_{t+1}\, x_{t+1}\,]\,, \tag{19}$$

where p_t is the asset price at time t, x_{t+1} its payoff, and m_{t+1} the SDF. In the absence of arbitrage opportunities, risk-adjusted excess returns must have zero mean and satisfy the pricing condition:

$$\mathbb{E}_t[\, m_{t+1}\, rx_{i,t+1}\,] \;=\; 0. \tag{20}$$

A linear SDF can be represented as

$$m_t = \xi\big[1 - (f_t - \mathbb{E}[f_t])b^{\top}\big], \tag{21}$$

where ξ is a scalar intercept, f_t is a $1 \times K$ vector of risk factors observed at time t, $\mathbb{E}[f_t]$ denotes the vector of their expected values, and b collects the factor loadings of the SDF. The element b_k measures how much the k-th factor contributes to the SDF, conditional on the presence of the other factors.

Combining the SDF representation with the pricing condition leads to the standard beta-pricing relation,

$$\mathbb{E}[rx_t] = \underbrace{\text{cov}(rx_t, f_t)\Sigma_f^{-1}}_{\beta}\underbrace{\Sigma_f b^\top}_{\lambda^\top}, \tag{22}$$

which shows that expected excess returns are determined by two components: the $1 \times K$ vector of risk prices, λ, and the $N \times K$ matrix of risk exposures, β. Therefore, the expected excess return for asset i is

$$\mathbb{E}[rx_i] = \beta_i \lambda^\top. \tag{23}$$

In empirical work, the risk exposures (β_i) and risk prices (λ) are typically estimated using the two-pass Fama–MacBeth procedure. The first pass estimates the risk exposures (or factor loadings), and the second pass recovers the prices of risk.

The first pass involves running time-series regressions of each asset's excess return on the factors:

$$rx_{i,t} = \alpha_i + f_t \beta_i^\top + \epsilon_{i,t}, \quad i = 1, \ldots, N, \quad t = 1, \ldots, T, \tag{24}$$

where α_i captures the risk-adjusted average excess return of asset i.

The second pass then runs a cross-sectional regression of average asset returns on the estimated betas. Since expected excess returns $\mathbb{E}[rx_i]$ are not observable, they are proxied by average realized excess returns, $\bar{rx}_i = \frac{1}{T}\sum_{t=1}^{T} rx_{i,t}$. The cross-sectional regression becomes

$$\bar{rx}_i = \hat{\beta}_i \lambda^\top + a_i, \quad i = 1, \ldots, N, \tag{25}$$

where a_i is the pricing error for asset i, and $\hat{\beta}_i \hat{\lambda}^\top$ is its model-implied risk premium. Therefore, the vector of risk prices is estimated as

$$\hat{\lambda} = \bar{rx}\hat{\beta}(\hat{\beta}^\top \hat{\beta})^{-1}, \tag{26}$$

where $\hat{\beta}$ is the $N \times K$ matrix of estimated factor loadings, and $\bar{rx}$ is the vector of average realized excess returns across assets. Then, the cross-sectional R^2 is computed as

$$R^2 = 1 - \frac{\frac{1}{N}\sum_{i=1}^{N} \hat{a}_i^2}{\text{var}\left(\bar{rx}_i\right)}. \tag{27}$$

The covariance matrix of the estimated risk prices, however, must be adjusted to account for the fact that the factor loadings are estimated rather than observed. I apply the Shanken (1992) correction to obtain consistent standard errors for the estimated prices of risk, in which the standard errors are inflated to account for the fact that $\hat{\beta}_i$ are estimated rather than known. This yields robust inference for the estimated risk premia under the two-pass Fama–MacBeth procedure.

Additionally, the overall pricing accuracy of the model is evaluated with the Gibbons et al., 1989 statistic, which tests whether all first-stage alphas are jointly zero. The test uses the time-series regression residuals and the factor covariance structure to assess whether the model leaves any systematic mispricing. Failure to reject the null means the alphas are statistically indistinguishable from zero, so the model is not leaving predictable errors. Rejection of the null means at least one asset shows a nonzero alpha, indicating the model cannot fully eliminate pricing errors.

4.2. Test assets: style-sorted portfolios and country excess returns

I use four different sets of currency test-assets, all of which I also include the factors themselves as test-assets. The first set includes the five Carry, five Momentum and five Value portfolios, while the second set includes only the five Momentum and five Value portfolios. The third and fourth cross-sections are country-level excess returns[10], where the third includes all available countries and the fourth restricts the sample to advanced countries only. The Carry portfolios are the five forward premium sorted portfolios, following Lustig et al., 2011. The Momentum portfolios are built by sorting currencies on previous-month excess return, as in Menkhoff et al., 2012, and the Value portfolios are built by sorting currencies on the 5-year log-return of the real exchange rate, as in Menkhoff et al., 2017. The foreign exchange literature tends to focus more on explaining the Carry sorted portfolios, I go a step further and add additional portfolio and country cross-sections to test the hypothesis.

4.3. Cross-sectional pricing results: estimated risk premia

The second stage results[11] for portfolio cross-sections are shown in Table 3, where columns (1) to (5) are for the Carry, Momentum and Value test assets, and columns (6) to (10) are for the Momentum and Value test assets. Table 4 show the second stage results for the country level analysis. The first noticeable thing is that there is a jump in the cross-sectional R^2 when

[10]Currencies with fewer than five years of data are excluded. Missing observations are replaced by the monthly cross-sectional mean.

[11]I use annualized monthly excess returns for the test assets and the risk factors.

any of the VIX term-structure currency risk factors are included, relative to the model only with the Carry and Dollar factors, showing that these additional factors provide valuable information to pricing the cross-section of returns. This pattern is even more pronounced in the country level cross-sections. It is also noticeable that when the VIX term-structure risk factors are included I cannot reject the null-hypothesis that all the pricing errors are null with 5% significance level, but it is rejected when only the Carry and Dollar risk factors are included in the all country cross-section.

I find that the level (PC1) and curvature (PC3) factors have a negative price of risk, while the slope (PC2) factor has a positive price of risk. A higher exposure to the level (PC1) and curvature (PC3) factors leads to a lower average excess return across assets, where an asset with a beta of one pays a risk premium of about −6% to −8% and −5% to −6% per year, respectively, acting as insurance.

In terms of the model in Section 2, experts often net supply long-horizon volatility ($h_L < 0$), for example due to long-run liability matching or duration-hedging needs. This implies that expert demand for long-horizon volatility is low relative to outstanding supply, so that arbitrageurs absorb a positive residual position. Being long level volatility is therefore akin to selling long-term volatility insurance, and insurance sellers earn a premium. In equilibrium, this corresponds to a negative price of long-horizon volatility risk, $\lambda_L < 0$, so currencies with positive level betas earn negative expected premia.

Experts may also exhibit demand for curvature exposure ($h_C > 0$) because their portfolios feature nonlinear sensitivities to volatility. However, if exogenous supply of curvature exposure is sufficiently large or if arbitrageurs face high costs of absorbing convexity risk (W_{CC} large), the residual demand absorbed by intermediaries remains negative. Supplying curvature exposure is akin to selling convexity insurance, and the buyer of such insurance pays a premium. As a result, the equilibrium price of curvature risk is negative, $\lambda_C < 0$, and currencies with positive curvature betas earn negative risk premia.

The slope (PC2) risk factor, by contrast, earns a positive risk premium of about 6% to 8% per year for an asset with a beta of one. Experts can exhibit strong demand for short-horizon volatility protection ($h_S > 0$), while arbitrageurs face relatively high costs of absorbing short-dated volatility exposure (W_{SS} large). When expert demand exceeds outstanding supply, the residual position $h_S - \Theta_S$ that intermediaries must absorb is positive, generating a positive price of risk, $\lambda_S > 0$. Currencies with positive slope betas therefore deliver payoffs precisely when short-horizon volatility spikes, states in which experts value protection most, and must earn positive risk premia in equilibrium.

This highlights that these risk prices are both statistically and economically significant across all four cross-sections, with a price of risk about half of the Carry factor in most specifications. Taken together, in terms of the model, preferred-habitat hedging motives across volatility horizons and maturity-dependent intermediary frictions generate the scarcity-based prices of risk. Currency premia arise naturally because currencies differ in their exposures to these global volatility term-structure shocks. These results show that volatility term-structures contain information about shocks to the global pricing kernel that are systematically priced in currency markets, reflecting horizon-specific expert hedging demand. Changes in the term-structure generate differential compensation across currencies. The empirical evidence confirms that these novel currency factors are priced in the cross-section of currency excess returns, and that a measure tied to global risk pricing across horizons captures information essential for explaining them.

[Table 3 about here.]

[Table 4 about here.]

5. Predictive power for currency excess returns

Having shown in the cross-sectional tests that the tradable VIX term-structure currency factors carry significant prices of risk, it is natural to ask whether the underlying innovations themselves contain information about the time-series dynamics of currency excess returns. The Fama–MacBeth analysis evaluates whether returns on factor-mimicking portfolios are compensated in the cross section, whereas the panel regressions examine whether the raw shocks to the VIX term-structure components forecast next-month currency excess returns. This provides a complementary perspective: the cross-sectional tests assess whether factor-mimicking portfolios span priced dimensions of global risk, while the panel regressions test whether the shocks generating those dimensions have predictive content for individual currencies, in light of the known difficulty in predicting exchange rates (Meese and Rogoff, 1983).

5.1. Panel data regressions

Table 5 reports the predictive regressions using pooled OLS and currency fixed-effects specifications. Across all models, innovations to the level ($\Delta PC1$), slope ($\Delta PC2$), and curvature ($\Delta PC3$) components display statistically significant predictive power for one-month-ahead currency excess returns. The coefficients are uniformly positive, indicating that upward shocks to any part of the VIX term-structure are followed, on average, by higher subsequent currency excess returns. Among the three, $\Delta PC2$ delivers the strongest and most precisely estimated effect, highlighting the importance of shocks to the slope for near-term currency payoffs. $\Delta PC3$ also yields economically meaningful forecasts, and $\Delta PC1$ contributes modest but significant predictive power in both pooled and fixed-effects specifications.

Including lagged currency returns shows that the predictive content of the VIX-based innovations is not absorbed by simple return persistence. The coefficient on rx_t is positive and significant in the pooled models, but adding it leaves the ΔPC coefficients largely unchanged, indicating that VIX term-structure shocks contain incremental information rather than merely proxying for autoregressive dynamics. The fixed-effects results further show that the predictive coefficients survive after removing all time-invariant currency characteristics, implying that the forecasting power reflects within-currency variation.

Although the adjusted R^2 values are modest, as typical in short-horizon currency forecasting, the results align closely with the economic interpretation developed earlier: the same shocks that produce priced, tradable risk factors in the cross-section also transmit forward-looking information about expected currency returns.

[Table 5 about here.]

6. Conclusion

This paper examines the pricing implications of the VIX futures term-structure on currency excess returns. In a simple two-agent model where experts have preferred-habitat demand for volatility-linked contracts and arbitrageurs supply these contracts subject to maturity-dependent capital costs, I show that each country's stochastic discount factor is driven by local and global shocks, and that sorting currencies by their exposure to innovations in the volatility term-structure yields pricing factors related to hedging demand through volatility-linked contracts.

I build currency risk factors from long-short beta-sorted portfolios based on each currency's exposure to the principal components of the VIX futures term-structure. These term-structures summarize market-implied pricing of volatility risk across different horizons and exhibit a strong factor structure: the first three principal components (level, slope, and curvature) account for over 99% of their variation. Sorting currencies by their conditional beta with respect to these components, as suggested by the model, captures the global component of volatility term-structure risk.

Finally, I provide empirical evidence across four sets of test assets showing that these currency factors explain a significant portion of the cross-sectional variation in excess returns. This explanatory power is incremental to the well-established carry and dollar factors of Lustig et al., 2011, and supports the view that innovations in VIX futures term-structure reflect priced global SDF shocks. Additionally, the innovations in the components of the VIX futures term-structure provide incremental information for currency excess return predictability.

References

Aït-Sahalia, Yacine, Mustafa Karaman, and Loriano Mancini (2020). "The term structure of equity and variance risk premia". In: *Journal of Econometrics* 219.2. Annals Issue: Econometric Estimation and Testing: Essays in Honour of Maxwell King, pp. 204–230. ISSN: 0304-4076. DOI: https://doi.org/10.1016/j.jeconom.2020.03.002. URL: https://www.sciencedirect.com/science/article/pii/S030440762030097X.

Backus, David K., Silverio Foresi, and Chris I. Telmer (Feb. 2001). "Affine Term Structure Models and the Forward Premium Anomaly". In: *The Journal of Finance* 56 (1), pp. 279–304. ISSN: 0022-1082. DOI: 10.1111/0022-1082.00325.

Bollerslev, Tim, George Tauchen, and Hao Zhou (Feb. 2009). "Expected Stock Returns and Variance Risk Premia". In: *The Review of Financial Studies* 22.11, pp. 4463–4492. ISSN: 0893-9454. DOI: 10.1093/rfs/hhp008. eprint: https://academic.oup.com/rfs/article-pdf/22/11/4463/24429122/hhp008.pdf. URL: https://doi.org/10.1093/rfs/hhp008.

Cochrane, John H. (2005). *Asset Pricing*. Revised Edition. Princeton, NJ: Princeton University Press.

Corte, Pasquale Della, Steven J. Riddiough, and Lucio Sarno (Aug. 2016). "Currency Premia and Global Imbalances". In: *Review of Financial Studies* 29 (8), pp. 2161–2193. ISSN: 0893-9454. DOI: 10.1093/rfs/hhw038.

Dew-Becker, Ian et al. (Feb. 2017). "The price of variance risk". In: *Journal of Financial Economics* 123 (2), pp. 225–250. ISSN: 0304405X. DOI: 10.1016/j.jfineco.2016.04.003.

Dobrynskaya, Victoria (Aug. 2014). "Downside Market Risk of Carry Trades*". In: *Review of Finance* 18 (5), pp. 1885–1913. ISSN: 1573-692X. DOI: 10.1093/rof/rfu004.

Drechsler, Itamar and Amir Yaron (Sept. 2010). "What's Vol Got to Do with It". In: *The Review of Financial Studies* 24.1, pp. 1–45. ISSN: 0893-9454. DOI: 10.1093/rfs/hhq085. eprint: https://academic.oup.com/rfs/article-pdf/24/1/1/24441046/hhq085.pdf. URL: https://doi.org/10.1093/rfs/hhq085.

Egloff, Daniel, Markus Leippold, and Liuren Wu (2010). "The Term Structure of Variance Swap Rates and Optimal Variance Swap Investments". In: *The Journal of Financial and Quantitative Analysis* 45.5, pp. 1279–1310. ISSN: 00221090, 17566916. URL: http://www.jstor.org/stable/27919564 (visited on 12/18/2025).

Fama, Eugene F and James D MacBeth (1973). "Risk, return, and equilibrium: Empirical tests". In: *Journal of Political Economy* 81.3, pp. 607–636.

Fan, Zhenzhen, Juan M. Londono, and Xiao Xiao (Jan. 2022). “Equity tail risk and currency risk premiums”. In: *Journal of Financial Economics* 143 (1), pp. 484–503. ISSN: 0304405X. DOI: 10.1016/j.jfineco.2021.05.020.

Gârleanu, Nicolae and Lasse Pedersen (2011). “Margin-based asset pricing and deviations from the law of one price”. In: *Review of Financial Studies*.

Gibbons, Michael R., Stephen A. Ross, and Jay Shanken (1989). “A Test of the Efficiency of a Given Portfolio”. In: *Econometrica* 57.5, pp. 1121–1152. ISSN: 00129682, 14680262. URL: http://www.jstor.org/stable/1913625 (visited on 06/14/2025).

He, Zhiguo and Arvind Krishnamurthy (2013). “Intermediary Asset Pricing”. In: *Annual Review of Financial Economics*.

Johnson, Travis L. (2017). “Risk Premia and the VIX Term Structure”. In: *Journal of Financial and Quantitative Analysis* 52.6, pp. 2461–2490. DOI: 10.1017/S0022109017000825.

Joslin, Scott, Kenneth Singleton, and Haoxiang Zhu (2011). “A New Perspective on Gaussian Dynamic Term Structure Models”. In: *Review of Financial Studies* 24.3, pp. 926–970. URL: https://EconPapers.repec.org/RePEc:oup:rfinst:v:24:y:2011:i:3:p:926-970.

Kroencke, Tim A., Felix Schindler, and Andreas Schrimpf (Aug. 2014). “International Diversification Benefits with Foreign Exchange Investment Styles*”. In: *Review of Finance* 18 (5), pp. 1847–1883. ISSN: 1573-692X. DOI: 10.1093/rof/rft047.

Lettau, Martin, Matteo Maggiori, and Michael Weber (Nov. 2014). “Conditional risk premia in currency markets and other asset classes”. In: *Journal of Financial Economics* 114 (2), pp. 197–225. ISSN: 0304405X. DOI: 10.1016/j.jfineco.2014.07.001.

Londono, Juan M. and Hao Zhou (May 2017). “Variance risk premiums and the forward premium puzzle”. In: *Journal of Financial Economics* 124 (2), pp. 415–440. ISSN: 0304405X. DOI: 10.1016/j.jfineco.2017.02.002.

Lustig, Hanno, Nikolai Roussanov, and Adrien Verdelhan (Nov. 2011). “Common Risk Factors in Currency Markets”. In: *Review of Financial Studies* 24 (11), pp. 3731–3777. ISSN: 0893-9454. DOI: 10.1093/rfs/hhr068.

Lustig, Hanno and Adrien Verdelhan (2007). “The Cross Section of Foreign Currency Risk Premia and Consumption Growth Risk”. en. In: *The American Economic Review* 97.1.

Meese, Richard A. and Kenneth Rogoff (1983). “Empirical exchange rate models of the seventies: Do they fit out of sample?” In: *Journal of International Economics* 14.1, pp. 3–24. ISSN: 0022-1996. DOI: https://doi.org/10.1016/0022-1996(83)90017-X. URL: https://www.sciencedirect.com/science/article/pii/002219968390017X.

Menkhoff, Lukas et al. (Dec. 2012). "Currency momentum strategies". In: *Journal of Financial Economics* 106 (3), pp. 660–684. ISSN: 0304405X. DOI: 10.1016/j.jfineco.2012.06.009.

— (Feb. 2017). "Currency Value". In: *Review of Financial Studies* 30 (2), pp. 416–441. ISSN: 0893-9454. DOI: 10.1093/rfs/hhw067.

Shanken, Jay (1992). "On the Estimation of Beta-Pricing Models". In: *The Review of Financial Studies* 5.1, pp. 1–33. ISSN: 08939454, 14657368. URL: http://www.jstor.org/stable/2962011 (visited on 06/14/2025).

Vayanos, Dimitri and Jean-Luc Vila (2021). "A Preferred-Habitat Model of the Term Structure of Interest Rates". In: *Econometrica* 89.1, pp. 77–112. DOI: https://doi.org/10.3982/ECTA17440. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA17440. URL: https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA17440.

Verdelhan, Adrien (Feb. 2018). "The Share of Systematic Variation in Bilateral Exchange Rates". In: *The Journal of Finance* 73 (1), pp. 375–418. ISSN: 0022-1082. DOI: 10.1111/jofi.12587.

Zhou, Hao (2018). "Variance Risk Premia, Asset Predictability Puzzles, and Macroeconomic Uncertainty". In: *Annual Review of Financial Economics* 10.Volume 10, 2018, pp. 481–497. ISSN: 1941-1375. DOI: https://doi.org/10.1146/annurev-financial-110217-022737. URL: https://www.annualreviews.org/content/journals/10.1146/annurev-financial-110217-022737.

Tables and figures

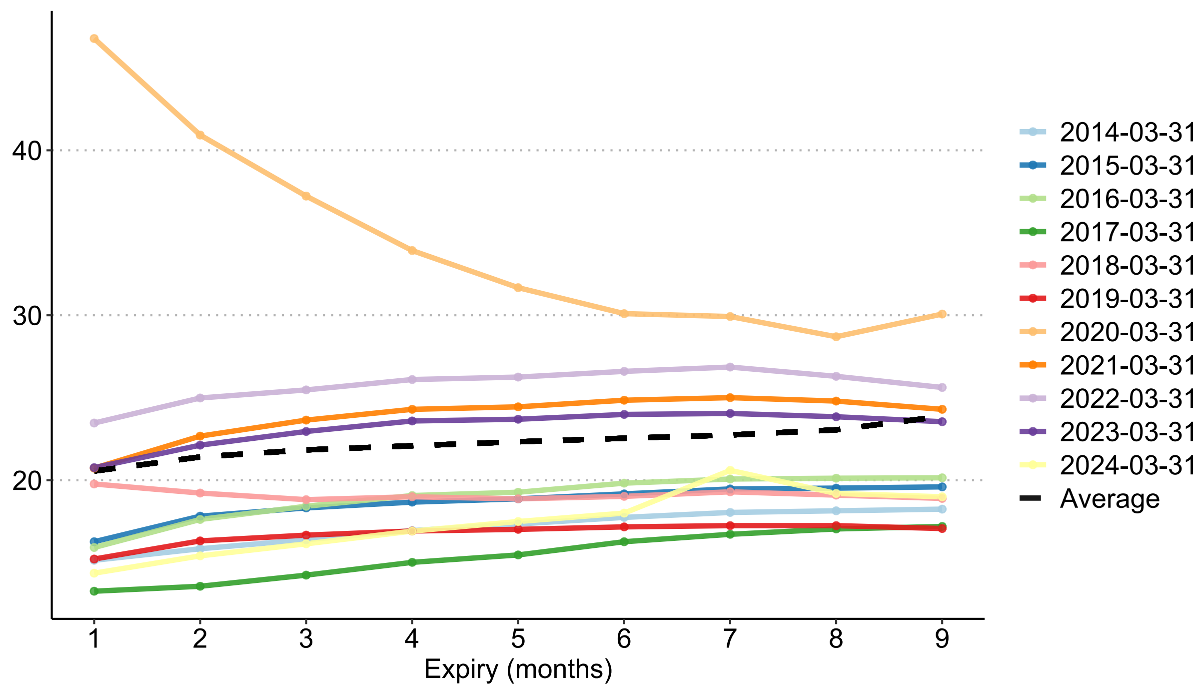

Figure 1: Selected end-of-month VIX futures term-structures

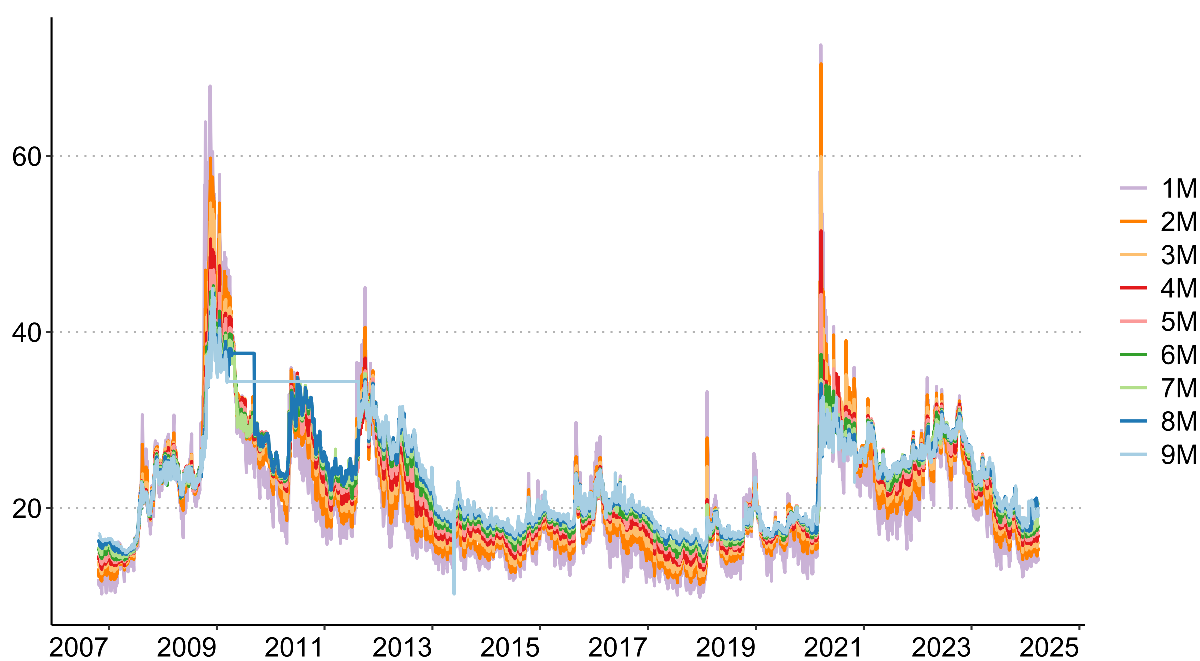

Figure 2: Time-series of VIX futures, end-of-month

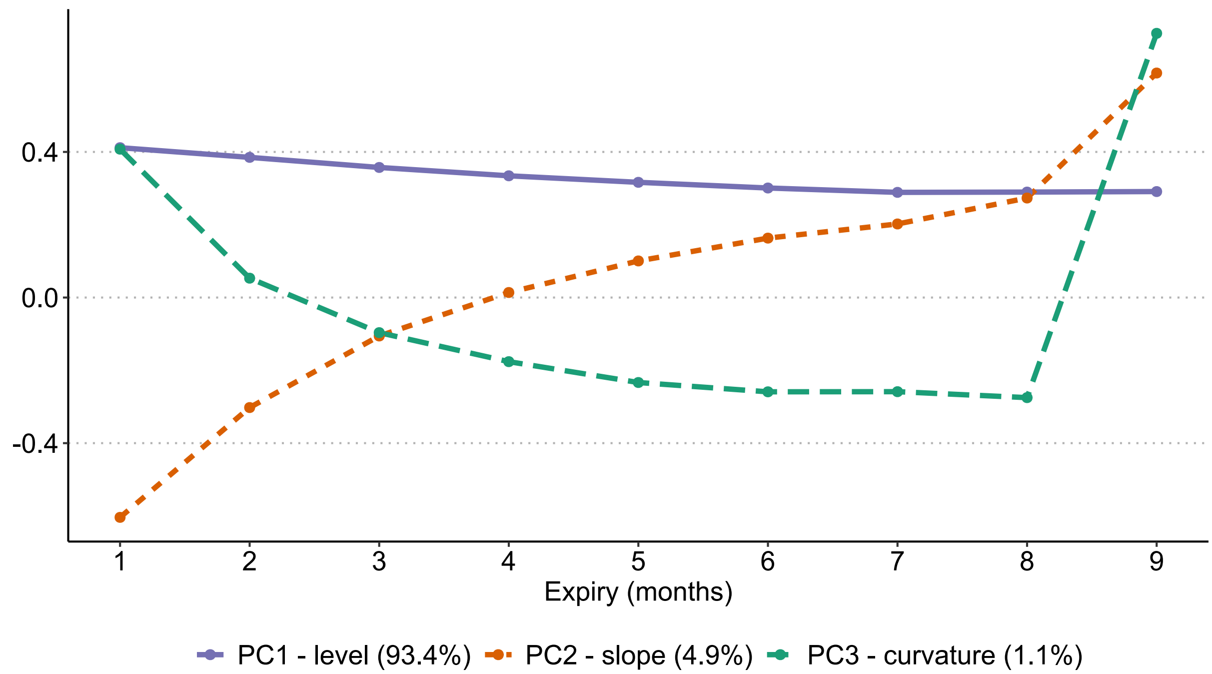

Figure 3: Factor loadings of VIX futures PCA. This figure shows the factor loadings of the full-sample principal component analysis on the level of VIX futures. Loadings are scaled such that they correspond to common measures of level, slope, and curvature. Percentage of variance explained by each principal component is between parenthesis.

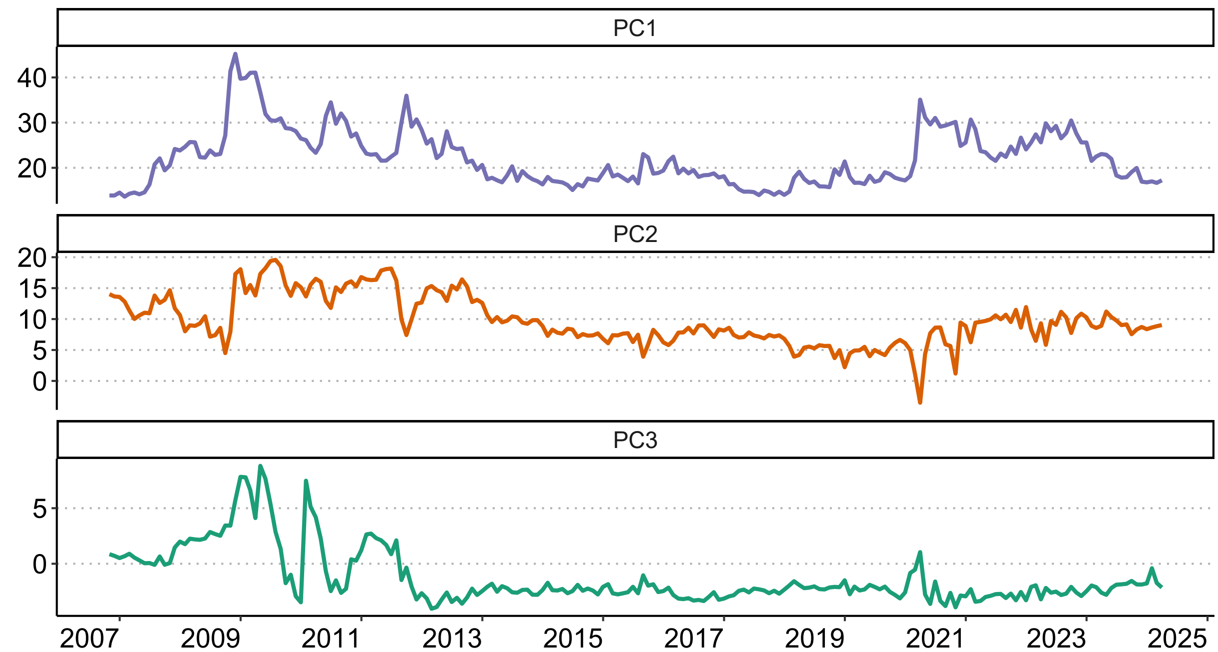

Figure 4: Principal components of expanding window PCA. This figure shows the time-series of the first three principal components extracted from VIX futures using an expanding window PCA, with a startup period of 252 days. PCA is performed on daily data, with the resulting factors sampled at month-end. The signs of the loadings are such that PC1 captures the level of the curve, PC2 captures the slope, and PC3 captures curvature. Increases in PC1 reflect upward shifts in the entire curve—typically associated with deteriorating market conditions—while increases in PC2 reflect steepening, often interpreted as improving conditions. PC3 captures changes in curvature, though its interpretation is less directly linked to macro-financial states.

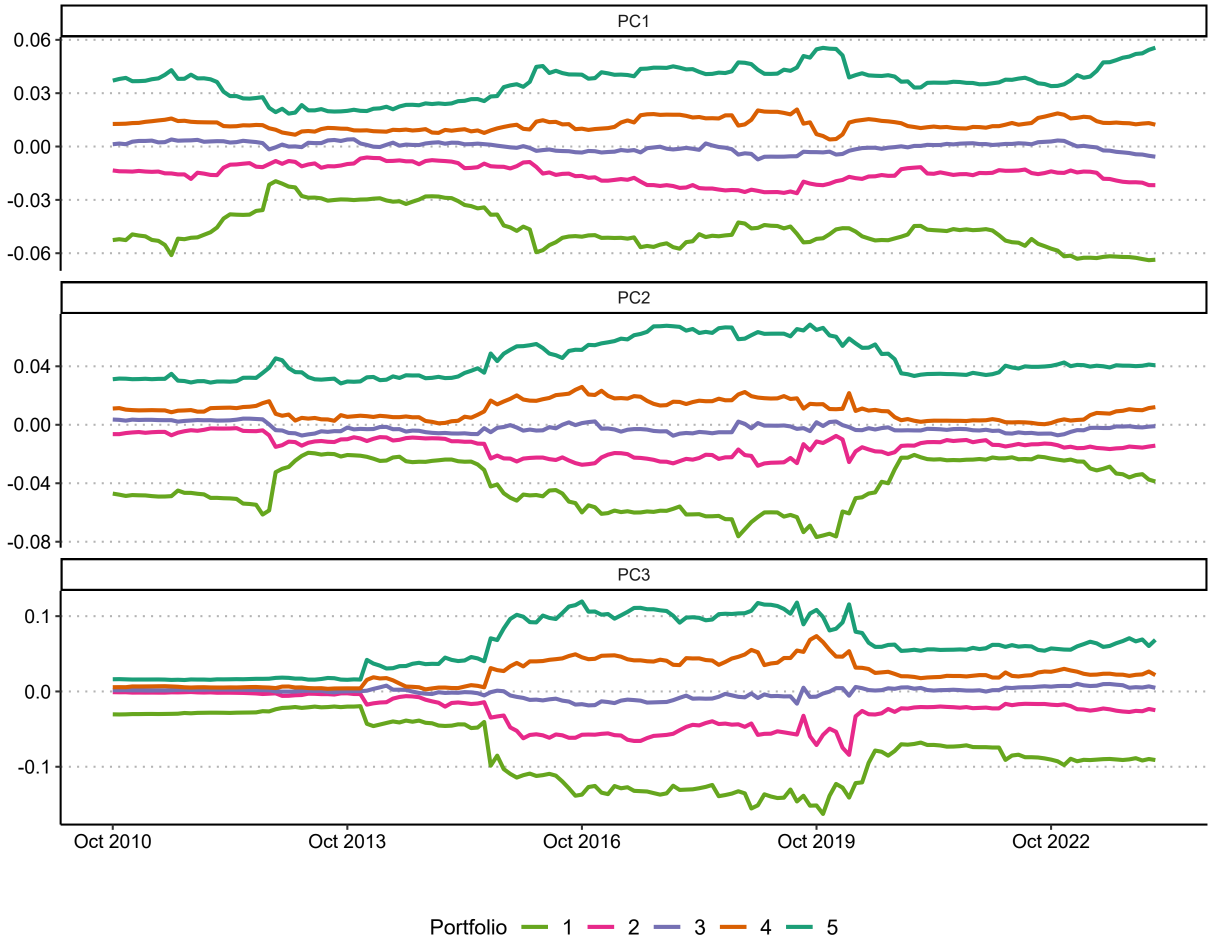

Figure 5: Rolling betas to VIX principal component innovations by currency portfolio. This figure shows the time-series of average betas for five currency portfolios sorted each month by their estimated exposure to innovations in VIX term-structure components. The principal components—capturing the level (PC1), slope (PC2), and curvature (PC3) of the VIX futures curve—are extracted monthly using an expanding-window PCA. Innovations in each component are computed as the monthly change relative to the previous observation. For each currency, rolling 48-month regressions of excess returns on these innovations are estimated following Equation 18, controlling for the dollar factor. The resulting betas measure time-varying sensitivities to shocks in the VIX term-structure. The figure reports portfolio-average betas across the five beta-sorted groups, illustrating persistent cross-sectional heterogeneity and cyclical variation in exposures to volatility-level (PC1), slope (PC2), and curvature (PC3) innovations.

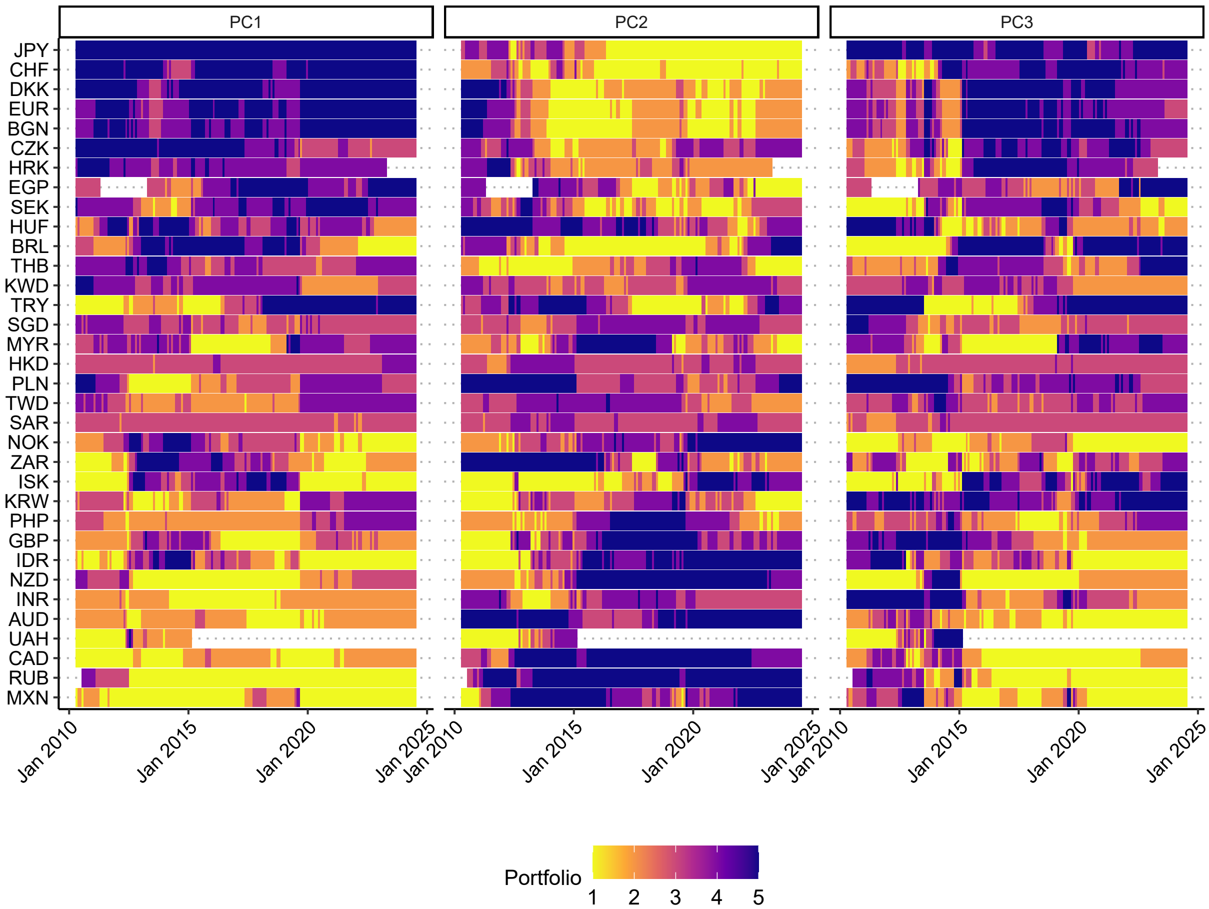

Figure 6: Composition of currency portfolios sorted on VIX principal component betas. This figure displays the time-varying portfolio assignments of individual currencies based on their rolling 48-month betas with respect to innovations in the first three VIX term-structure principal components: level (PC1), slope (PC2), and curvature (PC3). Each color represents the portfolio (from 1 to 5) to which a currency belongs in a given month, determined by its beta ranking from the previous month's estimation. Left corresponds to portfolios sorted on PC1 betas, middle panel on PC2, and right on PC3. Lighter shades indicate lower-exposure portfolios, while darker shades represent higher exposures. The figure highlights persistent clustering of currencies with similar sensitivities to volatility term-structure shocks, as well as transitions across portfolios reflecting evolving exposure patterns over time.

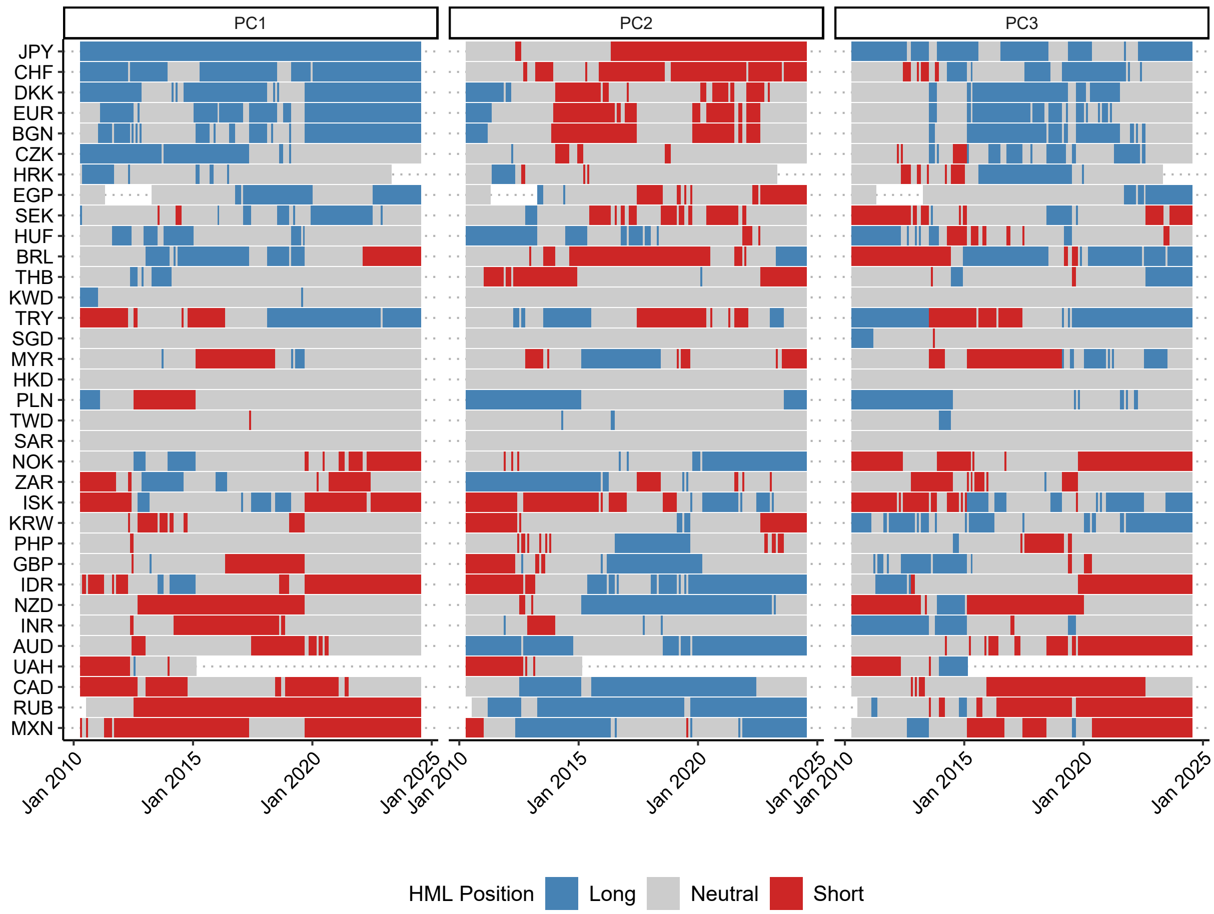

Figure 7: Composition of currency risk factors constructed from VIX principal component betas. This figure shows the time-varying composition of the high-minus-low (HML) currency risk factors associated with exposures to innovations in the first three VIX term-structure principal components: level (PC1), slope (PC2), and curvature (PC3). For each month, currencies are sorted into five portfolios based on their rolling 48-month betas with respect to the respective principal component innovation, and the factor is constructed as the return difference between the highest- and lowest-beta portfolios. The heatmap indicates, for each currency and month, whether it enters the factor with a long position (blue), short position (red), or is excluded from the factor (gray). Each panel corresponds to one VIX component. The figure highlights persistent cross-sectional asymmetries in how currencies load onto the VIX futures term-structure level, slope, and curvature shocks, as well as shifts in factor composition reflecting changing sources of volatility risk exposure across time.

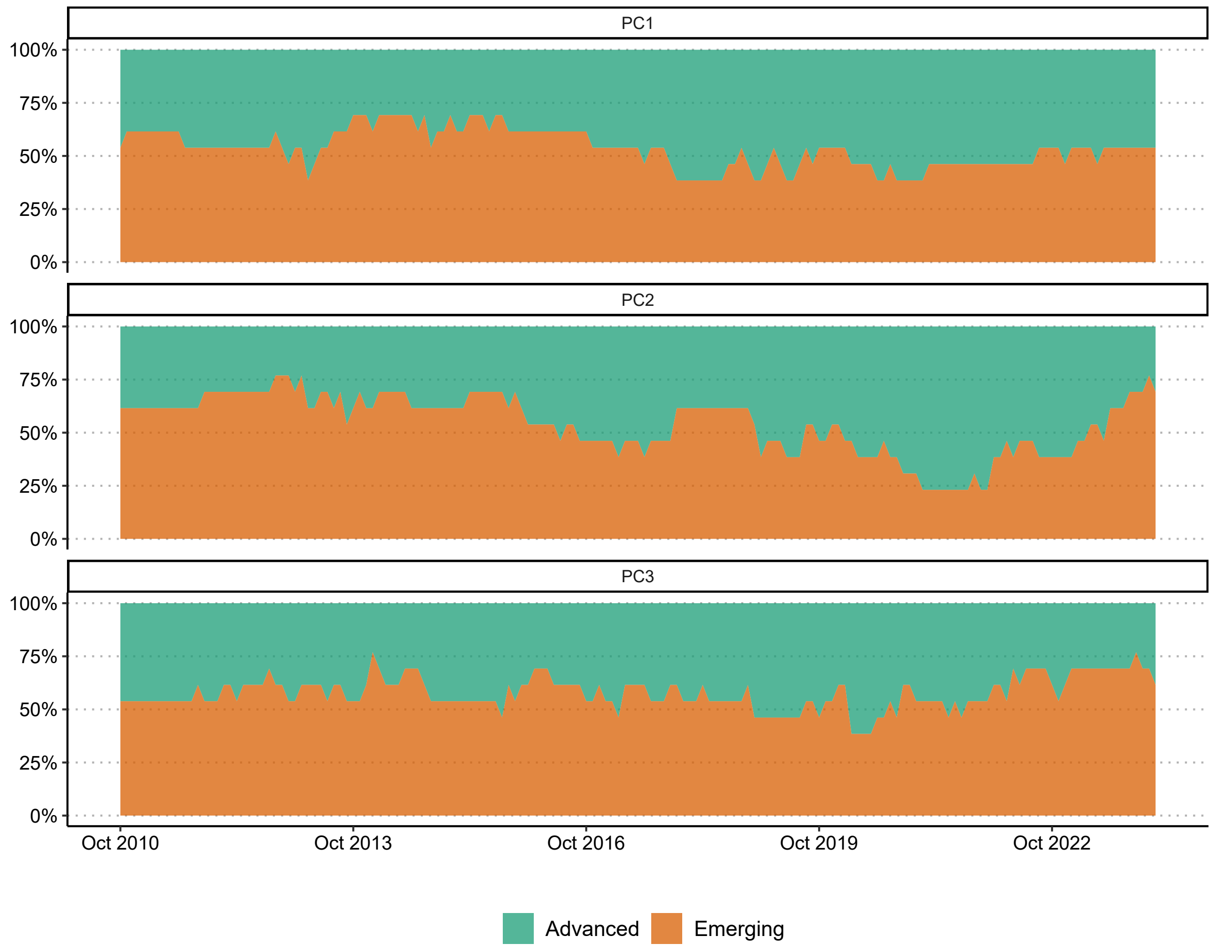

Figure 8: Advanced and emerging market composition of VIX principal component factors. This figure presents the share of advanced and emerging market currencies in the aggregate high-minus-low (HML) portfolios constructed from exposures to innovations in the first three VIX term-structure principal components: level (PC1), slope (PC2), and curvature (PC3). Each panel shows, over time, the proportion of advanced (green) and emerging (orange) currencies included in the combined factor portfolio, computed monthly based on the underlying beta-sorted portfolio assignments. The composition reveals that there is no consistent dominance of either country group across factors: the relative shares of advanced and emerging currencies fluctuate over time and differ across components, suggesting that exposure to volatility term-structure shocks is not systematically concentrated in one market type but varies with global conditions.

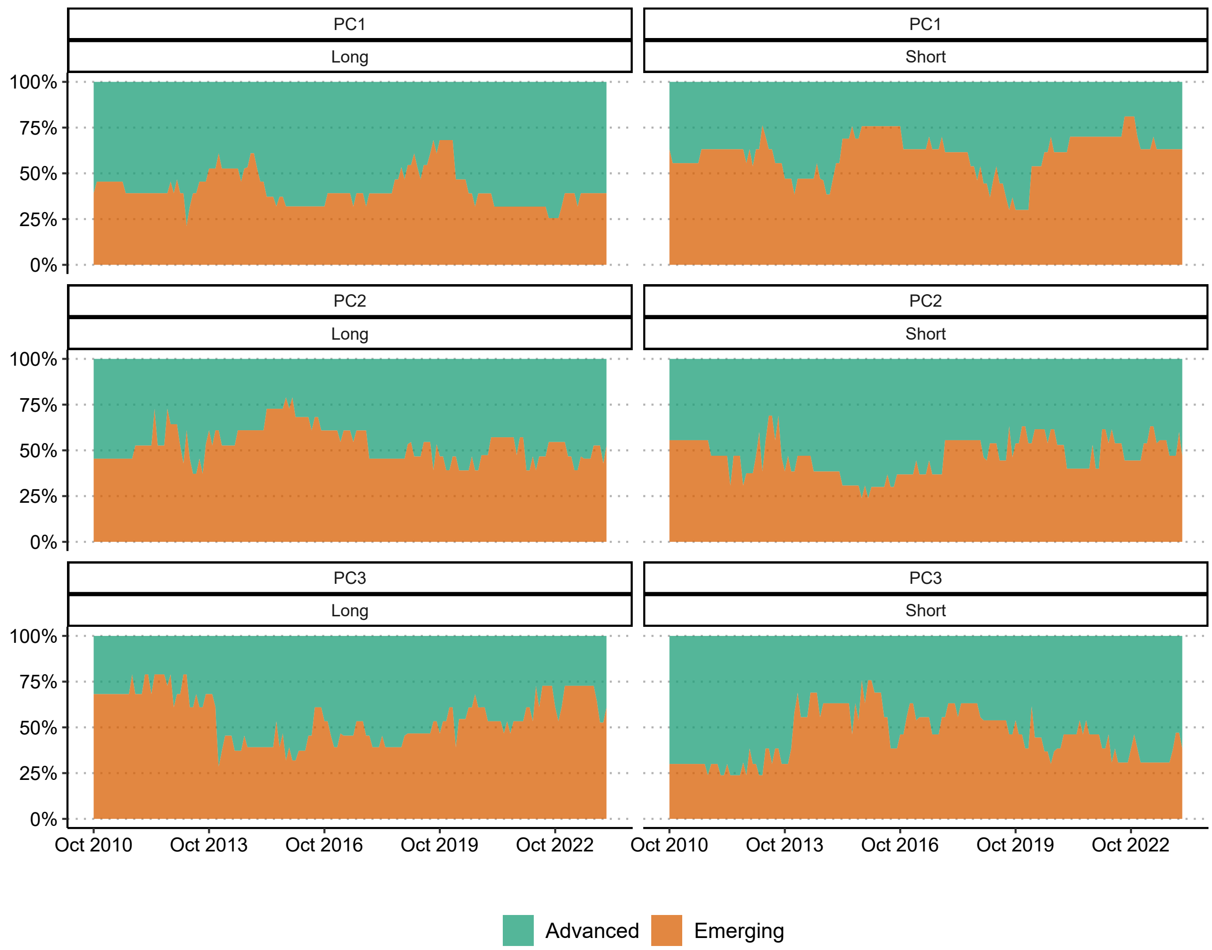

Figure 9: Advanced and emerging market composition of VIX principal component factors. This figure shows the share of advanced and emerging market currencies included in the long and short sides of the high-minus-low (HML) currency risk factors constructed from exposures to innovations in the first three VIX term-structure principal components: level (PC1), slope (PC2), and curvature (PC3). Each panel plots, over time, the proportion of advanced (green) and emerging (orange) currencies composing the long (left) and short (right) portfolios of each factor. The proportions are computed monthly based on the portfolio assignments underlying the factor construction. The figure highlights systematic differences in how advanced and emerging currencies load onto volatility term-structure shocks: exposures to the level (PC1) are typically dominated by advanced currencies on the long side, while slope (PC2) and curvature (PC3) factors exhibit more balanced or alternating compositions, reflecting shifts in global volatility transmission across market groups.

Table 1: VIX futures beta-sorted portfolios summary statistics. This table reports summary statistics for monthly currency excess returns of portfolios sorted by estimated betas with respect to VIX futures innovations. Panel A sorts currencies on β_{PC1} (level), Panel B on β_{PC2} (slope), and Panel C on β_{PC3} (curvature). P1 contains currencies with the lowest beta exposures, while P5 contains those with the highest. HML denotes the high-minus-low portfolio (P5 – P1). The Sharpe Ratio is computed as the ratio of average annualized return to annualized volatility. *t-stat* refers to the t-statistic testing whether the average monthly excess return is significantly different from zero.

Portfolio	P1	P2	P3	P4	P5	HML
Panel A: sorts on β_{PC1}						
Avg Ann. Return	0.34	-1.54	-0.85	-2.12	-6.56	-6.91
t-stat	[0.16]	[-0.82]	[-0.56]	[-1.10]	[-2.77]	[-3.37]
Ann. Volatility	7.73	6.92	5.59	7.05	8.67	7.50
Sharpe Ratio	**0.04**	**-0.22**	**-0.15**	**-0.30**	**-0.76**	**-0.92**
Skewness	-0.01	-0.56	-0.49	-0.40	-0.63	-0.69
Panel B: sorts on β_{PC2}						
Avg Ann. Return	-7.43	-1.44	-1.17	-1.39	-0.47	6.96
t-stat	[-3.09]	[-0.72]	[-0.83]	[-0.78]	[-0.20]	[3.16]
Ann. Volatility	8.81	7.33	5.19	6.56	8.78	8.06
Sharpe Ratio	**-0.84**	**-0.20**	**-0.23**	**-0.21**	**-0.05**	**0.86**
Skewness	-0.96	-0.35	-0.18	-0.98	-0.53	0.63
Panel C: sorts on β_{PC3}						
Avg Ann. Return	0.11	-1.90	-1.33	-1.18	-6.57	-6.67
t-stat	[0.05]	[-0.99]	[-0.93]	[-0.62]	[-2.83]	[-3.61]
Ann. Volatility	8.52	7.00	5.22	7.00	8.50	6.77
Sharpe Ratio	**0.01**	**-0.27**	**-0.25**	**-0.17**	**-0.77**	**-0.99**
Skewness	-0.27	-0.20	-0.84	-0.28	-0.60	-0.80

Table 2: Time-series regression of VIX futures beta-sorted portfolios on the Carry factor. This table reports the results of time-series regressions of monthly excess returns of VIX futures beta-sorted currency portfolios on the Carry factor of Lustig et al., 2011. Panels A–C correspond to portfolios sorted on betas with respect to innovations to the first three principal components of VIX futures: level (PC1), slope (PC2), and curvature (PC3), respectively. P1 includes currencies with the lowest beta exposure and P5 includes those with the highest. HML denotes the high-minus-low portfolio (P5 – P1). Newey-West t-statistics are reported in brackets.

Portfolio	P1	P2	P3	P4	P5	HML
Panel A: sorts on β_{PC1}						
Intercept	-0.03	-0.03	-0.02	-0.03	-0.04	-0.01
	[-1.70]	[-1.66]	[-1.13]	[-1.38]	[-1.63]	[-0.58]
Carry Factor	0.32	0.17	0.10	0.09	-0.21	-0.54
	[3.38]	[1.86]	[1.31]	[1.00]	[-1.14]	[-5.89]
R^2	0.147	0.053	0.024	0.014	0.052	0.433
Panel B: sorts on β_{PC2}						
Intercept	-0.05	-0.02	-0.02	-0.03	-0.04	0.02
	[-1.97]	[-0.83]	[-1.06]	[-1.65]	[-1.60]	[0.66]
Carry Factor	-0.20	0.04	0.05	0.19	0.30	0.50
	[-0.75]	[0.43]	[0.96]	[1.87]	[2.84]	[3.15]
R^2	0.043	0.002	0.008	0.071	0.096	0.319
Panel C: sorts on β_{PC3}						
Intercept	-0.03	-0.03	-0.03	-0.02	-0.05	-0.02
	[-1.31]	[-1.31]	[-1.91]	[-1.05]	[-1.87]	[-1.24]
Carry Factor	0.28	0.08	0.15	0.08	-0.13	-0.41
	[2.48]	[1.06]	[2.21]	[0.86]	[-0.58]	[-3.81]
R^2	0.088	0.011	0.068	0.010	0.020	0.307

Table 3: Asset pricing tests of VIX futures term-structure factors - second stage. This table reports the second-stage cross-sectional asset pricing results for the portfolio-level test-assets using the VIX futures term-structure currency risk factors. The dependent variables are the annualized average excess returns of the test assets, while the regressors are the betas estimated in the first-stage time-series regressions. Columns (1)–(5) correspond to a cross-section of Carry, Momentum, and Value portfolios, while columns (6)–(10) use only the Momentum and Value portfolios. Each specification includes the Dollar and Carry factors from the literature, along with different combinations of the three VIX-based currency risk factors: level ($fPC1$), slope ($fPC2$), and curvature ($fPC3$). The λ estimates represent the price of risk associated with each factor. I also report the p-values from the test of Gibbons et al., 1989, which tests the joint hypothesis that all pricing errors (alphas) are equal to zero—i.e., that the model prices all assets correctly. A failure to reject this null suggests that the pricing model performs well in explaining cross-sectional variation in returns. Standard errors are corrected following Shanken, 1992, and statistical significance is denoted by *, **, and *** for the 10%, 5%, and 1% levels, respectively.

	Test assets									
	Carry + Momentum + Value					Momentum + Value				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
β_{fPC1}		−0.064***			−0.061***		−0.063***			−0.062***
		(0.019)			(0.019)		(0.020)			(0.019)
β_{fPC2}			0.064***		0.063***			0.063***		0.062***
			(0.020)		(0.020)			(0.021)		(0.020)
β_{fPC3}				−0.061***	−0.060***				−0.060***	−0.059***
				(0.017)	(0.017)				(0.018)	(0.017)
β_{CAR}	0.114***	0.110***	0.110***	0.110***	0.110***	0.161**	0.112***	0.112***	0.112***	0.112***
	(0.028)	(0.026)	(0.026)	(0.026)	(0.026)	(0.065)	(0.025)	(0.025)	(0.025)	(0.025)
β_{DOL}	−0.022	−0.022	−0.022	−0.023	−0.023	−0.023	−0.022	−0.022	−0.023	−0.023
	(0.017)	(0.017)	(0.017)	(0.017)	(0.017)	(0.018)	(0.017)	(0.017)	(0.017)	(0.017)
R^2	0.890	0.972	0.977	0.974	0.981	0.792	0.973	0.977	0.973	0.982
GRS	0.871	0.892	0.928	0.913	0.919	0.808	0.861	0.899	0.849	0.867

Table 4: Asset pricing tests of VIX futures term-structure factors - second stage. This table reports the second-stage cross-sectional asset pricing results for the country-level test-assets using the VIX futures term-structure currency risk factors. The dependent variables are the annualized average excess returns of the test assets, while the regressors are the betas estimated in the first-stage time-series regressions. Columns (1)–(5) correspond to a cross-section of all country excess returns, while columns (6)–(10) use only advanced countries currency excess returns. Each specification includes the Dollar and Carry factors from the literature, along with different combinations of the three VIX-based currency risk factors: level ($fPC1$), slope ($fPC2$), and curvature ($fPC3$). The λ estimates represent the price of risk associated with each factor. I also report the p-values from the test of Gibbons et al., 1989, which tests the joint hypothesis that all pricing errors (alphas) are equal to zero—i.e., that the model prices all assets correctly. A failure to reject this null suggests that the pricing model performs well in explaining cross-sectional variation in returns. Standard errors are corrected following Shanken, 1992, and statistical significance is denoted by *, **, and *** for the 10%, 5%, and 1% levels, respectively.

	Test assets									
	Country level - all					Country level - advanced				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
β_{fPC1}		-0.088^{**}			-0.079^{***}		-0.067^{***}			-0.072^{***}
		(0.037)			(0.022)		(0.021)			(0.020)
β_{fPC2}			0.081^{**}		0.062^{**}			0.060^{***}		0.057^{***}
			(0.038)		(0.031)			(0.021)		(0.021)
β_{fPC3}				-0.072^{**}	-0.060^{**}				-0.058^{***}	-0.053^{***}
				(0.033)	(0.028)				(0.019)	(0.018)
β_{CAR}	0.075^{**}	0.065^{**}	0.063^{**}	0.064^{**}	0.063^{**}	0.114^{**}	0.111^{***}	0.111^{***}	0.113^{***}	0.112^{***}
	(0.032)	(0.029)	(0.028)	(0.028)	(0.028)	(0.052)	(0.026)	(0.026)	(0.026)	(0.026)
β_{DOL}	−0.021	−0.021	−0.021	−0.021	−0.021	−0.014	−0.017	−0.016	−0.017	−0.018
	(0.017)	(0.017)	(0.017)	(0.017)	(0.017)	(0.018)	(0.018)	(0.018)	(0.018)	(0.018)
R^2	0.669	0.842	0.833	0.833	0.859	0.520	0.917	0.905	0.895	0.930
GRS	0.018	0.072	0.049	0.080	0.126	0.196	0.317	0.23	0.250	0.259

Table 5: Predictive panel regressions of one-month-ahead currency excess returns (rx_{t+1}) on month-t innovations to the VIX futures term-structure common components: level (PC1), slope (PC2), and curvature (PC3). Columns (1)–(3) report pooled OLS estimates for all currencies; Columns (4)–(6) report corresponding within estimations with individual (currency) fixed effects. Newey–West standard errors (one lag) for pooled models and cluster-robust (by currency) standard errors for fixed-effects models are shown in parentheses. Constant terms are included in the pooled OLS specifications but their coefficients are omitted from the table. Adjusted R^2 values are reported in percent. Statistical significance is denoted by *, **, and *** for the 10%, 5%, and 1% levels, respectively.

	rx_{t+1}					
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta PC1$	0.005*		0.015***	0.005*		0.015*
	(0.003)		(0.005)	(0.003)		(0.008)
$\Delta PC2$	0.018***		0.018***	0.018***		0.018***
	(0.003)		(0.003)	(0.003)		(0.003)
$\Delta PC3$	0.009**		0.011***	0.009***		0.011***
	(0.004)		(0.004)	(0.002)		(0.003)
rx_t		0.155***	0.179***		0.149	0.173
		(0.056)	(0.064)		(0.109)	(0.124)
FE	No	No	No	Yes	Yes	Yes
R^2 (%)	0.54	2.24	3.12	0.05	1.60	2.47

Appendix

Appendix A. Downward sloping hedging demand

The baseline model assumes that expert hedging demand h is inelastic at the horizon of interest. This is consistent with preferred-habitat models (Vayanos and Vila, 2021) where institutional mandates, liability structures, and regulation impose near-vertical demand schedules for specific risk exposures. In practice, however, hedging demand may exhibit some degree of price sensitivity—for example due to funding constraints, risk-budget adjustments, or partial substitution across hedging instruments. In this appendix we develop a simple extension in which expert demand slopes downward with respect to the price of volatility exposure. This extension nests the baseline model and illustrates how the interaction between supply frictions and demand elasticities amplifies the price of volatility term-structure risk.

In the baseline model, expert demand for the three volatility buckets is represented by a fixed vector $h \in \mathbb{R}^3$. Arbitrageur supply is $x = \kappa^{-1} W^{-1} \lambda_{\Delta TS}$, leading to the scarcity-pricing condition

$$\lambda_{\Delta TS} = \kappa W(\Theta - h).$$

To introduce price-sensitive demand, I allow expert demand to depend (linearly) on the price vector $\lambda_{\Delta TS}$. Let

$$h(\lambda_{\Delta TS}) \;=\; h_0 - H\lambda_{\Delta TS}, \tag{A.1}$$

where $h_0 \in \mathbb{R}^3$ is the experts' target hedge (the inelastic quantity demanded in the baseline) and $H \succeq 0$ is a 3×3 demand-elasticity matrix. Diagonal elements of H govern own-price elasticities for the level, slope, and curvature buckets, while off-diagonal elements capture cross-horizon substitution or interaction effects. When $H = 0$, I recover the baseline model exactly.

Market clearing remains

$$x + h(\lambda_{\Delta TS}) = \Theta, \qquad x = \kappa^{-1} W^{-1} \lambda_{\Delta TS}.$$

Inputting Equation A.1 and rearranging gives us the price of volatility term-structure risk

$$\lambda_{\Delta TS} = \big(I - \kappa W H\big)^{-1} \kappa W(\Theta - h_0), \tag{A.2}$$

provided the stability condition $\rho(\kappa WH) < 1$ holds, where $\rho(\cdot)$ denotes the spectral radius.

Compared to the baseline case, the elasticity matrix H alters the mapping between net scarcity $(\Theta - h_0)$ and the equilibrium price of risk. The term $(I - \kappa WH)^{-1}$ acts as an amplification matrix: when prices rise, expert demand falls, reducing scarcity; however, because arbitrageurs face maturity-dependent costs, the adjustment of $\lambda_{\Delta TS}$ feeds back into demand and supply. If $H = 0$, I recover the baseline scarcity-pricing equation $\lambda_{\Delta TS} = \kappa W(\Theta - h_0)$. When H has positive diagonal elements, expert demand becomes less rigid, making the equilibrium price of risk more elastic with respect to supply shocks. Nonzero off-diagonal elements generate cross-horizon interactions, allowing shocks to one maturity bucket to propagate into others. Finally, as $\rho(\kappa WH)$ approaches unity, the amplification becomes large, reflecting strong feedback effects between hedging pressure and intermediary frictions.

Hence, downward-sloping demand curves naturally generate stronger and more interconnected pricing responses to shocks in net supply or hedging motives. This mechanism is present in the linear-quadratic preferred-habitat literature, the margin-based intermediary models of Gârleanu and Pedersen, 2011, and in option-market equilibrium frameworks where hedging demand and capital constraints jointly determine volatility risk premia.

Appendix B. FX data

Table B.1: Sample coverage for foreign exchange series. The table lists all countries included in the sample, their respective currency codes, start and end dates of data availability, and the economic classification of each country (Advanced or Emerging).

Country	Currency code	Start	End	Group
Austria	ATS	Dec-95	Dec-98	Advanced
Australia	AUD	Dec-95	Feb-24	Advanced
Belgium	BEF	Dec-95	Dec-98	Advanced
Bulgaria	BGN	Mar-04	Feb-24	Emerging
Brazil	BRL	Mar-04	Feb-24	Emerging
Canada	CAD	Dec-95	Feb-24	Advanced
Switzerland	CHF	Dec-95	Feb-24	Advanced
Cyprus	CYP	Mar-04	Nov-07	Emerging
Czech Republic	CZK	Dec-96	Sep-13	Emerging
Germany	DEM	Dec-95	Dec-98	Advanced
Denmark	DKK	Dec-95	Feb-24	Advanced
Egypt	EGP	Mar-04	Feb-24	Emerging
Spain	ESP	Dec-95	Dec-98	Advanced
Euro Area	EUR	Jan-99	Feb-24	Advanced
Finland	FIM	Dec-96	Dec-98	Advanced
France	FRF	Dec-95	Dec-98	Advanced
United Kingdom	GBP	Dec-95	Feb-24	Advanced
Greece	GRD	Dec-96	Nov-00	Advanced
Hong Kong	HKD	Dec-95	Feb-24	Advanced
Croatia	HRK	Mar-04	Nov-22	Emerging
Hungary	HUF	Oct-97	Feb-24	Emerging
Indonesia	IDR	Dec-96	Feb-24	Emerging
Ireland	IEP	Dec-95	Dec-98	Advanced
India	INR	Oct-97	Feb-24	Emerging
Iceland	ISK	Mar-04	Feb-24	Advanced
Italy	ITL	Dec-95	Dec-98	Advanced
Japan	JPY	Dec-95	Feb-24	Advanced
South Korea	KRW	Feb-03	Feb-24	Emerging
Kuwait	KWD	Dec-95	Feb-24	Emerging
Mexico	MXN	Dec-96	Feb-24	Emerging
Malaysia	MYR	Dec-96	Feb-24	Emerging
Netherlands	NLG	Dec-95	Dec-98	Advanced
Norway	NOK	Dec-95	Feb-24	Advanced
New Zealand	NZD	Dec-95	Feb-24	Advanced
Philippines	PHP	Dec-96	Feb-24	Emerging
Poland	PLN	Aug-96	Feb-24	Emerging
Portugal	PTE	Dec-95	Dec-98	Emerging
Russia	RUB	Mar-04	Feb-24	Emerging
Saudi Arabia	SAR	Dec-95	Feb-24	Emerging
Sweden	SEK	Dec-95	Feb-24	Advanced
Singapore	SGD	Dec-95	Feb-24	Advanced
Slovenia	SIT	Mar-04	Nov-06	Advanced
Slovakia	SKK	Feb-03	Nov-08	Advanced
Thailand	THB	Dec-95	Feb-24	Emerging
Turkey	TRY	Dec-96	Feb-24	Emerging
Taiwan	TWD	Dec-96	Feb-24	Emerging
Ukraine	UAH	Mar-04	Sep-14	Emerging
South Africa	ZAR	Dec-95	Feb-24	Emerging

Table B.2: FX Spot and Forward Rates. This table reports the Datastream mnemonics of the exchange rates spot and forward data used for each country. Data is sourced from Barclays Bank International (BBI), World Markets Company/Refinitiv (WMR) and Refinitiv, accessed via Datastream.

Country	Spot	1M Forward
Australia	BBAUDSP, AUSTDO$	BBAUD1F, USAUD1F, TDAUD1M
Austria	AUSTSC$	USATS1F, NBBEC1F
Belgium	BELGLU$	USBEF1F, NBBEC1F
Brazil	BRACRU$	USBRL1F, BRL$1MF
Bulgaria	BULGLV$	USBGN1F
Canada	BBCADSP, CNDOLL$	BBCAD1F, USCAD1F, TDCAD1M
Croatia	CROATK$	USHRK1F
Cyprus	CYPRUS$	USCYP1F
Czech	CZECHC$	USCZK1F, TDCZK1M
Denmark	BBDKKSP, DANISH$	BBDKK1F, USDKK1F, TDDKK1M
Egypt	EGYPTN$	USEGP1F
Euro	BBEURSP, USEURSP	BBEUR1F, USEUR1F, TDEUR1M
Finland	FINMAR$	USFIM1F
France	BBFRFSP, FRENFR$	BBFRF1F, USFRF1F, NBFRF1F
Germany	BBDEMSP, DMARKE$	BBDEM1F, USDEM1F, NBDEM1F
Greece	GREDRA$	USGRD1F
Hong Kong	BBHKDSP, HKDOLL$	BBHKD1F, USHKD1F, TDHKD1M
Hungary	HUNFOR$	USHUF1F, TDHUF1M
Iceland	ICEKRO$	USISK1F, TDISK1M
India	INDRUP$	USINR1F, TDINR1M
Indonesia	INDORU$	USIDR1F, TDIDR1M
Ireland	BBIEPSP, IPUNTE$	USIEP1F
Italy	BBITLSP, ITALIR$	BBITL1F, USITL1F, NBITL1F
Japan	BBJPYSP, JAPAYE$	BBJPY1F, USJPY1F, TDJPY1M
Kuwait	KUWADI$	USKWD1F, TDKWD1M
Malaysia	MALADL$	USMYR1F, TDMYR1M
Mexico	MEXPES$	USMXN1F, TDMXN1M
Netherlands	BBNLGSP, GUILDE$	BBNLG1F
New Zealand	BBNZDSP, NZDOLL$	BBNZD1F, USNZD1F, TDNZD1M
Norway	BBNOKSP, NORKRO$	BBNOK1F, USNOK1F, TDNOK1M
Philippines	PHILPE$	USPHP1F, TDPHP1M
Poland	POLZLO$	USPLN1F, TDPLN1M
Portugal	PORTES$	USPTE1F
Russia	CISRUB$	USRUB1F, TDRUB1M
Saudi Arabia	SAUDRI$	USSAR1F, TDSAR1M
Singapore	BBSGDSP, SINGDO$	BBSGD1F, USSGD1F, TDSGD1M
Slovakia	SLOVKO$	USSKK1F
Slovenia	SLOVTO$	USSIT1F
South Africa	BBZARSP, COMRAN$	BBZAR1F, USZAR1F, TDZAR1M
South Korea	KORSWO$	USKRW1F, TDKRW1M
Spain	SPANPE$	USESP1F
Sweden	BBSEKSP, SWEKRO$	BBSEK1F, USSEK1F, TDSEK1M
Switzerland	BBCHFSP, SWISSF$	BBCHF1F, USCHF1F, TDCHF1M
Taiwan	TAIWDO$	USTWD1F, TDTWD1M
Thailand	THABAH$	USTHB1F, TDTHB1M
Turkey	TURKLI$	USTRY1F, TDTRY1M
United Kingdom	BBGBPSP, USDOLLR	BBGBP1F, USGBP1F, TDGBP1M
Ukraine	UKRAHY$	USUAH1F, TDUAH1M

Table B.3: Periods of large deviations from covered interest rate parity. This table reports the time intervals during which covered interest rate parity exhibited substantial deviations for specific countries. Following Kroencke et al., 2014 and Corte et al., 2016, among others, these currencies are excluded from the empirical analysis during these periods.

Country	Start Date	End Date
Egypt	01/01/2011	30/08/2013
	03/10/2016	28/02/2017
Indonesia	01/12/1997	31/07/1998
	01/02/2001	31/05/2005
Malaysia	01/05/1998	30/06/2005
Russia	01/12/2008	30/01/2009
	03/11/2014	27/02/2015
South Africa	01/08/1985	30/08/1985
	01/01/2002	31/05/2005
Turkey	01/11/2000	30/11/2001
Ukraine	03/11/2014	12/04/2024

Appendix C. Summary statistics of test-assets

Appendix C.1. Individual currencies

Table C.1: Currencies test-assets summary statistics. This table reports summary statistics for monthly currency excess returns of test-asset currencies. The Sharpe Ratio is computed as the ratio of average annualized return to annualized volatility. *t-stat* refers to the t-statistic testing whether the average monthly excess return is significantly different from zero. The sample corresponds to the baseline analysis period using traded VIX futures.

Currency	Avg Ann. Return	t-stat	Ann. Volatility	Sharpe Ratio	Skewness
AUD	3.09	[0.29]	11.37	0.27	-0.01
BGN	-4.05	[-0.65]	6.76	-0.60	-0.11
BRL	4.58	[0.41]	12.16	0.38	-1.00
CAD	-4.20	[-0.63]	7.21	-0.58	-0.34
CHF	0.07	[0.01]	8.94	0.01	0.99
CZK	-11.36	[-1.41]	8.68	-1.31	-0.10
DKK	-4.13	[-0.67]	6.64	-0.62	-0.10
EGP	6.32	[2.59]	2.64	2.39	-0.21
EUR	-4.08	[-0.66]	6.71	-0.61	-0.13
GBP	-0.31	[-0.07]	5.10	-0.06	0.13
HKD	-0.13	[-0.45]	0.31	-0.42	-0.77
HRK	-3.83	[-0.54]	7.65	-0.50	-0.15
HUF	-3.86	[-0.40]	10.56	-0.37	-0.37
IDR	2.91	[0.32]	9.86	0.30	-0.12
INR	12.69	[2.23]	6.15	2.06	0.38
ISK	5.13	[0.79]	7.00	0.73	0.36
JPY	-8.90	[-1.14]	8.41	-1.06	-0.22
KRW	4.42	[0.77]	6.17	0.72	-0.78
KWD	-0.58	[-0.37]	1.72	-0.34	-0.78
MXN	1.82	[0.39]	5.09	0.36	-0.41
MYR	4.10	[0.63]	7.08	0.58	-0.35
NOK	-2.83	[-0.26]	11.65	-0.24	0.60
NZD	1.36	[0.14]	10.87	0.13	-0.29
PHP	-1.98	[-0.45]	4.73	-0.42	-0.45
PLN	-0.00	[-0.00]	9.30	-0.00	0.38
RUB	-13.88	[-1.14]	13.21	-1.05	-0.27
SAR	-0.06	[-1.84]	0.04	-1.50	-1.34
SEK	-7.68	[-1.07]	7.74	-0.99	0.94
SGD	0.62	[0.15]	4.44	0.14	0.62
THB	-1.48	[-0.34]	4.65	-0.32	-0.79
TRY	-1.34	[-0.13]	10.96	-0.12	-0.46
TWD	0.63	[0.13]	5.15	0.12	1.66
UAH	-31.28	[-1.90]	17.77	-1.76	-1.31
ZAR	3.72	[0.33]	12.32	0.30	-0.04

Appendix C.2. Portfolios

Table C.2: Test-assets portfolios summary statistics. This table reports summary statistics for monthly currency excess returns of test-asset portfolios. Panel A sorts currencies on the previous month forward premium as in Lustig et al., 2011, Panel B on previous month excess return as in Menkhoff et al., 2012, and Panel C on the 5-year-log-return of the real exchange rate as in Menkhoff et al., 2017. P1 contains currencies with the lowest beta exposures, while P5 contains those with the highest. The Sharpe Ratio is computed as the ratio of average annualized return to annualized volatility. *t-stat* refers to the t-statistic testing whether the average monthly excess return is significantly different from zero. The sample corresponds to the baseline analysis period using traded VIX futures.

Portfolio	P1	P2	P3	P4	P5
Panel A: Carry					
Avg Ann. Return	-8.85	-1.71	-1.38	-1.26	2.13
t-stat	[-3.83]	[-1.01]	[-0.73]	[-0.70]	[0.88]
Ann. Volatility	8.45	6.16	6.88	6.57	8.84
Sharpe Ratio	-1.05	-0.28	-0.20	-0.19	0.24
Skewness	-0.99	0.12	-0.39	-0.55	-0.94
Panel B: Momentum					
Avg Ann. Return	-5.92	-1.45	-1.45	-1.64	-0.21
t-stat	[-2.16]	[-0.76]	[-0.81]	[-0.92]	[-0.11]
Ann. Volatility	10.04	6.98	6.49	6.51	7.19
Sharpe Ratio	-0.59	-0.21	-0.22	-0.25	-0.03
Skewness	-0.68	0.28	-0.31	-0.73	-0.48
Panel C: Value					
Avg Ann. Return	0.63	-1.86	-2.65	-1.17	-1.80
t-stat	[0.24]	[-0.79]	[-1.27]	[-0.56]	[-0.91]
Ann. Volatility	9.66	8.61	7.62	7.53	7.24
Sharpe Ratio	0.07	-0.22	-0.35	-0.16	-0.25
Skewness	0.05	-0.26	-0.31	-0.54	-0.70

Appendix D. Alternative number of portfolios in factor construction

Appendix D.1. Tercile sorts

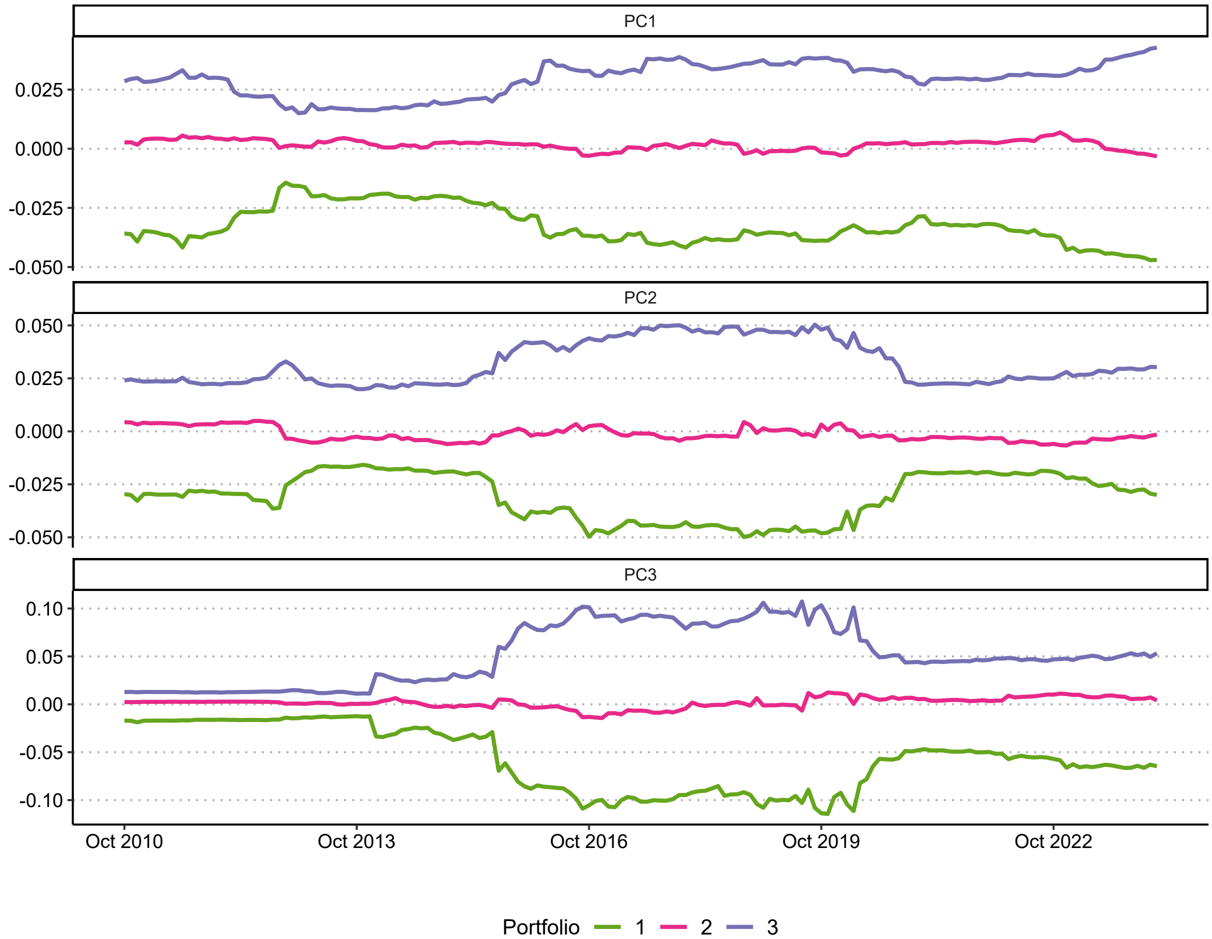

Figure D.1: Rolling betas to VIX principal component innovations by currency portfolio. This figure shows the time-series of average betas for five currency portfolios sorted each month by their estimated exposure to innovations in VIX term-structure components. The principal components—capturing the level (PC1), slope (PC2), and curvature (PC3) of the VIX futures curve—are extracted monthly using an expanding-window PCA. Innovations in each component are computed as the monthly change relative to the previous observation. For each currency, rolling 48-month regressions of excess returns on these innovations are estimated following Equation 18, controlling for the dollar factor. The resulting betas measure time-varying sensitivities to shocks in the VIX term-structure. The figure reports portfolio-average betas across the five beta-sorted groups, illustrating persistent cross-sectional heterogeneity and cyclical variation in exposures to volatility-level (PC1), slope (PC2), and curvature (PC3) innovations.

Table D.3: VIX futures beta-sorted portfolios summary statistics. This table reports summary statistics for monthly currency excess returns of portfolios sorted by estimated betas with respect to VIX futures innovations. Panel A sorts currencies on β_{PC1} (level), Panel B on β_{PC2} (slope), and Panel C on β_{PC3} (curvature). P1 contains currencies with the lowest beta exposures, while P5 contains those with the highest. HML denotes the high-minus-low portfolio (P5 – P1). The Sharpe Ratio is computed as the ratio of average annualized return to annualized volatility. *t-stat* refers to the t-statistic testing whether the average monthly excess return is significantly different from zero.

Portfolio	P1	P2	P3	HML
Panel A: sorts on β_{PC1}				
Avg Ann. Return	-0.39	-1.00	-5.37	-4.97
t-stat	[-0.21]	[-0.64]	[-2.49]	[-3.27]
Ann. Volatility	6.96	5.77	7.90	5.57
Sharpe Ratio	**-0.06**	**-0.17**	**-0.68**	**-0.89**
Skewness	-0.13	-0.58	-0.37	-0.69
Panel B: sorts on β_{PC2}				
Avg Ann. Return	-4.97	-0.87	-1.22	3.75
t-stat	[-2.39]	[-0.61]	[-0.58]	[2.55]
Ann. Volatility	7.63	5.21	7.66	5.38
Sharpe Ratio	**-0.65**	**-0.17**	**-0.16**	**0.70**
Skewness	-0.30	-0.41	-0.62	0.46
Panel C: sorts on β_{PC3}				
Avg Ann. Return	-0.53	-1.28	-4.96	-4.43
t-stat	[-0.26]	[-0.87]	[-2.39]	[-3.35]
Ann. Volatility	7.42	5.43	7.61	4.85
Sharpe Ratio	**-0.07**	**-0.24**	**-0.65**	**-0.91**
Skewness	-0.25	-0.59	-0.35	-0.42

Table D.4: Time-series regression of VIX futures beta-sorted portfolios on the Carry factor. This table reports the results of time-series regressions of monthly excess returns of VIX futures beta-sorted currency portfolios on the Carry factor of Lustig et al., 2011. Panels A–C correspond to portfolios sorted on betas with respect to innovations to the first three principal components of VIX futures: level (PC1), slope (PC2), and curvature (PC3), respectively. P1 includes currencies with the lowest beta exposure and P5 includes those with the highest. HML denotes the high-minus-low portfolio (P5 – P1). Newey-West t-statistics are reported in brackets.

Portfolio	P1	P2	P3	HML
Panel A: sorts on β_{PC1}				
Intercept	-0.03	-0.02	-0.05	-0.02
	[-1.65]	[-1.35]	[-1.99]	[-1.32]
Carry Factor	0.40	0.17	-0.09	-0.49
	[3.36]	[1.57]	[-0.43]	[-4.81]
R^2	0.148	0.040	0.005	0.342
Panel B: sorts on β_{PC2}				
Intercept	-0.04	-0.02	-0.04	0.01
	[-1.98]	[-1.21]	[-1.84]	[0.44]
Carry Factor	-0.08	0.15	0.39	0.47
	[-0.37]	[1.66]	[2.80]	[4.40]
R^2	0.005	0.038	0.116	0.338
Panel C: sorts on β_{PC3}				
Intercept	-0.03	-0.02	-0.05	-0.02
	[-1.27]	[-1.59]	[-2.10]	[-1.64]
Carry Factor	0.31	0.18	-0.00	-0.31
	[2.64]	[1.76]	[-0.01]	[-3.19]
R^2	0.079	0.048	0.000	0.187

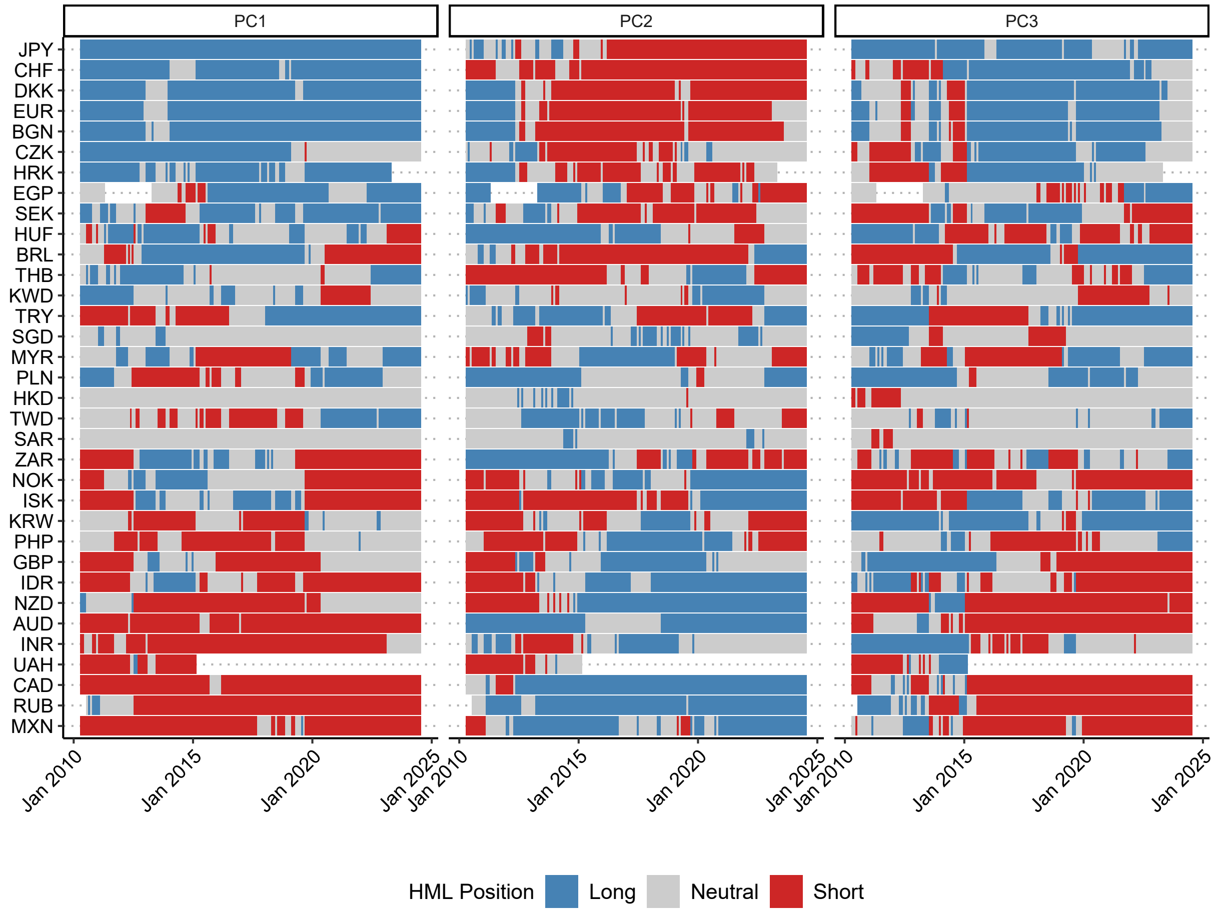

Figure D.2: Composition of currency risk factors constructed from VIX principal component betas. This figure shows the time-varying composition of the high-minus-low (HML) currency risk factors associated with exposures to innovations in the first three VIX term-structure principal components: level (PC1), slope (PC2), and curvature (PC3). For each month, currencies are sorted into five portfolios based on their rolling 48-month betas with respect to the respective principal component innovation, and the factor is constructed as the return difference between the highest- and lowest-beta portfolios. The heatmap indicates, for each currency and month, whether it enters the factor with a long position (blue), short position (red), or is excluded from the factor (gray). Each panel corresponds to one VIX component. The figure highlights persistent cross-sectional asymmetries in how currencies load onto the VIX futures term-structure level, slope, and curvature shocks, as well as shifts in factor composition reflecting changing sources of volatility risk exposure across time.

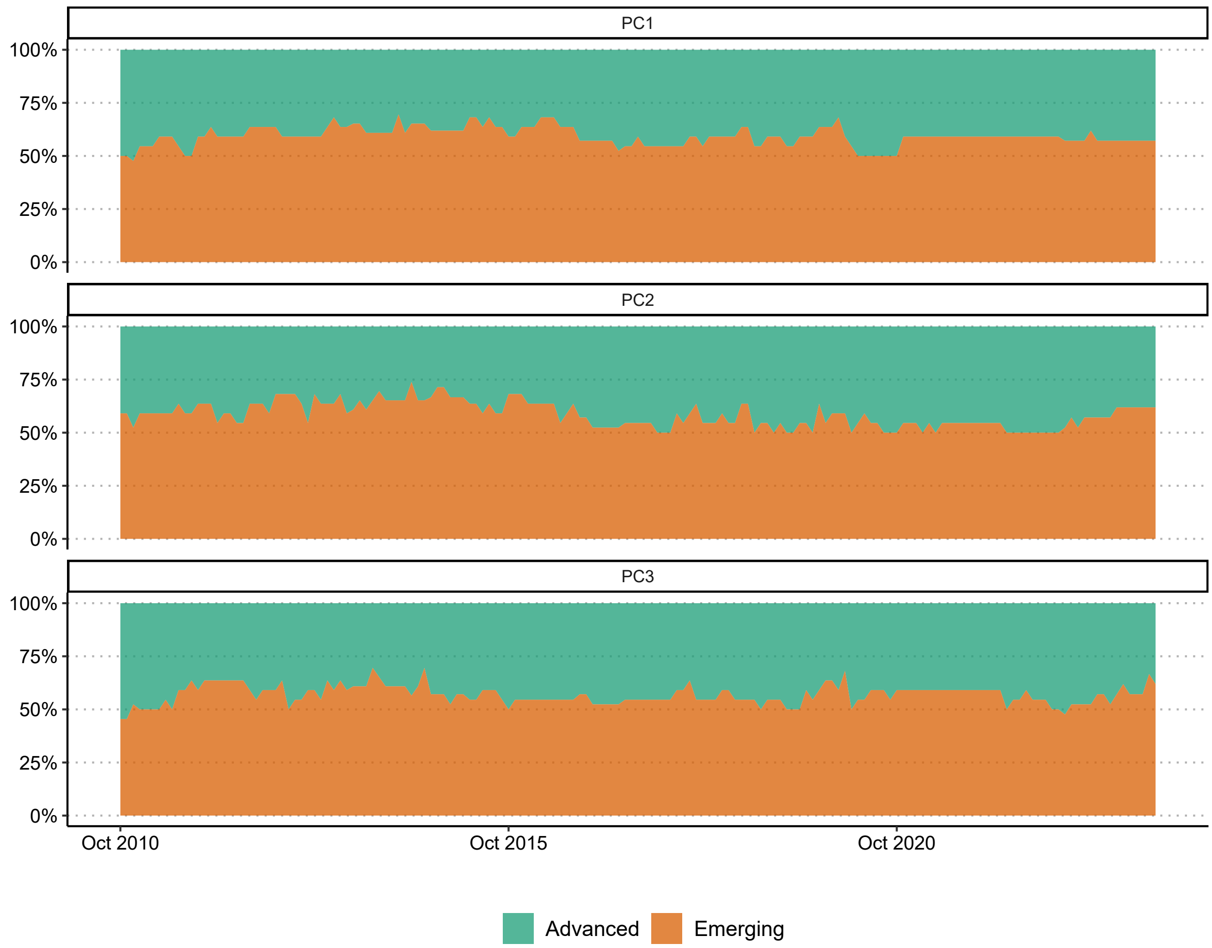

Figure D.3: Advanced and emerging market composition of VIX principal component factors. This figure presents the share of advanced and emerging market currencies in the aggregate high-minus-low (HML) portfolios constructed from exposures to innovations in the first three VIX term-structure principal components: level (PC1), slope (PC2), and curvature (PC3). Each panel shows, over time, the proportion of advanced (green) and emerging (orange) currencies included in the combined factor portfolio, computed monthly based on the underlying beta-sorted portfolio assignments. The composition reveals that there is no consistent dominance of either country group across factors: the relative shares of advanced and emerging currencies fluctuate over time and differ across components, suggesting that exposure to volatility term-structure shocks is not systematically concentrated in one market type but varies with global conditions.

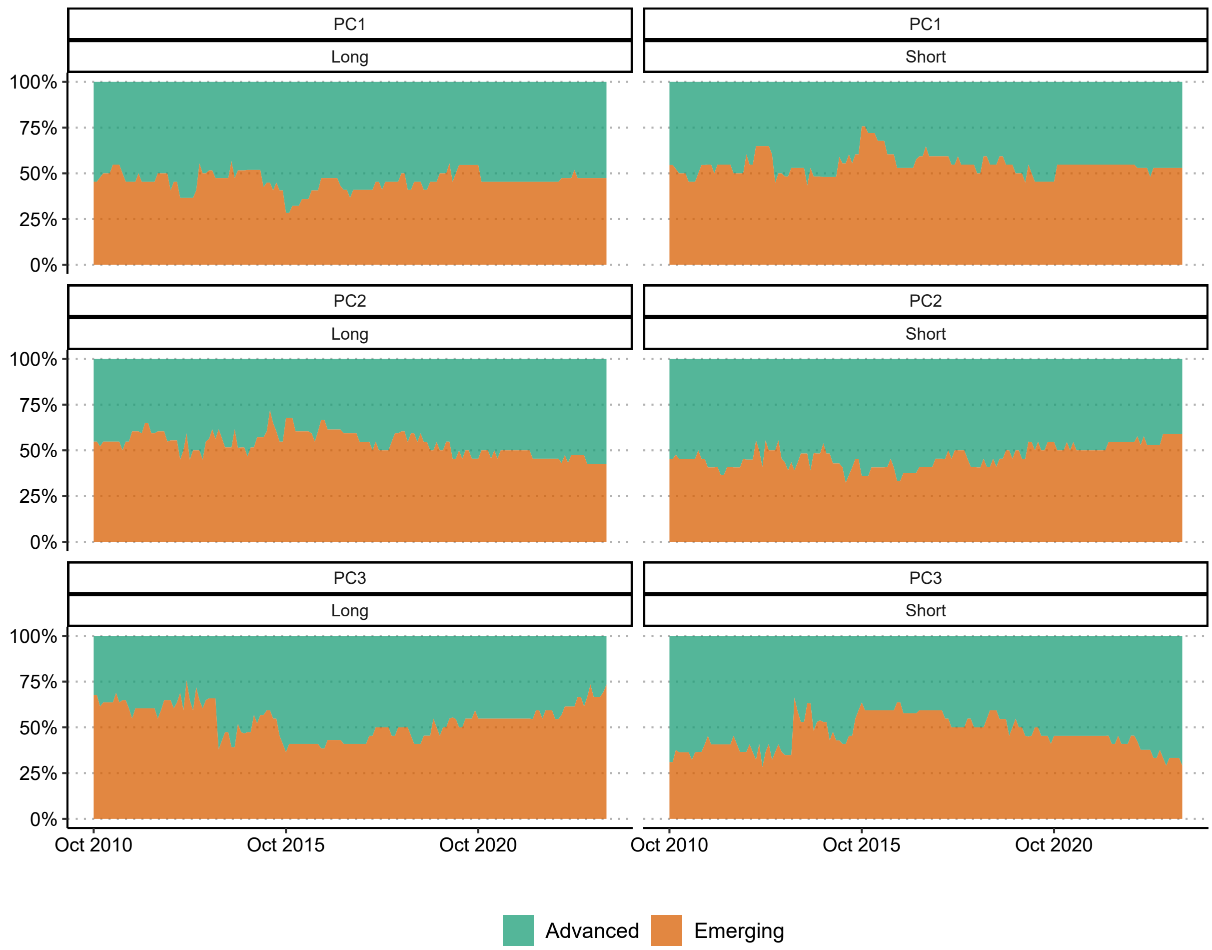

Figure D.4: Advanced and emerging market composition of VIX principal component factors. This figure shows the share of advanced and emerging market currencies included in the long and short sides of the high-minus-low (HML) currency risk factors constructed from exposures to innovations in the first three VIX term-structure principal components: level (PC1), slope (PC2), and curvature (PC3). Each panel plots, over time, the proportion of advanced (green) and emerging (orange) currencies composing the long (left) and short (right) portfolios of each factor. The proportions are computed monthly based on the portfolio assignments underlying the factor construction. The figure highlights systematic differences in how advanced and emerging currencies load onto volatility term-structure shocks: exposures to the level (PC1) are typically dominated by advanced currencies on the long side, while slope (PC2) and curvature (PC3) factors exhibit more balanced or alternating compositions, reflecting shifts in global volatility transmission across market groups.

Table D.5: Asset pricing tests of VIX futures term-structure factors - second stage. This table reports the second-stage cross-sectional asset pricing results for the portfolio-level test-assets using the VIX futures term-structure currency risk factors. The dependent variables are the annualized average excess returns of the test assets, while the regressors are the betas estimated in the first-stage time-series regressions. Columns (1)–(5) correspond to a cross-section of Carry, Momentum, and Value portfolios, while columns (6)–(10) use only the Momentum and Value portfolios. Each specification includes the Dollar and Carry factors from the literature, along with different combinations of the three VIX-based currency risk factors: level ($fPC1$), slope ($fPC2$), and curvature ($fPC3$). The λ estimates represent the price of risk associated with each factor. I also report the p-values from the test of Gibbons et al., 1989, which tests the joint hypothesis that all pricing errors (alphas) are equal to zero—i.e., that the model prices all assets correctly. A failure to reject this null suggests that the pricing model performs well in explaining cross-sectional variation in returns. Standard errors are corrected following Shanken, 1992, and statistical significance is denoted by *, **, and *** for the 10%, 5%, and 1% levels, respectively.

	Test assets									
	Carry + Momentum + Value					Momentum + Value				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
β_{fPC1}		−0.045***			−0.044***		−0.045***			−0.044***
		(0.014)			(0.014)		(0.014)			(0.014)
β_{fPC2}			0.034**		0.033**			0.033**		0.033**
			(0.014)		(0.014)			(0.014)		(0.014)
β_{fPC3}				−0.041***	−0.041***				−0.040***	−0.040***
				(0.013)	(0.013)				(0.013)	(0.013)
β_{CAR}	0.071***	0.068***	0.068***	0.068***	0.068***	0.130**	0.069***	0.069***	0.069***	0.069***
	(0.020)	(0.019)	(0.019)	(0.019)	(0.019)	(0.065)	(0.018)	(0.018)	(0.018)	(0.018)
β_{DOL}	−0.020	−0.021	−0.021	−0.021	−0.021	−0.021	−0.020	−0.020	−0.021	−0.021
	(0.017)	(0.017)	(0.017)	(0.017)	(0.017)	(0.018)	(0.018)	(0.017)	(0.017)	(0.017)
R^2	0.882	0.967	0.971	0.972	0.976	0.819	0.968	0.971	0.970	0.976
GRS	0.787	0.833	0.83	0.723	0.751	0.641	0.734	0.718	0.589	0.627

Table D.6: Asset pricing tests of VIX futures term-structure factors - second stage. This table reports the second-stage cross-sectional asset pricing results with country-level currency excess returns as test assets using the VIX futures term-structure currency risk factors. The dependent variables are the annualized average excess returns of the test assets, while the regressors are the betas estimated in the first-stage time-series regressions. Each specification includes the Dollar and Carry factors from the literature, along with different combinations of the three VIX-based currency risk factors: level ($fPC1$), slope ($fPC2$), and curvature ($fPC3$). The λ estimates represent the price of risk associated with each factor. I also report the p-values from the test of Gibbons et al., 1989, which tests the joint hypothesis that all pricing errors (alphas) are equal to zero—i.e., that the model prices all assets correctly. A failure to reject this null suggests that the pricing model performs well in explaining cross-sectional variation in returns. Standard errors are corrected following Shanken, 1992, and statistical significance is denoted by *, **, and *** for the 10%, 5%, and 1% levels, respectively.

	Test assets				
	Country level				
	(1)	(2)	(3)	(4)	(5)
β_{fPC1}		-0.072^{**}			-0.027
		(0.034)			(0.018)
β_{fPC2}			0.068^{**}		0.054^{**}
			(0.031)		(0.022)
β_{fPC3}				-0.056^{**}	-0.048
				(0.028)	(0.030)
β_{CAR}	0.055^{**}	0.045^{**}	0.041^{*}	0.042^{**}	0.041^{*}
	(0.024)	(0.021)	(0.021)	(0.021)	(0.021)
β_{DOL}	-0.021	-0.021	-0.021	-0.021	-0.021
	(0.017)	(0.017)	(0.017)	(0.017)	(0.017)
R^2	0.653	0.803	0.854	0.859	0.875
GRS	0.04	0.06	0.054	0.045	0.049

Appendix D.2. Decile sorts

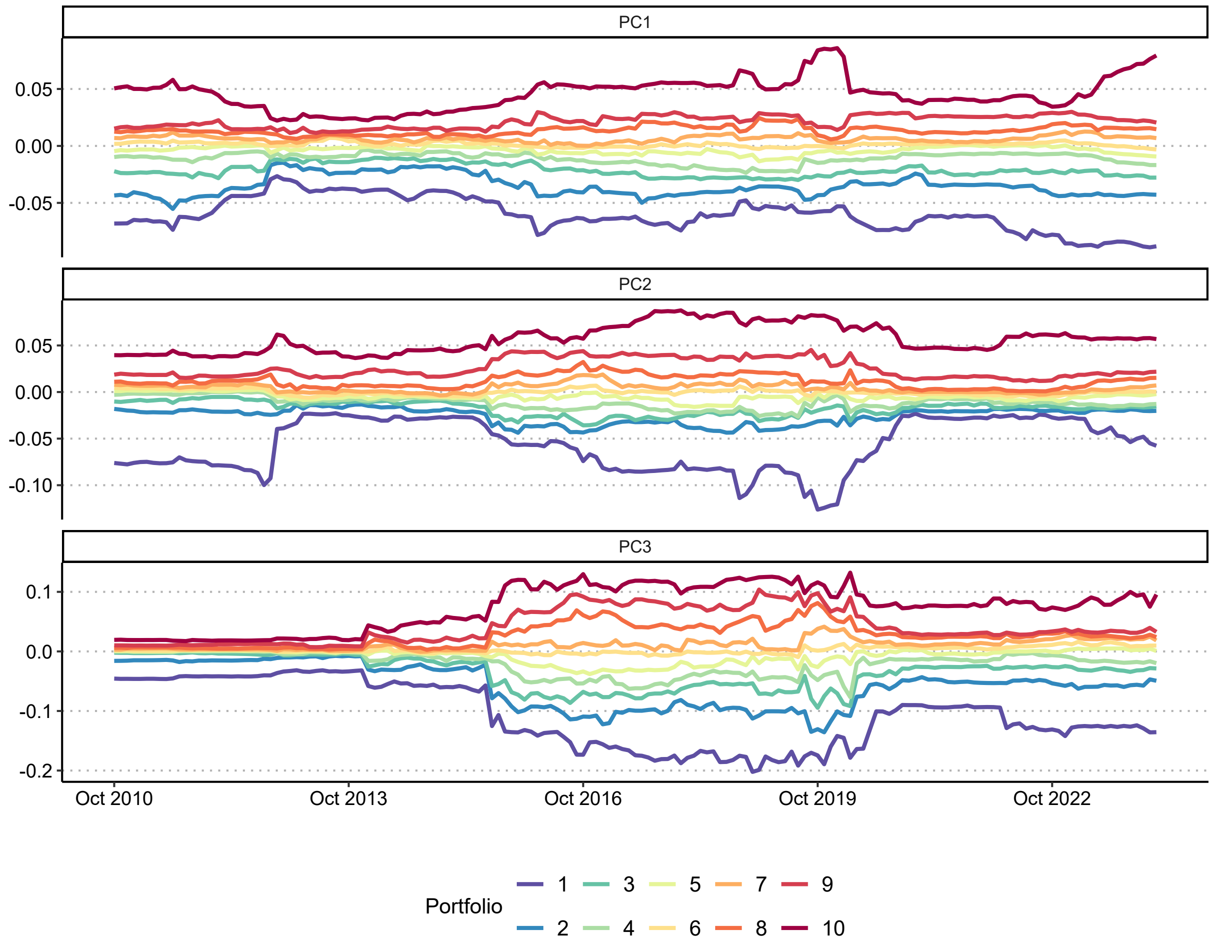

Figure D.5: Rolling betas to VIX principal component innovations by currency portfolio. This figure shows the time-series of average betas for five currency portfolios sorted each month by their estimated exposure to innovations in VIX term-structure components. The principal components—capturing the level (PC1), slope (PC2), and curvature (PC3) of the VIX futures curve—are extracted monthly using an expanding-window PCA. Innovations in each component are computed as the monthly change relative to the previous observation. For each currency, rolling 48-month regressions of excess returns on these innovations are estimated following Equation 18, controlling for the dollar factor. The resulting betas measure time-varying sensitivities to shocks in the VIX term-structure. The figure reports portfolio-average betas across the five beta-sorted groups, illustrating persistent cross-sectional heterogeneity and cyclical variation in exposures to volatility-level (PC1), slope (PC2), and curvature (PC3) innovations.

Table D.7: VIX futures beta-sorted portfolios summary statistics. This table reports summary statistics for monthly currency excess returns of portfolios sorted by estimated betas with respect to VIX futures innovations. Panel A sorts currencies on β_{PC1} (level), Panel B on β_{PC2} (slope), and Panel C on β_{PC3} (curvature). P1 contains currencies with the lowest beta exposures, while P5 contains those with the highest. HML denotes the high-minus-low portfolio (P5 – P1). The Sharpe Ratio is computed as the ratio of average annualized return to annualized volatility. *t-stat* refers to the t-statistic testing whether the average monthly excess return is significantly different from zero.

Portfolio	P1	P2	P3	P4	P5	P6	P7	P8	P9	P10	HML
Panel A: sorts on β_{PC1}											
Avg Ann. Return	1.33	-0.77	-1.26	-1.84	1.07	-2.39	-1.21	-3.37	-2.30	-9.58	-10.91
t-stat	[0.51]	[-0.35]	[-0.58]	[-0.95]	[0.60]	[-1.56]	[-0.57]	[-1.49]	[-1.11]	[-3.18]	[-3.33]
Ann. Volatility	9.49	8.04	7.98	7.09	6.58	5.62	7.77	8.31	7.62	11.02	11.98
Sharpe Ratio	**0.14**	**-0.10**	**-0.16**	**-0.26**	**0.16**	**-0.43**	**-0.16**	**-0.41**	**-0.30**	**-0.87**	**-0.91**
Skewness	-0.05	-0.05	-0.34	-0.66	-0.84	-0.40	-0.91	-0.84	-0.19	-1.37	-1.11
Panel B: sorts on β_{PC2}											
Avg Ann. Return	-11.04	-3.22	-1.63	-1.49	-1.31	-0.28	-1.93	-1.66	-0.18	-0.95	10.10
t-stat	[-3.21]	[-1.53]	[-0.67]	[-0.77]	[-0.74]	[-0.17]	[-1.12]	[-0.86]	[-0.08]	[-0.34]	[2.69]
Ann. Volatility	12.62	7.74	8.96	7.06	6.43	5.84	6.31	7.08	8.14	10.32	13.75
Sharpe Ratio	**-0.87**	**-0.42**	**-0.18**	**-0.21**	**-0.20**	**-0.05**	**-0.31**	**-0.23**	**-0.02**	**-0.09**	**0.73**
Skewness	-2.14	-0.29	-0.72	-0.44	-0.14	0.11	-0.50	-0.66	-0.56	-0.45	1.45
Panel C: sorts on β_{PC3}											
Avg Ann. Return	0.17	0.07	-0.58	-3.68	-1.98	-0.51	0.54	-2.66	-1.30	-10.17	-10.34
t-stat	[0.06]	[0.03]	[-0.26]	[-2.02]	[-1.16]	[-0.32]	[0.29]	[-1.26]	[-0.64]	[-3.39]	[-3.38]
Ann. Volatility	10.35	8.80	8.34	6.67	6.23	5.76	6.81	7.74	7.38	10.99	11.21
Sharpe Ratio	**0.02**	**0.01**	**-0.07**	**-0.55**	**-0.32**	**-0.09**	**0.08**	**-0.34**	**-0.18**	**-0.93**	**-0.92**
Skewness	-0.37	-0.22	-0.07	-0.60	-1.15	-0.76	-0.41	-0.48	-0.47	-1.30	-1.14

Table D.8: Time-series regression of VIX futures beta-sorted portfolios on the Carry factor. This table reports the results of time-series regressions of monthly excess returns of VIX futures beta-sorted currency portfolios on the Carry factor of Lustig et al., 2011. Panels A–C correspond to portfolios sorted on betas with respect to innovations to the first three principal components of VIX futures: level (PC1), slope (PC2), and curvature (PC3), respectively. P1 includes currencies with the lowest beta exposure and P5 includes those with the highest. HML denotes the high-minus-low portfolio (P5 – P1). Newey-West t-statistics are reported in brackets.

Portfolio	P1	P2	P3	P4	P5	P6	P7	P8	P9	P10	HML
Panel A: sorts on β_{PC1}											
Intercept	-0.03	-0.02	-0.02	-0.03	0.01	-0.03	-0.03	-0.03	-0.02	-0.04	-0.01
	[-1.45]	[-0.85]	[-0.94]	[-1.71]	[0.26]	[-1.97]	[-1.43]	[-1.19]	[-1.08]	[-1.20]	[-0.26]
Carry Factor	0.25	0.08	0.05	0.08	0.03	0.05	0.10	-0.02	-0.01	-0.33	-0.59
	[2.04]	[1.27]	[1.05]	[1.57]	[0.67]	[1.10]	[1.78]	[-0.31]	[-0.11]	[-1.38]	[-5.32]
R^2	0.139	0.018	0.008	0.027	0.004	0.016	0.030	0.001	0.000	0.176	0.465
Panel B: sorts on β_{PC2}											
Intercept	-0.04	-0.03	-0.02	-0.02	-0.02	-0.01	-0.03	-0.03	-0.01	-0.04	-0.00
	[-1.18]	[-1.28]	[-0.75]	[-0.84]	[-0.95]	[-0.51]	[-1.66]	[-1.40]	[-0.63]	[-1.71]	[-0.10]
Carry Factor	-0.41	-0.01	0.03	0.01	0.02	0.03	0.06	0.07	0.07	0.20	0.60
	[-1.23]	[-0.29]	[0.39]	[0.33]	[0.72]	[0.95]	[1.32]	[1.21]	[1.41]	[2.11]	[3.44]
R^2	0.203	0.001	0.002	0.001	0.003	0.006	0.017	0.018	0.016	0.071	0.375
Panel C: sorts on β_{PC3}											
Intercept	-0.04	-0.00	-0.01	-0.04	-0.04	-0.01	-0.00	-0.03	-0.02	-0.05	-0.01
	[-1.61]	[-0.23]	[-0.53]	[-2.04]	[-1.74]	[-0.75]	[-0.11]	[-1.44]	[-0.71]	[-1.74]	[-0.39]
Carry Factor	0.25	0.03	0.04	0.02	0.09	0.04	0.04	0.02	0.03	-0.29	-0.54
	[2.11]	[0.56]	[0.89]	[0.36]	[1.69]	[0.95]	[0.87]	[0.33]	[0.49]	[-0.65]	[-5.77]
R^2	0.113	0.002	0.006	0.001	0.043	0.008	0.008	0.001	0.002	0.137	0.453

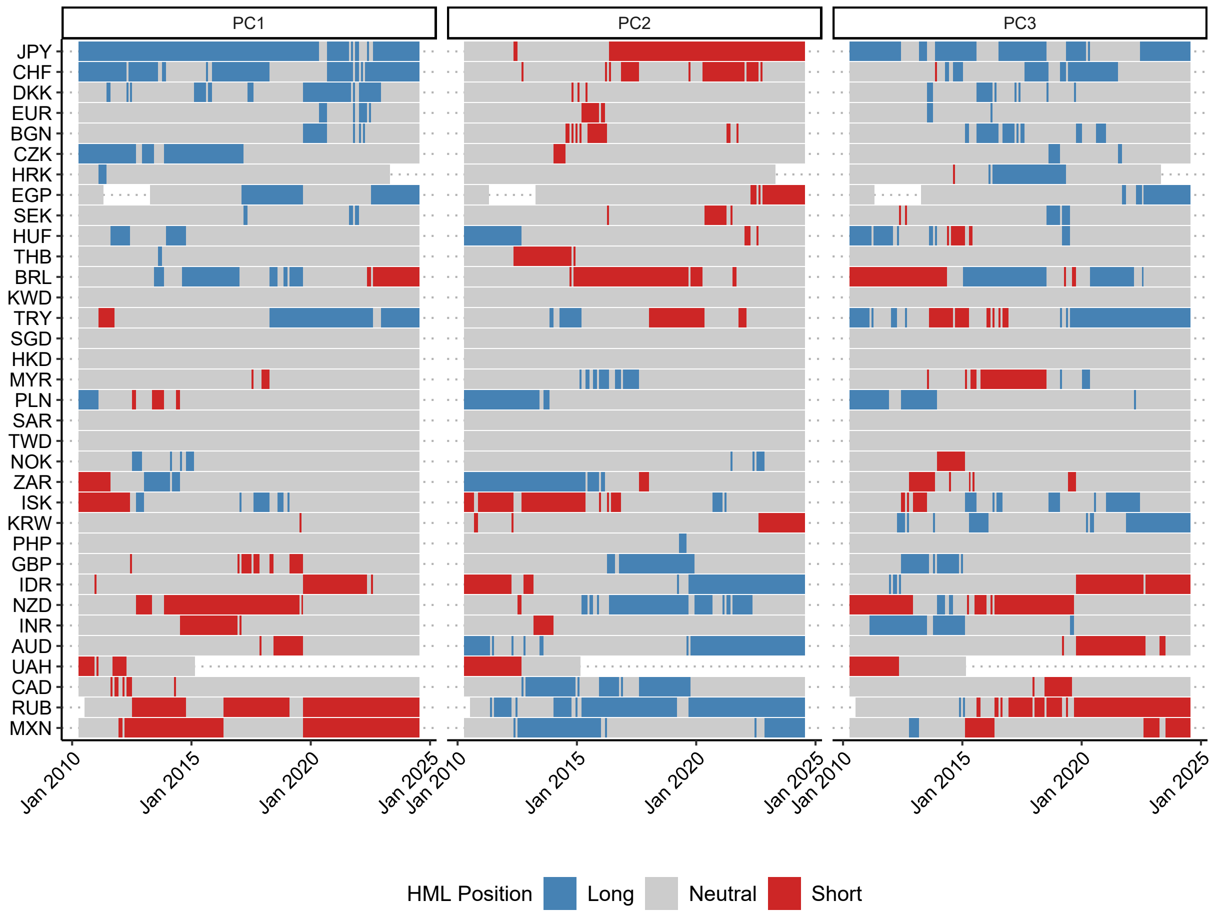

Figure D.6: Composition of currency risk factors constructed from VIX principal component betas. This figure shows the time-varying composition of the high-minus-low (HML) currency risk factors associated with exposures to innovations in the first three VIX term-structure principal components: level (PC1), slope (PC2), and curvature (PC3). For each month, currencies are sorted into five portfolios based on their rolling 48-month betas with respect to the respective principal component innovation, and the factor is constructed as the return difference between the highest- and lowest-beta portfolios. The heatmap indicates, for each currency and month, whether it enters the factor with a long position (blue), short position (red), or is excluded from the factor (gray). Each panel corresponds to one VIX component. The figure highlights persistent cross-sectional asymmetries in how currencies load onto the VIX futures term-structure level, slope, and curvature shocks, as well as shifts in factor composition reflecting changing sources of volatility risk exposure across time.

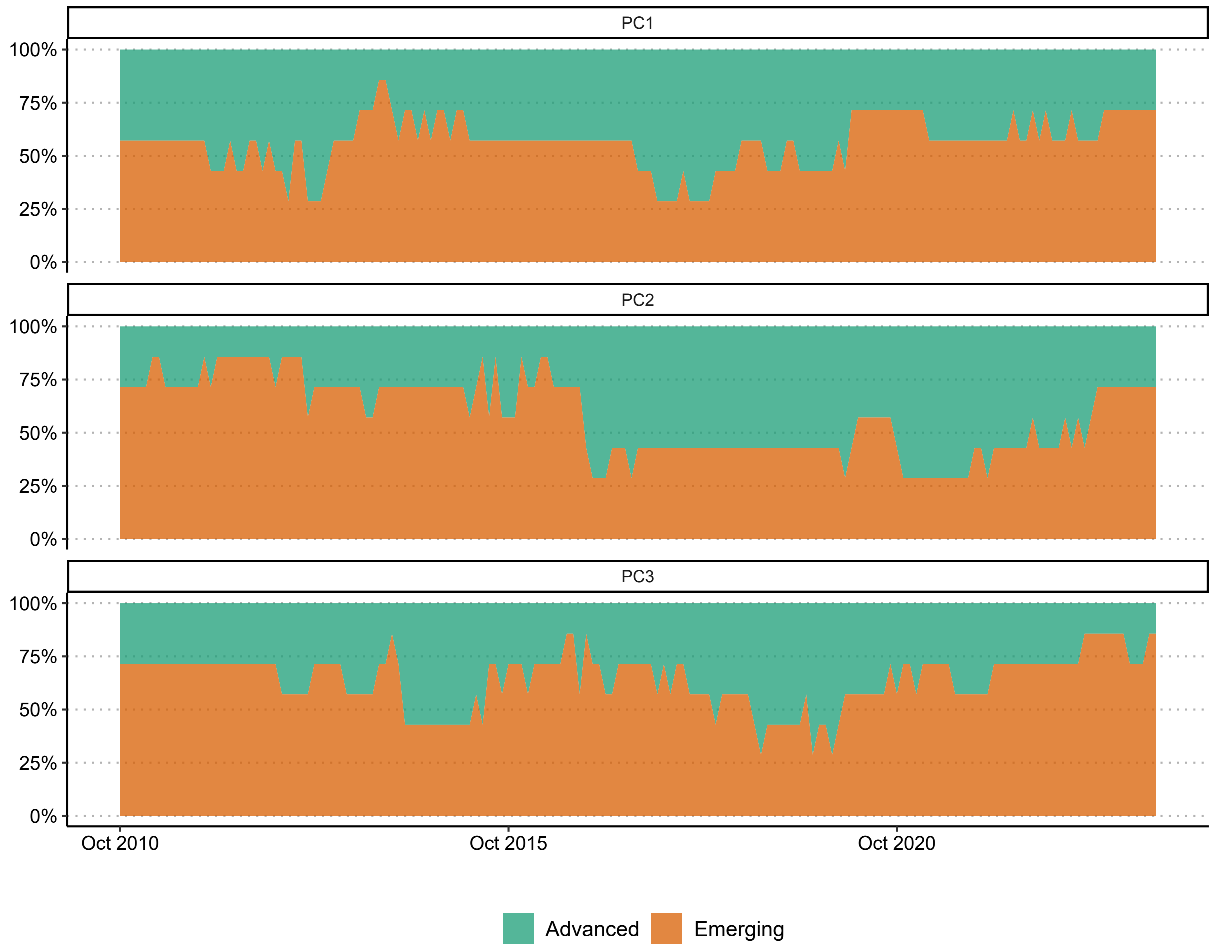

Figure D.7: Advanced and emerging market composition of VIX principal component factors. This figure presents the share of advanced and emerging market currencies in the aggregate high-minus-low (HML) portfolios constructed from exposures to innovations in the first three VIX term-structure principal components: level (PC1), slope (PC2), and curvature (PC3). Each panel shows, over time, the proportion of advanced (green) and emerging (orange) currencies included in the combined factor portfolio, computed monthly based on the underlying beta-sorted portfolio assignments. The composition reveals that there is no consistent dominance of either country group across factors: the relative shares of advanced and emerging currencies fluctuate over time and differ across components, suggesting that exposure to volatility term-structure shocks is not systematically concentrated in one market type but varies with global conditions.

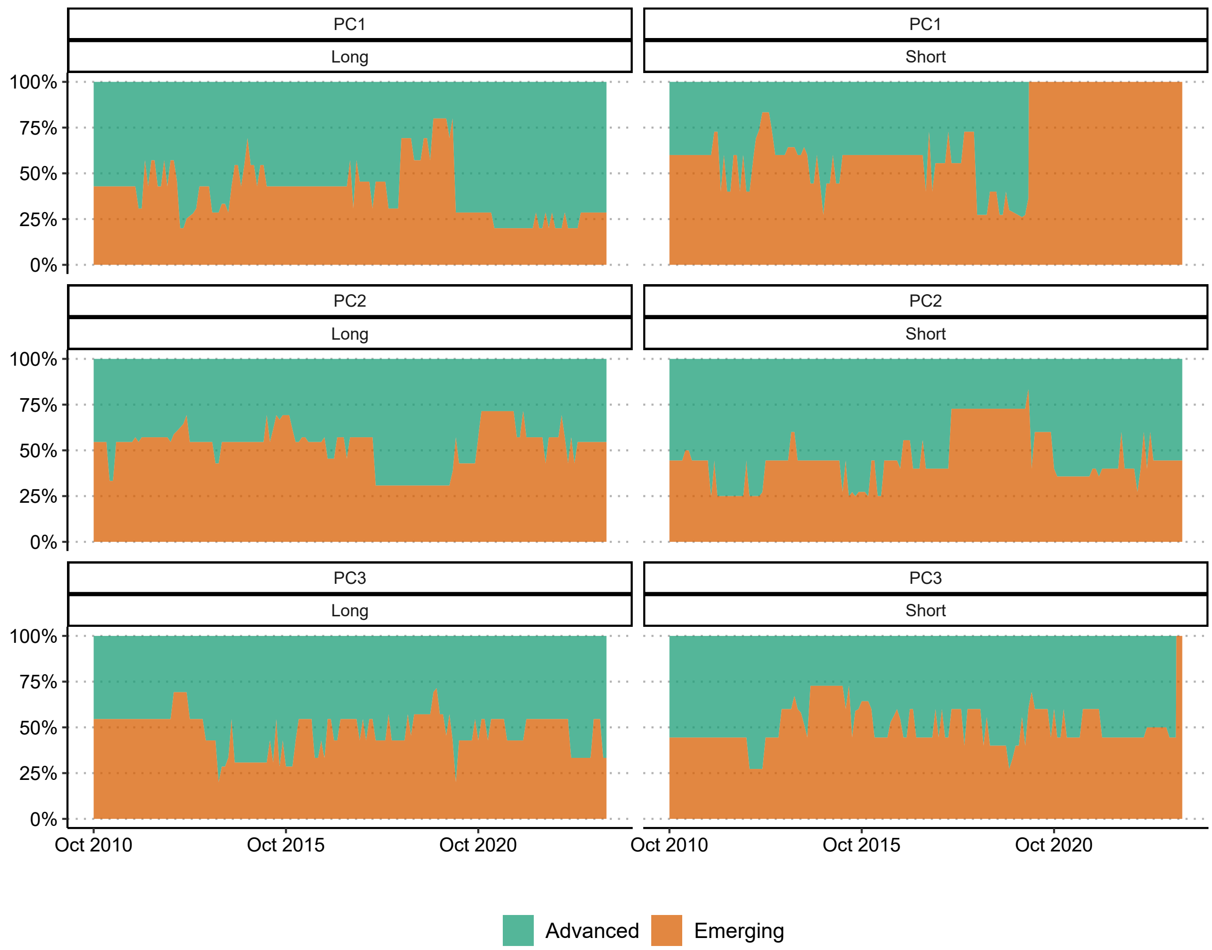

Figure D.8: Advanced and emerging market composition of VIX principal component factors. This figure shows the share of advanced and emerging market currencies included in the long and short sides of the high-minus-low (HML) currency risk factors constructed from exposures to innovations in the first three VIX term-structure principal components: level (PC1), slope (PC2), and curvature (PC3). Each panel plots, over time, the proportion of advanced (green) and emerging (orange) currencies composing the long (left) and short (right) portfolios of each factor. The proportions are computed monthly based on the portfolio assignments underlying the factor construction. The figure highlights systematic differences in how advanced and emerging currencies load onto volatility term-structure shocks: exposures to the level (PC1) are typically dominated by advanced currencies on the long side, while slope (PC2) and curvature (PC3) factors exhibit more balanced or alternating compositions, reflecting shifts in global volatility transmission across market groups.

Table D.9: Asset pricing tests of VIX futures term-structure factors - second stage. This table reports the second-stage cross-sectional asset pricing results for the portfolio-level test-assets using the VIX futures term-structure currency risk factors. The dependent variables are the annualized average excess returns of the test assets, while the regressors are the betas estimated in the first-stage time-series regressions. Columns (1)–(5) correspond to a cross-section of Carry, Momentum, and Value portfolios, while columns (6)–(10) use only the Momentum and Value portfolios. Each specification includes the Dollar and Carry factors from the literature, along with different combinations of the three VIX-based currency risk factors: level ($fPC1$), slope ($fPC2$), and curvature ($fPC3$). The λ estimates represent the price of risk associated with each factor. I also report the p-values from the test of Gibbons et al., 1989, which tests the joint hypothesis that all pricing errors (alphas) are equal to zero—i.e., that the model prices all assets correctly. A failure to reject this null suggests that the pricing model performs well in explaining cross-sectional variation in returns. Standard errors are corrected following Shanken, 1992, and statistical significance is denoted by *, **, and *** for the 10%, 5%, and 1% levels, respectively.

	Test assets									
	Carry + Momentum + Value					Momentum + Value				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
β_{fPC1}		−0.094***			−0.098***		−0.095***			−0.099***
		(0.030)			(0.030)		(0.030)			(0.030)
β_{fPC2}			0.087***		0.087***			0.087***		0.088***
			(0.034)		(0.034)			(0.034)		(0.034)
β_{fPC3}				−0.084***	−0.081***				−0.086***	−0.081***
				(0.028)	(0.028)				(0.028)	(0.028)
β_{CAR}	0.175***	0.173***	0.173***	0.173***	0.173***	0.181***	0.174***	0.174***	0.174***	0.173***
	(0.040)	(0.039)	(0.039)	(0.039)	(0.039)	(0.046)	(0.039)	(0.039)	(0.039)	(0.039)
β_{DOL}	−0.023	−0.023	−0.023	−0.023	−0.023	−0.023	−0.023	−0.023	−0.023	−0.023
	(0.017)	(0.017)	(0.017)	(0.017)	(0.017)	(0.017)	(0.017)	(0.017)	(0.017)	(0.017)
R^2	0.879	0.952	0.954	0.947	0.958	0.758	0.950	0.951	0.944	0.959
GRS	0.87	0.894	0.878	0.883	0.869	0.798	0.831	0.807	0.803	0.769

Table D.10: Asset pricing tests of VIX futures term-structure factors - second stage. This table reports the second-stage cross-sectional asset pricing results with country-level currency excess returns as test assets using the VIX futures term-structure currency risk factors. The dependent variables are the annualized average excess returns of the test assets, while the regressors are the betas estimated in the first-stage time-series regressions. Each specification includes the Dollar and Carry factors from the literature, along with different combinations of the three VIX-based currency risk factors: level ($fPC1$), slope ($fPC2$), and curvature ($fPC3$). The λ estimates represent the price of risk associated with each factor. I also report the p-values from the test of Gibbons et al., 1989, which tests the joint hypothesis that all pricing errors (alphas) are equal to zero—i.e., that the model prices all assets correctly. A failure to reject this null suggests that the pricing model performs well in explaining cross-sectional variation in returns. Standard errors are corrected following Shanken, 1992, and statistical significance is denoted by *, **, and *** for the 10%, 5%, and 1% levels, respectively.

	Test assets				
	Country level				
	(1)	(2)	(3)	(4)	(5)
β_{fPC1}		-0.112^{***}			-0.111^{***}
		(0.040)			(0.032)
β_{fPC2}			0.097^{**}		0.081^{**}
			(0.043)		(0.038)
β_{fPC3}				-0.099^{**}	-0.092^{***}
				(0.040)	(0.035)
β_{CAR}	0.112^{**}	0.115^{***}	0.117^{***}	0.116^{***}	0.116^{***}
	(0.046)	(0.042)	(0.042)	(0.042)	(0.042)
β_{DOL}	-0.022	-0.021	-0.022	-0.022	-0.021
	(0.017)	(0.017)	(0.017)	(0.017)	(0.017)
R^2	0.726	0.853	0.838	0.844	0.873
GRS	0.033	0.045	0.036	0.045	0.055

Appendix E. Option-implied VIX futures

Given the shorter availability of exchange-traded VIX futures, I extend the sample by combining the option-replicated VIX futures from Johnson, 2017 with the actual traded contracts. The option-implied series, constructed from S&P500 index options following the CBOE VIX methodology, provides a daily synthetic term-structure of one- to nine-month maturities from January 1996 to July 2019. I then append the corresponding exchange-traded VIX futures from August 2019 through 2024. As shown in Johnson, 2017, the synthetic and traded contracts are highly correlated, supporting the continuity of the combined series.

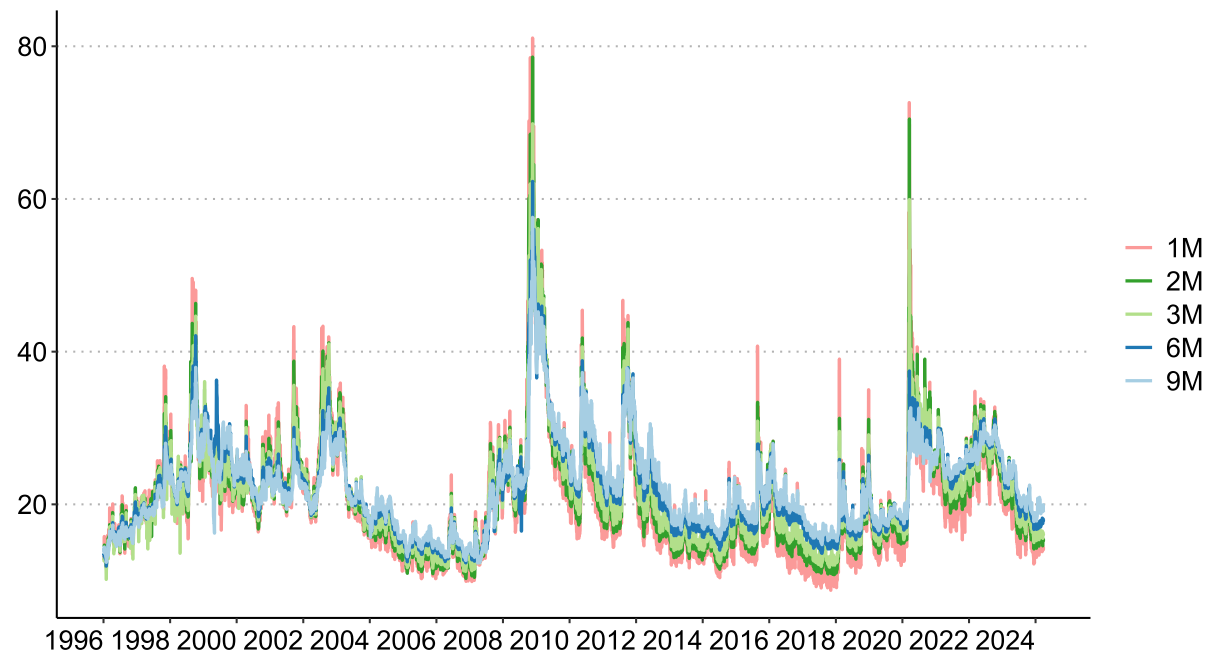

Figure E.1: Time-series of option-implied VIX futures, daily

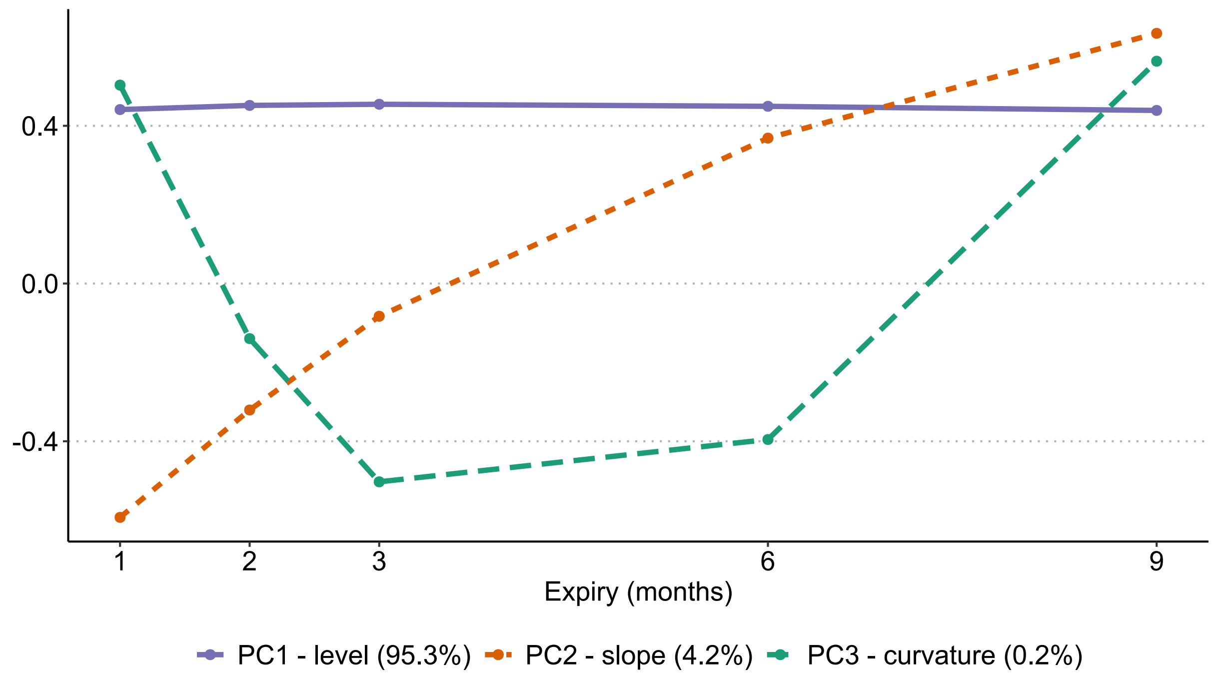

Figure E.2: Factor loadings of option-implied VIX futures PCA. This figure shows the factor loadings of the full-sample principal component analysis on the level of option-implied VIX futures. Loadings are scaled such that they correspond to common measures of level, slope, and curvature. Percentage of variance explained by each principal component is between parenthesis.

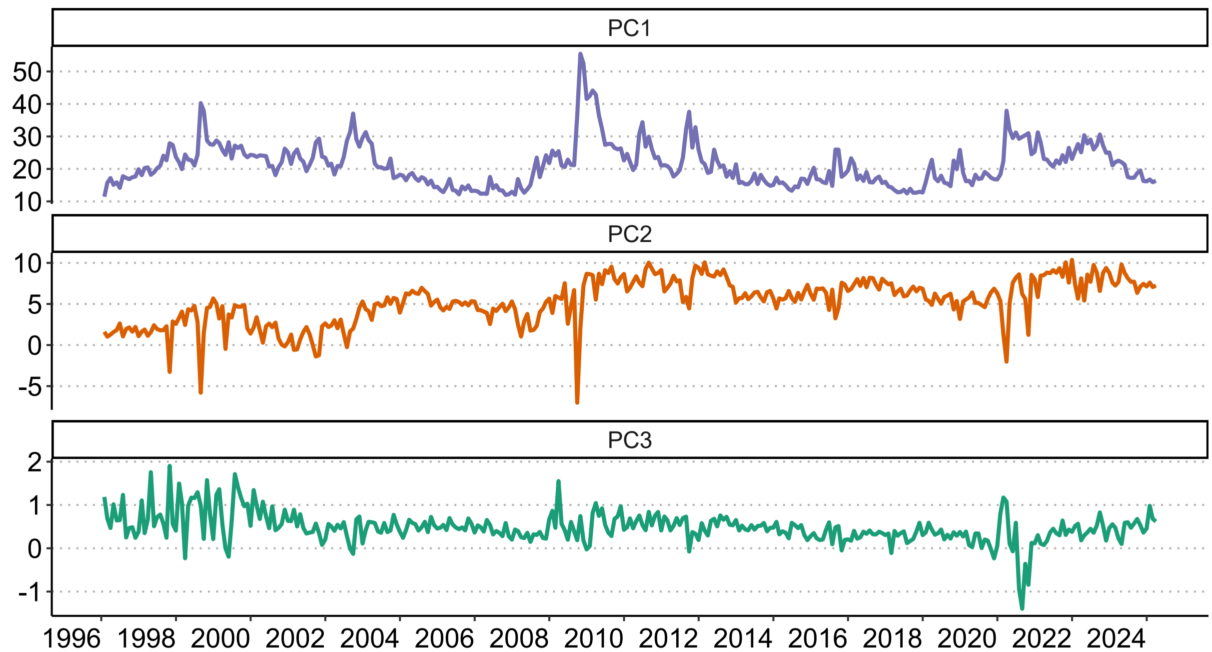

Figure E.3: Principal components of expanding window PCA. This figure shows the time-series of the first three principal components extracted from option-implied VIX futures using an expanding window PCA, with a startup period of 252 days. PCA is performed on daily data, with the resulting factors sampled at month-end. The signs of the loadings are such that PC1 captures the level of the curve, PC2 captures the slope, and PC3 captures curvature. Increases in PC1 reflect upward shifts in the entire curve—typically associated with deteriorating market conditions—while increases in PC2 reflect steepening, often interpreted as improving conditions. PC3 captures changes in curvature, though its interpretation is less directly linked to macro-financial states.

Figure E.4: Rolling betas to option-implied VIX principal component innovations by currency portfolio. This figure shows the time-series of average betas for five currency portfolios sorted each month by their estimated exposure to innovations in option-implied VIX term-structure components. The principal components—capturing the level (PC1), slope (PC2), and curvature (PC3) of the option-implied VIX futures curve—are extracted monthly using an expanding-window PCA. Innovations in each component are computed as the monthly change relative to the previous observation. For each currency, rolling 48-month regressions of excess returns on these innovations are estimated following Equation 18, controlling for the dollar factor. The resulting betas measure time-varying sensitivities to shocks in the option-implied VIX term-structure. The figure reports portfolio-average betas across the five beta-sorted groups, illustrating persistent cross-sectional heterogeneity and cyclical variation in exposures to volatility-level (PC1), slope (PC2), and curvature (PC3) innovations.

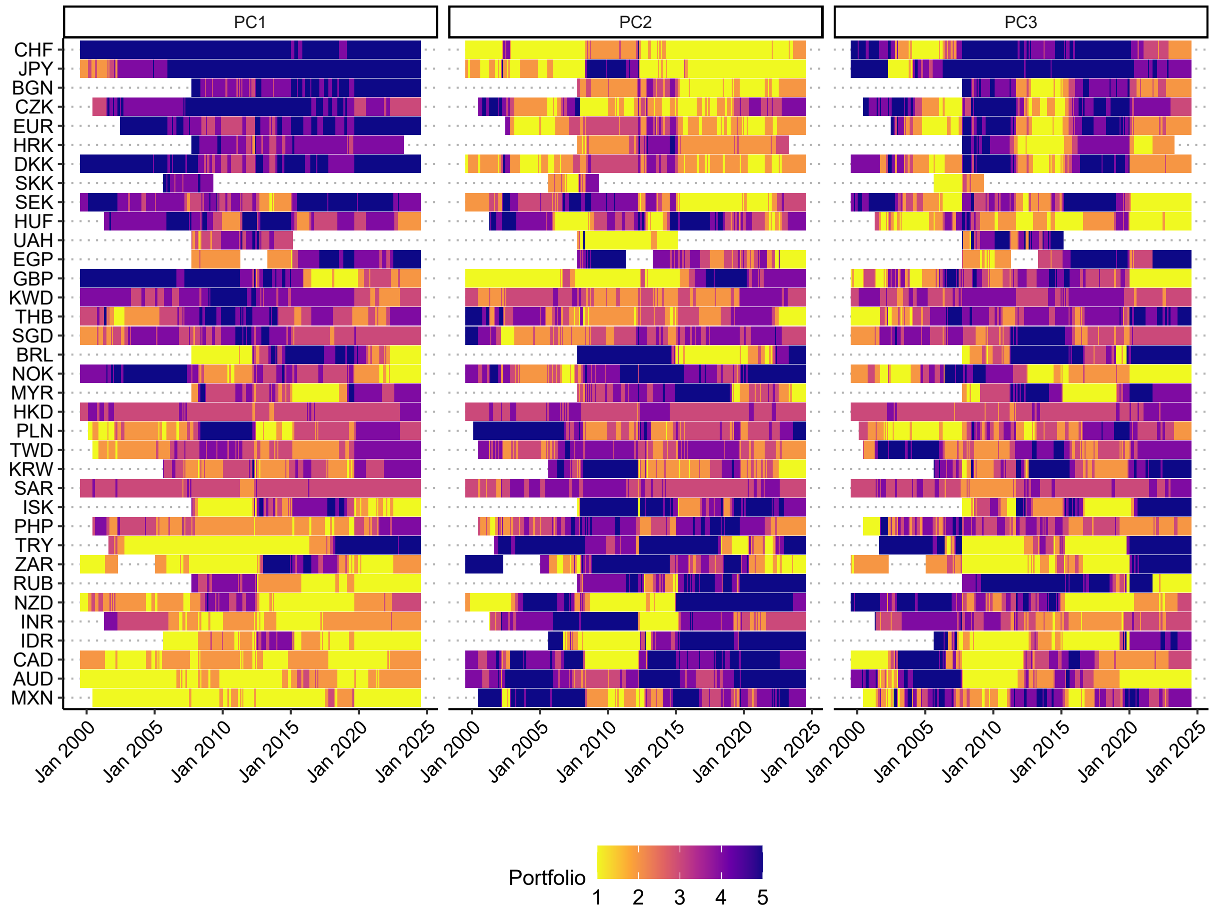

Figure E.5: Composition of currency portfolios sorted on option-implied VIX futures principal component betas. This figure displays the time-varying portfolio assignments of individual currencies based on their rolling 48-month betas with respect to innovations in the first three option-implied VIX term-structure principal components: level (PC1), slope (PC2), and curvature (PC3). Each color represents the portfolio (from 1 to 5) to which a currency belongs in a given month, determined by its beta ranking from the previous month's estimation. Left corresponds to portfolios sorted on PC1 betas, middle panel on PC2, and right on PC3. Lighter shades indicate lower-exposure portfolios, while darker shades represent higher exposures. The figure highlights persistent clustering of currencies with similar sensitivities to volatility term-structure shocks, as well as transitions across portfolios reflecting evolving exposure patterns over time.

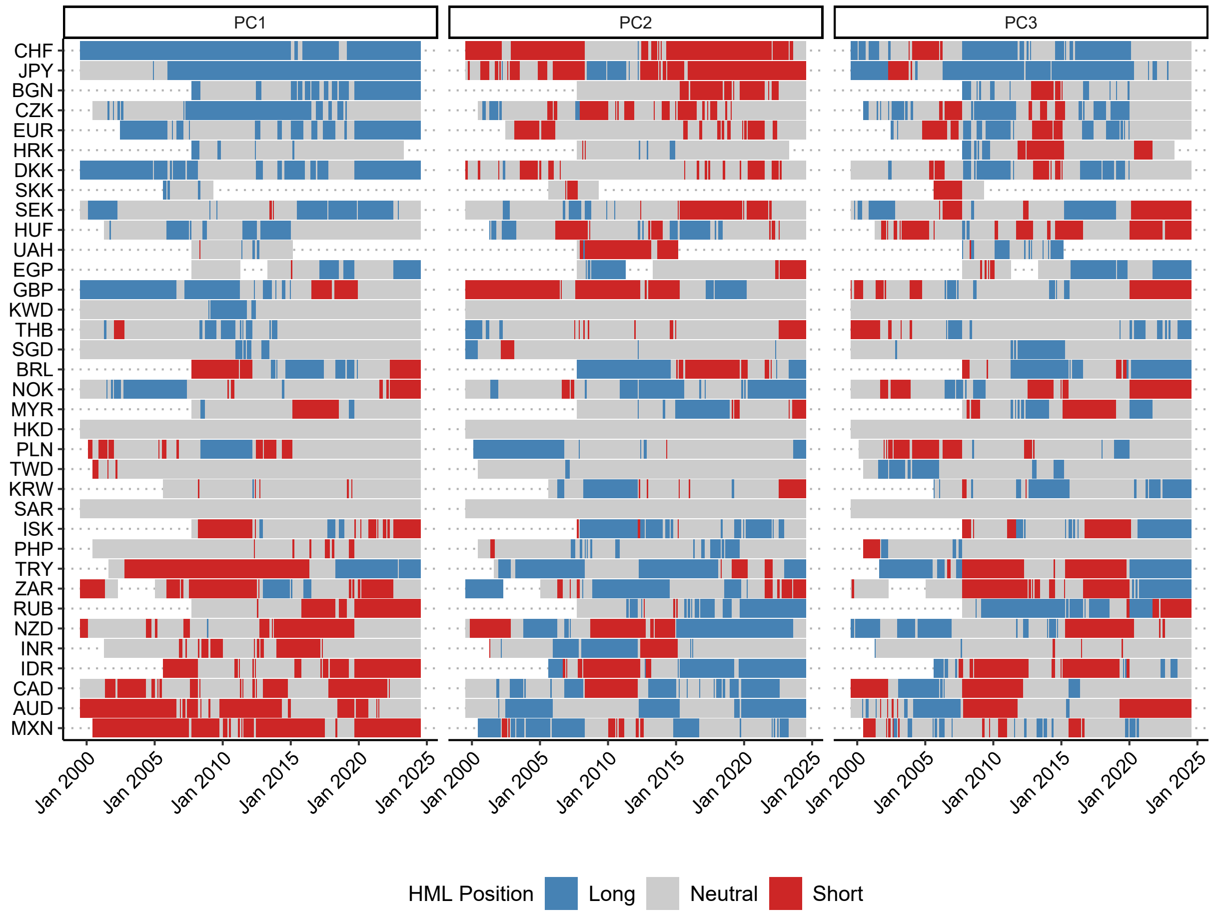

Figure E.6: Composition of currency risk factors constructed from option-implied VIX futures principal component betas. This figure shows the time-varying composition of the high-minus-low (HML) currency risk factors associated with exposures to innovations in the first three option-implied VIX term-structure principal components: level (PC1), slope (PC2), and curvature (PC3). For each month, currencies are sorted into five portfolios based on their rolling 48-month betas with respect to the respective principal component innovation, and the factor is constructed as the return difference between the highest- and lowest-beta portfolios. The heatmap indicates, for each currency and month, whether it enters the factor with a long position (blue), short position (red), or is excluded from the factor (gray). Each panel corresponds to one VIX component. The figure highlights persistent cross-sectional asymmetries in how currencies load onto the VIX futures term-structure level, slope, and curvature shocks, as well as shifts in factor composition reflecting changing sources of volatility risk exposure across time.

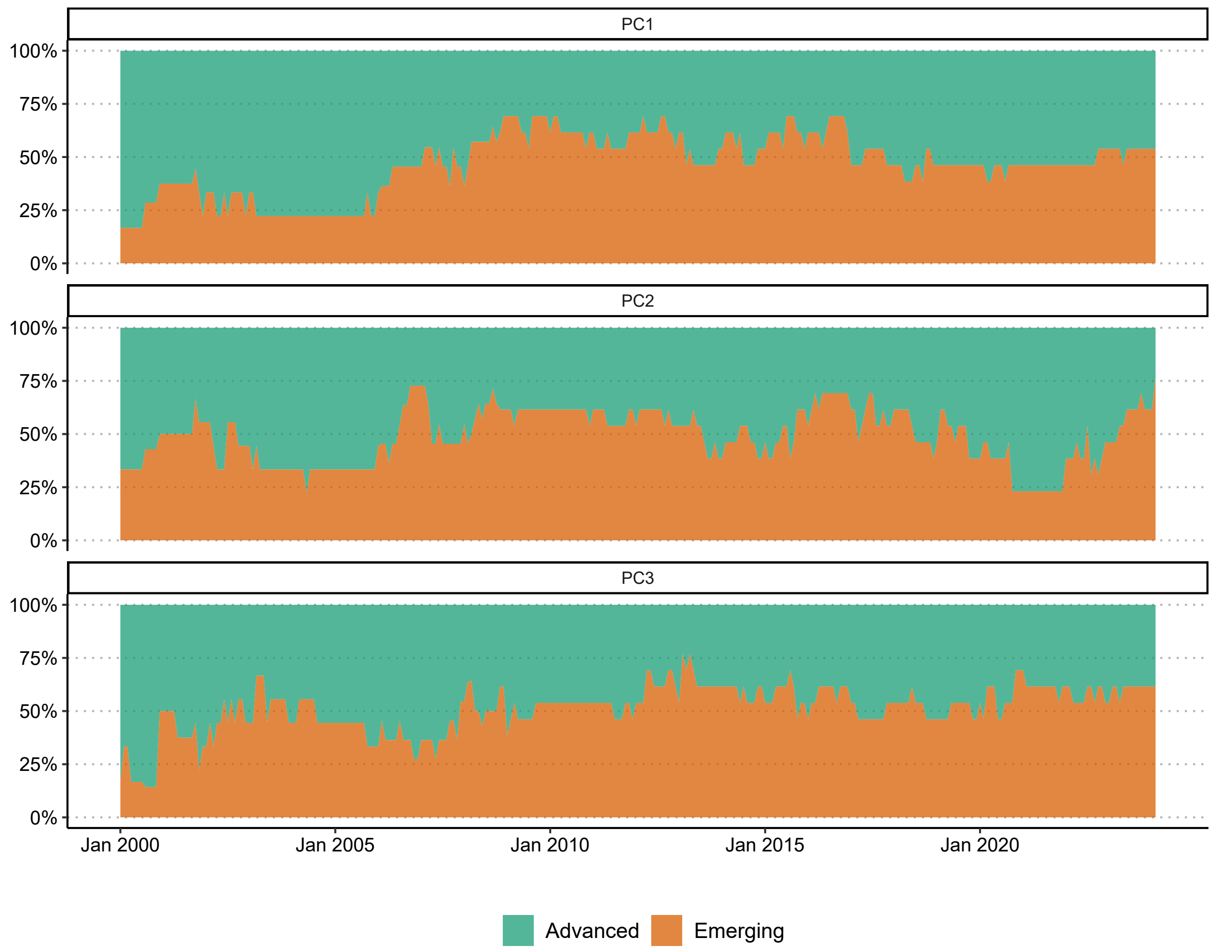

Figure E.7: Advanced and emerging market composition of option-implied VIX futures principal component factors. This figure presents the share of advanced and emerging market currencies in the aggregate high-minus-low (HML) portfolios constructed from exposures to innovations in the first three option-implied VIX term-structure principal components: level (PC1), slope (PC2), and curvature (PC3). Each panel shows, over time, the proportion of advanced (green) and emerging (orange) currencies included in the combined factor portfolio, computed monthly based on the underlying beta-sorted portfolio assignments. The composition reveals that there is no consistent dominance of either country group across factors: the relative shares of advanced and emerging currencies fluctuate over time and differ across components, suggesting that exposure to volatility term-structure shocks is not systematically concentrated in one market type but varies with global conditions.

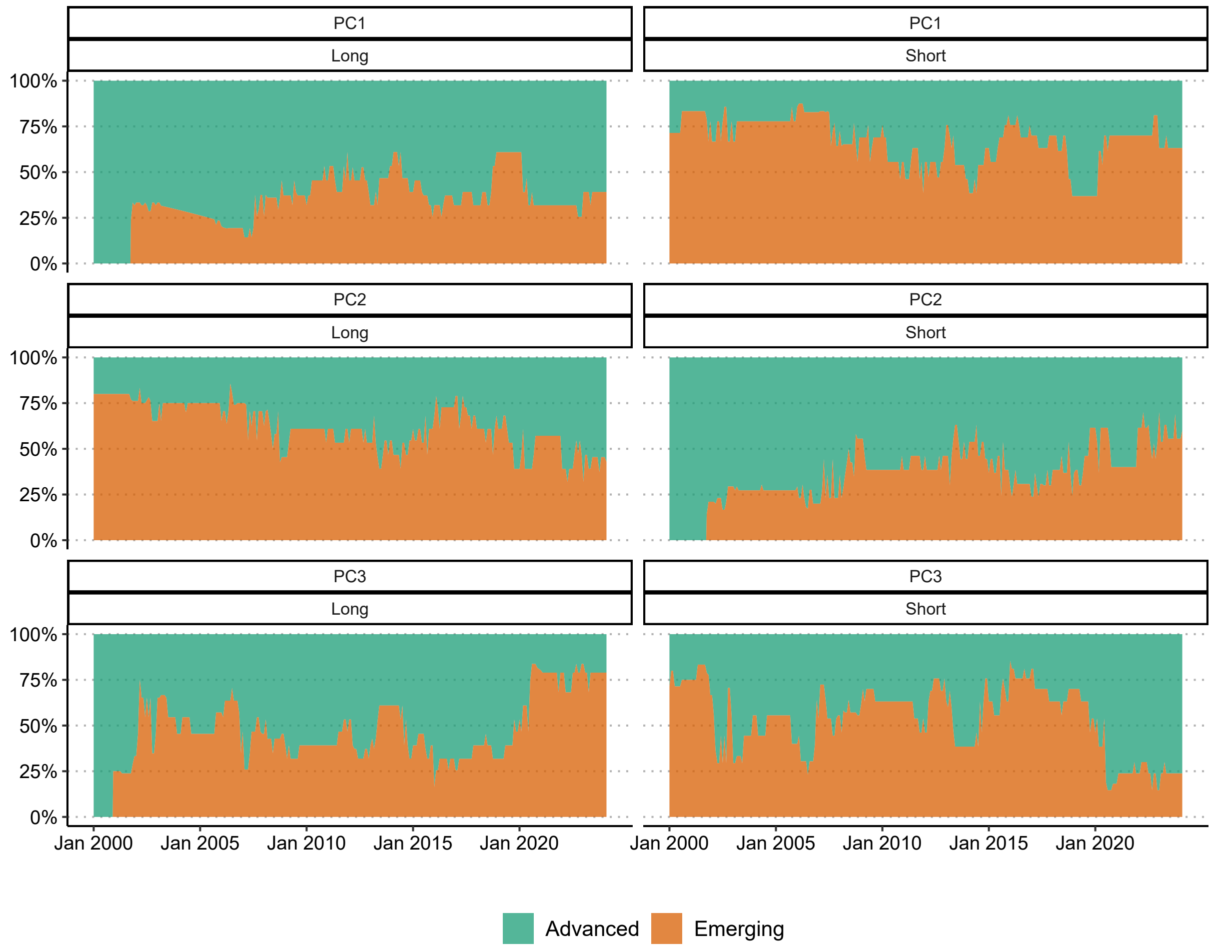

Figure E.8: Advanced and emerging market composition of option-implied VIX futures principal component factors. This figure shows the share of advanced and emerging market currencies included in the long and short sides of the high-minus-low (HML) currency risk factors constructed from exposures to innovations in the first three option-implied VIX term-structure principal components: level (PC1), slope (PC2), and curvature (PC3). Each panel plots, over time, the proportion of advanced (green) and emerging (orange) currencies composing the long (left) and short (right) portfolios of each factor. The proportions are computed monthly based on the portfolio assignments underlying the factor construction. The figure highlights systematic differences in how advanced and emerging currencies load onto volatility term-structure shocks: exposures to the level (PC1) are typically dominated by advanced currencies on the long side, while slope (PC2) and curvature (PC3) factors exhibit more balanced or alternating compositions, reflecting shifts in global volatility transmission across market groups.

Table E.1: Option-implied VIX futures beta-sorted portfolios summary statistics. This table reports summary statistics for monthly currency excess returns of portfolios sorted by estimated betas with respect to option-implied VIX futures innovations. Panel A sorts currencies on β_{PC1} (level), Panel B on β_{PC2} (slope), and Panel C on β_{PC3} (curvature). P1 contains currencies with the lowest beta exposures, while P5 contains those with the highest. HML denotes the high-minus-low portfolio (P5 − P1). The Sharpe Ratio is computed as the ratio of average annualized return to annualized volatility. *t-stat* refers to the t-statistic testing whether the average monthly excess return is significantly different from zero.

Portfolio	P1	P2	P3	P4	P5	HML
Panel A: sorts on β_{PC1}						
Avg Ann. Return	3.27	0.65	0.13	0.67	-2.08	-5.35
t-stat	[1.82]	[0.43]	[0.12]	[0.45]	[-1.16]	[-3.35]
Ann. Volatility	8.87	7.49	5.53	7.25	8.81	7.86
Sharpe Ratio	**0.37**	**0.09**	**0.02**	**0.09**	**-0.24**	**-0.68**
Skewness	-0.64	-0.76	-0.26	-0.31	-0.33	-0.00
Panel B: sorts on β_{PC2}						
Avg Ann. Return	-2.50	-0.57	0.87	1.38	2.58	5.08
t-stat	[-1.43]	[-0.36]	[0.70]	[1.06]	[1.46]	[3.55]
Ann. Volatility	8.60	7.81	6.12	6.37	8.69	7.03
Sharpe Ratio	**-0.29**	**-0.07**	**0.14**	**0.22**	**0.30**	**0.72**
Skewness	-0.50	-0.25	-0.83	-0.50	-0.79	0.65
Panel C: sorts on β_{PC3}						
Avg Ann. Return	1.31	-0.32	0.84	1.03	-0.58	-1.89
t-stat	[0.69]	[-0.21]	[0.78]	[0.74]	[-0.33]	[-1.25]
Ann. Volatility	9.35	7.49	5.30	6.78	8.73	7.46
Sharpe Ratio	**0.14**	**-0.04**	**0.16**	**0.15**	**-0.07**	**-0.25**
Skewness	-0.70	-0.58	-0.76	-0.25	-0.44	0.03

Table E.2: Time-series regression of option-implied VIX futures beta-sorted portfolios on the Carry factor. This table reports the results of time-series regressions of monthly excess returns of option-implied VIX futures beta-sorted currency portfolios on the Carry factor of Lustig et al., 2011. Panels A–C correspond to portfolios sorted on betas with respect to innovations to the first three principal components of VIX futures: level (PC1), slope (PC2), and curvature (PC3), respectively. P1 includes currencies with the lowest beta exposure and P5 includes those with the highest. HML denotes the high-minus-low portfolio (P5 – P1). Newey-West t-statistics are reported in brackets.

Portfolio	P1	P2	P3	P4	P5	HML
Panel A: sorts on β_{PC1}						
Intercept	-0.01	-0.02	-0.01	0.00	0.00	0.02
	[-0.62]	[-0.92]	[-0.90]	[0.16]	[0.12]	[1.15]
Carry Factor	0.42	0.26	0.13	0.03	-0.22	-0.64
	[3.74]	[2.42]	[2.03]	[0.40]	[-1.48]	[-10.49]
R^2	0.163	0.084	0.039	0.001	0.044	0.476
Panel B: sorts on β_{PC2}						
Intercept	-0.00	-0.01	-0.00	-0.00	-0.02	-0.01
	[-0.16]	[-0.62]	[-0.31]	[-0.00]	[-0.68]	[-1.17]
Carry Factor	-0.19	0.07	0.13	0.13	0.39	0.58
	[-1.02]	[0.84]	[1.78]	[1.73]	[3.17]	[8.54]
R^2	0.036	0.005	0.030	0.029	0.144	0.492
Panel C: sorts on β_{PC3}						
Intercept	-0.02	-0.02	-0.00	0.00	-0.00	0.01
	[-0.50]	[-0.92]	[-0.30]	[0.05]	[-0.12]	[0.60]
Carry Factor	0.26	0.14	0.12	0.09	-0.03	-0.29
	[1.91]	[1.68]	[1.86]	[1.12]	[-0.17]	[-2.40]
R^2	0.056	0.026	0.037	0.012	0.001	0.109

Table E.3: Asset pricing tests of option-implied VIX futures term-structure factors - second stage. This table reports the second-stage cross-sectional asset pricing results for the portfolio-level test-assets using the option-implied VIX futures term-structure currency risk factors. The dependent variables are the annualized average excess returns of the test assets, while the regressors are the betas estimated in the first-stage time-series regressions. Columns (1)–(5) correspond to a cross-section of Carry, Momentum, and Value portfolios, while columns (6)–(10) use only the Momentum and Value portfolios. Each specification includes the Dollar and Carry factors from the literature, along with different combinations of the three VIX-based currency risk factors: level ($fPC1$), slope ($fPC2$), and curvature ($fPC3$). The λ estimates represent the price of risk associated with each factor. I also report the p-values from the test of Gibbons et al., 1989, which tests the joint hypothesis that all pricing errors (alphas) are equal to zero—i.e., that the model prices all assets correctly. A failure to reject this null suggests that the pricing model performs well in explaining cross-sectional variation in returns. Standard errors are corrected following Shanken, 1992, and statistical significance is denoted by *, **, and *** for the 10%, 5%, and 1% levels, respectively.

	Test assets									
	Carry + Momentum + Value					Momentum + Value				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
β_{fPC1}		−0.050***			−0.048***		−0.050***			−0.049***
		(0.016)			(0.016)		(0.016)			(0.016)
β_{fPC2}			0.048***		0.049***			0.047***		0.048***
			(0.014)		(0.014)			(0.014)		(0.014)
β_{fPC3}				−0.017	−0.016				−0.016	−0.016
				(0.015)	(0.015)				(0.015)	(0.015)
β_{CAR}	0.107***	0.108***	0.108***	0.107***	0.108***	0.123**	0.109***	0.109***	0.108***	0.109***
	(0.020)	(0.018)	(0.018)	(0.018)	(0.018)	(0.052)	(0.017)	(0.017)	(0.017)	(0.017)
β_{DOL}	0.005	0.005	0.005	0.005	0.005	0.004	0.004	0.004	0.005	0.004
	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)
R^2	0.869	0.954	0.952	0.946	0.961	0.456	0.946	0.936	0.931	0.954
GRS	0.175	0.143	0.15	0.175	0.147	0.152	0.146	0.131	0.153	0.137

Table E.4: Asset pricing tests of option-implied VIX futures term-structure factors - second stage. This table reports the second-stage cross-sectional asset pricing results for the country-level test-assets using the option-implied VIX futures term-structure currency risk factors. The dependent variables are the annualized average excess returns of the test assets, while the regressors are the betas estimated in the first-stage time-series regressions. Columns (1)–(5) correspond to a cross-section of all country excess returns, while columns (6)–(10) use only advanced countries currency excess returns. Each specification includes the Dollar and Carry factors from the literature, along with different combinations of the three VIX-based currency risk factors: level ($fPC1$), slope ($fPC2$), and curvature ($fPC3$). The λ estimates represent the price of risk associated with each factor. I also report the p-values from the test of Gibbons et al., 1989, which tests the joint hypothesis that all pricing errors (alphas) are equal to zero—i.e., that the model prices all assets correctly. A failure to reject this null suggests that the pricing model performs well in explaining cross-sectional variation in returns. Standard errors are corrected following Shanken, 1992, and statistical significance is denoted by *, **, and *** for the 10%, 5%, and 1% levels, respectively.

	Test assets									
	Country level - all					Country level - advanced				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
β_{fPC1}		−0.056***			−0.060***		−0.049***			−0.053***
		(0.016)			(0.016)		(0.016)			(0.016)
β_{fPC2}			0.039***		0.039***			0.042***		0.040***
			(0.014)		(0.015)			(0.015)		(0.015)
β_{fPC3}				−0.018	−0.010				−0.013	−0.014
				(0.015)	(0.016)				(0.015)	(0.015)
β_{CAR}	0.058***	0.064***	0.063***	0.065***	0.064***	0.056**	0.090***	0.090***	0.090***	0.090***
	(0.018)	(0.018)	(0.018)	(0.018)	(0.018)	(0.026)	(0.019)	(0.019)	(0.019)	(0.019)
β_{DOL}	0.007	0.007	0.006	0.007	0.007	0.008	0.009	0.007	0.009	0.008
	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)
R^2	0.674	0.766	0.731	0.721	0.775	0.550	0.850	0.818	0.819	0.860
GRS	0.004	0.012	0.007	0.009	0.015	0.007	0.036	0.036	0.027	0.077

Table E.5: Predictive panel regressions of one-month-ahead currency excess returns (rx_{t+1}) on month-t innovations to the option-implied VIX futures term-structure common components: level (PC1), slope (PC2), and curvature (PC3). Columns (1)–(3) report pooled OLS estimates for all currencies; Columns (4)–(6) report corresponding within estimations with individual (currency) fixed effects. Newey–West standard errors (one lag) for pooled models and cluster-robust (by currency) standard errors for fixed-effects models are shown in parentheses. Constant terms are included in the pooled OLS specifications but their coefficients are omitted from the table. Adjusted R^2 values are reported in percent. Statistical significance is denoted by *, **, and *** for the 10%, 5%, and 1% levels, respectively.

	rx_{t+1}					
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta PC1$	-0.002		0.003	-0.002		0.003
	(0.002)		(0.002)	(0.002)		(0.003)
$\Delta PC2$	0.007**		0.008**	0.007***		0.008***
	(0.003)		(0.003)	(0.003)		(0.003)
$\Delta PC3$	-0.031***		-0.024**	-0.031***		-0.024***
	(0.010)		(0.010)	(0.008)		(0.009)
rx_t		0.150***	0.153***		0.144*	0.147*
		(0.045)	(0.049)		(0.081)	(0.088)
FE	No	No	No	Yes	Yes	Yes
R^2 (%)	0.30	2.17	2.34	-0.12	1.58	1.75

Appendix F. S&P500 variance swaps

The variance-forward data from Dew-Becker et al., 2017 are constructed from variance-swap quotes on the S&P 500. Each forward claim F_t^n represents the market's risk-neutral expectation of realized variance $E_t^Q[RV_{t+n}]$ at horizon n months ahead, obtained as the difference between successive variance-swap values:

$$F_t^n = VS_t^n - VS_t^{n-1},$$

where VS_t^n is the price of an n-month variance swap. A one-month forward claim (F_t^1) is exactly equivalent to a one-month variance swap. Because a variance forward is a claim that delivers realized variance in the future, an investor holding it is short volatility: the claim's value falls when expected future variance rises. This payoff orientation is the mirror image of a VIX future, whose price rises when expected variance increases.

Consequently, the principal components extracted from the variance-forward term-structure have the opposite economic sign to those from VIX-futures data. In the raw variance-forward loadings, the "level" factor (PC1) decreases in periods such as 2008 when expected variance jumps, whereas the VIX-based level factor increases. To maintain a consistent interpretation—so that positive innovations in PC1 and PC3 correspond to higher expected variance (bad times) and negative innovations in PC2 correspond to flattening or normalization of the volatility term-structure—I multiply the variance-forward principal-component loadings (and scores) by -1. This adjustment aligns the forward-claim factors with the economic meaning of the VIX-futures factors, ensuring that all subsequent betas and risk-price estimates reflect the same "volatility-up = bad-times" convention.

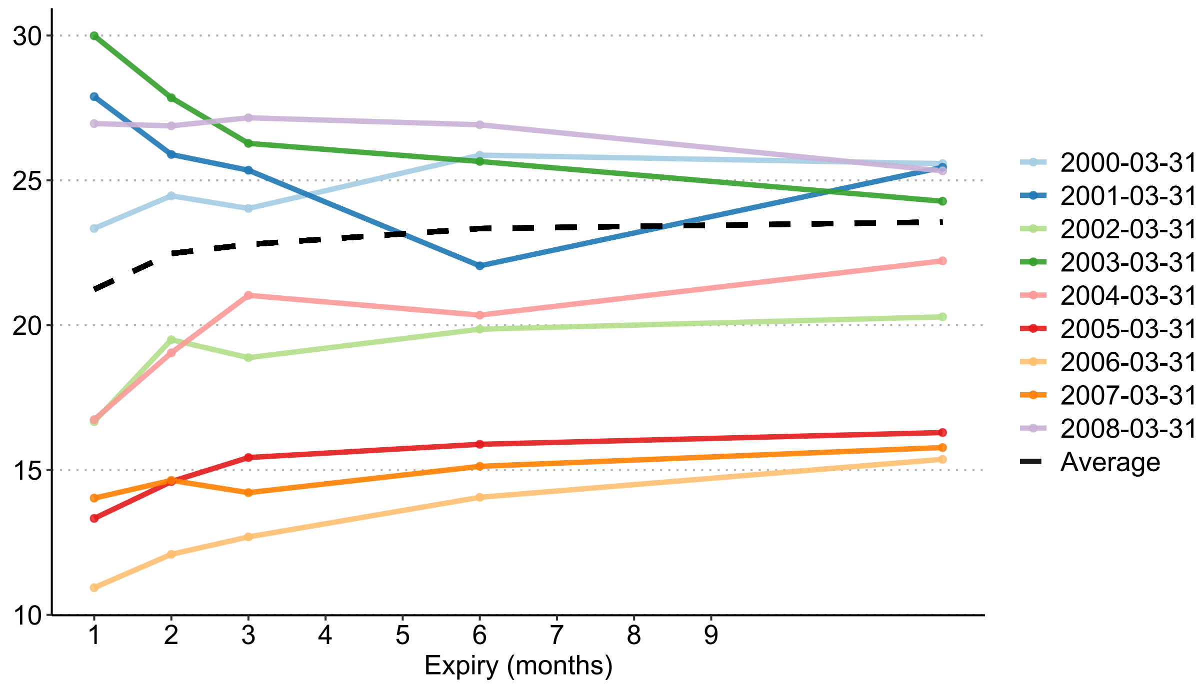

Figure F.1: Selected end-of-month variance forward claims term-structures

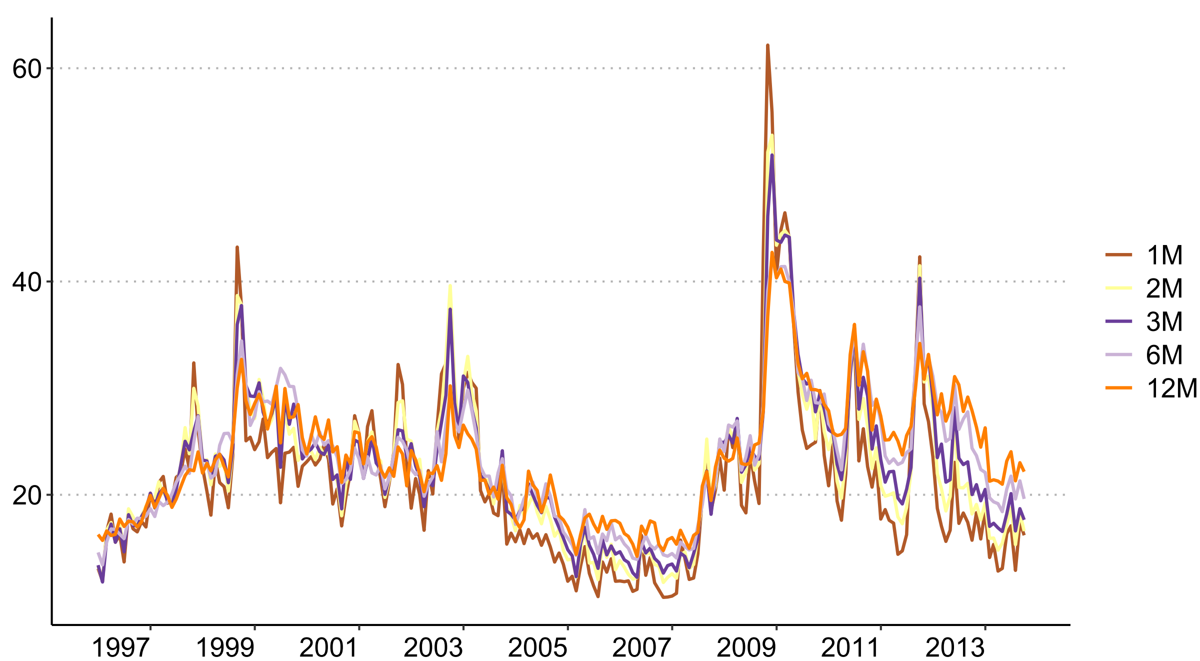

Figure F.2: Time-series variance forward claims term-structures, end-of-month

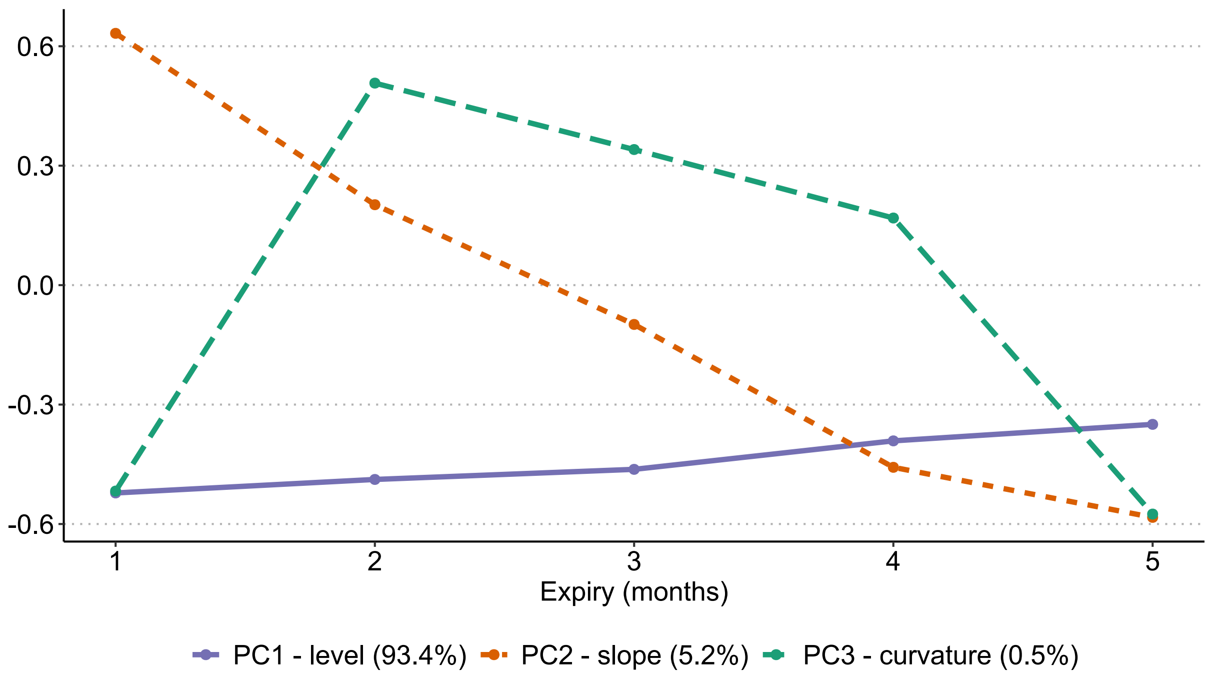

Figure F.3: Factor loadings of variance forward claims PCA. This figure shows the factor loadings of the full-sample principal component analysis on the level of S&P500 variance forward claims. Loadings are scaled such that they correspond to common measures of level, slope, and curvature. Percentage of variance explained by each principal component is between parenthesis.

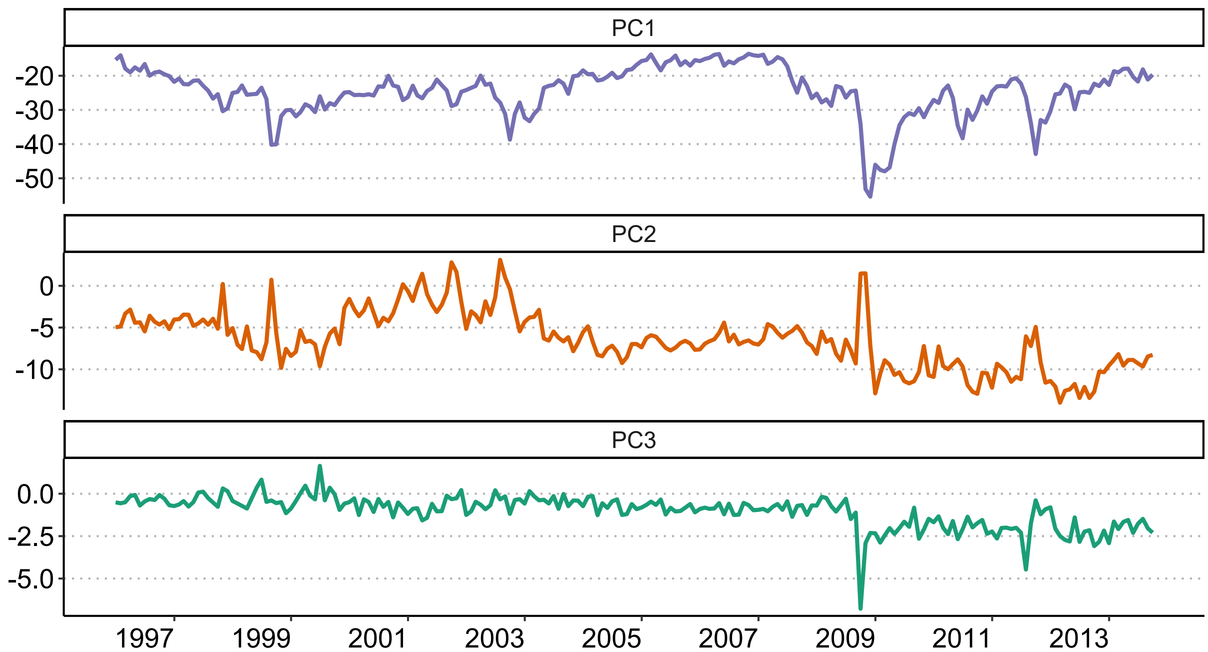

Figure F.4: Principal components of expanding window PCA. This figure shows the time-series of the first three principal components extracted from S&P 500 variance forward claims using an expanding window PCA, with a startup period of 50 months. The signs of the loadings are such that PC1 captures the level of the curve, PC2 captures the slope, and PC3 captures curvature. Increases in PC1 reflect upward shifts in the entire curve—typically associated with deteriorating market conditions—while increases in PC2 reflect steepening, often interpreted as improving conditions. PC3 captures changes in curvature, though its interpretation is less directly linked to macro-financial states.

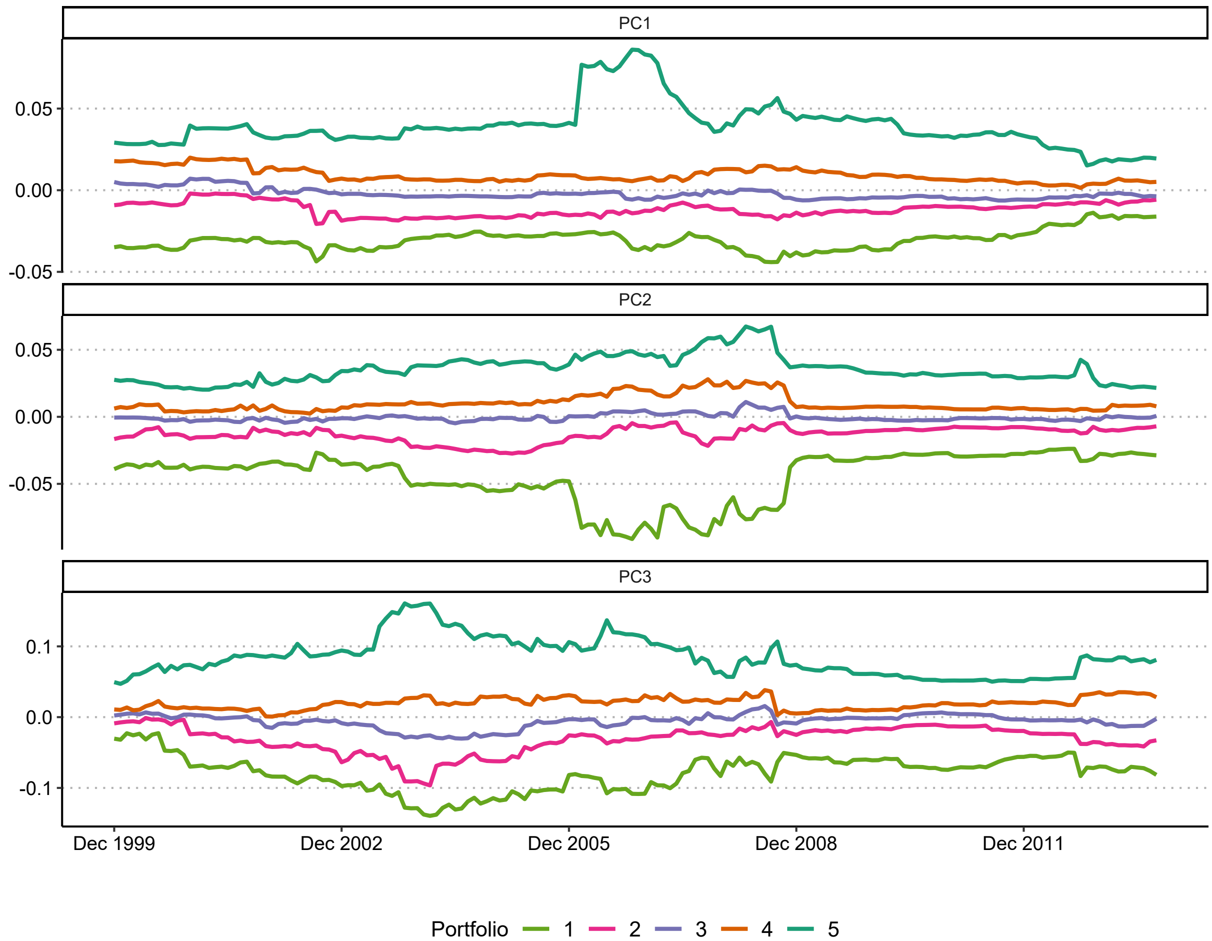

Figure F.5: Rolling betas to variance forward claims principal component innovations by currency portfolio. This figure shows the time-series of average betas for five currency portfolios sorted each month by their estimated exposure to innovations in variance forward claims term-structure components. The principal components—capturing the level (PC1), slope (PC2), and curvature (PC3) of the variance forward claims curve—are extracted monthly using an expanding-window PCA. Innovations in each component are computed as the monthly change relative to the previous observation. For each currency, rolling 48-month regressions of excess returns on these innovations are estimated following Equation 18, controlling for the dollar factor. The resulting betas measure time-varying sensitivities to shocks in the variance forward claims term-structure. The figure reports portfolio-average betas across the five beta-sorted groups, illustrating persistent cross-sectional heterogeneity and cyclical variation in exposures to volatility-level (PC1), slope (PC2), and curvature (PC3) innovations.

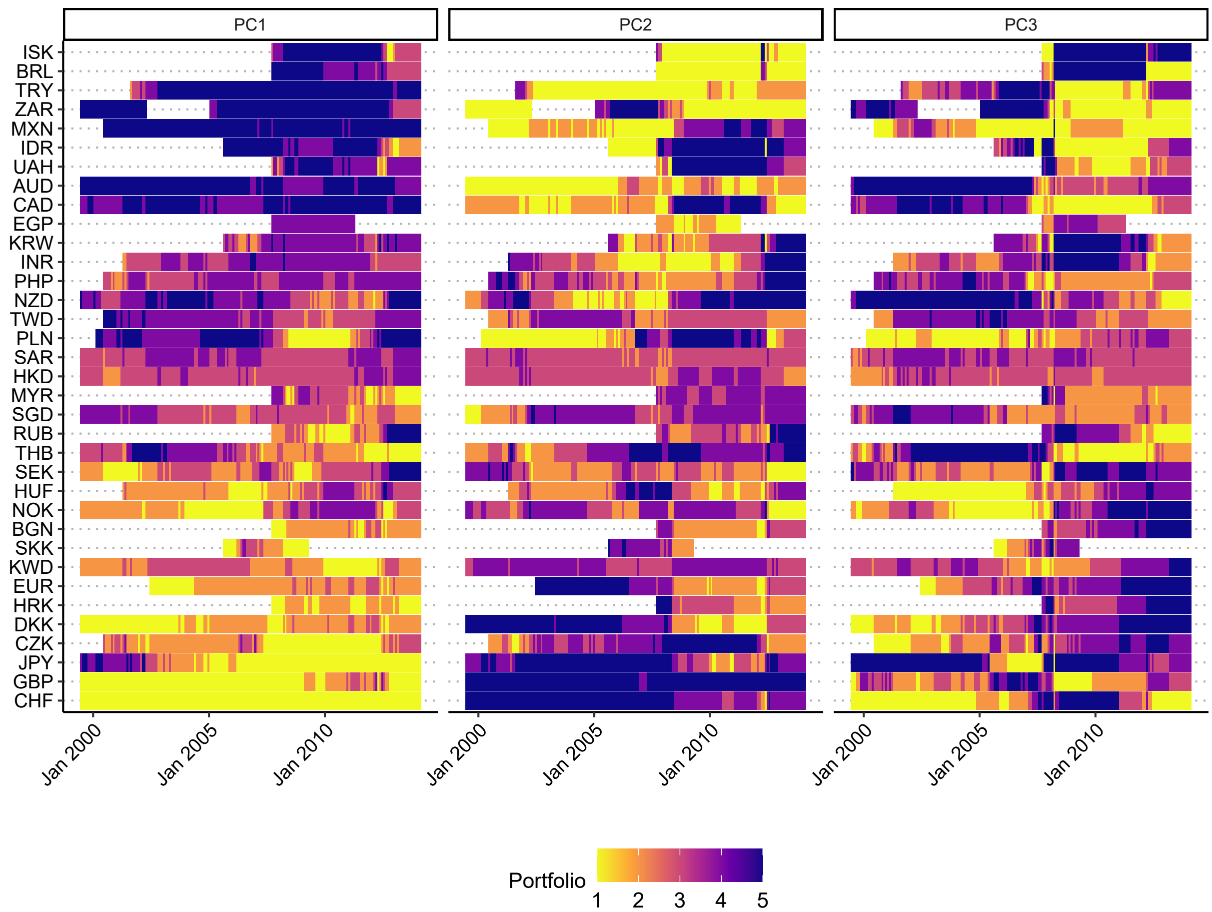

Figure F.6: Composition of currency portfolios sorted on variance forward claims principal component betas. This figure displays the time-varying portfolio assignments of individual currencies based on their rolling 48-month betas with respect to innovations in the first three VIX term-structure principal components: level (PC1), slope (PC2), and curvature (PC3). Each color represents the portfolio (from 1 to 5) to which a currency belongs in a given month, determined by its beta ranking from the previous month's estimation. Left corresponds to portfolios sorted on PC1 betas, middle panel on PC2, and right on PC3. Lighter shades indicate lower-exposure portfolios, while darker shades represent higher exposures. The figure highlights persistent clustering of currencies with similar sensitivities to volatility term-structure shocks, as well as transitions across portfolios reflecting evolving exposure patterns over time.

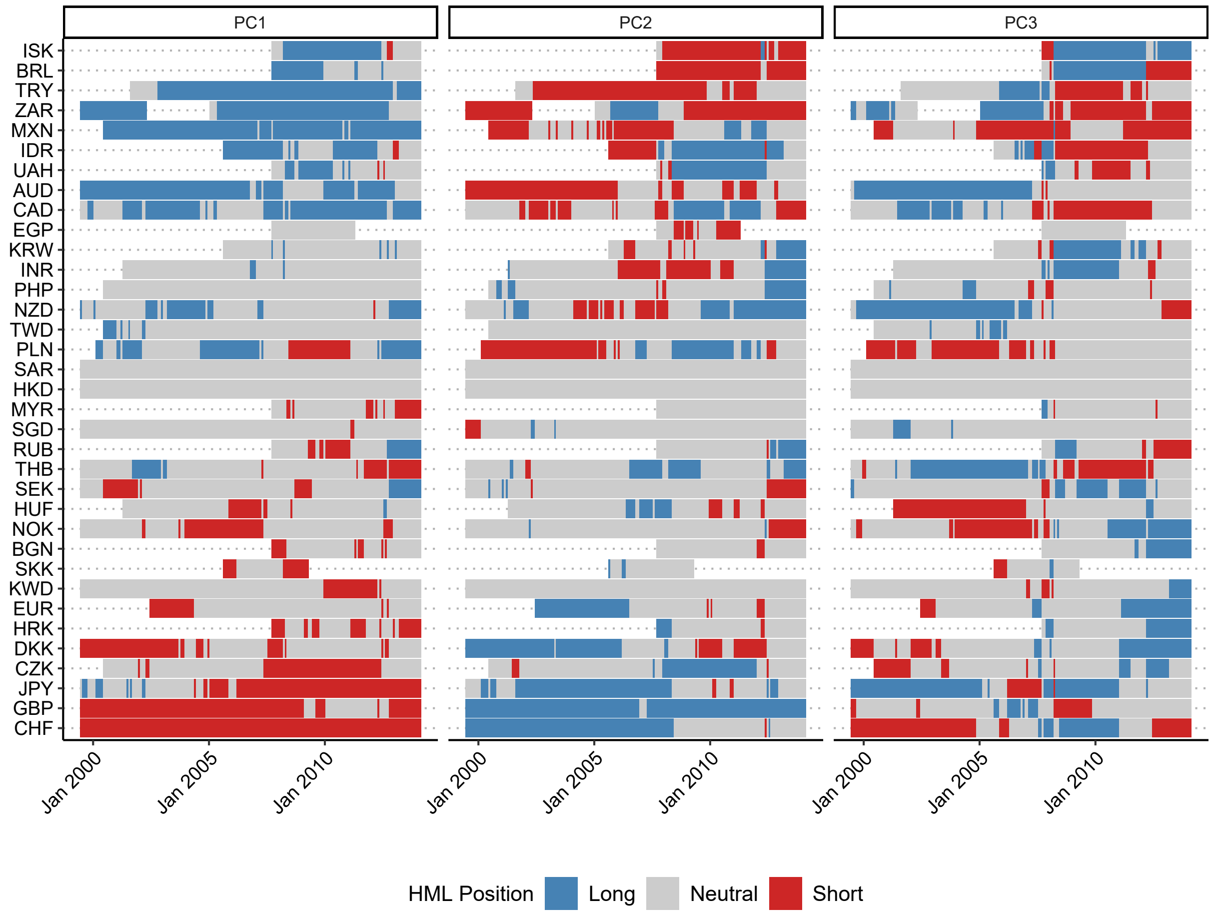

Figure F.7: Composition of currency risk factors constructed from variance forward claims principal component betas. This figure shows the time-varying composition of the high-minus-low (HML) currency risk factors associated with exposures to innovations in the first three variance forward claims term-structure principal components: level (PC1), slope (PC2), and curvature (PC3). For each month, currencies are sorted into five portfolios based on their rolling 48-month betas with respect to the respective principal component innovation, and the factor is constructed as the return difference between the highest- and lowest-beta portfolios. The heatmap indicates, for each currency and month, whether it enters the factor with a long position (blue), short position (red), or is excluded from the factor (gray). Each panel corresponds to one VIX component. The figure highlights persistent cross-sectional asymmetries in how currencies load onto the VIX futures term-structure level, slope, and curvature shocks, as well as shifts in factor composition reflecting changing sources of volatility risk exposure across time.

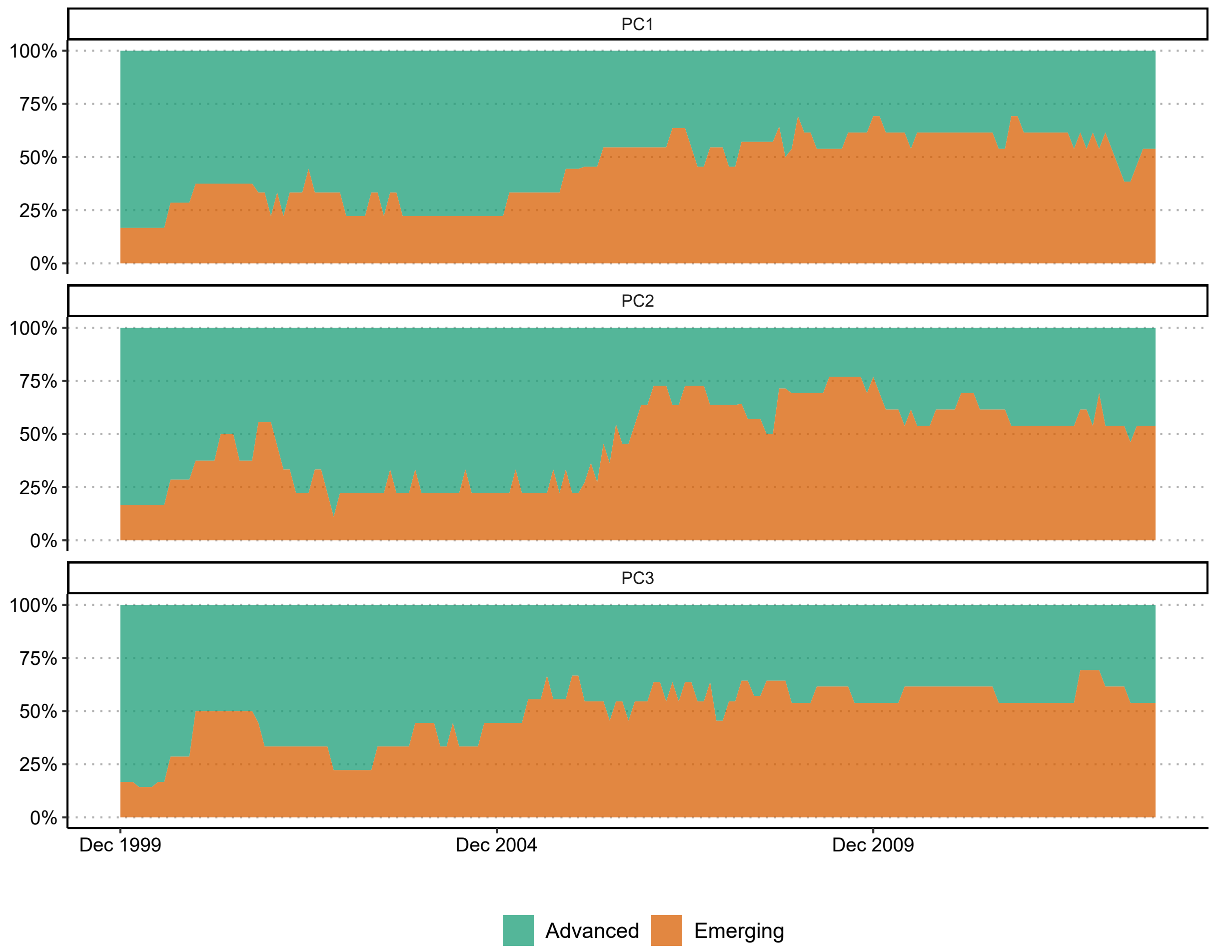

Figure F.8: Advanced and emerging market composition of variance forward claims principal component factors. This figure presents the share of advanced and emerging market currencies in the aggregate high-minus-low (HML) portfolios constructed from exposures to innovations in the first three variance forward claims term-structure principal components: level (PC1), slope (PC2), and curvature (PC3). Each panel shows, over time, the proportion of advanced (green) and emerging (orange) currencies included in the combined factor portfolio, computed monthly based on the underlying beta-sorted portfolio assignments. The composition reveals that there is no consistent dominance of either country group across factors: the relative shares of advanced and emerging currencies fluctuate over time and differ across components, suggesting that exposure to volatility term-structure shocks is not systematically concentrated in one market type but varies with global conditions.

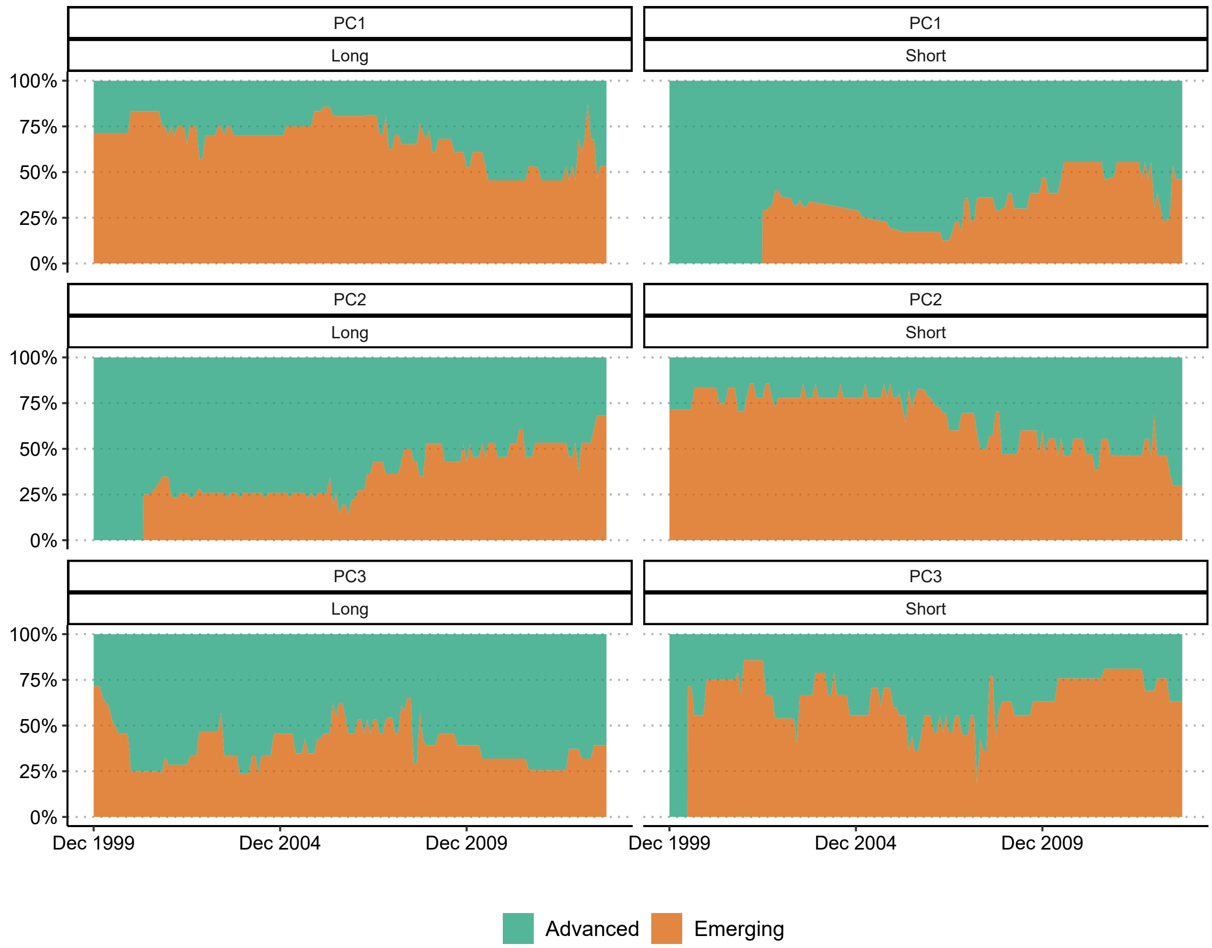

Figure F.9: Advanced and emerging market composition of variance forward claims principal component factors. This figure shows the share of advanced and emerging market currencies included in the long and short sides of the high-minus-low (HML) currency risk factors constructed from exposures to innovations in the first three variance forward claims term-structure principal components: level (PC1), slope (PC2), and curvature (PC3). Each panel plots, over time, the proportion of advanced (green) and emerging (orange) currencies composing the long (left) and short (right) portfolios of each factor. The proportions are computed monthly based on the portfolio assignments underlying the factor construction. The figure highlights systematic differences in how advanced and emerging currencies load onto volatility term-structure shocks: exposures to the level (PC1) are typically dominated by advanced currencies on the long side, while slope (PC2) and curvature (PC3) factors exhibit more balanced or alternating compositions, reflecting shifts in global volatility transmission across market groups.

Table F.1: Variance forward claims beta-sorted portfolios summary statistics. This table reports summary statistics for monthly currency excess returns of portfolios sorted by estimated betas with respect to innovations in S&P500 variance forward claims. Panel A sorts currencies on β_{PC1} (level), Panel B on β_{PC2} (slope), and Panel C on β_{PC3} (curvature). P1 contains currencies with the lowest beta exposures, while P5 contains those with the highest. HML denotes the high-minus-low portfolio (P5 – P1). The Sharpe Ratio is computed as the ratio of average annualized return to annualized volatility. *t-stat* refers to the t-statistic testing whether the average monthly excess return is significantly different from zero.

Portfolio	P1	P2	P3	P4	P5	HML
Panel A: sorts on β_{PC1}						
Avg Ann. Return	3.06	3.94	2.60	2.46	2.81	-0.25
t-stat	[1.23]	[2.18]	[1.41]	[1.04]	[1.21]	[-0.13]
Ann. Volatility	9.22	6.72	6.80	8.72	8.64	7.13
Sharpe Ratio	**0.33**	**0.59**	**0.38**	**0.28**	**0.33**	**-0.04**
Skewness	-1.39	-0.46	-0.46	-0.69	-0.76	0.55
Panel B: sorts on β_{PC2}						
Avg Ann. Return	1.06	2.30	1.86	3.81	5.21	4.14
t-stat	[0.57]	[1.35]	[0.86]	[1.50]	[2.05]	[2.08]
Ann. Volatility	6.86	6.32	8.01	9.39	9.40	7.40
Sharpe Ratio	**0.15**	**0.36**	**0.23**	**0.41**	**0.55**	**0.56**
Skewness	-1.51	-1.51	-0.64	-0.31	-0.74	-0.22
Panel C: sorts on β_{PC3}						
Avg Ann. Return	5.41	1.69	1.71	3.05	3.03	-2.38
t-stat	[2.28]	[0.78]	[0.85]	[1.53]	[1.24]	[-1.26]
Ann. Volatility	8.79	8.00	7.46	7.38	9.04	7.01
Sharpe Ratio	**0.62**	**0.21**	**0.23**	**0.41**	**0.34**	**-0.34**
Skewness	-0.60	-0.43	-1.04	-0.83	-0.56	0.04

Table F.2: Time-series regression of variance forwards beta-sorted portfolios on the Carry factor. This table reports the results of time-series regressions of monthly excess returns of S&P500 variance forward claims beta-sorted currency portfolios on the Carry factor of Lustig et al., 2011. Panels A–C correspond to portfolios sorted on betas with respect to innovations to the first three principal components of S&P500 variance forward claims: level (PC1), slope (PC2), and curvature (PC3), respectively. P1 includes currencies with the lowest beta exposure and P5 includes those with the highest. HML denotes the high-minus-low portfolio (P5 – P1). Newey-West t-statistics are reported in brackets.

Portfolio	P1	P2	P3	P4	P5	HML
Panel A: sorts on β_{PC1}						
Intercept	-0.02	0.01	0.01	0.00	0.01	0.03
	[-0.41]	[0.47]	[0.26]	[0.10]	[0.52]	[1.28]
Carry Factor	0.52	0.32	0.23	0.24	0.16	-0.37
	[2.57]	[2.19]	[2.06]	[1.40]	[1.12]	[-3.00]
R^2	0.179	0.122	0.064	0.043	0.018	0.147
Panel B: sorts on β_{PC2}						
Intercept	-0.01	-0.00	-0.01	0.01	0.02	0.03
	[-0.31]	[-0.07]	[-0.23]	[0.60]	[0.43]	[1.22]
Carry Factor	0.23	0.28	0.27	0.26	0.39	0.16
	[1.47]	[1.84]	[2.13]	[1.92]	[1.84]	[1.57]
R^2	0.059	0.111	0.065	0.043	0.095	0.027
Panel C: sorts on β_{PC3}						
Intercept	0.01	-0.01	-0.01	0.01	0.01	0.00
	[0.47]	[-0.57]	[-0.32]	[0.58]	[0.37]	[0.15]
Carry Factor	0.48	0.34	0.30	0.18	0.18	-0.30
	[3.99]	[2.80]	[1.76]	[1.39]	[0.86]	[-3.13]
R^2	0.165	0.101	0.090	0.034	0.022	0.103

Table F.3: Asset pricing tests of variance forward claims term-structure factors - second stage. This table reports the second-stage cross-sectional asset pricing results for the portfolio-level test-assets using the variance forwards term-structure. The dependent variables are the annualized average excess returns of the test assets, while the regressors are the betas estimated in the first-stage time-series regressions. Columns (1)–(5) correspond to a cross-section of Carry, Momentum, and Value portfolios, while columns (6)–(10) use only the Momentum and Value portfolios. Each specification includes the Dollar and Carry factors from the literature, along with different combinations of the three S&P500 variance forward claims-based currency risk factors: level ($fPC1$), slope ($fPC2$), and curvature ($fPC3$). The λ estimates represent the price of risk associated with each factor. I also report the p-values from the test of Gibbons et al., 1989, which tests the joint hypothesis that all pricing errors (alphas) are equal to zero—i.e., that the model prices all assets correctly. A failure to reject this null suggests that the pricing model performs well in explaining cross-sectional variation in returns. Standard errors are corrected following Shanken, 1992, and statistical significance is denoted by *, **, and *** for the 10%, 5%, and 1% levels, respectively.

	Test assets									
	Carry + Momentum + Value					Momentum + Value				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
β_{fPC1}		0.002			0.002		0.002			0.002
		(0.019)			(0.019)		(0.019)			(0.019)
β_{fPC2}			0.041**		0.040**			0.041**		0.041**
			(0.020)		(0.020)			(0.020)		(0.020)
β_{fPC3}				−0.024	−0.024				−0.025	−0.025
				(0.019)	(0.019)				(0.019)	(0.019)
β_{CAR}	0.079***	0.084***	0.084***	0.083***	0.084***	0.045	0.084***	0.084***	0.084***	0.084***
	(0.023)	(0.021)	(0.021)	(0.021)	(0.021)	(0.049)	(0.020)	(0.020)	(0.020)	(0.020)
β_{DOL}	0.033*	0.033*	0.033*	0.033*	0.033*	0.032*	0.033*	0.033*	0.033*	0.033*
	(0.019)	(0.019)	(0.019)	(0.019)	(0.019)	(0.019)	(0.019)	(0.019)	(0.019)	(0.019)
R^2	0.740	0.852	0.840	0.878	0.907	0.009	0.757	0.694	0.830	0.880
GRS	0.309	0.327	0.433	0.277	0.426	0.133	0.144	0.229	0.134	0.274

Table F.4: Asset pricing tests of variance forward claims term-structure factors - second stage. This table reports the second-stage cross-sectional asset pricing results for the country-level test-assets using variance forwards term-structure. The dependent variables are the annualized average excess returns of the test assets, while the regressors are the betas estimated in the first-stage time-series regressions. Columns (1)–(5) correspond to a cross-section of all country excess returns, while columns (6)–(10) use only advanced countries currency excess returns. Each specification includes the Dollar and Carry factors from the literature, along with different combinations of the three S&P500 variance forward claims-based currency risk factors: level ($fPC1$), slope ($fPC2$), and curvature ($fPC3$). The λ estimates represent the price of risk associated with each factor. I also report the p-values from the test of Gibbons et al., 1989, which tests the joint hypothesis that all pricing errors (alphas) are equal to zero—i.e., that the model prices all assets correctly. A failure to reject this null suggests that the pricing model performs well in explaining cross-sectional variation in returns. Standard errors are corrected following Shanken, 1992, and statistical significance is denoted by *, **, and *** for the 10%, 5%, and 1% levels, respectively.

	Test assets									
	Country level - all					Country level - advanced				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
β_{fPC1}		0.001			−0.001		0.005			0.001
		(0.021)			(0.021)		(0.020)			(0.020)
β_{fPC2}			0.057**		0.054**			0.035*		0.033
			(0.022)		(0.022)			(0.020)		(0.020)
β_{fPC3}				−0.015	−0.013				−0.020	−0.018
				(0.021)	(0.021)				(0.020)	(0.019)
β_{CAR}	0.048**	0.056***	0.055***	0.056***	0.057***	0.029	0.054**	0.051**	0.053**	0.054**
	(0.021)	(0.021)	(0.021)	(0.021)	(0.021)	(0.026)	(0.022)	(0.022)	(0.022)	(0.022)
β_{DOL}	0.034*	0.034*	0.034*	0.034*	0.034*	0.030	0.033*	0.032	0.033*	0.032
	(0.019)	(0.019)	(0.019)	(0.019)	(0.019)	(0.020)	(0.020)	(0.020)	(0.020)	(0.020)
R^2	0.433	0.461	0.515	0.485	0.564	0.664	0.633	0.572	0.668	0.670
GRS	0.026	0.029	0.039	0.028	0.041	0.011	0.012	0.016	0.013	0.021

Table F.5: Predictive panel regressions of one-month-ahead currency excess returns (rx_{t+1}) on month-t innovations to the variance forwards term-structure common components: level (PC1), slope (PC2), and curvature (PC3). Columns (1)–(3) report pooled OLS estimates for all currencies; Columns (4)–(6) report corresponding within estimations with individual (currency) fixed effects. Newey–West standard errors (one lag) for pooled models and cluster-robust (by currency) standard errors for fixed-effects models are shown in parentheses. Constant terms are included in the pooled OLS specifications but their coefficients are omitted from the table. Adjusted R^2 values are reported in percent. Statistical significance is denoted by *, **, and *** for the 10%, 5%, and 1% levels, respectively.

	rx_{t+1}					
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta PC1$	0.004*		0.001	0.003*		0.001
	(0.002)		(0.002)	(0.002)		(0.002)
$\Delta PC2$	−0.004		−0.002	−0.004		−0.002
	(0.003)		(0.003)	(0.003)		(0.002)
$\Delta PC3$	0.043***		0.046***	0.043***		0.046***
	(0.009)		(0.009)	(0.006)		(0.007)
rx_t		0.088***	0.092***		0.078***	0.081***
		(0.020)	(0.019)		(0.018)	(0.018)
FE	No	No	No	Yes	Yes	Yes
R^2 (%)	1.04	0.78	1.79	0.35	-0.11	0.92