Survey Expectations Meet Option Prices: New Insights from the FX Market*

Pasquale Della Corte
Imperial College London & CEPR

Can Gao
University of St. Gallen & SFI

Alexandre Jeanneret
University of New South Wales

August 28, 2026

*We are grateful to Svetlana Bryzgalova (Discussant), Mikhail Chernov, Xiang Fang, Niels Gormsen, Ralph Koijen, Lukas Kremens (discussant), Yang Liu, Yueran Ma, Ian Martin, Thomas Maurer, Philippe Mueller (discussant), Piotr Orlowski, Tarun Ramadorai, Valeri Sokolovski, Andreas Stathopoulos (discussant), Ine Van Robays (discussant), Irina Zviadadze (discussant), conference participants at the Vienna Symposium on Foreign Exchange Markets, Doctoral Consortium of the Asian Finance Association, QRFE Asset Pricing Workshop at Durham University, ICEEE Congress at the University of Cagliari, RAPS conference at CUHK, Swiss Finance Institute Research Day, Economics of Risk by BIS-ESM-EUI-CCA-HEC, 14th Workshop on Exchange Rates in Brussels, Canadian Derivative Institute Conference, SGF Conference, seminar participants at HEC Montréal, Imperial College London, Université Laval, University of Hong Kong, University of New South Wales, University of St. Gallen, and Fulcrum Asset Management for comments and suggestions. We gratefully acknowledge financial support from the Canadian Derivatives Institute (CDI) in Montréal. An earlier version of this paper circulated under the title "Expected Currency Returns and the Term Structure of Risk Preferences". All remaining errors are our own. E-mail contacts: `p.dellacorte@imperial.ac.uk`, `can.gao@unisg.ch`, and `a.jeanneret@unsw.edu.au`.

Survey Expectations Meet Option Prices: New Insights from the FX Market

Abstract

We develop an ex-ante measure of effective risk preferences in the FX market, linking consensus exchange-rate forecasts to option-implied currency risk premia through a no-arbitrage framework. The resulting parameter, which we refer to as RORO, for risk-off and risk-on, is the scaling that reconciles the two; higher values correspond to less willingness to bear risk. Estimated over nearly thirty years across a broad cross-section of currencies, RORO lies far above the log-utility benchmark recent work assumes. Four results follow. (i) Its term structure is steeply downward-sloping. (ii) That pattern reverses with market conditions, almost entirely at the short end, where the estimate falls to risk neutrality in calm markets. (iii) Level and slope both track the dispersion of risk and compensation across currencies more than global macro-financial conditions. (iv) No single value fits every currency: RORO is highest in high-yielding currencies. The effective price of currency risk is therefore horizon-dependent, state-dependent, and currency-specific.

Keywords: Exchange rate, risk premium, preferences, term structure, business cycle.
JEL Classification: F31, F37, F47, G12, G15.

1 Introduction

Expected currency returns combine two economically distinct objects: the risk of a currency payoff and the compensation investors require for bearing it. Realized returns cannot separate them. The distinction is central to interpreting violations of uncovered interest parity (UIP), cross-country differences in currency premia, and carry-trade returns.[1] Yet forward-looking measures of both expected returns and the price of currency risk have rarely been brought together across maturities and currencies. How much risk are currency investors willing to bear, and does the answer depend on when they must bear it?

In this paper, we address this gap by measuring effective risk preferences in the FX market. We use a straightforward no-arbitrage framework that draws on two sources of forward-looking information: (i) survey data from professional forecasters, which directly measure expected currency returns (e.g., Kremens, Martin, and Varela, 2025; Nagel and Xu, 2023; Piatti and Whelan, 2026); and (ii) FX derivatives, which capture risk premia and contribute to exchange rate predictability (e.g., Della Corte, Ramadorai, and Sarno, 2016; Kremens and Martin, 2019; Della Corte, Jeanneret, and Patelli, 2023). We exploit an identity that matches expected excess returns, derived from professional forecasts, with currency risk premia extracted from option prices. The two sides come from different data, and the parameter that reconciles them measures risk preferences.

We denote this parameter by ϕ and refer to it as RORO, for risk-off and risk-on: higher values correspond to less willingness to bear risk, and lower values to more. Because forecasts and options are both available at several horizons, we can measure RORO at each of them and trace its term structure. In traditional asset pricing models, risk preferences are either time-invariant (e.g., models with long-run risk or recursive preferences) or time-varying (e.g., models with habit preferences), but in both cases they are independent of the forecast horizon. A single stochastic discount factor further implies one value common to every currency.

[1] On UIP, see Hansen and Hodrick (1980) and Fama (1984); on the cross-section, Lustig, Roussanov, and Verdelhan (2011) and Verdelhan (2018); and on carry, Brunnermeier, Nagel, and Pedersen (2008), Burnside, Eichenbaum, Kleshchelski, and Rebelo (2011), and Menkhoff, Sarno, Schmeling, and Schrimpf (2012). Hassan and Zhang (2021) survey this literature.

Our framework is structurally parsimonious: it requires no specific utility function, no stationarity, and no restriction on the underlying stochastic process. It exploits an equivalent version of the no-arbitrage condition in asset pricing, the existence of the so-called *numéraire portfolio* (Long, 1990).[2] For a given currency i over a holding period $T-t$, the identity is

$$\underbrace{\mathbb{E}_t\left[\frac{S_{i,T}}{S_{i,t}}\right] - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}}}_{\text{Survey forecasts}} = \underbrace{\frac{1}{R^{\$}_{f,t}}\mathbb{C}\text{ov}^{\star}_t\left(X_T, \frac{S_{i,T}}{S_{i,t}}\right)}_{\text{Option implied}}, \tag{1}$$

where X_T is the gross return on the numéraire portfolio, inversely related to the stochastic discount factor (SDF), $S_{i,T}/S_{i,t}$ is the gross exchange rate return, $\mathbb{C}\text{ov}^{\star}_t$ is their risk-neutral covariance, and $R^{\$}_{f,t}/R^{i}_{f,t}$ captures the interest rate differential. The identity is exact. Its empirical content comes from what we assume about X_T: following Martin and Shi (2025), we specify it as a power function of the market return, $X_T = R^{\phi}_{m,T}$, with $R_{m,T}$ proxied by the S&P 500 index.[3] The currency risk premium is therefore a risk-neutral covariance with a nonlinear transformation of the market return. The exponent ϕ controls the curvature of that transformation, and hence the sensitivity of the pricing benchmark to market fluctuations.

We estimate ϕ in a daily panel of up to 33 currencies against the U.S. dollar over 1996–2025, matching options on the S&P 500 and on individual exchange rates to professional forecasts at horizons of one to 24 months. The breadth of the cross-section and the length of the sample ensure that the estimates reflect robust features of global FX markets rather than idiosyncratic episodes.

A natural concern is whether survey forecasts reliably measure expected exchange rates, since they may be biased or fail to reflect investors' true beliefs. Consensus forecasts measure the average expectation of professional forecasters, not the belief of the marginal investor who prices the option. Identification requires less than equality between the two: it requires that forecast errors

[2] The numéraire portfolio coincides with the growth-optimal portfolio of the earlier literature, considered in Kelly (1956), Roll (1973), Fama and MacBeth (1974), and Markowitz (1976), and more recently in Alvarez and Jermann (2005), Martin (2012), Martin and Wagner (2019), and Martin and Shi (2025) for the equity market.

[3] A third assumption enters here: we replace the risk-neutral correlation with its realized counterpart, which we test directly against traded quanto contracts.

are unrelated to the option-implied premium. Farmer, Nakamura, and Steinsson (2024) show that survey forecasts are consistent with rational updating under model uncertainty. Stavrakeva and Tang (2025) find that average exchange rate forecasts, Consensus Economics among them, closely mirror the average positions of major FX market participants. Kremens, Martin, and Varela (2025) provide direct evidence that survey forecasts help predict exchange rates. The evidence supports survey forecasts as a measure of expected exchange rates.

In the general framework, ϕ is an effective curvature parameter, and its magnitude carries no structural reading. One special case supplies one. Suppose a representative investor with power utility is the marginal trader in currency markets, faces no constraints, and holds the market portfolio as their entire wealth. Then ϕ is exactly the coefficient of relative risk aversion (RRA). The estimation requires none of this.[4] Read through that lens, our unconditional estimate of 7.17 is an RRA far above, and statistically distinct from, the log-utility benchmark often adopted in the recent FX predictability literature (e.g., Kremens and Martin, 2019; Della Corte, Jeanneret, and Patelli, 2023).

The level of RORO also governs how heavily a currency's exposure to higher-order market risk is priced; Backus, Chernov, and Martin (2011) show that sensitivity to higher-order cumulants rises with risk aversion once one departs from the lognormal benchmark. When $\phi = 2$, for instance, investors require compensation for comovement with market variance as well as with market returns. Our framework thus generalizes Kremens and Martin (2019), Della Corte, Jeanneret, and Patelli (2023), and Kremens, Martin, and Varela (2025), which are the special case $\phi = 1$. Our estimates therefore point to a large price on higher-order equity risk in currency markets, consistent with the evidence in Fan, Londono, and Xiao (2022).

Most important, risk preferences vary with the horizon over which risk is borne. Estimated separately at each maturity, ϕ declines from 8.01 at one month to 1.99 at two years. To our knowledge,

[4]Under partial market participation, the curvature reflects both preferences and the share of wealth exposed to market risk, so ϕ and the participation share are not separately identified. Alternatively, ϕ can be endogenously time varying in models with heterogeneous agents (Chan and Kogan, 2002), thereby generating business-cycle variation in ϕ. Under recursive preferences, ϕ also captures the preference for early resolution of uncertainty. Since our methodology nests these different cases, it can be applied in a wide range of contexts and asset pricing applications.

this is the first such term structure estimated from market prices. The pattern matches evidence that the price of higher-order risk is concentrated at short horizons: Dew-Becker, Giglio, Le, and Rodriguez (2017) show that the spot variance risk premium in the S&P 500 index is large while forward premia vanish beyond one or two months, Della Corte, Kozhan, and Neuberger (2021) report the same in FX markets, and Lustig, Stathopoulos, and Verdelhan (2019) and Zviadadze (2017) show that carry-trade compensation declines with horizon.[5] Crash risk offers a natural reading: over long horizons a currency has time to recover from a severe depreciation, so the possibility is priced less heavily at the long end. Horizon-specific measurement error, in the forecasts or in the correlation proxy, would produce the same unconditional decline. It would not produce a slope that changes sign with the state of the market.

The slope does change sign. We split the sample on prevailing financial conditions and re-estimate at each horizon,[6] using the CBOE equity-option implied volatility index (VIX) and, alternatively, the one-month realized volatility of the S&P 500 index. At the one-month horizon ϕ rises to 8.07 in adverse times but collapses to 0.15 in favorable ones, a value indistinguishable from risk neutrality, while the two-year estimate moves only from 1.52 to 1.64. The term structure is therefore steeply downward-sloping in stressed markets and upward-sloping in calm ones.

Market conditions move the price of imminent risk far more than the price of distant risk, and the reversal in the slope comes almost entirely from the short end. The pattern matters for anyone pricing currency exposure at more than one horizon: a single price of risk cannot fit both ends of the term structure, and because the short end moves with market conditions while the long end does not, the gap between them is widest exactly when markets are stressed. Because the estimates rest on forecasts rather than realized returns, which are poor proxies for currency risk premia at short horizons, their precision is greatest exactly where the action is.

What moves RORO over time? Every indicator we examine points the same way: RORO rises

[5] This pattern is consistent with the downward-sloping term structure of risk premia in equity markets, documented in Binsbergen, Brandt, and Koijen (2012), Weber (2018), Gonçalves (2021), and Gormsen (2021), among others.

[6] Because we impose no utility function, our approach addresses the critique of Bekaert, Engstrom, and Xu (2022), who argue that many measures of time-varying risk aversion are derived from models that assume a constant risk aversion coefficient and are therefore internally inconsistent.

as conditions deteriorate, whether conditions are read from implied volatility, credit spreads, intermediary capital, or the real economy. What differs is how much each accounts for. Conditions observed inside currency markets explain more than twice as much of the monthly variation as global macro-financial series, and the dispersion of risk and compensation across currencies accounts for more of that than the dollar does. Volatility lines up with the term structure and carry with the level, so the two have distinct sources.

Does one value fit every currency? It does not. The estimates range from -1.4 to 9.8 across currencies, and the five negative ones all belong to currencies with managed or pegged exchange rates. What orders the spread is how much a currency pays: ϕ is lowest in the currencies that yield least and highest in those that yield most, about 0.9 for the Swiss franc against 8.1 for the Brazilian real. High-yield currencies are the ones that depreciate in bad states of the world (e.g., Brunnermeier, Nagel, and Pedersen, 2008; Menkhoff, Sarno, Schmeling, and Schrimpf, 2012), so compensation per unit of risk is highest exactly where a sudden loss is most likely.

Our paper makes three contributions. First, we measure effective risk preferences without imposing a utility function. The measure prices both linear and higher-order exposures to market risk, and we estimate it by regression from ex-ante data, where existing methods back preferences out of structural models (e.g., Bekaert, Engstrom, and Xu, 2022; Orłowski, Sokolovski, and Sverdrup, 2025). Kremens and Martin (2019) document that realized excess returns move more than one-for-one with the Quanto Risk Premium, which points to preferences beyond log utility.[7] We formalize that observation, deriving the exact mapping between RORO and higher-order risk exposures. The expected-return side comes from professional forecasts, where the literature has substituted realized returns; the risk-premium side comes from option prices, with a realized-correlation proxy for the unobservable risk-neutral covariance.

Second, we estimate the term structure that existing models assume. Individuals do show greater

[7] Kremens and Martin (2019) Result 1 holds for *any* arbitrary return R and *any* stochastic discount factor: it does not require log utility. Log utility is a sufficient condition for the residual term in their decomposition to be zero, not a maintained assumption of their paper. The choice $X_T = R^{\phi}_{m,T}$ in the present paper is one specific implementation of their identity that eliminates the residual by construction, enabling direct OLS estimation of ϕ.

aversion to near-term risk than to distant risk, as field and laboratory evidence documents (e.g., Holt and Laury, 2002; Coble and Lusk, 2010; Abdellaoui, Diecidue, and Öncüler, 2011), and Eisenbach and Schmalz (2016) and Andries, Eisenbach, and Schmalz (2024) show that horizon-dependent risk aversion helps resolve asset pricing puzzles. Neither literature observes it in market prices. The quanto risk premium studied by Kremens, Martin, and Varela (2025) is observable only monthly, for a few currencies, and at the two-year horizon, which leaves every other maturity unobserved.

Third, we recover variation in preferences that a large literature infers indirectly. Guiso, Sapienza, and Zingales (2018) find that investors' risk aversion rises after the 2008 crisis, Cohn, Engelmann, Fehr, and Maréchal (2015) that fear rises with financial risk in the laboratory, Baker and Wurgler (2006) that sentiment varies over time, Pflueger, Siriwardane, and Sunderam (2020) that perceived risk moves with the business cycle, and Bekaert, Engstrom, and Xu (2022) and Orłowski, Sokolovski, and Sverdrup (2025) that aggregate risk aversion and the SDF can be recovered from macro data and asset prices. Consistent with these studies, RORO rises in bad times. Our estimates, by contrast, rest only on information available at the time of the forecast, which is what allows the horizon dimension to be measured at all.

The remainder of the paper is organized as follows. Section 2 presents our conceptual framework. Section 3 describes the data, the construction of the main variables, and the estimation approach. Section 4 reports the results. Section 5 asks which financial conditions move RORO, and Section 6 turns to the cross-section. We conclude in Section 7. The Internet Appendix contains technical details and additional results.

2 Our Framework

This section develops our framework for linking currency risk premia to investor preferences through the numéraire portfolio. Starting from a standard asset-pricing relation, we derive a tractable representation of expected currency returns that connects an unobservable pricing benchmark to observable market returns. We then show that this framework admits several economic interpretations

and clarifies how higher-order market risk enters currency pricing.

2.1 Asset Pricing Setup

Consider a currency strategy that converts a dollar into foreign currency at time t, lends at the foreign riskless rate between times t and T, and then exchanges the proceeds denominated in foreign currency for dollars at time T. Using the euro (€) as the foreign currency to ease the notation, this strategy's gross return is equivalent to

$$R_T = \frac{S_T}{S_t} R^{€}_{f,t}, \tag{2}$$

where S_t is the spot exchange rate at time t in units of dollars per unit of euro such that an increase reflects an appreciation of the euro, and $R^{€}_{f,t}$ denotes the euro gross riskless rate observed at time t with maturity $T-t$.

The standard asset pricing equation applied to Equation (2) implies the following expression of currency excess return

$$\mathbb{E}_t\left[\frac{S_T}{S_t}\right] - \frac{R^{\$}_{f,t}}{R^{€}_{f,t}} = \mathbb{C}\text{ov}_t\left(\frac{-M_T}{\mathbb{E}_t[M_T]}, \frac{S_T}{S_t}\right), \tag{3}$$

where $\mathbb{E}_t$ and $\mathbb{C}\text{ov}_t$ are the real-world expectation and covariance operators conditional on all information available at time t, respectively, $R^{\$}_{f,t}$ denotes the dollar gross riskless rate observed at time t with maturity $T-t$, and M_T refers to a stochastic discount factor (SDF) that prices future dollar payoffs, with conditional expectation equal to the inverse of the dollar gross riskless rate.

The identity in Equation (3) states that investors demand a time-varying risk premium determined by the conditional covariance between the SDF and the gross exchange rate return. As shown by the recent literature (e.g., Fama, 1984; Lustig, Roussanov, and Verdelhan, 2011), accounting for this risk adjustment is central to understanding the dynamics of currency excess returns. In theory, this covariance should be informative about risk preferences because it determines the compensation an investor requires for holding currencies whose payoffs are exposed to adverse states of the world.

In practice, however, it is difficult to measure because it is unobservable ex ante and may change both over time and across investment horizons. These features make it challenging to isolate precisely how preferences shape this covariance. To address this challenge, we adopt a structurally parsimonious approach based on the numéraire portfolio.

2.2 Numéraire Portfolio Representation

Our framework is anchored in the existence of the numéraire portfolio formalized by Long (1990). It is a strictly positive, self-financing benchmark portfolio that values all traded assets in the absence of arbitrage opportunities. When the gross return on a traded asset is measured relative to the gross return on the numéraire portfolio, its conditional expectation is exactly one. Combining the property of the numéraire portfolio with the standard asset pricing condition implies that the stochastic discount factor is simply the inverse of the gross return on the numéraire portfolio. Formally, the defining property of the numéraire portfolio gives $\mathbb{E}_t[R_T X_T^{-1}] = 1$, while standard asset pricing yields $\mathbb{E}_t[R_T M_T] = 1$. Combining these relations yields the exact identity

$$M_T X_T = 1, \tag{4}$$

where X_T denotes the gross return on the numéraire portfolio. As shown by Long (1990), the numéraire portfolio also coincides with the growth-optimal portfolio, namely the portfolio that maximizes the geometric mean of portfolio returns. Often associated with the Kelly criterion, it can also be viewed as the portfolio that maximizes the probability of outperforming any other portfolio over a sufficiently long investment horizon (e.g., Kelly, 1956; Latané, 1959). We use the numéraire notation X_T throughout, as in Equation (1).

The numéraire portfolio is a theoretical benchmark. To bring this framework to the data, we assume that its gross return can be expressed as a power function of the gross return on the market portfolio. Specifically, we consider the following specification

$$X_T = R_{m,T}^{\phi},$$

where $R_{m,T}$ is the gross return on the market portfolio and ϕ is a curvature parameter that governs the sensitivity of the numéraire portfolio to aggregate market fluctuations. A useful intuition for this specification comes from Long (1990), who discusses the numéraire portfolio as being more exposed to aggregate market risk than the market portfolio itself. In this spirit, $R^{\phi}_{m,T}$ provides a parsimonious reduced-form way to allow the numéraire benchmark to scale market exposure non-linearly. The parameter ϕ plays a dual role. First, it governs the sensitivity of the numéraire portfolio payoff to aggregate market fluctuations. Second, it determines the curvature of the associated pricing kernel. This specification provides a parsimonious yet flexible one-parameter family that allows the pricing benchmark to respond nonlinearly to market movements, while nesting the market portfolio as a special case. Values of $\phi > 1$ imply a benchmark that is more sensitive to market movements than the market portfolio itself, whereas values of $0 < \phi < 1$ imply a more attenuated response. Throughout, we refer to ϕ as a measure of risk preferences rather than as risk aversion, a label we reserve for the power-utility case below in which the two coincide. A higher ϕ means less willingness to bear risk. A negative ϕ corresponds to a negative effective price of risk, which need not imply a literal preference for risk. More broadly, this formulation offers a tractable and economically interpretable way to connect an unobservable benchmark portfolio to an observable market return, while remaining agnostic about the precise structural origins of ϕ. At the same time, the specification admits natural interpretations in several prominent asset-pricing environments, which we discuss in Section 2.3.

Building on this representation, we derive an alternative version of Equation (3) that makes explicit how the curvature parameter ϕ shapes currency excess returns. This result is stated in the following proposition.

Proposition 1. *Starting from the existence of the numéraire portfolio, so that $M_T X_T = 1$, and imposing the assumption $X_T = R^{\phi}_{m,T}$, we can rewrite Equation (3) as follows*

$$\mathbb{E}_t\left[\frac{S_T}{S_t}\right] - \frac{R^{\$}_{f,t}}{R^{€}_{f,t}} = \underbrace{\mathbb{C}ov^{\star}_t\left(\frac{R^{\phi}_{m,T}}{\mathbb{E}^{\star}_t[R^{\phi}_{m,T}]}, \frac{S_T}{S_t}\right)}_{ERP^{\phi}_t}. \tag{5}$$

Here, $\mathbb{E}_t^\star$ and $\mathbb{C}ov_t^\star$ denote the risk-neutral conditional expectation and covariance operators, respectively, and $\mathbb{E}_t^\star[R_{m,T}^\phi] = R_{f,t}^\$$.

Proof. See Internet Appendix A.1. □

We hereafter refer to the above risk-neutral covariance term as the currency's expected risk premium or simply ERP_t^ϕ. To help interpret the expected risk premium, Proposition 2 presents an explicit expression as a function of investors' risk preferences ϕ, based on a second-order Taylor expansion.

Proposition 2. *The expected risk premium, for a given level of ϕ, can be expressed as follows:*

$$ERP_t^\phi = \phi ERP_t + \frac{\phi(\phi-1)}{2R_{f,t}^\$}\mathbb{C}ov_t^\star\left(\xi_T^{\phi-2}(R_{m,T}-1)^2, \frac{S_T}{S_t}\right), \tag{6}$$

where

$$ERP_t = \frac{1}{R_{f,t}^\$}\mathbb{C}ov_t^\star\left(R_{m,T}, \frac{S_T}{S_t}\right) \tag{7}$$

is the expected risk premium when $\phi = 1$.

Proof. See Internet Appendix A.2. □

Proposition 2 shows that ERP_t^ϕ is equivalent to ϕ times ERP_t, where ERP_t reflects the linear exposure of a currency to the market return, plus a quadratic term that captures the exposure to a nonlinear change in the market return. Specifically, this second term reflects the curvature of the SDF, which increases with ϕ.

2.3 Economic Interpretations of the Curvature Parameter ϕ

Although our specification is intentionally parsimonious, the parameter ϕ has a rich economic content. Rather than imposing a single microfoundation, we show that the same reduced-form representation can be rationalized in a range of models commonly used in asset pricing. This

mapping helps clarify how the curvature parameter should be interpreted and why it provides a useful empirical summary of investor preferences.

Power Utility. Our framework admits a direct interpretation in an economy with a representative investor endowed with power utility, who acts as the marginal trader in FX markets and allocates their wealth entirely to the equity market. In this setting, the investor's marginal rate of substitution is governed by standard constant-relative-risk-aversion preferences, so that the pricing kernel inherits a simple power form in aggregate wealth. As shown in Internet Appendix B.1, solving the agent's static portfolio choice problem implies that the numéraire portfolio payoff satisfies $X_T \propto R_{m,T}^{\phi}$, where ϕ acts as the coefficient of relative risk aversion. Intuitively, higher values of ϕ imply that the benchmark payoff becomes more sensitive to market outcomes, reflecting a stronger state dependence of marginal utility. This mapping provides a simple microfoundation for our empirical specification: under power utility, the curvature parameter ϕ has a clear structural interpretation as a preference parameter, linking the shape of the numéraire portfolio directly to the degree of aggregate risk aversion.

Ambiguity Aversion. A similar result obtains under ambiguity aversion (Hansen, 2007). Consider an agent with logarithmic utility who fears model misspecification, penalizing belief deviations via a robustness penalty $\theta > 0$, with a lower θ denoting stronger ambiguity aversion. As shown in Internet Appendix B.2, this optimal belief distortion reduces the robust agent's problem to that of a power utility agent with an effective risk aversion of $\phi = 1 + \theta^{-1}$. This concern for robustness mathematically guarantees $\phi > 1$, even without baseline curvature in the agent's utility. The specification $X_T = R_{m,T}^{\phi}$ is therefore a unified reduced form. It captures both traditional risk aversion and aversion to model uncertainty in a single parameter, allowing ϕ to be taken to the data without requiring a stand on its precise structural origins.

Partial Market Participation. Our framework extends naturally to environments in which the representative agent invests only a fraction $\omega \in (0, 1]$ of their wealth in the market portfolio, with the remainder allocated to the risk-free asset. In this case, the effective curvature of the numéraire portfolio is governed jointly by ϕ and ω rather than by ϕ alone, as shown in Internet Appendix B.3.

Economically, partial market participation dampens the representative agent's effective exposure to aggregate risk, so that the elasticity of the numéraire with respect to the market return reflects both risk preferences and the share of wealth intermediated through risky assets. This specification nests intermediary asset-pricing models, in which the agents that price risk are specialized investors facing leverage or balance-sheet constraints. In such environments, ϕ can be interpreted more broadly as summarizing both the risk aversion of marginal investors and the constraints that shape their effective market exposure.

Heterogeneous Agents. The parameter ϕ need not be constant and may instead vary endogenously over time. In models with heterogeneous agents (Chan and Kogan, 2002; Longstaff and Wang, 2012), aggregate risk aversion is given by the wealth-weighted average of individual agents' risk aversion coefficients. As wealth shifts over the business cycle toward agents with different preferences, the implied ϕ for the representative agent also changes endogenously. As formally shown in Internet Appendix B.4, this mechanism carries over directly to our theoretical benchmark, implying a numéraire portfolio of the form $X_T \propto R_{m,T}^{\phi(W_T)}$, where the exponent varies counter-cyclically. This provides a rigorous microfoundation for a time-varying ϕ and allows our generalized specification to capture macroeconomic shifts in risk tolerance, which we subsequently examine in the data.

Recursive Preferences. Finally, our framework also accommodates Epstein-Zin recursive preferences, which break the tight link between risk aversion and the elasticity of intertemporal substitution imposed by standard power utility. In this setting, when investors value the early resolution of uncertainty, the aggregate market return enters the pricing kernel directly and independently of current consumption growth. As a result, ϕ is no longer interpretable as a simple coefficient of standard risk aversion. Rather, it captures the effective curvature of the pricing kernel, reflecting both attitudes toward risk and preferences over the timing of uncertainty resolution, the central mechanism in long-run risk models (e.g., Bansal and Yaron, 2004).

These interpretations show that our specification, $X_T = R_{m,T}^{\phi}$, is best viewed as a flexible reduced-form benchmark that nests a broad range of economic mechanisms within a common empirical representation. Depending on the underlying environment, ϕ may reflect standard risk aversion,

ambiguity aversion, limited market participation, endogenous variation in aggregate risk tolerance, or preferences over the timing of uncertainty resolution. This unifying formulation allows us to connect the numéraire portfolio to observable market returns in a tractable and economically interpretable way, without imposing a restrictive ex ante commitment to any single structural model.

2.4 Implications for the Expected Risk Premium

Equation (6) also shows that a higher ϕ increases the expected risk premium, but not necessarily in a linear way, unlike in the case of lognormal shocks. Under a more general distribution of market returns, a higher ϕ also makes investors more sensitive to exposures to higher-order risk, which further raises the currency risk premium.

Special Cases. To better grasp the intuition, it is useful to compare two cases: $\phi = 1$ and $\phi = 2$. When $\phi = 1$, the expected risk premium exactly corresponds to the *quanto risk premium* considered in Kremens and Martin (2019). In this scenario, investors require compensation only for the comovement between the exchange rate and the market return. This situation arises when a log-utility representative agent fully invests in the market, as also analyzed in Della Corte, Jeanneret, and Patelli (2023). Therefore, our approach generalizes the currency risk premium examined in these two studies. Another simple but informative case is when $\phi = 2$, as the expected risk premium simplifies to

$$ERP_t^2 = 2ERP_t + \frac{1}{R_{f,t}^{\$}}\mathbb{C}\text{ov}_t^{\star}\left((R_{m,T}-1)^2, \frac{S_T}{S_t}\right), \tag{8}$$

where the first part amounts to twice ERP_t, and the second one captures a nonlinear term reflecting how the exchange rate additionally comoves with *market variance*.[8]

[8]The risk-neutral covariance in Equation (8) could be synthetically priced as a quanto contract on the payoff of a variance swap written on the market return, and interpreted as the co-skewness between exchange rate returns and the market implied variance (e.g., the square of the VIX index).

Role of High-Order Risk. In the example above, the currency risk premium was only a compensation for exposure to market returns and market variance, given the considered second-order Taylor expansion. Exposures to higher-order market risk also matter. In a more general case, we can explore how higher-order risk drives the expected risk premium by expanding the risk-neutral covariance in Equation (5) as follows, now written for a generic currency i so that $S_{i,T}/S_{i,t}$ denotes the gross exchange rate return,

$$\mathbb{E}_t\left[\frac{S_{i,T}}{S_{i,t}}\right] - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \frac{1}{R^{\$}_{f,t}}\left(e^{\phi}-1\right)\sum_{n=1}^{\infty} w^{\star}_{\phi,n}\mathbb{C}\text{ov}^{\star}_t\left(r^n_{m,T}, \frac{S_{i,T}}{S_{i,t}}-1\right), \tag{9}$$

where $r_{m,T} \approx R_{m,T}-1$ is the continuously compounded market return, and $w^{\star}_{\phi,n} = \left(e^{\phi}-1\right)^{-1}\phi^n/n!$ denotes a set of weights that sum to one. Each weight is a bell-shaped function of n with its maximum value around ϕ, as illustrated by Panel A of Figure 1. In addition, we can see that, as ϕ increases, more weight is shifted to the (risk-neutral) higher-order terms. The decomposition follows from expanding the exponential function in the covariance term as a power series. It shows that the level of ϕ is intrinsically related to an aversion to higher-order risk, as in Backus, Chernov, and Martin (2011).

FIGURE 1 ABOUT HERE

3 Data and Empirical Design

This section first presents the key datasets before turning to the construction of expected currency excess returns based on professional foreign exchange (FX) forecasts and expected currency risk premia inferred from FX options for a broad cross-section of countries and across several maturities. After reporting summary statistics for the main variables, we describe the empirical methods used to estimate ϕ within the framework developed in Section 2.

3.1 FX Forecasts

We use historical exchange rate forecasts from *Consensus Economics*. The database draws on monthly surveys of about 250 professional forecasters, including leading commercial and investment banks, economic consulting firms, credit rating agencies, and independent research institutions, and reports aggregate forecasts for bilateral exchange rates at horizons of one month, three months, one year, and two years. The survey date usually falls on the second Monday of each month. We use the consensus forecast as our proxy for exchange rate expectations, and high and low forecasts to capture heterogeneity in expectations. Since forecasts are observed monthly, we construct daily observations by forward filling the most recently available forecasts until the next survey becomes available. This approach is consistent with evidence of information rigidities in exchange rate expectations, whereby forecasters update their expectations only gradually as new information arrives, consistent with sticky-information models of expectation formation (e.g., Coibion and Gorodnichenko, 2015; Beckmann and Reitz, 2020).[9]

After matching these forecasts with other data described below, our sample runs daily from January 1996 to September 2025 and includes forecasts between one month and two years for 33 currency pairs against the US dollar, i.e., the Australian dollar (AUD), Canadian dollar (CAD), Swiss franc (CHF), Danish krone (DKK), euro (EUR), British pound sterling (GBP), Japanese yen (JPY), Norwegian krone (NOK), New Zealand dollar (NZD), Swedish krona (SEK), Brazilian real (BRL), Czech koruna (CZK), Hungarian forint (HUF), South Korean won (KRW), Mexican peso (MXN), Polish zloty (PLN), Singapore dollar (SGD), Turkish lira (TRY), New Taiwan dollar (TWD), South African rand (ZAR), Chilean peso (CLP), Chinese yuan (CNY), Colombian peso (COP), Hong Kong dollar (HKD), Indonesian rupiah (IDR), Israeli new shekel (ILS), Indian rupee (INR), Malaysian ringgit (MYR), Peruvian sol (PEN), Philippine peso (PHP), Romanian leu (RON), Russian ruble (RUB), and Thai baht (THB). We refer to this set of currencies as the *broad sample*. We also consider a subsample of liquid pairs, consisting of the first 20 currencies listed above, which

[9] Internet Appendix Table A.9 reports a Monte Carlo simulation showing that forward filling does not materially affect the coefficient estimates. This is consistent with forward filling introducing classical measurement error in the dependent variable. Provided this measurement error is uncorrelated with the regressors, it inflates the residual variance rather than biasing the slope.

we refer to as the *liquid sample*. We express exchange rate forecasts in units of US dollars per unit of foreign currency, so that an increase corresponds to an appreciation of the foreign currency against the dollar.

3.2 FX Options

We also use daily implied volatilities for over-the-counter (OTC) options on exchange rates against the US dollar from JP Morgan and Bloomberg. The sample covers the same 33 currency pairs listed above over the period from January 1996 to September 2025, with maturities ranging from one month to two years. These FX options are European style options quoted as implied volatilities on baskets of options for fixed deltas and fixed maturities. From these data, we first recover implied volatility on plain-vanilla options (10-delta call and put, 25-delta call and put, and at-the money options). We then combine these quotes with the spot exchange rates and interest rates described below to map deltas into strike prices and implied volatilities into option prices using the Garman and Kohlhagen (1983) valuation formula.

3.3 Other Data

We complement our data on FX forecasts and FX options with daily spot exchange rates, interest rates, and equity index option data. First, we obtain daily WMR spot exchange rates from Datastream for the same set of currency pairs and sample period, and express them in units of US dollars per unit of foreign currency, consistent with our exchange rate forecasts (see Della Corte and Riddiough (2026) for further details on the specific tickers and data cleaning process). Second, for each currency pair in our sample, we construct daily zero-coupon yield curves by bootstrapping money market rates and interest rate swap rates collected from Bloomberg, Datastream, JP Morgan, and LSEG Workspace, ranging from one month to two years. This gives a consistent panel of daily zero-coupon rates across currencies and over time. Finally, we use the daily implied volatility surface of SPX options, i.e., European style options on the S&P 500 index, from OptionMetrics and convert implied volatilities into option prices for maturities between one month and two years. This equity option sample spans the same period as the rest of our data.

3.4 Measuring Expected Excess Returns

We measure expected excess returns by combining consensus exchange rate forecasts with spot exchange rates and interest rates. First, we construct the expected exchange rate return on day t for currency i and maturity τ as

$$EFX_{i,t,T} = \frac{\mathbb{E}_t\left[S_{i,T}\right]}{S_{i,t}} - 1 \tag{10}$$

where $\mathbb{E}_t\left[S_{i,T}\right]$ is the exchange rate forecast formed on day t for currency i over the horizon τ and $S_{i,t}$ is the corresponding spot exchange rate. The terminal date $T = t + \tau$ varies across forecast horizons, with τ ranging from one month to two years. Recall that both exchange rates are defined in units of dollars per unit of foreign currency, so that an increase denotes a foreign currency appreciation. Next, we calculate the interest rate differential on day t between the dollar and foreign currency i for maturity τ as

$$IRD_{i,t,T} = \frac{R^{\$}_{f,t,T}}{R^{i}_{f,t,T}} - 1, \tag{11}$$

where $R^{\$}_{f,t,T}$ is the gross dollar interest rate on day t and maturity τ, and $R^{i}_{f,t,T}$ is the corresponding gross interest rate for currency i. Finally, the expected excess return on day t for currency i and maturity τ is simply the difference between the expected exchange rate return and the interest rate differential as

$$ERX_{i,t,T} = \frac{\mathbb{E}_t\left[S_{i,T}\right]}{S_{i,t}} - \frac{R^{\$}_{f,t,T}}{R^{i}_{f,t,T}}. \tag{12}$$

3.5 Measuring the Expected Risk Premia

Constructing the expected risk premium for a given ϕ on day t for currency i and maturity τ

$$ERP^{\phi}_{i,t,T} = \mathbb{C}\text{ov}^{\star}_t\left(\frac{R^{\phi}_{m,T}}{\mathbb{E}^{\star}_t[R^{\phi}_{m,T}]}, \frac{S_{i,T}}{S_{i,t}}\right), \tag{13}$$

requires the risk-neutral covariance at time t between the normalized powered gross market return and the gross spot exchange rate return for currency i between times t and T. This risk-neutral covariance is not directly observable from market prices, except when ϕ equals one. In this case, it corresponds to the quanto risk premium inferred from quanto contracts studied by Kremens and Martin (2019). However, those data are available only for a limited cross-section of currency pairs, a single maturity, and at the monthly frequency.

To overcome this challenge, decompose the expected risk premium into the product of three components as

$$ERP_{i,t,T}^{\phi} = \mathbb{C}\text{or}_t^{\star}\left(\frac{R_{m,T}^{\phi}}{\mathbb{E}_t^{\star}[R_{m,T}^{\phi}]}, \frac{S_{i,T}}{S_{i,t}}\right)\sqrt{\mathbb{V}\text{ar}_t^{\star}\left(\frac{R_{m,T}^{\phi}}{\mathbb{E}_t^{\star}[R_{m,T}^{\phi}]}\right)}\sqrt{\mathbb{V}\text{ar}_t^{\star}\left(\frac{S_{i,T}}{S_{i,t}}\right)}, \tag{14}$$

where the first term captures the risk-neutral correlation between the normalized powered gross market return and the gross exchange rate return on currency i, the second term measures the risk-neutral variance of the normalized powered gross market return, and the last term quantifies the square root of the risk-neutral variance of the gross exchange rate return on currency i. These quantities are measured at time t for a given maturity τ. While the risk-neutral variances can be constructed using the model-free approach applied to equity index and currency options, as shown in Internet Appendix C.2, the risk-neutral correlation is not directly observable from traded instruments. As a proxy, we instead use the backward-looking realized correlation between the powered gross market return and the gross exchange rate return for currency i between times $t-\tau$ to t. Suppose, for example, that we want to construct the one-year expected risk premium at time t. We use one-year SPX and FX options traded on day t to construct the risk-neutral variances, and daily returns over the previous year to construct the realized correlation.

FIGURE 2 ABOUT HERE

Using realized correlation as a proxy for the unobservable risk-neutral correlation is admittedly a simple approach. When ϕ equals one, the expected risk premium implied by Equation (14) yields

a synthetic quanto risk premium that can be compared with the observed measure of Kremens and Martin (2019). For the overlapping currency pairs and dates, we regress the observed quanto premium on its synthetic counterpart. The estimated slope coefficients, reported in Internet Appendix Table A.1, are statistically indistinguishable from one. Figure 2 provides a visual comparison of the observed and synthetic quanto premia, complementing the regression evidence. These results support the use of realized correlations as proxies for unobservable risk-neutral correlations.

3.6 Descriptive Statistics

Table 1 reports descriptive statistics for the main variables across maturities from one month to two years. Panel A covers the broad sample of 33 currencies, while Panel B reports the corresponding moments for the 20 liquid currencies. Expected exchange rate returns are close to zero on average at all maturities. Their volatility declines substantially with the horizon: in the broad sample, the standard deviation falls from 10.81% at one month to 4.54% at two years. The pattern is similar in the liquid-currency sample. The short-horizon distributions are strongly right-skewed and fat-tailed, whereas both features weaken at longer maturities. Expected exchange rate returns are also highly persistent, with first-order autocorrelations increasing from about 0.95 at one month to 0.99 at two years.

TABLE 1 ABOUT HERE

Expected excess returns are positive on average at all maturities. Their means range from 1.83% to 2.08% in the broad sample and from 1.22% to 1.80% in the liquid-currency sample. Their volatility declines with maturity, while persistence rises from about 0.95 at one month to 0.99 at two years. Expected currency risk premia are also positive throughout. Their average level increases with ϕ: in the broad sample, the mean ranges from 0.29%–0.30% for ERP^1, 0.51%–0.58% for ERP^2, and 0.68%–0.83% for ERP^3. These premia are more volatile at longer maturities and for higher values of ϕ, but their positive skewness and excess kurtosis are concentrated at short maturities. Their persistence is close to one from the one-year horizon onward.

Overall, the table shows that exchange rate expectations are volatile at short horizons and that both expected excess returns and expected risk premia are positive and highly persistent. These features motivate our analysis of how survey expectations and option-implied premia jointly identify the term structure of risk preferences.

4 Empirical Analysis

In this section, we provide evidence on the estimation of ϕ using the two approaches presented in Section 2, both unconditionally and conditionally, as well as for different horizons. Our analysis allows us to quantify how strongly FX market participants price higher-order risk, explore how that pricing varies over time, and characterize the term structure of risk preferences.

4.1 Unconditional Estimates of ϕ

We start our empirical investigation by estimating the unconditional value of ϕ using panel regressions subsumed by the following specification

$$ERX_{i,t,T} = \alpha_i + \alpha_\tau + \alpha_t + \beta_\phi ERP^{\phi}_{i,t,T} + \varepsilon_{i,t,T}, \tag{15}$$

where $ERX_{i,t,T}$ denotes the expected excess return for currency i on day t over the horizon τ (described in Section 3.4), and $ERP^{\phi}_{i,t,T}$ represents the expected risk premium on day t for the same currency pair i and maturity τ (detailed in Section 3.5). We use daily observations and the maturity/horizon ranges between 1 month and 2 years. We include currency fixed effects α_i to absorb time-invariant differences across expected returns, maturity fixed effects α_τ to absorb average differences across horizons, and calendar-time fixed effects α_t to control for shocks simultaneously affecting expected returns and option-implied risk premia across all currencies.

We recover ϕ from the data using two complementary approaches, named *search* and *simple* method. The first method builds on Proposition 1 and imposes the restriction that β_ϕ must equal one. Empirically, we compute the expected risk premium over a fine grid of candidate values of ϕ,

estimate the corresponding panel regressions, and select the value for which the estimated slope β_ϕ is closest to one. We leave the intercept unrestricted, since a nonzero average pricing error may reflect trading frictions or transaction costs that the model does not capture. The second method follows the approximation in Proposition 2. When the nonlinear component in ϕ is small, we simply compute the expected risk premium under ϕ equal one, estimate the corresponding panel regression, and interpret the estimated slope β_ϕ directly as an estimate of ϕ.

TABLE 2 ABOUT HERE

We report our estimates of ϕ in Table 2. Panel A uses the broad sample of currency pairs, while Panel B uses the liquid sample. We first present estimates based on the simple method, with standard errors clustered by currency and calendar time, and then turn to estimates based on the search method, with standard errors for ϕ computed using the delta method. The first column in each panel pools all maturities. In the broad sample, the estimate of ϕ is 6.01 under the simple method and 7.17 under the search method. In the liquid sample, the corresponding estimates are 6.11 and 7.27. All four estimates are statistically significant and virtually identical across the broad and liquid samples, suggesting that our estimates are not driven by the inclusion of less liquid currencies. The simple estimates are slightly lower than the search estimates, since the simple method ignores the nonlinear component in ϕ and therefore provides only a first-order approximation. Overall, both methods point to values of ϕ well above one, the value corresponding to log utility and often used in the exchange rate predictability literature (e.g., Kremens and Martin, 2019; Della Corte, Jeanneret, and Patelli, 2023). Interpreted through the lens of power utility, these estimates imply substantial risk aversion. More generally, they indicate that FX market participants place meaningful weight on higher-order risk-neutral moments when pricing currency risk premia.

Since expected excess returns and expected risk premia are both available at multiple horizons, we can also recover the term structure of ϕ by estimating the panel regression in Equation (15) separately for each maturity. The remaining columns in each panel of Table 2 presents these estimates separately for each maturity and reveal a clear downward-sloping term structure of ϕ. In

the broad sample, the simple estimates decline from 7.25 at the one-month horizon to 3.59 at three months, 2.10 at one year, and 1.75 at two years. The search estimates display the same pattern, falling from 8.01 at one month to 4.02 at three months, 2.40 at one year, and 1.99 at two years. The liquid sample delivers nearly identical results: estimates are above 7 at the one-month horizon and decline to about 2 at the two-year horizon. Hence, the panel evidence points to a steep term structure of risk preferences, with higher-order risk priced most heavily at short maturities and ϕ remaining above the log-utility benchmark even at the longest horizons.

This downward-sloping term structure is consistent with broader evidence that compensation for higher-order risk is concentrated at short maturities. In equity markets, Dew-Becker, Giglio, Le, and Rodriguez (2017) show that the spot variance risk premium in the S&P 500 index is large, whereas forward premia become small beyond one or two months. Della Corte, Kozhan, and Neuberger (2021) report similar evidence in FX markets. Relatedly, Lustig, Stathopoulos, and Verdelhan (2019) and Zviadadze (2017) show that the compensation embedded in carry-trade returns declines with horizon. Our estimates provide complementary evidence from survey-based expected excess returns and option-implied risk premia. The evidence also connects to models and experiments in which risk attitudes depend on the horizon of the payoff. Laboratory and field studies find that people are more risk averse toward near-term risks than toward risks realized further in the future (e.g., Holt and Laury, 2002; Coble and Lusk, 2010; Abdellaoui, Diecidue, and Öncüler, 2011; Eisenbach and Schmalz, 2016). Consistent with this evidence, Eisenbach and Schmalz (2016) and Andries, Eisenbach, and Schmalz (2024) develop asset pricing models with horizon-dependent risk aversion. Our estimates provide direct empirical support for this class of models by showing that the price of higher-order currency risk is largest at short horizons and declines with maturity.

We further show that these estimates are robust to alternative sample and specification choices. First, Internet Appendix Table A.2 shows that the results remain similar when we remove the fixed effects and estimate pooled regressions. The estimates are close to those in the baseline panel specification, indicating that the main relation is not driven mechanically by the fixed-effect structure. Second, Internet Appendix Table A.3 restricts the analysis to the traditional 10 major

currency pairs. The estimates again remain close to the baseline results, suggesting that the evidence is not driven by less liquid currencies or by the broader cross-section. These robustness checks show that the estimated values of ϕ are stable across both regression specifications and currency samples.

4.2 Accounting for Heterogeneity in Expectations

We next examine whether our estimates of ϕ are robust to incorporating heterogeneity in survey expectations. The consensus forecasts used in our core analysis capture the average view, but our database also reports the high and low forecasts across respondents. Following Section 3.4, we construct expected excess returns separately for each forecast measure and estimate panel regressions based on the following specification:

$$ERX_{i,f,t,T} = \alpha_i + \alpha_f + \alpha_t + \alpha_\tau + \beta_\phi ERP^{\phi}_{i,f,t,T} + \varepsilon_{i,f,t,T}, \tag{16}$$

where f indexes the forecast measure, namely the consensus, high, or low forecast, and α_f absorbs systematic differences across these measures. All other elements of the specification follow the baseline regression in Equation (15).

TABLE 3 ABOUT HERE

Table 3 shows that the results are very close to those obtained using only the consensus forecast. When all maturities are pooled, the simple estimate of ϕ is 6.25 in the broad sample and 6.39 in the liquid sample, compared with 6.01 and 6.11 in the baseline specification. The corresponding search estimates are 7.54 and 7.70, compared with 7.17 and 7.27 in the baseline. The maturity-specific estimates also preserve the downward-sloping term structure documented above. Estimates are largest at the one-month horizon and decline steadily as the maturity increases. Overall, incorporating the high and low forecasts alongside the consensus forecast leaves the main conclusions unchanged. The estimates are slightly larger than in the baseline specification, but the economic message remains the same.

4.3 Using Expected Exchange Rate Returns

Our definition of expected excess returns in Section 3.4 combines expected exchange rate returns with interest rate differentials matched over the same horizon/maturity. A natural question is whether our estimates are driven by the interest rate differential component. To address this issue, we use the expected exchange rate return as the dependent variable and include the interest rate differential as an additional regressor as follows

$$EFX_{i,t,T} = \alpha_i + \alpha_t + \alpha_\tau + \beta_\phi ERP^{\phi}_{i,t,T} + \gamma IRD_{i,t,\tau} + \varepsilon_{i,t,T}, \tag{17}$$

where $EFX_{i,t,T}$ denotes the expected exchange rate return for currency i on day t over the horizon τ and $IRD_{i,t,\tau}$ represents the interest rate differential between the dollar and foreign currency i on day t for maturity τ. All other terms are defined as in the baseline specification in Equation (15).

TABLE 4 ABOUT HERE

Table 4 shows that the estimates of ϕ are almost unchanged. When all maturities are pooled, the simple estimate is 5.97 in the broad sample and 6.11 in the liquid sample, while the corresponding search estimates are 7.10 and 7.27. These values are very close to the baseline estimates reported in Table 2. The maturity-specific estimates also display the same downward-sloping pattern, with the largest values at the one-month horizon and substantially lower values at longer maturities. The coefficient on IRD is positive and statistically significant across all specifications, but controlling for it directly leaves the estimated price of higher-order risk essentially unchanged.

4.4 Time Variation in ϕ

We next examine whether the term structure of ϕ varies over time. In particular, we start from the specification in Equation (15) and estimate ϕ over 252-day rolling windows using the simple method. Because each window needs a full year of data, the rolling estimates begin in January 1997. Within each window, we estimate the regressions separately for each maturity and without fixed effects. While the baseline estimates summarize the average relation between expected excess

returns and expected risk premia, after absorbing currency, time, and maturity fixed effects, here we focus on the evolution of the maturity profile of ϕ. We then summarize the maturity-specific estimates using level, slope, and curvature. The level is the average of ϕ across maturities. The slope is the difference between the 2-year and 3-month ϕ estimates. The curvature is the average of the inner maturities, 3 months and 1 year, minus the average of the outer maturities, 1 month and 2 years.

FIGURES 3–5 ABOUT HERE

Figure 3 displays the rolling level component of ϕ, which captures common movements across maturities. The level is far from constant, and rises around the Asian currency crisis, the Russian stress episode in 1998, the global financial crisis, the euro-area debt crises, the taper episode and the Fragile Five sell-off, the yuan devaluation, the Covid-19 pandemic, the Russia–Ukraine war, and the Israel–Iran war. These episodes are typically associated with sharp repricing, funding stress, or sudden shifts in global risk appetite. The level component therefore shows that ϕ rises when market conditions worsen. Figure 4 reports the rolling slope component, which measures how long-maturity estimates of ϕ move relative to short-maturity estimates. The slope is often negative, consistent with the downward-sloping term structure in the full-sample estimates. It becomes more negative during stress episodes such as the global financial crisis, the Greek debt crisis, the taper tantrum, the SNB intervention, the Volmageddon episode, and the Russia–Ukraine war. This indicates that the short end of the term structure rises more than the long end during periods of stress. In calmer periods, the slope moves closer to zero and can temporarily turn positive. Figure 5 reports the rolling curvature component, which captures movements in the middle of the term structure that are not explained by changes in level or slope. Intermediate maturities sometimes move independently of the one-month and two-year estimates. This is visible around the US yield-curve inversion, the Bear Stearns and Lehman shocks, the Italian fiscal crisis, the taper guidance period, the Volmageddon episode, the Covid-19 pandemic, and the Russia–Ukraine war. In these episodes, the term structure does not merely shift up or steepen, and its middle segment also changes shape.

The rolling estimates show that ϕ varies over time along several dimensions. Stress periods are associated with a higher level of ϕ, a steeper term structure, and occasional changes in curvature. The full-sample estimates therefore mask substantial time variation in how FX markets price risk across horizons.

4.5 Conditional Term Structure

We next examine whether the term structure of ϕ changes with market conditions. We first split the sample into good and bad states using two measures of volatility. The first is the VIX index, which captures option-implied volatility on the S&P 500. The second is one-month realized volatility of the S&P 500. In each case, good states are days below the sample mean and bad states are days at or above the sample mean. We then estimate ϕ within each state using the simple method on Equation (15), both pooling all maturities and separately by maturity.

TABLES 5–6 ABOUT HERE

Table 5 shows that the term structure of ϕ changes sharply across market conditions. The estimate is modest in low-VIX states, ranging from 0.15 at the one-month horizon to 1.64 at the two-year horizon, but much larger at short maturities in high-VIX states, equal to 8.07 at one month, 3.76 at three months, 1.90 at one year, and 1.52 at two years. Table 6 delivers a similar result when market conditions are measured using realized equity volatility. In good states, the estimate ranges from 0.46 to 2.04, whereas in bad states it starts at 8.49 at one month and declines to 3.71 at three months, 1.69 at one year, and 1.42 at two years. The corresponding search-method estimates, reported in Internet Appendix Table A.4 and Internet Appendix Table A.5, lead to the same conclusion.[10]

FIGURE 6 ABOUT HERE

[10]The same pattern also holds for the 10 major currencies, as shown in Internet Appendix Table A.6 and Internet Appendix Table A.7.

Figure 6 summarizes these estimates visually. In bad states, ϕ is much higher at short maturities and declines sharply with the horizon. In good states, the term structure is flatter and, in some cases, mildly upward sloping. This pattern is present whether market conditions are measured with option-implied volatility or realized volatility. Hence, the average downward-sloping term structure documented above is driven mainly by periods in which equity-market risk is elevated. The conditional estimates therefore show that the term structure of ϕ is substantially steeper in bad states, with the response concentrated at short maturities, indicating that near-term risks become especially important when market volatility is high.

5 What Moves Risk Preferences in the FX Market?

Earlier sections establish that ϕ is large, that it varies over time, and that its term structure steepens in bad times. Those sections do not say what the variation lines up with. "Bad times" names several conditions that arrive together, and showing that all of them correlate with risk preferences does not say which of them ϕ tracks. We set out four explanations, ask which the data favor and in what proportion, and put the question to the level and the slope together.

We work with the level of the ϕ term structure and its slope, estimated over a 252-day rolling window on 33 currencies against the dollar. The level is the average of the maturity-specific estimates and the slope is $\phi(2\text{y}) - \phi(3\text{m})$, both daily from January 1997 to September 2025, and all the analysis is monthly, at month-end.[11]

5.1 Why Risk-Taking in Currency Markets Should Vary

A first explanation is that investors respond to how risky the world is (Bloom, 2009), or to how much they dislike risk. Bekaert, Engstrom, and Xu (2022) separate the two and show they come apart in important episodes. Menkhoff, Sarno, Schmeling, and Schrimpf (2012) document that

[11] The level averages 5.16 with a standard deviation of 4.19 and a half-life of 10.3 months; the slope averages 1.02 with a standard deviation of 5.22 and is negative 45% of the time. Because each observation is a 252-day rolling estimate, a month-end value is a backward-looking one-year average, so a monthly surprise in any driver moves it only slowly.

global FX volatility is priced, and that high-yield currencies lose when it spikes. Investors should then be less willing to take risk when volatility and uncertainty rise.

A second explanation locates the variation in the balance sheets of the intermediaries who warehouse the risk, making it a property of who bears the risk rather than of the environment everyone faces. He and Krishnamurthy (2013) and He, Kelly, and Manela (2017) show that the marginal investor in many markets is a constrained intermediary whose capital moves the price of risk. Gabaix and Maggiori (2015) make the mechanism explicit for currencies: financiers with limited risk-bearing capacity absorb currency order flow, and the premium they require rises as their capacity shrinks. The cross-currency basis is the observable price of that balance sheet. Investors should then pull back when intermediary capital is depleted and the basis widens.

A third explanation is macroeconomic, and unlike the second it turns on conditions every investor faces rather than on the balance sheet of whoever happens to hold the risk. Lustig, Roussanov, and Verdelhan (2014) show that currency risk premia are countercyclical, and Pflueger, Siriwardane, and Sunderam (2020) link measured risk perceptions to credit conditions and real investment. Risk-taking should then follow the global business cycle, credit spreads, and equity markets.

A fourth explanation is specific to currencies. Brunnermeier, Nagel, and Pedersen (2008) show that carry positions unwind in crashes, with the size of the trade and its funding conditions governing the unwind, and Lustig, Roussanov, and Verdelhan (2011) that a common factor prices the cross-section of currency returns. When interest differentials and risks are widely spread across currencies, the carry trade is larger and more concentrated in a few high-yield names, and its crash exposure is greater. Risk should then be priced more heavily when carry and risk are widely spread, whatever the macro-financial aggregates do.

Volatility rises as credit spreads widen and intermediaries lose capital, so the four cannot be told apart by inspection. Which accounts for the variation in ϕ, and how much, is the empirical question.

5.2 What ϕ Measures

Eleven series enter the joint regression, one group for each explanation, and Table 7 reports them.[12] In levels, ϕ moves with all of them. Every driver clears $|t| = 2.8$ with the sign its explanation predicts: risk is priced more heavily when volatility, uncertainty, credit spreads, the basis, the dollar, and cross-currency dispersion rise, and when activity and equities fall. The estimate of ϕ is recovered from survey forecasts and option prices, and the drivers come from entirely separate sources, so the two need not have moved together at all.

In monthly changes, in Table 7, the set thins to six: the option-implied volatility of volatility at $t = 4.96$, the equity market return at -3.48, financial uncertainty at 3.38, the cross-sectional dispersion of realized volatility at 2.95, the credit spread at 2.78, and the dollar at 2.56. Global activity, intermediary capital, and the cross-currency basis drop out. The strongest of the six is the option-implied volatility of volatility. A one-standard-deviation rise in it goes with a rise in ϕ of 0.44, a tenth of the standard deviation of the level. Since ϕ scales the option-implied premium, that is 16 basis points a year of expected currency risk premium.

TABLE 7 ABOUT HERE

Table 7 also shows the drivers lining up with the slope in the opposite direction and at similar magnitude: the option-implied volatility of volatility enters the level at $t = 4.96$ and the slope at -4.90, the dispersion of realized volatility at 2.95 and -2.62. Level and slope are close to mirror images, correlating -0.57 in levels and -0.80 in monthly changes, so accounting for one is most of the way to accounting for the other. Tighter conditions raise the near end of the term structure by more than the far end.

These drivers explain the fluctuations in the level and the slope of ϕ well. Figure 7 plots the fitted value from a joint regression of each series on the full set of drivers, estimated in levels at the monthly frequency, against the original series. For the level they account for 38.5% of the

[12] ϕ is stationary (augmented Dickey-Fuller -4.04, $p = 0.001$) but persistent enough that a regression in levels leaves residual autocorrelation of 0.88; year-on-year changes leave 0.85, since consecutive annual differences overlap in eleven months of twelve, and monthly changes bring it to 0.23. AR(2) innovations give the same ranking of groups. Inference is Newey-West throughout.

variation, a figure that also appears in the levels columns of Table 7, and the fitted series correlates 0.62 with the original. It follows the major turning points: the run-up through 2014 to 2016, the spike in 2020, and the 2022 tightening cycle. The same conditions reproduce the slope of the term structure, accounting for 21.3% of its variation at a correlation of 0.46. Observable financial conditions account for a substantial share of both how much risk investors take and the horizon over which they take it.

FIGURE 7 ABOUT HERE

5.3 Which Conditions Account for the Variation

Owen values allocate the joint R^2 across groups rather than individual series, so explanations with different numbers of variables are comparable (Owen, 1977; Huettner and Sunder, 2012). The drivers account for 15.7% of the monthly change in the level and 10.9% of the slope, allocated in Figure 8.

FIGURE 8 ABOUT HERE

Mapped onto the four explanations of Section 5.1 in Figure 8, the dollar and cross-currency dispersion group and risk and uncertainty take the largest shares of the level, 5.4 and 5.2 points, and credit and intermediary conditions the smallest, 2.3. A second cut is sharper. Five of the drivers are taken from currency markets and six are global macro-financial series; allocated between the two sets, the currency-market series take 11.1 points of the level against 4.6, and 8.6 against 2.3 for the slope. Panel D shows the currency-market set taking the larger share at three months, at the level, and for the slope. The global set leads only at two years. Panel C shows the two much closer in levels, with the global set ahead at the short end. The result is about innovations, not about the long swings. Within that set the dispersion measures do the work rather than the dollar.[13] What matters inside currency markets is the spread of conditions across them, not the dollar cycle.

Volatility lines up with the term structure; carry with the level. These conditions act more strongly at the short end. In Figure 8 the joint fit falls from 13.2% at three months to 7.9% at two years.

[13]The dispersion measures share no input with ϕ, being built from realized volatility and interest differentials rather than option prices or survey forecasts. Estimating ϕ on the ten developed currencies and building dispersion from the other twenty-three, so that no currency enters both, leaves the relation at $t = 2.13$.

What differs across them is how steeply each tilts the curve. In Table 7 the option-implied volatility of volatility enters the slope at $t = -4.90$ and the dispersion of realized volatility at -2.62, while the dispersion of carry and the dollar barely tilt it. Risk and uncertainty leads the slope with 5.2 points, as much as it takes of the level. How volatile currency markets are priced to be changes the horizon over which investors are unwilling to bear risk; how much carry is on offer changes the level at every horizon at once.

What counts as important depends on whether one looks at levels or innovations. Well under half the fit in levels survives differencing, and differencing also changes the ranking. Global activity ranks second in levels and third in monthly changes; risk and uncertainty ranks third in levels and second in monthly changes. In the Owen columns of Table 7 the intermediary channel shifts most, dominating the slope in levels at 8.8 points while contributing little in monthly changes, so balance-sheet conditions appear more relevant for low-frequency variation than for month-to-month movements.

At the monthly frequency, conditions observed inside currency markets account for more than twice as much variation in ϕ as global macro-financial ones, and do so through the dispersion of risk and compensation across currencies rather than through aggregates.

6 The Cross-Section of ϕ

Asset pricing normally assumes that one stochastic discount factor prices everything. In our framework that means a single ϕ: the same value reconciles survey expectations with option prices for every currency, and all cross-sectional variation in expected returns runs through covariances. This section estimates ϕ currency by currency and asks whether that assumption holds.

6.1 Risk Preferences Differ across Currencies

A single ϕ does not fit the cross-section. Panel C of Table 8 averages each currency's rolling estimate over the sample. The estimate ranges from -1.4 for the Singapore dollar to 9.8 for the Turkish lira,

with a cross-sectional standard deviation of 2.9 against a mean of 3.0, and five of the 31 currencies carry a negative estimate.[14] The effective price of risk is not common across currencies.

The spread is not sampling error. The ordering of currencies lines up with their average interest rate differential, which Section 6.2 reports, and sampling error cannot line up with a quantity measured entirely outside the estimation. Systematic differences across currencies, in option liquidity or in how tightly the exchange rate is managed, are a separate question; the negative estimates above are the clearest case.[15]

How much risk investors take in a given currency changes over time. Panel C of Table 8 decomposes the variance of currency-month ϕ: a common time effect accounts for 15%, permanent differences across currencies for 7%, and the remaining 78% varies by currency and over time. The rankings say the same thing. Splitting the sample in half, the ordering of currencies by ϕ correlates only 0.25 across the two halves, while the ordering by interest rate differential correlates 0.90. A currency can require a high ϕ in one decade and not in the next, whereas a high-yield currency tends to stay high-yield.

6.2 What the Differences Line Up With

Figure 9 plots the currency-level ϕ against six characteristics, and every panel slopes the same way. Panel A is the average interest rate differential, which correlates 0.73 with ϕ. Panels B and (c) show what holding the currency paid: high-ϕ currencies depreciate more, correlating -0.65 with the exchange rate return, and the total return still correlates 0.37, because the depreciation absorbs most but not all of the differential. Panels D to (f) add option-implied skew at 0.62, sovereign CDS at 0.65, and the cross-currency basis at 0.49. Risk is priced more heavily in the currencies that

[14] The cross-section covers 31 of the 33 currencies in our sample; the Chinese yuan and the Hong Kong dollar are excluded as pegs. The five negative estimates all belong to currencies with managed or pegged exchange rates: the Singapore dollar, the Philippine peso, the Romanian leu, the Danish krone and the Thai baht. Where the exchange rate is administered, the risk-neutral covariance with the equity market is a weaker measure of the currency's risk and survey forecasts are anchored by the policy regime rather than by risk, so the assumptions behind Equation (13) bind least well for these currencies.

[15] A formal test of a common ϕ rejects in every month of the sample, and Internet Appendix D.1 reports it. We do not lean on it. The first-stage regressions are estimated on overlapping currency-days, so the standard errors the rolling files report understate the true sampling uncertainty, inflating any statistic that scales deviations by them.

look risky on any of these measures.

FIGURE 9 ABOUT HERE

The Turkish lira sits at the extreme of the first three panels of Figure 9. It paid the widest differential in the sample, suffered the steepest depreciation against the dollar, and retained almost none of the difference. It also carries the highest ϕ of any currency. Investors in the currency that commanded the most were paid the most and gave back nearly all of it.

Panel A also shows where the ordering comes from, and Internet Appendix Table A.8 quantifies it. Among the ten currencies that yielded less than the dollar the slope is flat, 0.09 ($t = 0.15$). Among the twenty-one that yielded more it is 0.93 ($t = 6.86$). Risk preferences therefore shift precisely where carry compensation and crash exposure are extreme. The split is about yield rather than development status: within emerging and frontier currencies alone the slope is 0.88 ($t = 6.66$).

Table 8 puts carry against each of the others in Panel B. Carry holds at $t = 4.97$ against the cross-currency basis, 4.36 against implied skew, and 2.81 against sovereign CDS, while every rival falls below 1.4. High-yield currencies do trade at a wider basis and are issued by riskier sovereigns. Neither characteristic adds anything to what the yield already says.

TABLE 8 ABOUT HERE

Carry therefore dominates the cross-section, and it dominates out of sample as well. In Panel D of Internet Appendix Table A.8, carry averaged before 2011 explains 28% of the variation in ϕ averaged after 2011, and carry averaged after 2011 explains 33% of ϕ before it. Carry keeps its cross-sectional order across the two halves while ϕ does not. A currency's yield is a durable attribute, and it lines up with how much risk investors take in that currency a decade later. The ordering also holds at each maturity taken on its own, and is sharper at the long end than the short: in Panel E of Internet Appendix Table A.8 carry explains 36% of the cross-section of the one-month and three-month estimates and 53% of the one-year and two-year ones. What orders currencies is therefore not a property of the short-dated estimates that dominate the level factor. Internet Appendix D shows that the ranking is not an accounting artifact of the interest differential

sitting inside the survey expected excess return, and reports the thresholds and subsamples behind the concentration.

6.3 Crash Risk of the Carry Trade

Section 6.2 shows that carry orders ϕ across currencies, but not why. One reading is crash risk. A high-yield currency earns its differential in most months and loses it in a few that coincide with equity market declines, so reconciling a survey expectation of its excess return with the option-implied premium at $\phi = 1$ requires a higher ϕ. That reading makes a prediction about something we have not used. Realized returns play no part in constructing ϕ, which is built from survey forecasts and option prices. If the reading is right, portfolios holding more carry should lose more when risk-taking contracts, and should lose it in bad states.

We test it with a portfolio sort. We sort currencies into quintiles on the lagged interest rate differential, hold each portfolio for one month, and regress its excess return on the monthly change in ϕ, with the long-short portfolio buying the highest-carry quintile and selling the lowest. Table 9 asks three things of that regression: whether the exposure grows with carry, whether it concentrates in bad equity states, and whether it survives the dollar factor, the equity market return, and global FX volatility.

The prediction holds. In Panel A of Table 9, exposures become steadily more negative with carry, from -0.088 in the lowest quintile to -0.464 in the highest, and the long-short portfolio loads at -0.376 ($t = -2.94$). The magnitude is large: a one-standard-deviation monthly rise in ϕ goes with a carry return 0.54% lower that month, or 6.4% a year.

Those losses also fall in the states the reading names. In Panel C of Table 9 the long-short loading is -0.626 ($t = -3.75$) in the worst 30% of months for the equity market, against -0.202 in the rest. The carry trade gives back its differential when investors pull away from risk, not at other times.

TABLE 9 ABOUT HERE

The exposure is not, however, incremental. In Panels A and B of Table 9, the dollar factor, the

equity market return, and global FX volatility together leave the long-short loading at -0.130 ($t = -1.13$), and global FX volatility absorbs most of it on its own. The redundancy is with known risk factors rather than with the construction, since rebuilding ϕ from currencies that do not form the portfolios leaves the raw gradient intact. Realized returns therefore confirm that the cross-section of ϕ tracks crash exposure. Those returns do not make ϕ a priced factor beyond the global risk conditions already known to move carry returns.

Three findings stand out. A single ϕ cannot reconcile survey expectations with option prices across currencies: estimates range from -1.4 to 9.8, and a currency's estimate changes over time rather than being fixed by the currency. What orders that spread is carry, concentrated among the high-yield currencies, and the ordering holds when carry is measured in one decade and ϕ in another. Realized returns are consistent with that reading: carry portfolios lose when risk-taking contracts and lose most in the months when equities fall.

Two implications follow. The data reject a common value of ϕ across currencies, so the aggregate series of Section 5 is an average over currencies that price risk differently. And ϕ tracks crash exposure across currencies without pricing it beyond the global risk conditions already known to move carry returns.

7 Concluding Remarks

Risk preferences in the FX market are not a single number. We measure RORO from survey forecasts and option prices, and find that it falls steeply with the forecast horizon, reverses that slope with market conditions, tracks conditions inside currency markets rather than macro-financial aggregates, and varies across currencies by more than its own average level. What investors demand for bearing currency risk therefore depends on when the exposure falls due, on the state of the market, and on which currency carries it.

Two features of our approach bound what we can say, and both point to work worth doing. We proxy the unobservable risk-neutral correlation with its realized counterpart. The proxy passes the

one test available to us, since the synthetic premium it produces is indistinguishable from the traded quanto premium of Kremens and Martin (2019), but a genuinely forward-looking estimate of that correlation would sharpen every estimate here. We also work with a second-order expansion of the expected risk premium, which suffices unconditionally but leaves the nonlinear term unexamined precisely when it should matter most, in periods of stress.

More broadly, the framework needs only survey expectations and option prices at matching horizons. Those exist for equities, commodities, and interest rates, so the term structure of risk preferences can be measured well beyond the currency market. Whether it slopes the same way there, and reverses in the same states, would show whether what we document is a property of FX investors or of investors.

References

Abdellaoui, Mohammed, Enrico Diecidue, and Ayse Öncüler, 2011, Risk preferences at different time periods: An experimental investigation, *Management Science* 57, 975–987.

Alvarez, Fernando, and Urban J. Jermann, 2005, Using asset prices to measure the persistence of the marginal utility of wealth, *Econometrica* 73, 1977–2016.

Andries, Marianne, Thomas M Eisenbach, and Martin C Schmalz, 2024, Horizon-dependent risk aversion and the timing and pricing of uncertainty, *Review of Financial Studies* 37, 3272–3334.

Backus, David, Mikhail Chernov, and Ian Martin, 2011, Disasters implied by equity index options, *Journal of Finance* 66, 1969–2012.

Baker, Malcolm, and Jeffrey Wurgler, 2006, Investor sentiment and the cross-section of stock returns, *Journal of Finance* 61, 1645–1680.

Bakshi, Gurdip, and Dilip Madan, 2000, Spanning and derivative-security valuation, *Journal of Financial Economics* 55, 205–238.

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

Beckmann, Joscha, and Stefan Reitz, 2020, Information rigidities and exchange rate expectations, *Journal of International Money and Finance* 105, 102136.

Bekaert, Geert, Eric C. Engstrom, and Nancy R. Xu, 2022, The time variation in risk appetite and uncertainty, *Management Science* 68, 3975–4004.

Binsbergen, Jules van, Michael Brandt, and Ralph Koijen, 2012, On the timing and pricing of dividends, *American Economic Review* 102, 1596–1618.

Bloom, Nicholas, 2009, The impact of uncertainty shocks, *Econometrica* 77, 623–685.

Britten-Jones, M., and Anthony Neuberger, 2000, Option prices, implied price processes, and stochastic volatility, *Journal of Finance* 55, 839–866.

Brunnermeier, M. K., S. Nagel, and L. H. Pedersen, 2008, Carry trades and currency crashes, *NBER Macroeconomics Annual* 23, 313–348.

Burnside, Craig, Martin Eichenbaum, Isaac Kleshchelski, and Sergio Rebelo, 2011, Do peso problems explain the returns to the carry trade?, *Review of Financial Studies* 24, 853–891.

Chan, Yeung Lewis, and Leonid Kogan, 2002, Catching up with the joneses: Heterogeneous preferences and the dynamics of asset prices, *Journal of Political Economy* 110, 1255–1285.

Coble, Keith H, and Jayson L Lusk, 2010, At the nexus of risk and time preferences: An experimental investigation, *Journal of Risk and Uncertainty* 41, 67–79.

Cohn, Alain, Jan Engelmann, Ernst Fehr, and Michel André Maréchal, 2015, Evidence for countercyclical risk aversion: An experiment with financial professionals, *American Economic Review* 105, 860–85.

Coibion, Olivier, and Yuriy Gorodnichenko, 2015, Information rigidity and the expectations formation process: A simple framework and new facts, *American Economic Review* 105, 2644–2678.

Della Corte, Pasquale, Alexandre Jeanneret, and Ella Patelli, 2023, A credit-based theory of the currency risk premium, *Journal of Financial Economics,* 149, 473–496.

Della Corte, Pasquale, Roman Kozhan, and Anthony Neuberger, 2021, The cross-section of currency volatility premia, *Journal of Financial Economics* 139, 950–970.

Della Corte, Pasquale, Tarun Ramadorai, and Lucio Sarno, 2016, Volatility risk premia and exchange rate predictability, *Journal of Financial Economics* 120, 21–40.

Della Corte, Pasquale, and Steven J. Riddiough, 2026, Currency speculation, *Annual Review of Financial Economics* 18 Forthcoming.

Dew-Becker, Ian, Stefano Giglio, Anh Le, and Marius Rodriguez, 2017, The price of variance risk, *Journal of Financial Economics* 123, 225–250.

Eisenbach, Thomas M, and Martin C Schmalz, 2016, Anxiety in the face of risk, *Journal of Financial Economics* 121, 414–426.

Fama, Eugene F., 1984, Forward and spot exchange rates, *Journal of Monetary Economics* 14, 319–338.

Fama, Eugene F, and James D MacBeth, 1974, Long-term growth in a short-term market, *Journal of Finance* 29, 857–885.

Fan, Zhenzhen, Juan M Londono, and Xiao Xiao, 2022, Equity tail risk and currency risk premiums, *Journal of Financial Economics* 143, 484–503.

Farmer, Leland E, Emi Nakamura, and Jón Steinsson, 2024, Learning about the long run, *Journal of Political Economy* 132, 3334–3377.

Gabaix, Xavier, and Matteo Maggiori, 2015, International liquidity and exchange rate dynamics, *Quarterly Journal of Economics* 130, 1369–1420.

Garman, Mark B., and Steven W. Kohlhagen, 1983, Foreign currency option values, *Journal of International Money and Finance* 2, 231–237.

Gonçalves, Andrei S, 2021, The short duration premium, *Journal of Financial Economics* 141, 919–945.

Gormsen, Niels Joachim, 2021, Time variation of the equity term structure, *Journal of Finance* 76, 1959–1999.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales, 2018, Time varying risk aversion, *Journal of Financial Economics* 128, 403–421.

Hansen, Lars, and Robert Hodrick, 1980, Forward exchange rates as optimal predictors of future spot rates: An econometric analysis, *Journal of Political Economy* 88, 829–53.

Hansen, Lars Peter, 2007, Beliefs, doubts and learning: Valuing macroeconomic risk, *American Economic Review* 97, 1–30.

——— , and Thomas J Sargent, 2011, *Robustness* (Princeton university press).

Hassan, Tarek A, and Tony Zhang, 2021, The economics of currency risk, *Annual Reviews of Economics* 13, 281–307.

He, Zhiguo, Bryan Kelly, and Asaf Manela, 2017, Intermediary asset pricing: New evidence from many asset classes, *Journal of Financial Economics* 126, 1–35.

He, Zhiguo, and Arvind Krishnamurthy, 2013, Intermediary asset pricing, *American Economic Review* 103, 732–770.

Holt, Charles A, and Susan K Laury, 2002, Risk aversion and incentive effects, *American economic review* 92, 1644–1655.

Huettner, Frank, and Marco Sunder, 2012, Axiomatic arguments for decomposing goodness of fit according to Shapley and Owen values, *Electronic Journal of Statistics* 6, 1239–1250.

Jiang, George J., and Yisong S. Tian, 2005, The model-free implied volatility and its information content, *Review of Financial Studies* 18, 1305–1342.

Jurado, Kyle, Sydney C. Ludvigson, and Serena Ng, 2015, Measuring uncertainty, *American Economic Review* 105, 1177–1216.

Kelly, J.L., 1956, A new interpretation of the information rate, *Bell Systems Technical Journal* 35, 917–926.

Kremens, Lukas, and Ian Martin, 2019, The quanto theory of exchange rates, *American Economic Review* 109, 810–843.

——— , and Liliana Varela, 2025, Long-horizon exchange rate expectations, *Journal of Finance* 80, 3695–3724.

Latané, Henry A., 1959, Criteria for choice among risky ventures, *Journal of Political Economy* 67, 144–155.

Long, John B Jr., 1990, The numeraire portfolio, *Journal of Financial Economics* 26, 29–69.

Longstaff, Francis A, and Jiang Wang, 2012, Asset pricing and the credit market, *Review of Financial Studies* 25, 3169–3215.

Ludvigson, Sydney C., Sai Ma, and Serena Ng, 2021, Uncertainty and business cycles: Exogenous impulse or endogenous response?, *American Economic Journal: Macroeconomics* 13, 369–410.

Lustig, Hanno, Nikolai Roussanov, and Adrien Verdelhan, 2011, Common risk factors in currency markets, *Review of Financial Studies* 24, 3731–3777.

——— , 2014, Countercyclical currency risk premia, *Journal of Financial Economics* 111, 527–553.

Lustig, Hanno, Andreas Stathopoulos, and Adrien Verdelhan, 2019, The term structure of currency carry trade risk premia, *American Economic Review* 109, 4142–4177.

Markowitz, Harry M., 1976, Investment for the long run: New evidence for an old rule, *Journal of Finance* 31, 1273–1286.

Martin, Ian, 2012, On the valuation of long-dated assets, *Journal of Political Economy* 120, 346–358.

——— , 2017, What is the expected return on the market?, *Quarterly Journal of Economics* 132, 367–433.

——— , and Ran Shi, 2025, Forecasting crashes with a smile, *Working Paper*.

Martin, Ian WR, and Christian Wagner, 2019, What is the expected return on a stock?, *Journal of Finance* 74, 1887–1929.

Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf, 2012, Carry trades and global foreign exchange volatility, *Journal of Finance* 67, 681–718.

Nagel, Stefan, and Zhengyang Xu, 2023, Dynamics of subjective risk premia, *Journal of Financial Economics* 150, 103713.

Orłowski, Piotr, Valeri Sokolovski, and Erik Sverdrup, 2025, Constrained currency stochastic discount factors, *Working Paper*.

Owen, Guillermo, 1977, Values of games with a priori unions, in Rudolf Henn, and Otto Moeschlin, ed.: *Essays in Mathematical Economics and Game Theory* . pp. 76–88 (Springer: Berlin).

Pflueger, Carolin, Emil Siriwardane, and Adi Sunderam, 2020, Financial market risk perceptions and the macroeconomy, *Quarterly Journal of Economics* 135, 1443–1491.

Piatti, Ilaria, and Paul Whelan, 2026, Subjective risk premia in bond and FX markets, *Working Paper*.

Roll, Richard, 1973, Evidence on the "growth-optimum" model, *Journal of Finance* 28, 551–566.

Stavrakeva, Vania, and Jenny Tang, 2025, Individual beliefs, demand for currency, and exchange rate dynamics, *Working Paper* CEPR Discussion Paper 20321.

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

Weber, Michael, 2018, Cash flow duration and the term structure of equity returns, *Journal of Financial Economics* 128, 486–503.

Zviadadze, Irina, 2017, Term structure of consumption risk premia in the cross section of currency returns, *The Journal of Finance* 72, 1529–1566.

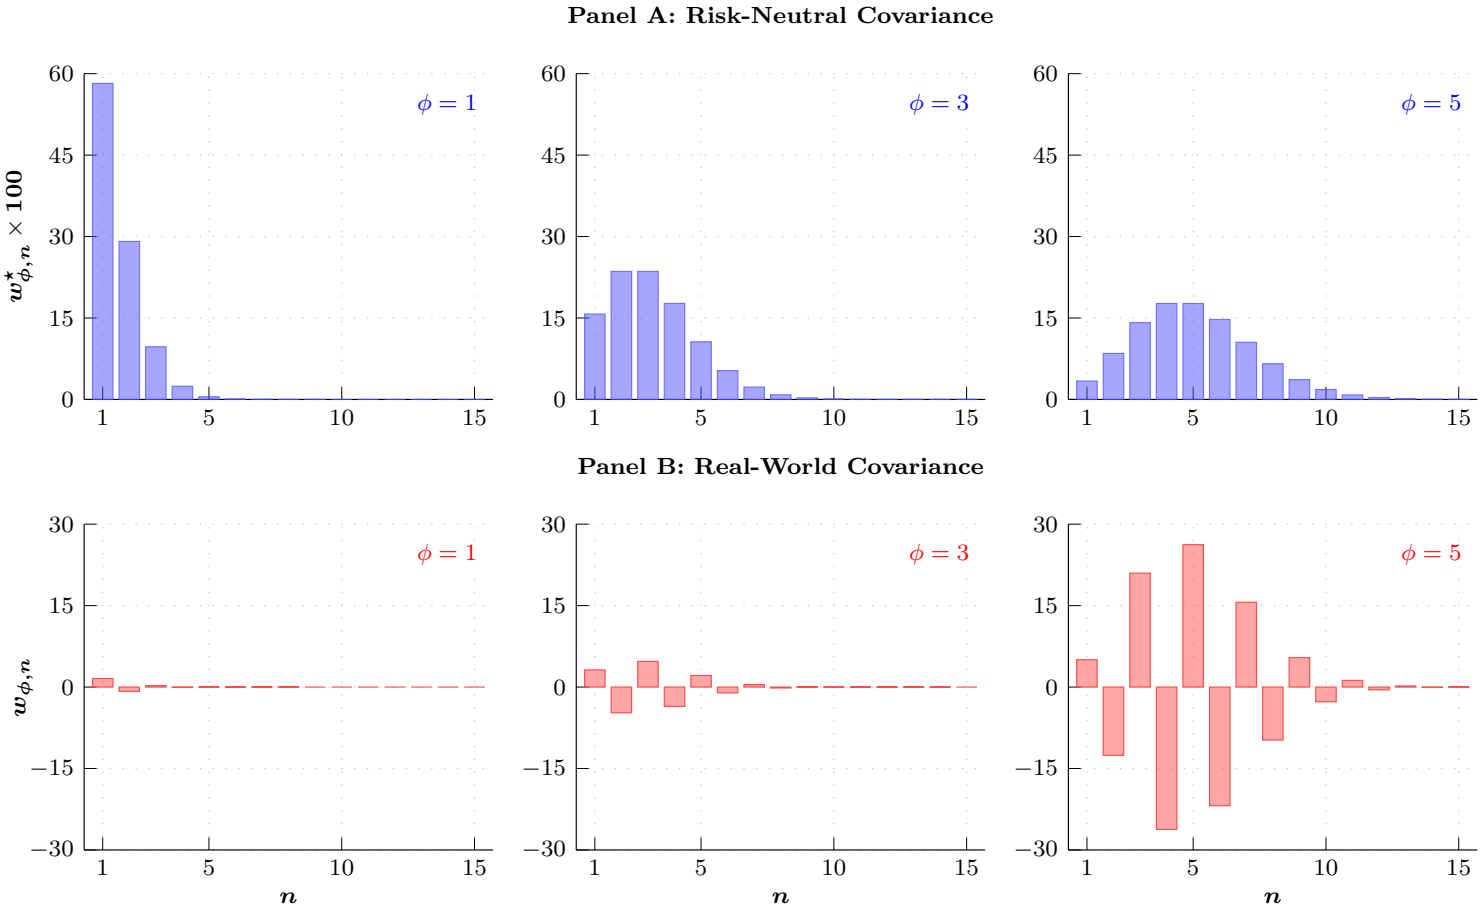


Figure 1. Higher-Order Risk and the Expected Risk Premium

This figure displays the weights associated with the higher-order risk decomposition of the expected risk premium for selected values of ϕ. Panel A plots the risk-neutral weights $w^{\star}_{\phi,n} = (e^{\phi} - 1)^{-1}\phi^n/n!$ from Equation (9) (scaled by 100 to ease comparison), whereas Panel B displays the real-world weights $w_{\phi,n} = (e^{-\phi} - 1)^{-1}(-\phi)^n/n!$ from Equation (A.7). The weights are associated with the risk-neutral and real-world covariances between the continuously compounded market return raised to power n and the exchange rate return. The risk-neutral weight takes its maximum value around $n = \phi$. The real-world weight switches signs between odd and even market return moments as FX investors require a positive risk premium when the covariance between exchange rate returns and odd (even) market return moments is positive (negative).

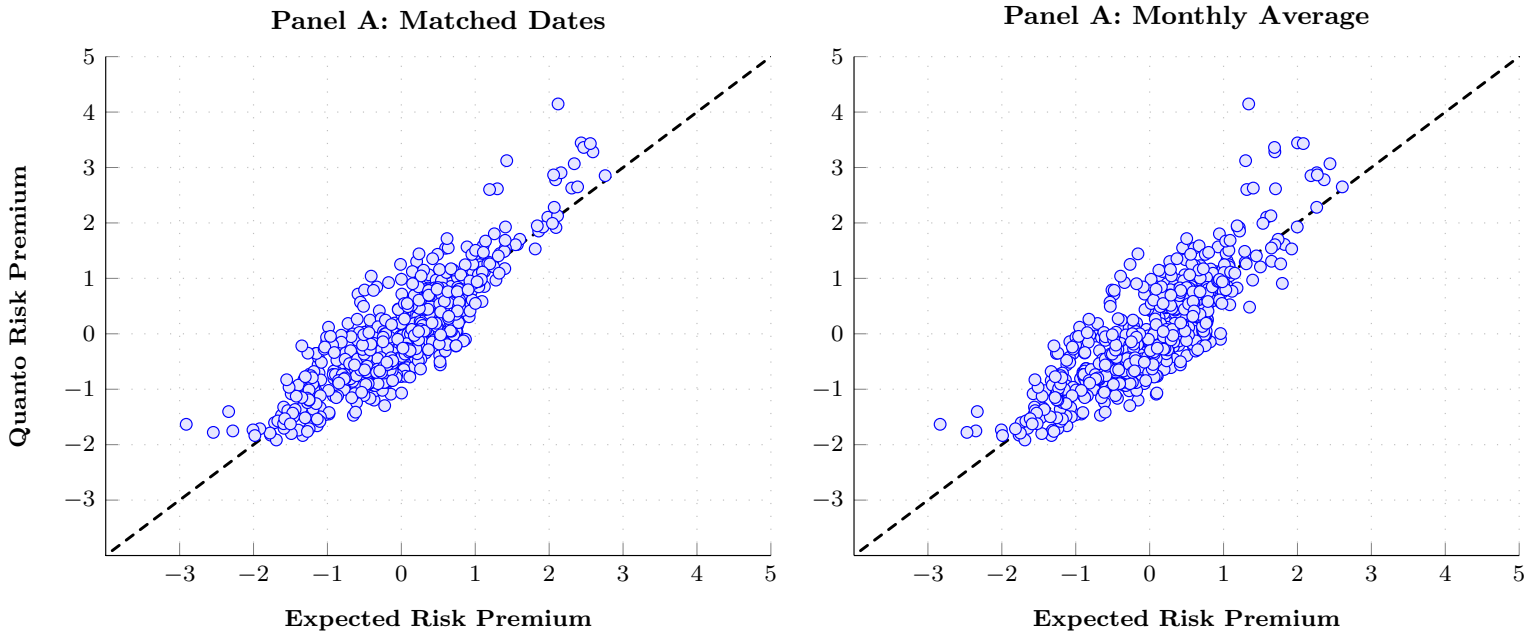


Figure 2. Quanto Risk Premium versus Expected Risk Premium

This figure displays scatter plots of the quanto risk premium (QRP) against the expected risk premium (ERP) for a maturity τ of two years. QRP is extracted from quanto contract prices and reflects the risk-neutral covariance between exchange rate movements and global equity market risk. ERP is constructed from SPX options, FX options, and backward-looking realized correlations between S&P 500 returns and FX returns, estimated over the window from $t - \tau$ to t. Since QRP is observed at a monthly frequency, Panel A synchronizes QRP and ERP by calendar date. In Panel B, for each QRP observation date, we compute the average daily ERP between the previous QRP date and the current QRP date. Both QRP and ERP are demeaned at the currency level. The sample covers monthly observations from December 2009 to October 2015 for 11 currencies against the US dollar at the two-year maturity. Data on QRP are from Kremens and Martin (2019); data used to construct ERP are from Bloomberg, JP Morgan, and OptionMetrics.

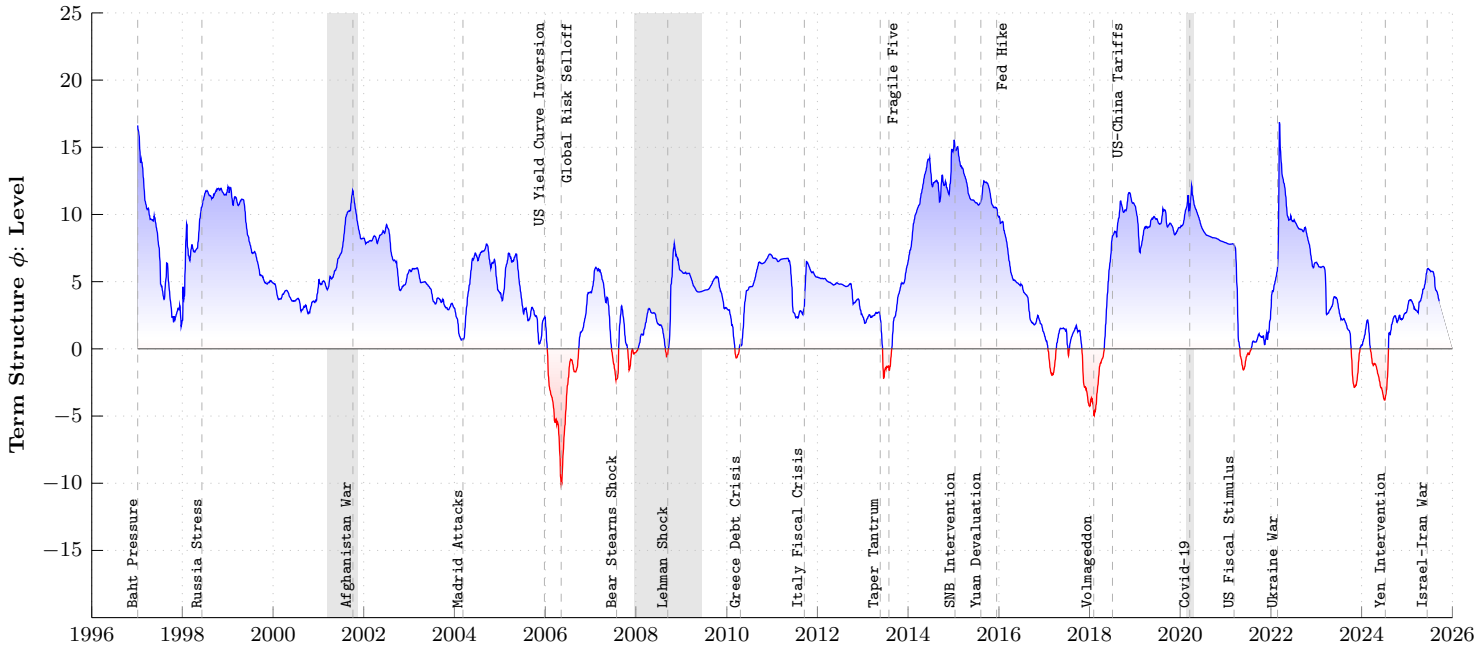


Figure 3. Level Component of the Term Structure of ϕ

This figure displays 252-day rolling estimates of the level component of the term structure of ϕ using the *Simple* method. Within each window, we estimate ϕ across different maturities by regressing the expected excess return for currency i on day t over the horizon τ on the corresponding expected risk premium computed under $\phi = 1$. The level is the average across maturity-specific estimates of ϕ. Gray shaded regions denote NBER-dated US recessions. The expected excess return is constructed from consensus exchange rate forecasts, spot exchange rates, and interest rate differentials, while the expected risk premium is derived from SPX options, FX options, and backward-looking realized correlations between S&P 500 index returns and FX spot returns, estimated over the window from $t - \tau$ to t. Positive estimates are shown in blue and negative estimates in red; the horizontal line marks zero. Vertical dashed lines indicate selected global macro, financial and geopolitical events. Daily forecasts are obtained by forward-filling monthly forecasts. The sample covers daily observations from January 1996 to September 2025 for 33 currencies in the broad sample (Panel A) and 20 currencies in the liquid sample (Panel B) against the US dollar, with maturities ranging from one month to two years. Data are from Bloomberg, Consensus Economics, JP Morgan, and OptionMetrics.

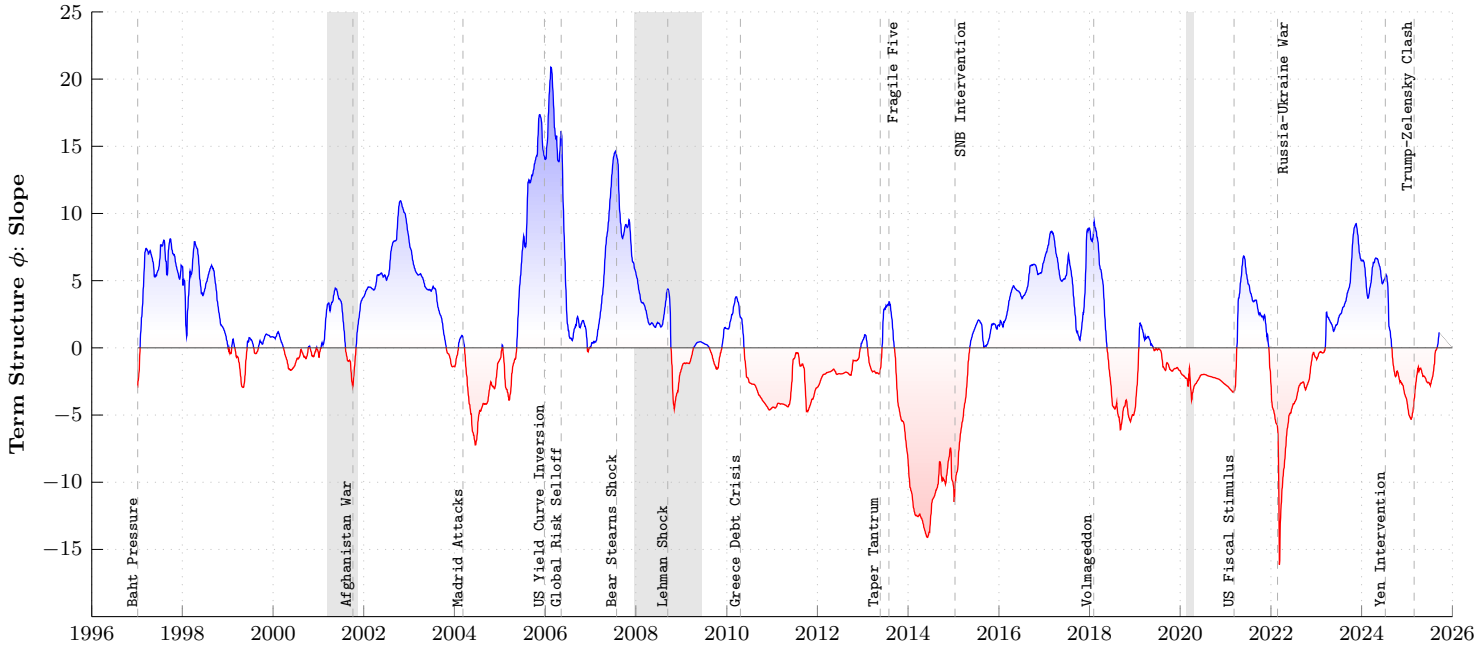


Figure 4. Slope Component of the Term Structure of ϕ

This figure displays 252-day rolling estimates of the slope component of the term structure of ϕ using the *Simple* method. Within each window, we estimate ϕ across different maturities by regressing the expected excess return for currency i on day t over the horizon τ on the corresponding expected risk premium computed under $\phi = 1$. The slope is the difference between the 2-year and 3-month estimates of ϕ. Gray shaded regions denote NBER-dated US recessions. The expected excess return is constructed from consensus exchange rate forecasts, spot exchange rates, and interest rate differentials, while the expected risk premium is derived from SPX options, FX options, and backward-looking realized correlations between S&P 500 index returns and FX spot returns, estimated over the window from $t - \tau$ to t. Positive estimates are shown in blue and negative estimates in red; the horizontal line marks zero. Vertical dashed lines indicate selected global macro, financial and geopolitical events. Daily forecasts are obtained by forward-filling monthly forecasts. The sample covers daily observations from January 1996 to September 2025 for 33 currencies in the broad sample (Panel A) and 20 currencies in the liquid sample (Panel B) against the US dollar, with maturities ranging from one month to two years. Data are from Bloomberg, Consensus Economics, JP Morgan, and OptionMetrics.

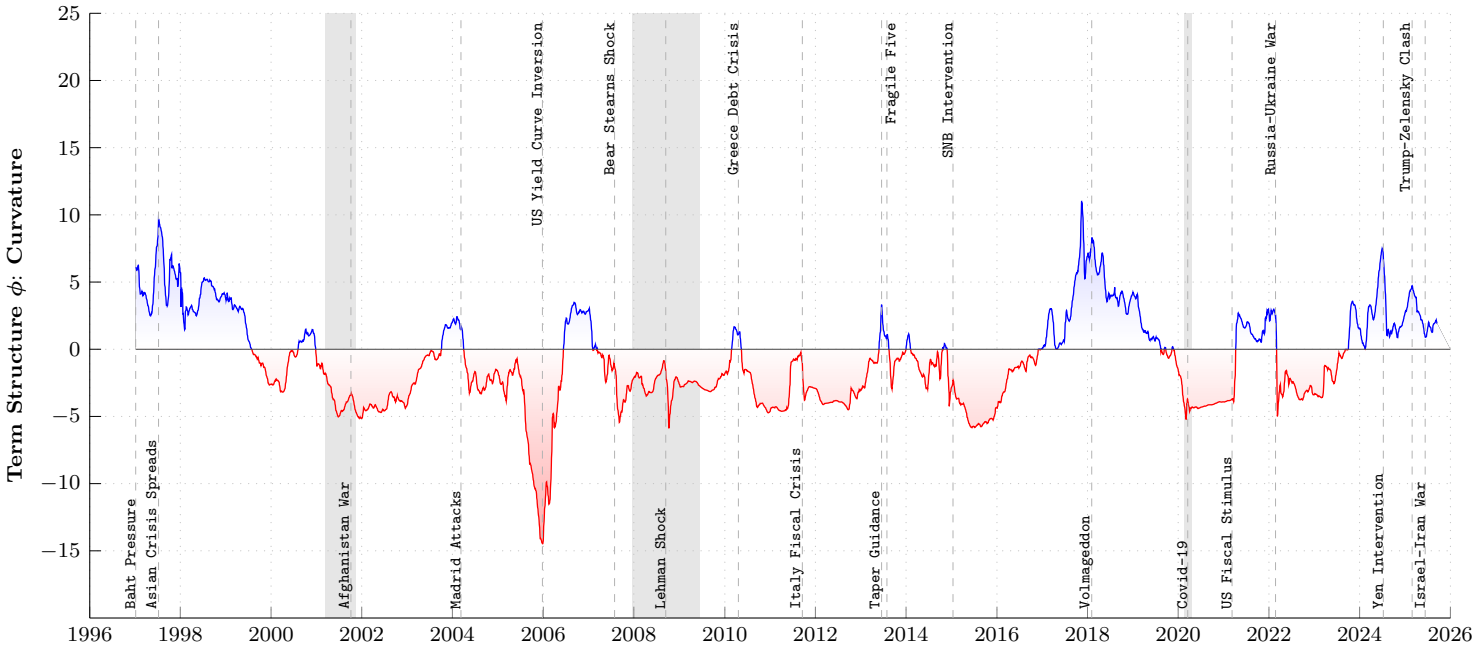


Figure 5. Curvature Component of the Term Structure of ϕ

This figure displays 252-day rolling estimates of the curvature component of the term structure of ϕ using the *Simple* method. Within each window, we estimate ϕ across different maturities by regressing the expected excess return for currency i on day t over the horizon τ on the corresponding expected risk premium computed under $\phi = 1$. The curvature is the average of the 3-month and 1-year estimates of ϕ minus the average of the 1-month and 2-year estimates of ϕ. Gray shaded regions denote NBER-dated US recessions. The expected excess return is constructed from consensus exchange rate forecasts, spot exchange rates, and interest rate differentials, while the expected risk premium is derived from SPX options, FX options, and backward-looking realized correlations between S&P 500 index returns and FX spot returns, estimated over the window from $t - \tau$ to t. Positive estimates are shown in blue and negative estimates in red; the horizontal line marks zero. Vertical dashed lines indicate selected global macro, financial and geopolitical events. Daily forecasts are obtained by forward-filling monthly forecasts. The sample covers daily observations from January 1996 to September 2025 for 33 currencies in the broad sample (Panel A) and 20 currencies in the liquid sample (Panel B) against the US dollar, with maturities ranging from one month to two years. Data are from Bloomberg, Consensus Economics, JP Morgan, and OptionMetrics.

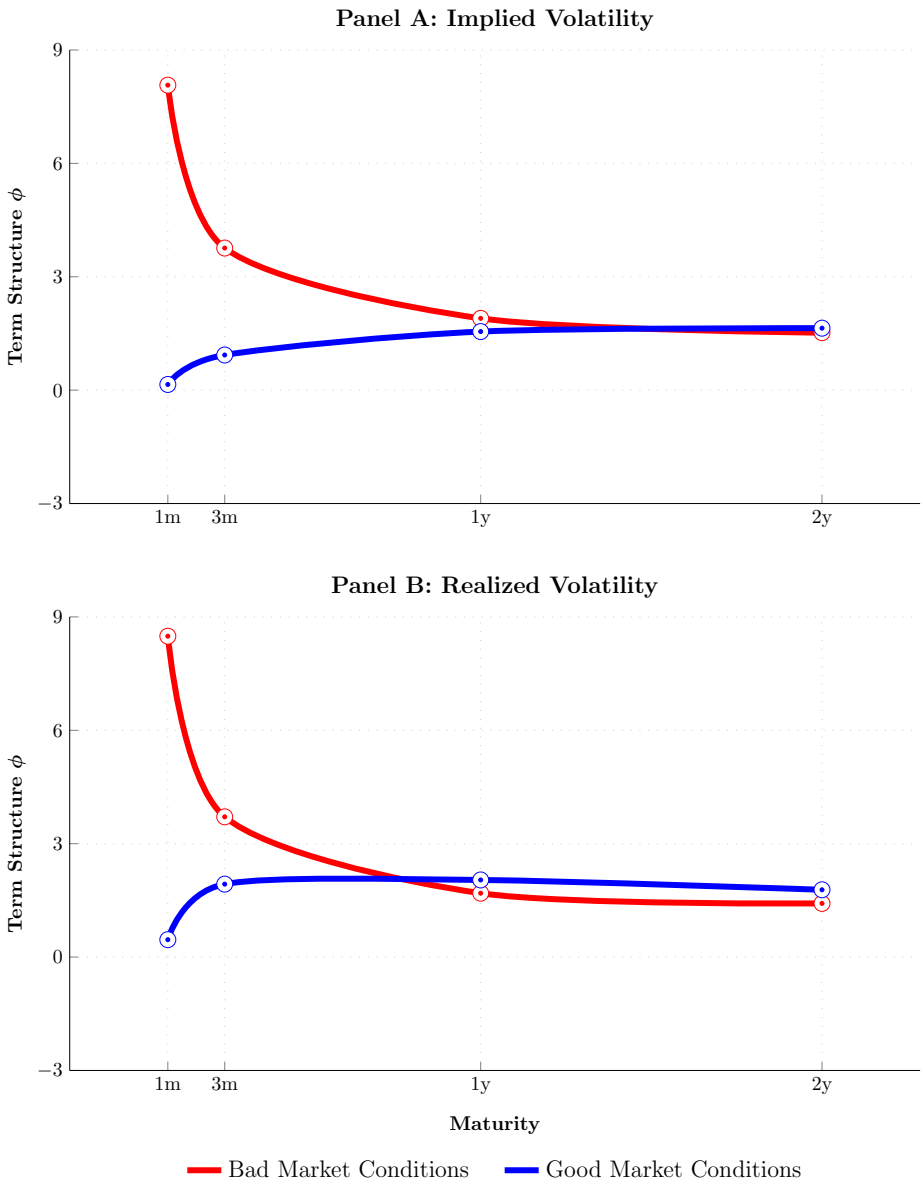


Figure 6. Term Structure of ϕ and Market Conditions

This figure displays the term structure of ϕ between *Good* and *Bad* market conditions. Good market conditions are defined as days on which volatility is below its sample mean, and bad market conditions as days on which volatility is at or above its sample mean. Panel A employs the one-month implied volatility of the S&P 500 index (VIX index), whereas Panel B uses the one-month realized volatility of the S&P 500 index (based on daily squared returns). We estimate ϕ using the *simple* method in Equation (15) for maturities between 1 month and 2 years. The sample covers daily observations from January 1996 to September 2025 for 33 currencies against the US dollar, with maturities ranging from one month to two years. Data are from Bloomberg, Consensus Economics, JP Morgan, and OptionMetrics.

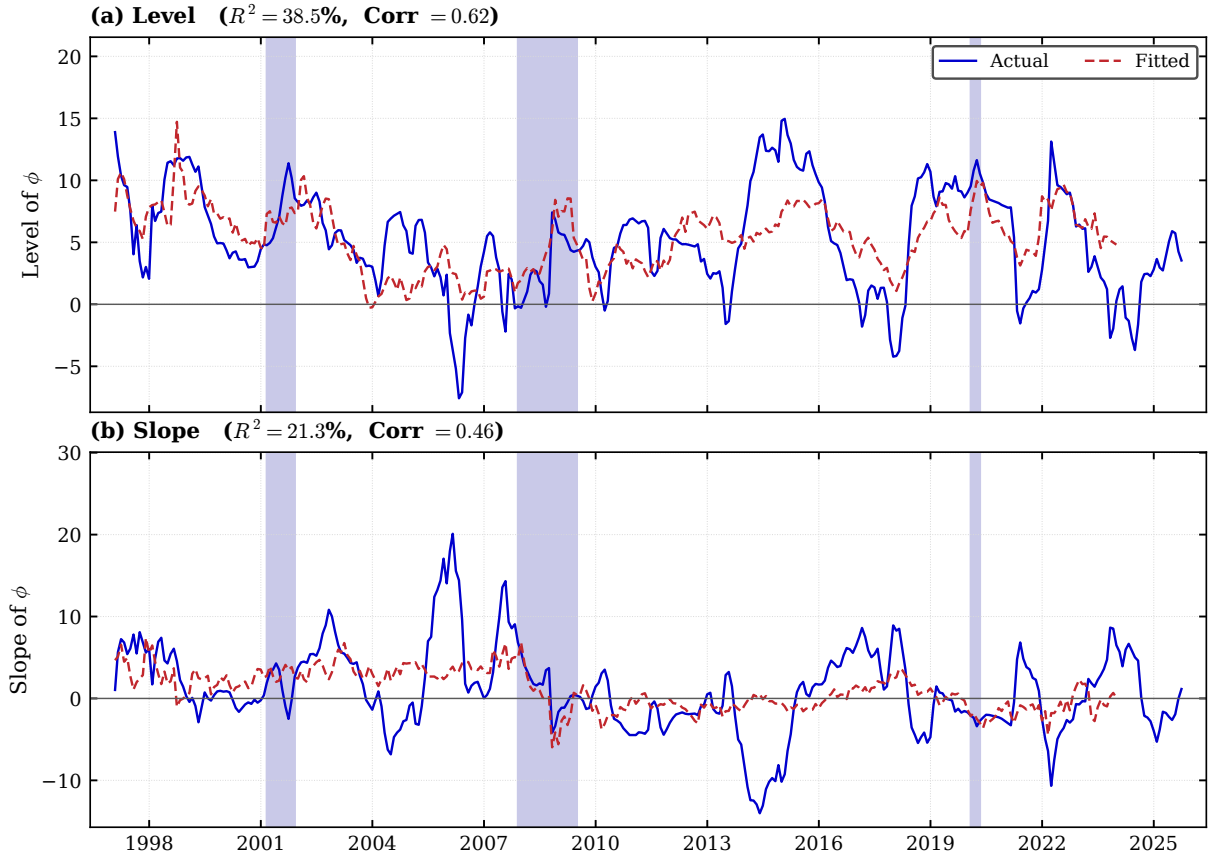


Figure 7. The Level and Slope of ϕ and Their Determinants

This figure displays the level and the slope of the ϕ term structure against their fitted values. Panel A plots the level and Panel B the slope, $\phi(2y) - \phi(3m)$, each against the fitted value from a joint regression on the eleven drivers of Table 7, estimated in levels at the monthly frequency and expressed in units of ϕ. The two regressions use the same drivers, so that the panels are directly comparable. The level factor is the average of the four maturity-specific estimates and the slope is the difference between the two-year and the three-month estimate. Gray shaded regions denote NBER-dated US recessions. The sample covers monthly observations, at month-end, from January 1997 to September 2025. Data are from Bloomberg, Consensus Economics, JP Morgan, and OptionMetrics.

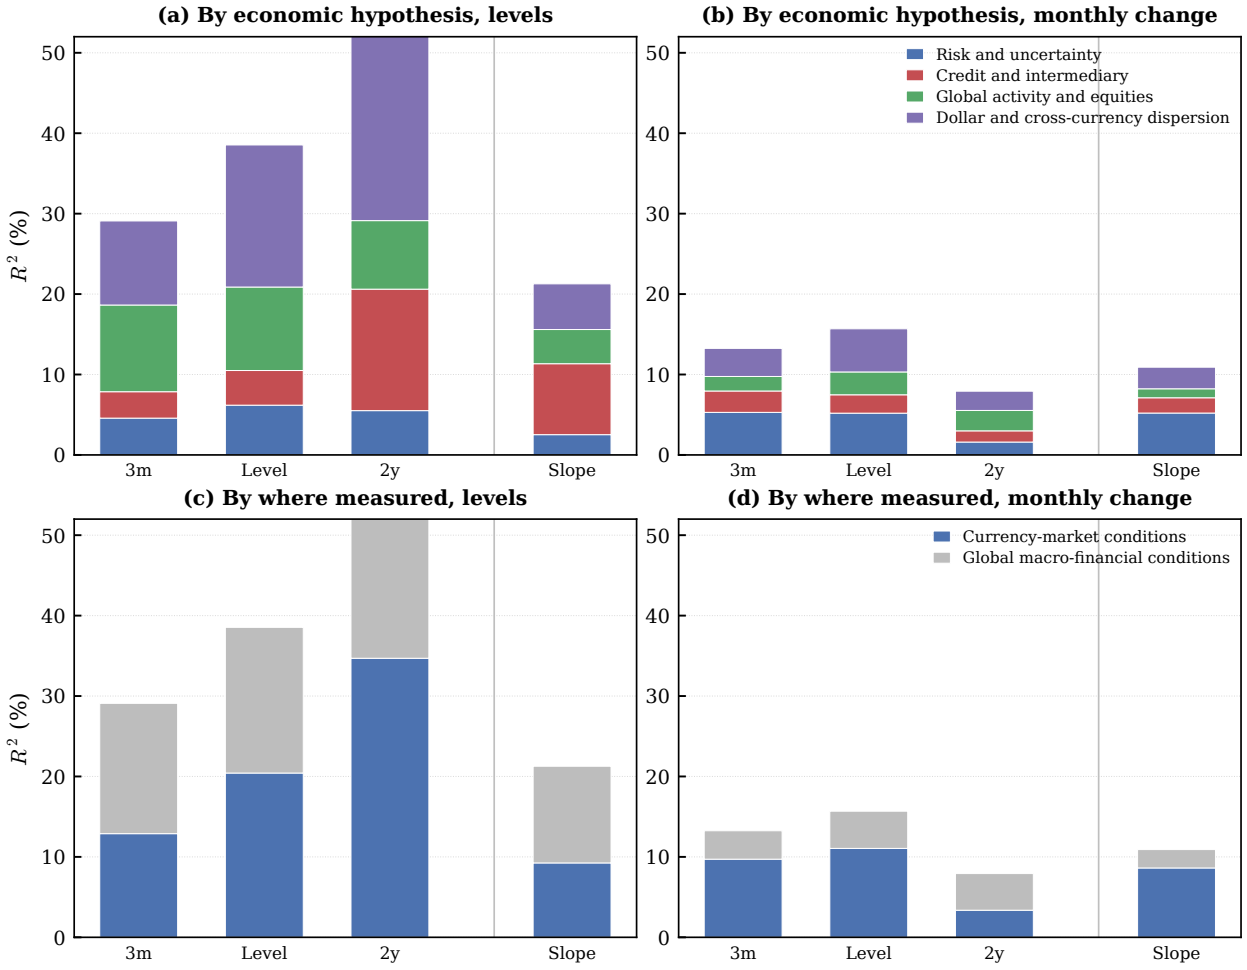


Figure 8. Decomposition of the Term Structure of ϕ across Drivers

This figure displays the Owen decomposition of the R^2 from a joint regression of ϕ on the eleven drivers of Table 7, expressed in percentage points of R^2. Bars are shown for the three-month estimate, the level factor, and the two-year estimate, which bracket the term structure, and for the slope, $\phi(2y) - \phi(3m)$, which is set apart by a rule. The height of each bar is the total explanatory power for that object and each segment is one set's share of it. The two rows apply the same regression and differ only in how the drivers are grouped. The top row groups them by economic hypothesis, following the four topic groups of Table 7. The bottom row groups the same series by where they are measured: five taken from currency markets, namely the option-implied vol of vol, the cross-currency basis, the two dispersion measures, and the dollar, against six global macro-financial series. Panels A and C use levels, whereas Panels B and D use monthly changes. The Owen value averages each set's marginal contribution over all orderings of the sets and each driver's over all orderings within its set, is computed exactly over all 2^{11} subsets, and sums to the R^2 of the full regression. The sample covers monthly observations, at month-end, from January 1997 to September 2025. Data are from Bloomberg, Consensus Economics, JP Morgan, and OptionMetrics.

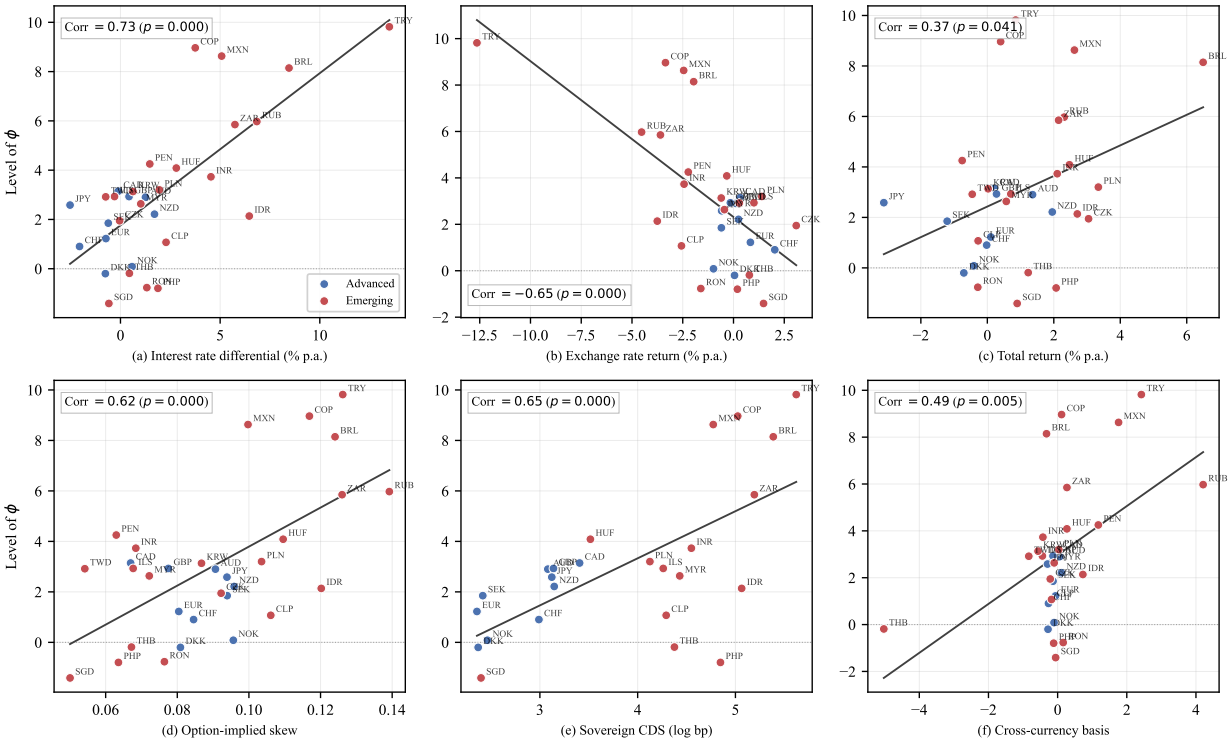


Figure 9. The Level of ϕ across Currencies

This figure displays the currency-level estimate of ϕ against six currency characteristics. The variable ϕ is the currency-level level factor, defined as the average of a currency's maturity-specific 252-day rolling estimates, averaged over the windows ending between January 1997 and September 2025. Characteristics are currency means over the same sample, taken from the currency-day control panel: the interest rate differential, signed so that a high-yield currency is positive (Panel A); the realized exchange rate return against the US dollar, positive when the currency appreciates (Panel B); the total return of a dollar-funded long position, which is the sum of the two preceding panels (Panel C); the option-implied skew (Panel D); the median five-year sovereign CDS par spread in log basis points (Panel E); and the cross-currency basis (Panel F). The interest differential and both returns are realized over the horizon of each contract and annualized by that horizon before averaging across maturities, so that Panels A and B sum to Panel C for every currency. Blue markers denote currencies classified as developed and red markers those classified as emerging or frontier. Solid lines are cross-sectional least-squares fits, and *Corr* is the Pearson correlation between the characteristic and ϕ across currencies, with its p-value in parentheses. Table 8 reports the same correlations together with their Spearman counterparts. Hong Kong and China are excluded because a currency peg fixes the interest differential and the exchange rate return by construction. The sample covers daily observations from January 1997 to September 2025 for 31 currencies against the US dollar, with maturities ranging from one month to two years. Data are from Bloomberg, Consensus Economics, JP Morgan, and OptionMetrics.

Table 1. Descriptive Statistics

This table reports descriptive statistics for the expected exchange rate return EFX, expected excess return ERX, and expected risk premium ERP^{ϕ} for selected $\phi \in \{1, 2, 3\}$, all measured for a given maturity τ. Expected exchange rate returns are constructed from consensus forecasts and spot exchange rates, forecast dispersion is measured as the log difference between the highest and lowest exchange rate forecasts, expected excess returns combine expected exchange rate returns with interest rate differentials, and expected risk premia are inferred from SPX options, FX options, and backward-looking realized correlations between S&P 500 index and FX returns, estimated over a window of length τ for each value of ϕ. We report the sample mean, standard deviation (*sdev*), skewness (*skew*), kurtosis (*kurt*), and first-order serial correlation (AC_1). Returns are expressed in percentage per annum. Daily forecasts are obtained by forward-filling monthly forecasts. The sample covers daily observations from January 1996 to September 2025 for 33 currencies in the broad sample (Panel A) and 20 currencies in the liquid sample (Panel B) against the US dollar, with maturities ranging from one month to two years. Data are from Bloomberg, Consensus Economics, JP Morgan, and OptionMetrics.

	maturity	*mean*	*sdev*	*skew*	*kurt*	AC_1	*mean*	*sdev*	*skew*	*kurt*	AC_1
		Panel A: Broad Sample					Panel B: Liquid Sample				
EFX	1 month	−0.15	10.81	3.37	96.92	0.95	−0.54	10.37	0.98	8.84	0.95
	3 months	−0.27	6.90	2.45	66.11	0.96	−0.32	6.75	0.64	6.40	0.96
	1 year	0.24	4.81	0.93	24.09	0.98	0.24	4.98	−0.14	2.61	0.98
	2 years	0.22	4.54	0.13	9.06	0.99	0.22	4.87	−0.37	2.34	0.99
ERX	1 month	1.90	10.82	3.83	102.82	0.95	1.22	10.34	1.20	9.76	0.95
	3 months	1.83	6.97	3.63	77.96	0.96	1.45	6.69	1.29	9.07	0.96
	1 year	2.08	5.21	2.62	31.38	0.98	1.80	5.09	1.16	4.37	0.98
	2 years	1.84	5.27	1.78	13.48	0.99	1.56	5.16	0.99	2.56	0.99
ERP^1	1 month	0.30	0.28	3.73	37.00	0.96	0.32	0.30	3.47	32.76	0.97
	3 months	0.30	0.36	2.39	14.16	0.99	0.31	0.40	2.08	11.41	0.99
	1 year	0.30	0.54	1.33	4.42	1.00	0.31	0.61	1.12	3.28	1.00
	2 years	0.29	0.68	1.03	3.30	1.00	0.29	0.76	0.87	2.49	1.00
ERP^2	1 month	0.58	0.54	3.56	34.32	0.96	0.61	0.58	3.32	30.37	0.97
	3 months	0.57	0.68	2.26	12.86	0.99	0.59	0.75	1.98	10.38	0.99
	1 year	0.53	0.98	1.26	4.14	1.00	0.55	1.10	1.06	3.05	1.00
	2 years	0.51	1.19	0.97	3.13	1.00	0.52	1.34	0.82	2.34	1.00
ERP^3	1 month	0.83	0.77	3.40	31.88	0.96	0.88	0.83	3.17	28.17	0.97
	3 months	0.80	0.95	2.16	11.79	0.99	0.84	1.05	1.88	9.52	0.99
	1 year	0.73	1.33	1.21	3.89	1.00	0.75	1.49	1.01	2.86	1.00
	2 years	0.68	1.58	0.94	2.99	1.00	0.68	1.77	0.79	2.22	1.00

Table 2. Panel Regression Estimates of ϕ

This table presents estimates based on the following specification

$$ERX_{i,t,T} = \alpha_i + \alpha_t + \alpha_\tau + \beta_\phi ERP^\phi_{i,t,T} + \varepsilon_{i,t,T},$$

where ERX is the expected excess return for currency i on day t over the horizon $\tau = T - t$ and ERP^ϕ is the corresponding expected risk premium for a given ϕ. The terms α_i, α_t, and α_τ denote currency, calendar-time, and maturity fixed effects (*fe*), respectively. ERX is constructed from consensus forecasts, spot exchange rates, and interest rate differentials, while ERP^ϕ is derived from SPX options, FX options, and backward-looking realized correlations between S&P 500 index returns and FX spot returns, estimated over the window from $t - \tau$ to t for different levels of ϕ. The column *All* pools observations across all maturities, whereas the remaining columns report estimates for a specific maturity τ. *Simple* reports estimates of ϕ, obtained from estimates of β after setting $\phi = 1$, as implied by Proposition 2. *Search* reports estimates obtained when the estimate of β is closest to one, using a fine grid search, as implied by Proposition 1. Standard errors, reported in parentheses, are clustered by currency and calendar-time dimensions. Standard errors for ϕ are obtained via the delta method. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. Daily forecasts are obtained by forward-filling monthly forecasts. The sample covers daily observations from January 1996 to September 2025 for 33 currencies in the broad sample (Panel A) and 20 currencies in the liquid sample (Panel B) against the US dollar, with maturities ranging from one month to two years. Data are from Bloomberg, Consensus Economics, JP Morgan, and OptionMetrics.

		Panel A: Broad Sample					Panel B: Liquid Sample				
		All	1 month	3 months	1 year	2 years	All	1 month	3 months	1 year	2 years
Simple	β	6.01***	7.25***	3.59***	2.10***	1.75***	6.11***	7.17***	3.56***	2.34***	1.93***
		(0.80)	(1.04)	(0.73)	(0.50)	(0.47)	(0.82)	(0.91)	(0.60)	(0.42)	(0.43)
	R^2	18.88	33.37	33.48	44.31	52.04	23.52	42.75	40.91	49.91	56.59
	N	824,750	204,355	206,452	206,798	207,145	538,076	134,519	134,519	134,519	134,519
Search	ϕ	7.17***	8.01***	4.02***	2.40***	1.99**	7.27***	8.01***	3.98***	2.75***	2.27***
		(1.20)	(1.41)	(1.00)	(0.78)	(0.75)	(1.21)	(1.32)	(0.85)	(0.68)	(0.70)
	R^2	18.42	32.92	33.29	44.27	52.04	22.96	42.29	40.70	49.84	56.57
	N	824,750	204,355	206,452	206,798	207,145	538,076	134,519	134,519	134,519	134,519
	Currency fe	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Maturity fe	✓					✓				
	Time fe	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Table 3. Panel Regression Estimates of ϕ: Adding High and Low Forecasts

This table presents estimates based on the following specification

$$ERX_{i,f,t,T} = \alpha_i + \alpha_f + \alpha_t + \alpha_\tau + \beta_\phi ERP^{\phi}_{i,f,t,T} + \varepsilon_{i,f,t,T},$$

where ERX is the expected excess return for currency i on day t over the horizon $\tau = T - t$ and ERP^{ϕ} is the corresponding expected risk premium for a given ϕ. The terms α_i, α_f, α_t, and α_τ denote currency, forecast-type, calendar-time, and maturity fixed effects (*fe*), respectively. ERX is constructed from exchange rate forecasts (high, low, and consensus), spot exchange rates, and interest rate differentials, while ERP^{ϕ} is derived from SPX options, FX options, and backward-looking realized correlations between S&P 500 index returns and FX spot returns, estimated over the window from $t - \tau$ to t for different levels of ϕ. The column *All* pools observations across all maturities, whereas the remaining columns report estimates for a specific maturity τ. *Simple* reports estimates of ϕ, obtained from estimates of β_ϕ after setting $\phi = 1$, as implied by Proposition 2. *Search* reports estimates obtained when the estimate of β_ϕ is closest to one, using a fine grid search, as implied by Proposition 1. Standard errors, reported in parentheses, are clustered by currency and calendar-time dimensions. Standard errors for ϕ are obtained via the delta method. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. Daily forecasts are obtained by forward-filling monthly forecasts. The sample covers daily observations from January 1996 to September 2025 for 33 currencies in the broad sample (Panel A) and 20 currencies in the liquid sample (Panel B) against the US dollar, with maturities ranging from one month to two years. Data are from Bloomberg, Consensus Economics, JP Morgan, and OptionMetrics.

		Panel A: Broad Sample					Panel B: Liquid Sample				
		All	1 month	3 months	1 year	2 years	All	1 month	3 months	1 year	2 years
Simple	β	6.25***	7.52***	3.82***	2.08***	1.82***	6.39***	7.51***	3.83***	2.26***	1.96***
		(0.83)	(1.05)	(0.79)	(0.54)	(0.53)	(0.84)	(0.98)	(0.67)	(0.46)	(0.44)
	R^2	28.97	49.24	53.89	61.39	58.64	32.98	56.73	60.17	66.70	64.07
	N	2,461,920	613,065	615,246	616,284	617,325	1,614,228	403,557	403,557	403,557	403,557
Search	ϕ	7.54***	8.42***	4.36***	2.38***	2.09**	7.70***	8.55***	4.41***	2.65***	2.31***
		(1.27)	(1.54)	(1.18)	(0.83)	(0.83)	(1.30)	(1.50)	(0.99)	(0.72)	(0.74)
	R^2	28.80	49.07	53.83	61.38	58.64	32.80	56.57	60.11	66.69	64.06
	N	2,461,920	613,065	615,246	616,284	617,325	1,614,228	403,557	403,557	403,557	403,557
	Currency fe	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Forecast fe	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Maturity fe	✓					✓				
	Time fe	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Table 4. Panel Regression Estimates of ϕ: Using Expected FX Returns

This table presents estimates based on the following specification

$$EFX_{i,t,T} = \alpha_i + \alpha_t + \alpha_\tau + \beta_\phi ERP^{\phi}_{i,t,T} + \gamma IRD_{i,t,T} + \varepsilon_{i,t,T},$$

where EFX is the expected exchange rate return for currency i on day t over the horizon $\tau = T - t$, ERP^{ϕ} is the corresponding expected risk premium for a given ϕ, and IRD is the interest rate differential relative to the US dollar. The terms α_i, α_t, and α_τ denote currency, calendar-time, and maturity fixed effects (*fe*), respectively. EFX is constructed from consensus forecasts and spot exchange rates, while ERP^{ϕ} is derived from SPX options, FX options, and backward-looking realized correlations between S&P 500 index returns and FX spot returns, estimated over the window from $t - \tau$ to t for different levels of ϕ. The column *All* pools observations across all maturities, whereas the remaining columns report estimates for a specific maturity τ. *Simple* reports estimates of ϕ, obtained from estimates of β_ϕ after setting $\phi = 1$, as implied by Proposition 2. *Search* reports estimates obtained when the estimate of β_ϕ is closest to one, using a fine grid search, as implied by Proposition 1. Standard errors, reported in parentheses, are clustered by currency and calendar-time dimensions. Standard errors for ϕ are obtained via the delta method. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. Daily forecasts are obtained by forward-filling monthly forecasts. The sample covers daily observations from January 1996 to September 2025 for 33 currencies in the broad sample (Panel A) and 20 currencies in the liquid sample (Panel B) against the US dollar, with maturities ranging from one month to two years. Data are from Bloomberg, Consensus Economics, JP Morgan, and OptionMetrics.

		Panel A: Broad Sample					Panel B: Liquid Sample				
		All	1 month	3 months	1 year	2 years	All	1 month	3 months	1 year	2 years
Simple	β	5.97***	7.25***	3.49***	2.06***	1.78***	6.11***	7.18***	3.56***	2.42***	2.14***
		(0.81)	(1.08)	(0.70)	(0.50)	(0.47)	(0.81)	(0.93)	(0.59)	(0.43)	(0.42)
	γ	0.75***	1.02***	0.69***	0.61***	0.52***	0.99***	1.35***	0.91***	0.78***	0.69***
		(0.20)	(0.34)	(0.18)	(0.10)	(0.08)	(0.16)	(0.31)	(0.15)	(0.09)	(0.10)
	R^2	18.32	33.25	32.84	38.50	43.65	23.96	43.18	41.98	48.54	53.65
	N	824,750	204,355	206,452	206,798	207,145	538,076	134,519	134,519	134,519	134,519
Search	ϕ	7.10***	8.01***	3.88***	2.34***	2.03***	7.27***	8.03***	3.99***	2.89***	2.62***
		(1.20)	(1.47)	(0.99)	(0.74)	(0.72)	(1.20)	(1.34)	(0.83)	(0.72)	(0.75)
	γ	0.75***	1.01***	0.69***	0.61***	0.52***	0.99***	1.35***	0.91***	0.78***	0.69***
		(0.20)	(0.34)	(0.18)	(0.10)	(0.08)	(0.16)	(0.31)	(0.15)	(0.09)	(0.10)
	R^2	17.86	32.79	32.65	38.45	43.64	23.41	42.72	41.77	48.46	53.61
	N	824,750	204,355	206,452	206,798	207,145	538,076	134,519	134,519	134,519	134,519
	Currency fe	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Maturity fe	✓					✓				
	Time fe	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Table 5. Panel Regression Estimates of ϕ: Implied Volatility and Market Conditions

This table presents estimates based on the following specification

$$ERX_{i,t,T} = \alpha_i + \alpha_t + \alpha_\tau + \beta_\phi ERP^{\phi}_{i,t,T} + \varepsilon_{i,t,T},$$

where ERX is the expected excess return for currency i on day t over the horizon $\tau = T - t$ and ERP^{ϕ} is the corresponding expected risk premium, evaluated at $\phi = 1$. The terms α_i, α_t, and α_τ denote currency, calendar-time, and maturity fixed effects (*fe*), respectively. ERX is constructed from consensus forecasts, spot exchange rates, and interest rate differentials, while ERP^{ϕ} is derived from SPX options, FX options, and backward-looking realized correlations between S&P 500 index returns and FX spot returns. The column *All* pools observations across all maturities, whereas the remaining columns report estimates for a specific maturity τ. The top panel reports estimates for *Good* market conditions, defined as days on which the VIX index is below its sample mean. The bottom panel reports estimates for *Bad* market conditions, defined as days on which the VIX index is at or above its sample mean. This table reports *Simple* estimates of ϕ, obtained from estimates of β after setting $\phi = 1$, as implied by Proposition 2. The corresponding *Search* estimates, obtained when the estimate of β is closest to one, using a fine grid search, are displayed in Internet Appendix Table A.4. Standard errors, reported in parentheses, are clustered by currency and calendar-time dimensions. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. Daily forecasts are obtained by forward-filling monthly forecasts. The sample covers daily observations from January 1996 to September 2025 for 33 currencies in the broad sample (Panel A) and 20 currencies in the liquid sample (Panel B) against the US dollar, with maturities ranging from one month to two years. Data are from Bloomberg, Consensus Economics, JP Morgan, and OptionMetrics.

		Panel A: Broad Sample					Panel B: Liquid Sample				
		All	1 month	3 months	1 year	2 years	All	1 month	3 months	1 year	2 years
Good	β	1.29*	0.15	0.93	1.55*	1.64**	2.53***	1.79	1.42	2.10*	2.19**
		(0.75)	(1.40)	(0.89)	(0.77)	(0.68)	(0.71)	(1.55)	(1.03)	(1.04)	(0.82)
	R^2	16.56	31.32	31.87	44.38	52.66	18.89	36.75	36.11	46.67	54.42
	N	509,847	126,433	127,596	127,876	127,942	327,388	81,847	81,847	81,847	81,847
Bad	β	6.93***	8.07***	3.76***	1.90***	1.52***	6.58***	7.25***	3.32***	1.86***	1.42***
		(1.03)	(1.19)	(0.93)	(0.59)	(0.54)	(1.08)	(0.95)	(0.79)	(0.44)	(0.42)
	R^2	21.13	35.60	35.48	45.77	53.26	27.79	48.47	45.93	54.29	60.45
	N	314,903	77,922	78,856	78,922	79,203	210,688	52,672	52,672	52,672	52,672
	Currency fe	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Maturity fe	✓					✓				
	Time fe	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Table 6. Panel Regression Estimates of ϕ: Realized Volatility and Market Conditions

This table presents estimates based on the following specification

$$ERX_{i,t,T} = \alpha_i + \alpha_t + \alpha_\tau + \beta_\phi ERP^{\phi}_{i,t,T} + \varepsilon_{i,t,T},$$

where ERX is the expected excess return for currency i on day t over the horizon $\tau = T - t$ and ERP^{ϕ} is the corresponding expected risk premium, evaluated at $\phi = 1$. The terms α_i, α_t, and α_τ denote currency, calendar-time, and maturity fixed effects (*fe*), respectively. ERX is constructed from consensus forecasts, spot exchange rates, and interest rate differentials, while ERP^{ϕ} is derived from SPX options, FX options, and backward-looking realized correlations between S&P 500 index returns and FX spot returns. The column *All* pools observations across all maturities, whereas the remaining columns report estimates for a specific maturity τ. The top panel reports estimates for *Good* market conditions, defined as days on which the one-month realized volatility of the S&P 500 index is below its sample mean. The bottom panel reports estimates for *Bad* market conditions, defined as days on which the one-month realized volatility of the S&P 500 index is at or above its sample mean. This table reports *Simple* estimates of ϕ, obtained from estimates of β after setting $\phi = 1$, as implied by Proposition 2. The corresponding *Search* estimates, obtained when the estimate of β is closest to one, using a fine grid search, are displayed in Internet Appendix Table A.5. Standard errors, reported in parentheses, are clustered by currency and calendar-time dimensions. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. Daily forecasts are obtained by forward-filling monthly forecasts. The sample covers daily observations from January 1996 to September 2025 for 33 currencies in the broad sample (Panel A) and 20 currencies in the liquid sample (Panel B) against the US dollar, with maturities ranging from one month to two years. Data are from Bloomberg, Consensus Economics, JP Morgan, and OptionMetrics.

		Panel A: Broad Sample					Panel B: Liquid Sample				
		All	1 month	3 months	1 year	2 years	All	1 month	3 months	1 year	2 years
Good	β	1.87**	0.46	1.93	2.04**	1.78**	2.80***	1.96	1.73*	2.29**	2.19***
		(0.77)	(1.16)	(1.38)	(0.84)	(0.65)	(0.65)	(1.27)	(0.88)	(0.84)	(0.68)
	R^2	15.94	29.82	30.54	42.74	50.86	19.35	37.55	36.77	46.43	53.53
	N	531,012	131,633	132,879	133,188	133,312	343,660	85,915	85,915	85,915	85,915
Bad	β	7.17***	8.49***	3.71***	1.69***	1.42***	6.80***	7.57***	3.48***	1.81***	1.37***
		(1.03)	(1.21)	(0.82)	(0.47)	(0.47)	(1.11)	(0.98)	(0.81)	(0.43)	(0.40)
	R^2	22.46	37.94	37.11	46.88	54.23	27.95	48.54	45.69	54.20	60.53
	N	293,738	72,722	73,573	73,610	73,833	194,416	48,604	48,604	48,604	48,604
	Currency fe	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Maturity fe	✓					✓				
	Time fe	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Table 7. The Level of ϕ and Financial Conditions

This table reports estimates based on the following specification

$$\Delta\phi_t = \alpha + \beta\, z(\Delta Z_t) + u_t,$$

where ϕ is the level factor of the ϕ term structure, defined as the average of the four maturity-specific estimates, and Z is a candidate driver, standardized so that β measures the association with a one-standard-deviation change in the driver, expressed in units of ϕ. The first four columns report this specification in monthly changes. The column *Slope* repeats it with the monthly change in the slope of the ϕ term structure, $\phi(2y) - \phi(3m)$, as the dependent variable. The remaining columns report each driver's Owen contribution, in percentage points of R^2, from a single regression on all eleven drivers, with the four topic groups as the coalition structure and group totals in italics, for the level and the slope in monthly changes and in levels. The average basis, the average implied vol of vol, and the two dispersion measures are cross-sectional aggregates of the currency-day control panel: the basis and the vol of vol are means across currencies, and the dispersion measures are cross-sectional standard deviations of realized volatility and of the interest rate differential. The remaining drivers are macro-financial series. The last two rows report the joint R^2 over the two halves of the sample. t-statistics, reported in parentheses, are based on Newey-West standard errors with 12 lags, and R^2 is reported in percent. The sample covers monthly observations, at month-end, from January 1997 to September 2025, with $N = 329$ in changes and 341 in levels. Data are from Bloomberg, Consensus Economics, JP Morgan, OptionMetrics, FRED, the Chicago Fed, the Federal Reserve Board's Senior Loan Officer Survey, the OECD, Jurado, Ludvigson, and Ng (2015), Ludvigson, Ma, and Ng (2021), and He, Kelly, and Manela (2017).

	Univariate, AR Innovations				Owen, Innovations		Owen, Levels	
	b	t	R^2	Slope t	Level	Slope	Level	Slope
Risk and uncertainty								
VIX	0.23	(2.24)	2.6	−1.75	0.4	0.3	1.2	0.3
Financial uncertainty	0.25	(3.38)	2.9	−1.70	0.6	0.2	2.1	1.4
Implied vol of vol	0.44	(4.96)	9.5	−4.90	4.2	4.7	2.9	0.8
group total					5.2	5.2	6.2	2.5
Credit and intermediary								
BAA–10Y spread	0.28	(2.78)	3.7	−2.52	0.8	0.7	1.3	1.7
Intermediary capital	−0.23	(−2.27)	2.6	1.57	0.5	0.2	2.2	3.8
CIP basis	0.20	(1.51)	1.9	−1.39	1.0	0.9	0.8	3.4
group total					2.3	1.9	4.3	8.8
Global activity and equities								
Global real activity	−0.16	(−1.64)	1.2	0.66	0.9	0.1	9.7	1.0
Equity market	−0.32	(−3.48)	4.9	2.50	2.0	1.0	0.7	3.2
group total					2.8	1.1	10.4	4.3
Dollar and cross-currency dispersion								
Broad dollar	0.27	(2.56)	3.5	−1.17	1.3	0.2	9.7	1.1
Dispersion of realized vol	0.35	(2.95)	6.0	−2.62	3.0	2.4	4.0	1.0
Dispersion of carry	0.24	(2.24)	2.7	−0.77	1.1	0.1	4.0	3.5
group total					5.4	2.7	17.7	5.7
All eleven jointly, R^2					15.7	10.9	38.5	21.3
R^2, 1997–2010					24.4	19.3	53.5	33.1
R^2, 2011–2025					18.4	13.0	51.8	32.2

Table 8. Currency Characteristics and the Cross-Section of ϕ

This table relates the currency-level level factor to currency characteristics across 31 currencies. Panel A reports Pearson and Spearman correlations between ϕ and each characteristic, together with the associated p-values and the number of currencies for which the characteristic is available. Panel B reports least-squares regressions of ϕ on standardized characteristics, so that each coefficient is the response to a one-standard-deviation move in the cross-section; blank cells denote regressors not included. Panel C reports the distribution of ϕ across currencies and a decomposition of the variance of currency-month ϕ into a common time effect, permanent differences across currencies, and a residual that varies by currency and over time, together with rank correlations between the first and second halves of the sample. Carry is the mean interest rate differential and sovereign CDS is the median five-year par spread over the sample, in logs. The variable ϕ and the sample are defined as in Figure 9, as are the first five characteristics and the last. Three appear only here: realized volatility is the mean of the currency's realized exchange rate volatility, the term spread differential is the foreign less the domestic slope of the yield curve, and the risk-premium component is the option-implied premium at $\phi = 1$, each averaged over the same windows. Heteroskedasticity-robust t-statistics are reported in parentheses beside each coefficient, and R^2 is reported in percent. The sample covers daily observations from January 1997 to September 2025 for 31 currencies against the US dollar, with maturities ranging from one month to two years. Data are from Bloomberg, Consensus Economics, JP Morgan, and OptionMetrics.

	Pearson		Spearman		
	r	p	ρ	p	N
Panel A: Cross-Sectional Correlations with ϕ					
Interest rate differential	0.73	0.000	0.60	0.000	31
Exchange rate return	−0.65	0.000	−0.50	0.004	31
Total return	0.37	0.041	0.36	0.048	31
Cross-currency basis	0.49	0.005	0.37	0.040	31
Option-implied skew	0.62	0.000	0.48	0.007	31
Realized volatility	0.56	0.001	0.45	0.012	31
Term spread differential	0.69	0.000	0.47	0.008	31
Risk-premium component	0.66	0.000	0.63	0.000	31
Sovereign CDS (log)	0.65	0.000	0.60	0.001	25

Panel B: Multivariate, Standardized Regressors						
	Carry	Basis	Skew	CDS	R^2	N
Carry alone	2.11 (6.99)				53	31
Carry and basis	1.86 (4.97)	0.48 (1.38)			55	31
Carry and implied skew	1.66 (4.36)		0.65 (1.22)		55	31
Carry and sovereign CDS	1.87 (2.81)			0.53 (0.64)	56	25
Carry, basis and skew	1.55 (4.12)	0.36 (0.83)	0.54 (0.89)		56	31

Panel C: The Cross-Section of ϕ		
Mean	3.03	31
Standard deviation	2.90	31
Minimum	−1.41	31
Maximum	9.82	31
Currencies with a negative estimate	5	31
Variance from a common time effect (%)	15	
Variance from permanent currency effects (%)	7	
Variance currency-specific and time-varying (%)	78	
Split-half rank correlation, ϕ	0.25	
the same for carry	0.90	

Table 9. Carry-Portfolio Exposure to ϕ

This table reports estimates based on the following specification

$$rx_{p,t} = \alpha_p + \beta_p \, \Delta\phi_t + \gamma_p' X_t + \varepsilon_{p,t},$$

where $rx_{p,t}$ is the monthly excess return of carry-sorted currency portfolio p, $\Delta\phi_t$ is the monthly change in the level factor of the ϕ term structure, and X_t collects the controls. Portfolios are formed each month on the lagged interest rate differential and held for one month; P1 holds the lowest-yielding fifth and P5 the highest. The ordered-trend portfolio weights the five quintiles $-2, -1, 0, 1, 2$ and divides by three. Controls are the dollar factor, the return on the S&P 500 index, and the change in global FX volatility, measured as the cross-currency average of realized exchange rate volatility. Panel B reports the coefficient on the change in ϕ with one control set at a time, together with the regression R^2. Panel C splits the sample on the equity market return and excludes 2008–09. Panel D replaces the level factor with an equal-weighted average of currency-level estimates over the stated set of currencies. t-statistics, reported in parentheses, are based on Newey-West standard errors with three lags, and R^2 is reported in percent. The sample covers 344 monthly observations from February 1997 to September 2025, over which the monthly change in the level factor has a standard deviation of 1.43. Data are from Bloomberg, Consensus Economics, JP Morgan, and OptionMetrics.

	No Controls		With Controls	
	β	t	β	t
Panel A: Exposure by Carry Rank				
Lowest carry	−0.088	(−1.13)	0.049	(1.15)
2	−0.162	(−1.92)	0.063	(1.47)
3	−0.247	(−2.00)	0.018	(0.43)
4	−0.342	(−2.42)	−0.048	(−1.09)
Highest carry	−0.464	(−3.00)	−0.082	(−0.99)
High minus low	−0.376	(−2.94)	−0.130	(−1.13)
Ordered trend across ranks	−0.310	(−2.97)	−0.124	(−1.39)
Panel B: Which Control Absorbs the High-Minus-Low Exposure				
Dollar factor	−0.282	(−2.51)	12.2	
Equity market return	−0.316	(−2.46)	6.2	
Global FX volatility	−0.211	(−1.78)	17.0	
Change in the VIX	−0.294	(−2.64)	12.1	
Dollar and equity	−0.237	(−2.11)	13.3	
Dollar, equity, and FX volatility	−0.130	(−1.13)	21.0	
Panel C: States and Subsamples, High Minus Low				
Worst 30% of equity months	−0.626	(−3.75)	−0.286	(−1.69)
Other months	−0.202	(−1.21)	−0.088	(−0.60)
Excluding 2008–09	−0.318	(−2.50)	−0.138	(−1.14)
Panel D: ϕ Rebuilt without the Currencies Being Tested				
All currencies, equal-weighted	−0.254	(−3.00)	−0.028	(−0.39)
Excluding the top carry quintile	−0.165	(−1.90)	0.036	(0.47)
Developed currencies only	−0.137	(−1.88)	0.007	(0.11)

Online Appendix to

"Survey Expectations Meet Option Prices: New Insights from the FX Market"

(not for publication)

Preliminary and Incomplete

Abstract

This Internet Appendix presents supplementary material and results not included in the main body of the paper.

A Proofs

A.1 Proof of Proposition 1

Proof. Using the property that $M_T X_T = 1$, we start from the following identity

$$\mathbb{E}_t\left[\frac{S_T}{S_t}\right] = \mathbb{E}_t\left[M_T X_T \frac{S_T}{S_t}\right] \tag{A.1}$$

$$= \frac{1}{R^{\$}_{f,t}} \mathbb{E}^{\star}_t\left[X_T \frac{S_T}{S_t}\right], \tag{A.2}$$

with $\mathbb{E}^{\star}_t$ reflecting the risk-neutral expectation, which can be decomposed as

$$\mathbb{E}^{\star}_t\left[X_T \frac{S_T}{S_t}\right] = \mathbb{E}^{\star}_t\left[X_T\right] \mathbb{E}^{\star}_t\left[\frac{S_T}{S_t}\right] + \mathbb{C}\text{ov}^{\star}_t\left(X_T, \frac{S_T}{S_t}\right) \tag{A.3}$$

$$= R^{\$}_{f,t} \times \frac{R^{\$}_{f,t}}{R^{€}_{f,t}} + \mathbb{C}\text{ov}^{\star}_t\left(X_T, \frac{S_T}{S_t}\right), \tag{A.4}$$

where $\mathbb{E}^{\star}_t\left[X_T\right] = R^{\$}_{f,t}$ follows from the relation between the risk-neutral probability and the SDF valuation and $\mathbb{E}^{\star}_t\left[S_T/S_t\right] = R^{\$}_{f,t}/R^{€}_{f,t}$ from the Uncovered Interest Rate Parity (UIP) condition.

Combining and rearranging the above equations yield

$$\mathbb{E}_t\left[\frac{S_T}{S_t}\right] - \frac{R^{\$}_{f,t}}{R^{€}_{f,t}} = \frac{1}{R^{\$}_{f,t}} \mathbb{C}\text{ov}^{\star}_t\left(X_T, \frac{S_T}{S_t}\right) = \mathbb{C}\text{ov}^{\star}_t\left(\frac{X_T}{\mathbb{E}^{\star}_t[X_T]}, \frac{S_T}{S_t}\right). \tag{A.5}$$

Finally, after replacing X_T by $R^{\phi}_{m,T}$ above, we obtain

$$\mathbb{E}_t\left[\frac{S_T}{S_t}\right] - \frac{R^{\$}_{f,t}}{R^{€}_{f,t}} = \mathbb{C}\text{ov}^{\star}_t\left(\frac{R^{\phi}_{m,T}}{\mathbb{E}^{\star}_t[R^{\phi}_{m,T}]}, \frac{S_T}{S_t}\right). \tag{A.6}$$

□

A.2 Proof of Proposition 2

Proof. Recall that $f(x) = f(c) + f'(c)(x-c) + \frac{1}{2}f''(\xi)(x-c)^2$ for $\xi \in [1, x]$ (or $[x, 1]$). Here, we use $f(x) = x^\phi$, $c = 1$, and $x = R_{m,T}$. This allows us to express $R^\phi_{m,T} = 1 + \phi(R_{m,T} - 1) + \frac{\phi(\phi-1)}{2}\xi_T^{\phi-2}(R_{m,T}-1)^2$. Taking the covariance between $R^\phi_{m,T}$ and S_T/S_t would eliminate the constants: $\mathbb{C}\text{ov}^\star(R^\phi_{m,T}, S_T/S_t) = \phi\mathbb{C}\text{ov}^\star(R_{m,T}, S_T/S_t) + \frac{\phi(\phi-1)}{2}\mathbb{C}\text{ov}^\star(\xi_T^{\phi-2}(R_{m,T}-1)^2, S_T/S_t)$. Note that ϕ does not need to be an integer. □

A.3 Role of Higher-Order Moments

We can write the analogous version of (9) under the real-world measure by simply replacing $M_T = 1/X_T = R^{-\phi}_{m,T}$ in (3). We then obtain:

$$\mathbb{E}_t\left[\frac{S_{i,T}}{S_{i,t}}\right] - \frac{R^\$_{f,t}}{R^i_{f,t}} = \frac{1}{R^\$_{f,t}}\left(1 - e^{-\phi}\right)\sum_{n=1}^{\infty} w_{\phi,n}\mathbb{C}\text{ov}_t\left(r^n_{m,T}, \frac{S_{i,T}}{S_{i,t}} - 1\right), \tag{A.7}$$

where the weights $w_{\phi,n} = \left(e^{-\phi} - 1\right)^{-1}\frac{(-\phi)^n}{n!}$ sum to one. Note that the weights switch signs between the odd and even market return moments, as investors require a positive risk premium when the covariance between exchange rate returns and odd (even) market return moments is positive (negative). As an example, a currency that comoves positively with market return (or skewness) is a risky currency that should offer a positive risk premium. In contrast, a currency that comoves positively with market variance (or kurtosis) is a hedge currency that should offer a negative risk premium. The right-hand side panel of Figure 1 shows the change in signs between odd (even) market return moments and how the weights vary with ϕ.

B The Numéraire Portfolio under Various Cases

This appendix shows that our specification, whereby the return of the optimal growth portfolio is proportional to a power function of the market return, is consistent with various theoretical models. Throughout, we denote the growth optimal (numéraire) portfolio return by $R_{g,T}$, corresponding to X_T in Section 2.

B.1 Power Utility

We start with a simple static portfolio choice problem of an unconstrained CRRA agent who invests in the market.[16] The maximization problem can be written as follows:

$$\max_{\mathbf{w}} \mathbb{E}_t \frac{\left(\sum_i w_i R_{i,T}\right)^{1-\phi}}{1-\phi}, \quad s.t. \quad \sum_i w_i = 1\,.$$

Taking the first order condition for each weight w_i, we have $\lambda = \mathbb{E}_t\left[R_{i,T}\left(\sum_i w_i^\star R_{i,T}\right)^{-\phi}\right]$, where λ is the Lagrangian multiplier, and $w_i^\star$ is the optimal weight of an asset with gross return $R_{i,T}$ in the representative agent's portfolio. Note that the quantity $\left(\sum_i w_i^\star R_{i,T}\right)^{-\phi} = R_{m,T}^{-\phi}$ is proportional to the SDF, which proves that $R_{g,T} = \lambda R_{m,T}^{\phi}$ is the growth optimal portfolio return. The scaling constant λ is innocuous: Proposition 1 normalizes X_T by $\mathbb{E}_t^\star[X_T]$, so we set $X_T = R_{m,T}^{\phi}$ in the main text without loss of generality.

B.2 Ambiguity Aversion

A similar result can be obtained in a setting with ambiguity aversion. The following example builds on the results in Hansen (2007). We consider a robust portfolio optimization problem for an unconstrained representative agent with log utility, who penalizes their modeling mistake, i.e. the distance (relative entropy) between their subjective belief and the rational expectation (through

[16] This portfolio choice problem builds on Martin (2017). See Result 8 of its Online Appendix.

the change of measure ξ_T)

$$\max_{\mathbf{w}} \min_{\xi_T>0, \xi_T=\frac{d\mathbb{H}}{d\mathbb{P}}} \mathbb{E}_t\left[\xi_T \log\left(\sum_i w_i R_{i,T}\right)\right] - \underbrace{\theta\,\mathrm{KL}(\mathbb{H}|\mathbb{P})}_{\text{penalty of choosing } \mathbb{H}}, \quad \sum_i w_i = 1.$$

Hansen and Sargent (2011) shows the optimal distortion for the minimization problem (assume the portfolio weights are given) should be the *exponential tilting* (also called the Esscher transform), $\xi_T = R_{m,T}^{-\frac{1}{\theta}} / \mathbb{E}_t\left[R_{m,T}^{-\frac{1}{\theta}}\right]$. This would reduce the robust agent's problem into a CRRA agent's portfolio choice problem under rational expectation

$$\max_{\mathbf{w}} -\mathbb{E}_t \theta \left(\sum_i w_i R_{i,T}\right)^{-\frac{1}{\theta}}, \quad \sum_i w_i = 1.$$

As shown in Internet Appendix B.1, the growth optimal portfolio is $R_{g,T} = R_{m,T}^{\theta^{-1}+1}$. The value of $\phi = 1 + \theta^{-1}$ captures a log agent's ambiguity aversion, i.e., the higher ϕ (the lower θ), the stronger the ambiguity aversion of the log agent.

B.3 Partial Market Participation

Consider now, as an extension, that the representative agent holds only part of their portfolio in the market. This agent has a portfolio weight $\omega \in (0, 1]$ in the market portfolio and $(1-\omega)$ in the risk-free asset. This would imply the growth optimal portfolio being $R_{g,T} = (\omega R_{m,T} + (1-\omega) R_{f,t})^{\phi}$. Under some reasonable conditions, the above binomial function can be expanded as a Maclaurin series

$$R_{g,T} = (1-\omega)^{\phi} R_{f,t}^{\phi} \left(1 + \phi \frac{\omega}{1-\omega} \frac{R_{m,T}}{R_{f,t}} + \frac{\phi(\phi-1)}{2!} \left(\frac{\omega}{1-\omega} \frac{R_{m,T}}{R_{f,t}}\right)^2 + ...\right),$$

which is essentially a sum of integer powers of the market return with constant coefficients. The value of ϕ should be a rational number, which is a sensible choice here since the set of rational numbers is *dense* in the set of real numbers. Also, the random variable has compact support $R_{m,T}/R_{f,t} \in [0, 2(\omega^{-1} - 1)]$.

B.4 Heterogeneous Agents

We now show that ϕ can be endogenously time varying in a model with heterogeneous agents. Intuitively, variation in ϕ arises from the change in the distribution of wealth among agents with different preferences as in Chan and Kogan (2002) and Longstaff and Wang (2012). Consider a two-period model with complete markets and two agents from country 1 and country 2 with homogeneous beliefs and power utility, but with differing coefficients of risk aversion, $\gamma_2 > \gamma_1 \geq 1$.[17] Agent i's problem is as follows:

$$\max \frac{W_{i,t}^{1-\gamma_i}}{1-\gamma_i} + \beta \mathbb{E}_t \frac{W_{i,T}^{1-\gamma_i}}{1-\gamma_i}.$$

As markets are complete and beliefs are homogeneous, the SDF is unique, so that

$$\beta \left(\frac{W_{1,T}}{W_{1,t}}\right)^{-\gamma_1} = \beta \left(\frac{e_T W_{2,T}}{e_t W_{2,t}}\right)^{-\gamma_2},$$

where e_t is the exchange rate of one unit currency in country 2 valued in the currency of country 1.

Assuming that $\gamma_1 = \gamma$ and $\gamma_2 = 2\gamma$ to ensure a closed form solution, as in Longstaff and Wang (2012), we have

$$\frac{W_{1,T}}{W_{1,t}} = \left(\frac{e_T W_{2,T}}{e_t W_{2,t}}\right)^2.$$

Writing $W_t = W_{1,t} + e_t W_{2,t}$ for aggregate wealth measured in currency 1 implies that $W_{1,T} = \frac{2}{a}\left(\sqrt{1+aW_T}-1\right)^2$ and $e_T W_{2,T} = \frac{2}{a}\left(\sqrt{1+aW_T}-1\right)$, where the constant $a = 4W_{1,t}/(e_t W_{2,t})^2$ reflects the relative wealth of the two agents. Although agents 1 and 2 are not representative, since neither invests only in the market, they have the same beliefs and SDF as a representative agent. Such representative agent has a wealth W_T (measured in currency 1) and marginal utility $v'(W_T)$ that is proportional to $e_T W_{2,T}^{-2\gamma}$. Integrating across agents, this representative agent's utility function at time T is

$$v(W_T) = \frac{\left(\sqrt{1+aW_T}-1\right)^{2(1-\gamma)}}{2(1-\gamma)} + \frac{\left(\sqrt{1+aW_T}-1\right)^{1-2\gamma}}{1-2\gamma},$$

[17] This example builds on Longstaff and Wang (2012), where each agent faces a portfolio choice problem. We extend to a two-country environment and, at the same time, simplify the initial model as far as possible.

such that their relative risk aversion, denoted by $\phi(W_T)$, can be written as:

$$\phi(W_T) = -\frac{W_T v''(W_T)}{v'(W_T)} = \gamma + \frac{\gamma}{\sqrt{1 + aW_T}}$$

with the limits $\lim_{W_T \to \infty} \phi = \gamma$ and $\lim_{W_T \to 0} = 2\gamma$. The coefficient ϕ therefore varies over time, as the relative wealth of the agents changes, in a range given by γ and 2γ. We could represent marginal utility of the representative agent as, $v'(W_T) = a\left(\sqrt{1 + aW_T} - 1\right)^{-2\gamma} = \kappa(W_T) W_T^{-\phi(W_T)}$, where $\kappa(W_T)$ is a state dependent constant, and it implies the growth optimal portfolio return can be written as

$$R_{g,T} \propto \frac{1}{\kappa(W_T)} R_{m,T}^{\phi(W_T)}.$$

The time varying constant is $\kappa(W_T) = a\left(aW_T^{1-\frac{\phi(W_T)}{\gamma}} + 2W_T^{-\frac{\phi(W_T)}{\gamma}} - \frac{2\gamma W_T^{1-\frac{\phi(W_T)}{\gamma}}}{\phi(W_T)-\gamma}\right)^{-\gamma}$. Since $\phi(W_T)/\gamma$ is between $[1, 2]$, the expression $\kappa(W_T)$ is a polynomial in powers of W_T with exponents between 0 and 2γ.

C Empirical methodology

C.1 Estimating the Risk-Neutral Correlation $\rho^{\star}_{\phi,i,t}$

We measure $\rho_{\phi,i,t}$ using a backward-looking sample correlation over a window that matches the maturity of the options. Suppose that on day t, for example, we compute $\mathbb{V}\text{ar}^{\star}[\,\cdot\,]$ and $\mathbb{E}^{\star}[\,\cdot\,]$ using options between t and T. We then calculate $\rho_{\phi,i,t}$ using daily returns between $t-T$ and t.

Since risk-neutral correlations are not observable, we used the realized empirical correlations instead. This choice is backed by three observations. First, when $\phi = 1$ the premium built from our realized correlations tracks the quanto risk premium of Kremens and Martin (2019), which embeds the risk-neutral correlation directly. Figure 2 shows the comparison. Second, we find the realized correlations are similar across different values of ϕ. Third, we run similar regressions as Kremens and Martin (2019) by using the $ERP^{1}_{i,t,T}$ that we constructed using the realized correlations and match the same sample period.

C.2 Risk-Neutral Moments

We explain how we compute the two risk-neutral variances in detail.

Risk Neutral Moments of Equity Index Returns. To compute the risk-neutral variance of $\frac{R^{\phi}_{m,T}}{\mathbb{E}^{\star}_{t}[R^{\phi}_{m,T}]}$, we recall the formula

$$\mathbb{V}\text{ar}^{\star}_{t}\left(\frac{R^{\phi}_{m,T}}{\mathbb{E}^{\star}_{t}\left[R^{\phi}_{m,T}\right]}\right) = \frac{\mathbb{E}^{\star}_{t}\left[R^{2\phi}_{m,T}\right]}{\mathbb{E}^{\star}_{t}\left[R^{\phi}_{m,T}\right]^{2}} - 1\,,$$

where we could compute the risk-neutral moments using the following formula

$$\begin{aligned}\mathbb{E}^{\star}_{t}[R^{\theta}_{m,T}] = R^{\theta}_{f,t} &+ R_{f,t}\int_{0}^{F_{t,T}} \frac{\theta(\theta-1)}{E^{\theta}_{t}}(E_t R_{f,t} - F_{t,T} + K)^{\theta-2} P_{t,T}(K)dK \\ &+ R_{f,t}\int_{F_{t,T}}^{\infty} \frac{\theta(\theta-1)}{E^{\theta}_{t}}(E_t R_{f,t} - F_{t,T} + K)^{\theta-2} C_{t,T}(K)dK\,, \qquad \text{(A.8)}\end{aligned}$$

where $R_{f,t}$ is the risk free rate from yield curve with maturity T, K is the strike of option, E_t is the current level of equity index, and $P_{t,T}(K)$ and $C_{t,T}(K)$ are the out-of-the-money put and call option prices with strike K and maturity T. The formula's proof could be found in the online appendix (result 9) of Martin (2017).

We take the assumption $E_t R_{f,t} = F_{t,T}$ in our computation to ignore the dividend payment, i.e. we assume the payment of dividends are not reinvested. The formula simplifies to

$$\mathbb{E}_t^\star[R_{m,T}^\theta] = R_{f,t}^\theta + R_{f,t}\int_0^{F_{t,T}} \frac{\theta(\theta-1)}{E_t^\theta} K^{\theta-2} P_{t,T}(K)dK + R_{f,t}\int_{F_{t,T}}^{\infty} \frac{\theta(\theta-1)}{E_t^\theta} K^{\theta-2} C_{t,T}(K)dK\,. \tag{A.9}$$

Risk Neutral Variances of FX returns. The risk-neutral variance of the gross exchange rate return between two dates t and T

$$\mathbb{V}\text{ar}_t^\star\left(\frac{S_T}{S_t}\right) = \mathbb{E}_t^\star\left(\frac{S_T}{S_t}\right)^2 - \left(\mathbb{E}_t^\star\frac{S_T}{S_t}\right)^2, \tag{A.10}$$

is computed by integrating over an infinite range of the strike prices from European call and put options expiring on these dates as

$$\mathbb{V}\text{ar}_t^\star\left(\frac{S_T}{S_t}\right) = \frac{2}{B_{t,T}S_t^2}\left(\int_0^{F_{t,T}} P_{t,T}(K)dK + \int_{F_{t,T}}^{\infty} C_{t,T}(K)dK\right), \tag{A.11}$$

where $P_{t,T}(K)$ and $C_{t,T}(K)$ are put and call option prices at time t with strike price K and maturity date T, respectively, and $B_{t,T}$ is the price of a domestic bond at time t with maturity date T. The above equation builds on Bakshi and Madan (2000) and Britten-Jones and Neuberger (2000) and is based on no-arbitrage conditions that require no specific option pricing model. In our implementation, we follow Jiang and Tian (2005) and use a cubic spline around the available implied volatility points. This interpolation method is standard in the literature and has the advantage that the implied volatility smile is smooth between the maximum and minimum available strikes. We compute the option values using the Garman and Kohlhagen (1983) valuation formula and

solve the integral in Equation (A.11) via trapezoidal integration.

D Supporting Evidence

D.1 Heterogeneity and the Carry Relation

Section 6.1 reports that a single ϕ cannot reconcile survey expectations with option prices across currencies. The formal version of that statement is a month-by-month homogeneity test. In each month we compute $\sum_i(\phi_i - \bar{\phi})^2/\text{se}_i^2$, which is χ^2 with $N-1$ degrees of freedom if every currency shares one ϕ; computing it within a month avoids the serial dependence that overlapping windows induce across months. It rejects at the 1% level in all 345 months, with a median statistic of 622 against a critical value of 43. On the same standard errors, the cross-sectional variance of the estimates averages 179.8 against an average estimation variance of 10.0, which would put 94% of the spread down to genuine differences.

Both figures should be read as upper bounds on the evidence rather than as the evidence. The rolling regressions are estimated on roughly a thousand overlapping currency-days inside each 252-day window, and the standard errors the files report are ordinary least squares errors that treat those observations as independent. They are therefore too small, which inflates the homogeneity statistic and understates the estimation share. A common estimation component across currencies within a month would push the statistic the other way, since the cross-sectional demeaning removes it while the standard errors still carry it, but that does not offset the first problem.

Two checks avoid the reported standard errors entirely. The first check is persistence. Estimation error in ϕ_i is autocorrelated for up to twelve months by construction, because consecutive windows share data; beyond twelve months the windows are disjoint. The currency-specific component of ϕ_i has a median autocorrelation of 0.91 at one month and 0.35 at six, and essentially none past the window length, between -0.04 and -0.08 from twelve to forty-eight months. The currency-specific variation is real but short-lived, which is the same message as the split-half rank correlation of 0.25 in Section 6.1. The second check is external. Time-averaged ϕ_i correlates 0.73 with average carry across the 31 currencies, and sampling error cannot reproduce a correlation of that size with a variable measured outside the estimation. The ordering of currencies survives without the standard

errors; the month-by-month statistic does not carry the argument on its own.

Section 6.2 reports that carry is the dominant cross-sectional correlate of ϕ and that the relation is concentrated among high-yield currencies. This appendix collects the evidence behind both statements.

Internet Appendix Table A.8 runs the four checks. Panel A cuts the cross-section at two transparent points, a zero differential and the median, rather than searching for a break. Panel B removes the high-yield tail in stages and splits on development status. Panel C races carry against each of the other characteristics. Panel D addresses the concern that carry is mechanically embedded in the survey expectation of the excess return, and asks whether carry measured in one part of the sample lines up with ϕ measured in the other.

INTERNET APPENDIX TABLE A.8 ABOUT HERE

Three points are worth drawing out. The relation is flat among funding currencies and steep among investment currencies, so it is the high-yield half that identifies it; a formal test of the difference in slopes returns $t = 1.79$, which is suggestive rather than decisive at 31 currencies. The concentration is about yield rather than development status: within emerging and frontier currencies alone the slope is 0.88 ($t = 6.66$), and an emerging-market indicator adds nothing once carry is included. And no rival characteristic displaces carry. Sovereign CDS, the basis, skew, and both volatility measures each leave carry between $t = 2.84$ and $t = 7.70$, while none of them survives on its own once carry is present.

Panel D addresses the construction. The expected excess return is the expected exchange rate return less the interest differential, so the differential sits inside the dependent variable of the first-stage regression. Because that regression includes a constant, a currency's average differential cannot move its slope; only the comovement of the differential with the option-implied premium can. That comovement, the interest-differential leg, explains 2.5% of the cross-section on its own ($t = -0.99$) and leaves the carry coefficient at 0.91 ($t = 7.50$) when both are included. The sort is not an accounting identity.

The last two rows of Panel D make the relation harder to read as contemporaneous comovement. Carry averaged before 2011 explains 28% of the cross-sectional variation in ϕ averaged after 2011, and carry averaged after 2011 explains 33% of ϕ before it. Carry is a stable currency attribute, with a rank correlation of 0.90 across the two halves, while ϕ is not, at 0.25. A stable attribute lines up with an unstable one measured in a different decade.

E More Robustness Concerns

We consider some robustness issues that we omitted in the paper that we think might distract the readers from our main discussions.

E.1 Consensus vs. Rational Expectation

One might think the consensus forecast is not exactly the physical measure. To reflect this view, we use $\widetilde{\mathbb{E}}[\cdot]$ to represent the subjective measure (consensus forecasts), which differs from $\mathbb{E}[\cdot]$. We now decompose the risk premium forecasts, i.e. the ex-ante expected excess return of exchange rate under subjective measure, into two parts

$$\underbrace{\widetilde{\mathbb{E}}_t\left[\frac{S_T}{S_t}\right]-\frac{R^{\$}_{f,t}}{R^{i}_{f,t}}}_{\text{forecast risk premium}}=\underbrace{\mathbb{E}_t\left[\frac{S_T}{S_t}\right]-\frac{R^{\$}_{f,t}}{R^{i}_{f,t}}}_{\text{UIP premium}}+\underbrace{\widetilde{\mathbb{E}}_t\left[\frac{S_T}{S_t}\right]-\mathbb{E}_t\left[\frac{S_T}{S_t}\right]}_{\text{non-risk-premium component}}, \tag{A.12}$$

where the first part is the expected risk premium under physical measure and the second part reflects the difference between the subjective and physical measures. Empirically, the left side of (A.12) is directly observable from the data of surveys. Comparing the above to (A.15), we could view the subjective expected risk premium as

$$\widetilde{\mathbb{E}}_t\left[\frac{S_T}{S_t}\right]-\frac{R^{\$}_{f,t}}{R^{i}_{f,t}}=ERP^{\phi}_{i,t,T}+\xi_{i,t}\,, \tag{A.13}$$

where the residual has two parts

$$\xi_{i,t}=\underbrace{-\frac{R^{\$}_{f,t}\mathbb{C}\text{ov}_t\left(M_T R^{\phi}_{m,T},\frac{S_T}{S_t}\right)}{\mathbb{E}^{\star}_t[R^{\phi}_{m,T}]}}_{\xi^{(1)}_{i,t}}+\underbrace{\widetilde{\mathbb{E}}_t\left[\frac{S_T}{S_t}\right]-\mathbb{E}_t\left[\frac{S_T}{S_t}\right]}_{\xi^{(2)}_{i,t}}. \tag{A.14}$$

The optimal choice of ϕ that minimize the residual $\xi_{i,t}$ could be seen as a joint test that the two components are both zero.

E.2 Residual Term for $ERP^{\phi}_{i,t,T}$

A potential misspecification of $R_{g,T}$ would generate an error term in (5). We could explicitly compute this term in the following result, which is an extension of result 1 in Kremens and Martin (2019).

Result 1. *Under no-arbitrage, for any random quantity X_T, we have the following decomposition of expected return of FX*

$$\mathbb{E}_t\left[\frac{S_T}{S_t}\right] - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \frac{\mathbb{C}ov^{\star}_t\left(X_T, \frac{S_T}{S_t}\right)}{\mathbb{E}^{\star}_t[X_T]} - \frac{R^{\$}_{f,t}\mathbb{C}ov_t\left(M_T X_T, \frac{S_T}{S_t}\right)}{\mathbb{E}^{\star}_t[X_T]}, \tag{A.15}$$

where M_T is the SDF and $\mathbb{E}^{\star}[\cdot]$ stands for the expectation under the risk-neutral measure.

Given the choice of $R_{g,T} = R^{\phi}_{m,T}$, which value of ϕ would minimize the residual term in (A.15) is an empirical question.

Table A.1. Quanto Risk Premium vs Expected Risk Premium

This table reports estimates from the following specification:

$$QRP_{i,t,T} = \alpha_i + \alpha_t + \beta\, ERP_{i,t,T} + \varepsilon_{i,t,T},$$

where QRP denotes the quanto risk premium for currency i on date t over the horizon $\tau = T - t$, and ERP is the corresponding expected risk premium. The terms α_i and α_t denote currency and calendar-time fixed effects (*fe*), respectively. The maturity τ is fixed at the two-year horizon. QRP is extracted from quanto contract prices and reflects the risk-neutral covariance between exchange rate movements and global equity market risk. ERP is constructed from SPX options, FX options, and backward-looking realized correlations between S&P 500 index returns and FX spot returns, estimated over the window from $t - \tau$ to t. Since QRP is observed at a monthly frequency, Panel A synchronizes QRP and ERP by calendar date. In Panel B, for each QRP observation date, we compute the daily average ERP between the previous QRP date and the current QRP date. The intercept α and R^2 are reported in percent. Standard errors, reported in parentheses, are clustered by currency and calendar time. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. The sample covers monthly observations from December 2009 to October 2015 for 11 currencies against the US dollar at the two-year maturity. Data on QRP are from Kremens and Martin (2019), whereas data used to construct ERP are from Bloomberg, JP Morgan, and OptionMetrics.

Panel A: Matched Dates		
β	1.14***	1.00***
	(0.11)	(0.13)
α	0.52	
	(0.30)	
R^2	74.83	95.49
N	656	656
Panel B: Monthly Average		
β	1.15***	1.01***
	(0.11)	(0.13)
α	0.52	
	(0.30)	
R^2	73.34	95.22
N	656	656
currency fe		✓
time fe		✓

Table A.2. Pooled Regression Estimates of ϕ

This table presents estimates of ϕ from the following specification

$$ERX_{i,t,T} = \alpha + \beta_\phi ERP^\phi_{i,t,T} + \varepsilon_{i,t,T},$$

where ERX is the expected excess return for currency i on day t over the horizon $\tau = T - t$ and ERP^ϕ is the corresponding expected risk premium for a given ϕ. The column *All* reports estimates from the regression pooled across all maturities, whereas the remaining columns report estimates for a specific maturity τ. ERX is constructed from consensus forecasts, spot exchange rates, and interest rate differentials, while ERP^ϕ is derived from SPX options, FX options, and backward-looking realized correlations between S&P 500 index returns and FX spot returns, estimated over the window from $t - \tau$ to t for different levels of ϕ. *Simple* displays estimates of β, obtained by setting $\phi = 1$. *Search* reports estimates of ϕ that satisfies $\beta = 1$, obtained via a fine grid search. Standard errors, reported in parentheses, are clustered by currency and calendar-time dimensions. The standard error for ϕ is obtained via the delta method. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. Daily forecasts are obtained by forward-filling monthly forecasts. The sample covers daily observations from January 1996 to September 2025 for 33 currencies in the broad sample (Panel A) and 20 currencies in the liquid sample (Panel B) against the US dollar, with maturities ranging from one month to two years. Data are from Bloomberg, Consensus Economics, JP Morgan, and OptionMetrics.

		Panel A: Broad Sample					Panel B: Liquid Sample				
		All	1 month	3 months	1 year	2 years	All	1 month	3 months	1 year	2 years
Simple	β	5.15***	7.18***	3.17***	2.89***	2.80***	4.98***	7.24***	3.13***	2.96***	2.79***
		(0.69)	(0.87)	(0.58)	(0.46)	(0.42)	(0.78)	(1.01)	(0.69)	(0.52)	(0.46)
	α	0.38	−0.36	0.82*	1.20***	1.01***	−0.03	−1.09	0.45	0.88**	0.73*
		(0.45)	(0.77)	(0.45)	(0.38)	(0.34)	(0.44)	(0.89)	(0.52)	(0.42)	(0.37)
	R^2	3.27	3.93	3.03	9.82	13.65	3.98	4.57	3.54	12.67	16.99
	N	824,750	204,355	206,452	206,798	207,145	538,076	134,519	134,519	134,519	134,519
Search	ϕ	6.46***	8.40***	3.63***	3.82***	4.06***	6.22***	8.28***	3.45***	3.85***	3.97***
		(1.17)	(1.37)	(0.80)	(0.88)	(0.99)	(1.28)	(1.51)	(0.95)	(1.01)	(1.07)
	α	0.57	−0.10	0.88*	1.22***	1.02***	0.15	−0.85	0.50	0.90**	0.74*
		(0.44)	(0.80)	(0.46)	(0.38)	(0.34)	(0.45)	(0.89)	(0.52)	(0.42)	(0.37)
	R^2	2.81	2.76	2.55	9.06	13.15	3.47	3.43	3.11	12.18	16.82
	N	824,750	204,355	206,452	206,798	207,145	538,076	134,519	134,519	134,519	134,519

Table A.3. Panel Regression Estimates of ϕ: Major Currencies

This table presents estimates of ϕ from the following specification

$$ERX_{i,t,T} = \alpha_i + \alpha_t + \alpha_\tau + \beta_\phi ERP^{\phi}_{i,t,T} + \varepsilon_{i,t,T},$$

where ERX is the expected excess return for currency i on day t over the horizon $\tau = T - t$ and ERP^{ϕ} is the corresponding expected risk premium for a given ϕ. The terms α_i, α_t, and α_τ denote currency, calendar-time, and maturity fixed effects (*fe*), respectively. ERX is constructed from consensus forecasts, spot exchange rates, and interest rate differentials, while ERP^{ϕ} is derived from SPX options, FX options, and backward-looking realized correlations between S&P 500 index returns and FX spot returns, estimated over the window from $t - \tau$ to t for different levels of ϕ. The column *All* pools observations across all maturities, whereas the remaining columns report estimates for a specific maturity τ. *Simple* displays estimates of β, obtained by setting $\phi = 1$. *Search* reports estimates of ϕ that satisfies $\beta = 1$, obtained via a fine grid search. Standard errors, reported in parentheses, are clustered by currency and calendar-time dimensions. The standard error for ϕ is obtained via the delta method. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. Daily forecasts are obtained by forward-filling monthly forecasts. The sample covers daily observations from January 1996 to September 2025 for the 10 major currencies (G10) against the US dollar, with maturities ranging from one month to two years. Data are from Bloomberg, Consensus Economics, JP Morgan, and OptionMetrics.

		All	1 month	3 months	1 year	2 years
Simple	β	3.97***	5.62***	2.04**	1.99***	1.68***
		(0.65)	(1.01)	(0.75)	(0.43)	(0.39)
	R^2	26.02	51.67	47.54	47.66	54.66
	N	289,472	72,368	72,368	72,368	72,368
Search	ϕ	4.29***	6.12***	2.12**	2.27***	1.89**
		(0.81)	(1.35)	(0.83)	(0.64)	(0.60)
	R^2	25.80	51.45	47.51	47.66	54.65
	N	289,472	72,368	72,368	72,368	72,368
	Currency fe	✓	✓	✓	✓	✓
	Maturity fe	✓				
	Time fe	✓	✓	✓	✓	✓

Table A.4. Panel Regression Estimates of ϕ: Implied Volatility and Market Conditions

This table presents estimates from the following specification

$$ERX_{i,t,T} = \alpha_i + \alpha_t + \alpha_\tau + \beta_\phi ERP^{\phi}_{i,t,T} + \varepsilon_{i,t,T},$$

where ERX is the expected excess return for currency i on day t over the horizon $\tau = T - t$ and ERP^{ϕ} is the corresponding expected risk premium, evaluated at $\phi = 1$. The terms α_i, α_t, and α_τ denote currency, calendar-time, and maturity fixed effects (*fe*), respectively. ERX is constructed from consensus forecasts, spot exchange rates, and interest rate differentials, while ERP^{ϕ} is derived from SPX options, FX options, and backward-looking realized correlations between S&P 500 index returns and FX spot returns. The column *All* pools observations across all maturities, whereas the remaining columns report estimates for a specific maturity τ. The top panel reports estimates for *Good States*, defined as days on which the VIX index is below its sample mean. The bottom panel reports estimates for *Bad States*, defined as days on which the VIX index is at or above its sample mean. We report estimates of ϕ that satisfies $\beta = 1$, obtained via a fine grid search. Estimates of β, obtained by setting $\phi = 1$, are displayed in Table 5. Standard errors, reported in parentheses, are clustered by currency and calendar-time dimensions. The standard error for ϕ is obtained via the delta method. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. Daily forecasts are obtained by forward-filling monthly forecasts. The sample covers daily observations from January 1996 to September 2025 for 33 currencies in the broad sample (Panel A) and 20 currencies in the liquid sample (Panel B) against the US dollar, with maturities ranging from one month to two years. Data are from Bloomberg, Consensus Economics, JP Morgan, and OptionMetrics.

		Panel A: Broad Sample					Panel B: Liquid Sample				
		All	1 month	3 months	1 year	2 years	All	1 month	3 months	1 year	2 years
Good	ϕ	1.25*	0.08	0.91	1.65*	1.81*	2.51***	1.74	1.43	2.38	2.66*
		(0.73)	(1.08)	(1.10)	(0.95)	(0.97)	(0.70)	(1.50)	(1.22)	(1.53)	(1.42)
	R^2	16.56	31.32	31.87	44.39	52.66	18.85	36.75	36.11	46.67	54.40
	N	509,847	126,433	127,596	127,876	127,942	327,388	81,847	81,847	81,847	81,847
Bad	ϕ	8.57***	9.40***	4.29***	2.11**	1.66**	7.96***	8.29***	3.68***	2.05***	1.51**
		(1.56)	(1.78)	(1.37)	(0.83)	(0.77)	(1.53)	(1.38)	(1.02)	(0.59)	(0.53)
	R^2	20.29	34.82	35.18	45.73	53.27	26.88	47.74	45.66	54.24	60.44
	N	314,903	77,922	78,856	78,922	79,203	210,688	52,672	52,672	52,672	52,672
	Currency fe	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Maturity fe	✓					✓				
	Time fe	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Table A.5. Panel Regression Estimates of ϕ: Realized Volatility and Market Conditions

This table presents estimates from the following specification

$$ERX_{i,t,T} = \alpha_i + \alpha_t + \alpha_\tau + \beta_\phi ERP^{\phi}_{i,t,T} + \varepsilon_{i,t,T},$$

where ERX is the expected excess return for currency i on day t over the horizon $\tau = T - t$ and ERP^{ϕ} is the corresponding expected risk premium, evaluated at $\phi = 1$. The terms α_i, α_t, and α_τ denote currency, calendar-time, and maturity fixed effects (*fe*), respectively. ERX is constructed from consensus forecasts, spot exchange rates, and interest rate differentials, while ERP^{ϕ} is derived from SPX options, FX options, and backward-looking realized correlations between S&P 500 index returns and FX spot returns. The column *All* pools observations across all maturities, whereas the remaining columns report estimates for a specific maturity τ. The top panel reports estimates for *Good States*, defined as days on which the one-month realized volatility of the S&P 500 is below its sample mean. The bottom panel reports estimates for *Bad States*, defined as days on which the one-month realized volatility of the S&P 500 is at or above its sample mean. We report estimates of ϕ that satisfies $\beta = 1$, obtained via a fine grid search. Estimates of β, obtained by setting $\phi = 1$, are displayed in Table 6. Standard errors, reported in parentheses, are clustered by currency and calendar-time dimensions. The standard error for ϕ is obtained via the delta method. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. Daily forecasts are obtained by forward-filling monthly forecasts. The sample covers daily observations from January 1996 to September 2025 for 33 currencies in the broad sample (Panel A) and 20 currencies in the liquid sample (Panel B) against the US dollar, with maturities ranging from one month to two years. Data are from Bloomberg, Consensus Economics, JP Morgan, and OptionMetrics.

		Panel A: Broad Sample					Panel B: Liquid Sample				
		All	1 month	3 months	1 year	2 years	All	1 month	3 months	1 year	2 years
Good	ϕ	1.85**	0.48	1.99	2.33*	2.03**	2.86***	1.92	1.76*	2.68*	2.68**
		(0.84)	(2.07)	(1.71)	(1.25)	(0.96)	(0.75)	(1.37)	(0.97)	(1.28)	(1.20)
	R^2	15.93	29.82	30.53	42.76	50.86	19.31	37.55	36.77	46.43	53.51
	N	531,012	131,633	132,879	133,188	133,312	343,660	85,915	85,915	85,915	85,915
Bad	ϕ	8.98***	10.04***	4.21***	1.82***	1.52**	8.27***	8.76***	3.90***	1.98***	1.44**
		(1.63)	(1.83)	(1.17)	(0.63)	(0.62)	(1.57)	(1.44)	(1.10)	(0.61)	(0.56)
	R^2	21.44	36.95	36.77	46.84	54.23	26.90	47.70	45.37	54.14	60.52
	N	293,738	72,722	73,573	73,610	73,833	194,416	48,604	48,604	48,604	48,604
	Currency fe	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Maturity fe	✓					✓				
	Time fe	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Table A.6. Panel Regression Estimates of ϕ: Implied Volatility and Major Currencies

This table presents estimates of ϕ from the following specification

$$ERX_{i,t,T} = \alpha_i + \alpha_t + \alpha_\tau + \beta_\phi ERP^{\phi}_{i,t,T} + \varepsilon_{i,t,T},$$

where ERX is the expected excess return for currency i on day t over the horizon $\tau = T - t$ and ERP^{ϕ} is the corresponding expected risk premium for a given ϕ. The terms α_i, α_t, and α_τ denote currency, calendar-time, and maturity fixed effects (*fe*), respectively. ERX is constructed from consensus forecasts, spot exchange rates, and interest rate differentials, while ERP^{ϕ} is derived from SPX options, FX options, and backward-looking realized correlations between S&P 500 index returns and FX spot returns, estimated over the window from $t - \tau$ to t for different levels of ϕ. Panel A reports results for days on which the VIX index is below its full-sample mean and Panel B for days on which it is above. The column *All* pools observations across all maturities, whereas the remaining columns report estimates for a specific maturity τ. *Simple* displays estimates of β, obtained by setting $\phi = 1$. *Search* reports estimates of ϕ that satisfies $\beta = 1$, obtained via a fine grid search; *n.a.* indicates that no solution was found within the search range. Standard errors, reported in parentheses, are clustered by currency and calendar-time dimensions. The standard error for ϕ is obtained via the delta method. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. Daily forecasts are obtained by forward-filling monthly forecasts. The sample covers daily observations from January 1996 to September 2025 for the 10 major currencies (G10) against the US dollar, with maturities ranging from one month to two years. Data are from Bloomberg, Consensus Economics, JP Morgan, OptionMetrics, and the CBOE.

		Panel A: Good States					Panel B: Bad States				
		All	1 month	3 months	1 year	2 years	All	1 month	3 months	1 year	2 years
Simple	β	2.82***	0.29	1.22	2.94***	2.50***	3.61***	5.21***	1.00	0.75	0.73
		(0.56)	(1.71)	(1.38)	(0.75)	(0.54)	(1.00)	(1.07)	(1.03)	(0.67)	(0.48)
	R^2	24.59	49.38	47.45	47.46	54.20	28.00	54.90	50.42	53.12	59.32
	N	170,980	42,745	42,745	42,745	42,745	118,492	29,623	29,623	29,623	29,623
Search	ϕ	2.72***	0.27	1.21	3.79**	3.23**	3.91***	5.68***	1.00	0.70	0.67
		(0.62)	(1.25)	(1.44)	(1.33)	(1.04)	(1.18)	(1.38)	(1.33)	(0.68)	(0.49)
	R^2	24.53	49.38	47.45	47.50	54.13	27.82	54.62	50.42	53.13	59.32
	N	170,980	42,745	42,745	42,745	42,745	118,492	29,623	29,623	29,623	29,623
	Currency fe	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Maturity fe	✓					✓				
	Time fe	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Table A.7. Panel Regression Estimates of ϕ: Realized Volatility and Major Currencies

This table presents estimates of ϕ from the following specification

$$ERX_{i,t,T} = \alpha_i + \alpha_t + \alpha_\tau + \beta_\phi ERP^{\phi}_{i,t,T} + \varepsilon_{i,t,T},$$

where ERX is the expected excess return for currency i on day t over the horizon $\tau = T - t$ and ERP^{ϕ} is the corresponding expected risk premium for a given ϕ. The terms α_i, α_t, and α_τ denote currency, calendar-time, and maturity fixed effects (*fe*), respectively. ERX is constructed from consensus forecasts, spot exchange rates, and interest rate differentials, while ERP^{ϕ} is derived from SPX options, FX options, and backward-looking realized correlations between S&P 500 index returns and FX spot returns, estimated over the window from $t - \tau$ to t for different levels of ϕ. Panel A reports results for days on which the one-month realized volatility of the S&P 500 is below its full-sample mean and Panel B for days on which it is above. The column *All* pools observations across all maturities, whereas the remaining columns report estimates for a specific maturity τ. *Simple* displays estimates of β, obtained by setting $\phi = 1$. *Search* reports estimates of ϕ that satisfies $\beta = 1$, obtained via a fine grid search; *n.a.* indicates that no solution was found within the search range. Standard errors, reported in parentheses, are clustered by currency and calendar-time dimensions. The standard error for ϕ is obtained via the delta method. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. Daily forecasts are obtained by forward-filling monthly forecasts. The sample covers daily observations from January 1996 to September 2025 for the 10 major currencies (G10) against the US dollar, with maturities ranging from one month to two years. Data are from Bloomberg, Consensus Economics, JP Morgan, OptionMetrics, and the CBOE.

		Panel A: Good States					Panel B: Bad States				
		All	1 month	3 months	1 year	2 years	All	1 month	3 months	1 year	2 years
Simple	β	3.53***	3.45**	2.53**	3.09***	2.40***	3.55***	4.94***	0.85	0.59	0.67
		(0.46)	(1.29)	(0.95)	(0.55)	(0.51)	(0.98)	(1.05)	(0.93)	(0.58)	(0.40)
	R^2	25.21	50.49	47.68	46.91	52.69	27.52	54.05	49.78	52.70	59.74
	N	182,336	45,584	45,584	45,584	45,584	107,136	26,784	26,784	26,784	26,784
Search	ϕ	3.64***	3.40**	2.69**	4.11***	3.07**	3.80***	5.30***	0.82	0.56	0.62
		(0.54)	(1.38)	(1.12)	(1.05)	(0.95)	(1.12)	(1.33)	(1.04)	(0.47)	(0.38)
	R^2	25.11	50.48	47.68	46.97	52.65	27.32	53.79	49.78	52.71	59.74
	N	182,336	45,584	45,584	45,584	45,584	107,136	26,784	26,784	26,784	26,784
	Currency fe	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Maturity fe	✓					✓				
	Time fe	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Table A.8. Where the Cross-Sectional Carry Relation Comes From

This table reports least-squares regressions of the currency-level ϕ on average carry across currencies. The variable ϕ and the characteristics are constructed as in Figure 9. Carry is the mean interest rate differential in percent per annum, signed so that a high-yield currency is positive. Panel A splits the cross-section at a zero differential and at the median. Panel B excludes the highest-yielding currencies in stages and splits on development status. Panel C adds one rival characteristic at a time, and only the carry coefficient is shown. Panel D reports the interest-differential leg, defined as the slope of the interest differential on the option-implied premium within each currency, together with regressions of ϕ averaged over one part of the sample on carry averaged over the other, split at June 2011. Panel E replaces the level factor with the currency's average estimate at each maturity taken separately. Coefficients are expressed in units of ϕ per percentage point of carry. Heteroskedasticity-robust t-statistics are reported in parentheses, N is the number of currencies, and R^2 is reported in percent. The sample covers daily observations from January 1997 to September 2025 for 31 currencies against the US dollar, with maturities ranging from one month to two years. Data are from Bloomberg, Consensus Economics, JP Morgan, and OptionMetrics.

	Coefficient	t	N	R^2
Panel A: Shape of the Relation				
Full cross-section, linear	0.80	(7.26)	31	53.6
Funding currencies (carry ≤ 0)	0.09	(0.15)	10	0.2
Investment currencies (carry > 0)	0.93	(6.86)	21	55.2
Below median carry	0.28	(0.76)	16	2.8
Above median carry	0.95	(5.36)	15	50.7
Panel B: Subsamples				
Excluding the 3 highest-yielding	0.75	(2.65)	28	29.5
Excluding the 5 highest-yielding	0.97	(2.82)	26	33.9
Excluding the 8 highest-yielding	0.24	(0.79)	23	2.5
Developed currencies only	0.23	(0.72)	10	4.5
Emerging and frontier only	0.88	(6.66)	21	53.9
excluding their own highest 30%	1.31	(1.74)	14	27.5
Panel C: Horse Races				
Carry and sovereign CDS (log)	0.78	(3.72)	24	60.1
Carry and cross-currency basis	0.71	(5.42)	31	55.4
Carry and option-implied skew	0.84	(7.70)	28	61.9
Carry and realized volatility	0.67	(5.25)	31	56.3
Carry and implied volatility	0.64	(4.60)	31	55.9
Carry and emerging-market indicator	0.83	(6.73)	31	53.9
Carry, CDS, basis, and skew	0.67	(2.84)	24	74.6
Panel D: The Interest Differential inside the Survey Expectation				
Interest-differential leg alone	−0.15	(−0.99)	31	2.5
Carry and the interest-differential leg	0.91	(7.50)	31	57.5
Carry before 2011, ϕ after	0.64	(3.12)	29	28.4
Carry after 2011, ϕ before	0.88	(3.76)	29	33.4
Panel E: By Maturity of the Currency-Level Estimate				
One month	0.95	(4.54)	31	36.1
Three months	0.73	(5.17)	31	36.2
One year	0.73	(6.77)	31	53.8
Two years	0.81	(8.75)	31	52.9

Table A.9. Simulation: Forward Filling

This table reports Monte Carlo evidence on the effect of using forward-filled monthly exchange rate forecasts in daily panel regressions. In each simulation, we generate a panel of N currencies observed over T days, where the latent expected excess return depends on a currency fixed effect, a persistent factor-structured regressor, and an idiosyncratic error, with the true slope coefficient β set equal to one. Exchange rate forecasts are observed every 21 trading days and forward-filled to the daily frequency. The imputed daily expected excess return is then constructed using the forward-filled forecast and the daily spot exchange rate. For each simulated panel, we estimate identical regressions using (i) the imputed daily expected excess return and (ii) the latent daily expected excess return as the dependent variable, including currency and calendar-date fixed effects. We perform 10,000 simulations and report the mean estimated slope, its standard deviation, and the root mean squared error (RMSE) relative to the true value of one. We also report the average slope difference, γ. The column "Reject 5%" reports the fraction of simulations in which the null hypothesis $\gamma = 0$ is rejected at the 5% significance level.

Sample length	mean (β)	sdev (β)	RMSE	mean (γ)	Reject 5%
		Panel A: 10 Currencies			
5 years	0.98	0.05	0.06	−0.02	0.17
10 years	0.99	0.04	0.04	−0.01	0.15
30 years	1.00	0.02	0.02	0.00	0.14
		Panel B: 20 Currencies			
5 years	0.98	0.04	0.05	−0.02	0.24
10 years	0.99	0.03	0.03	−0.01	0.22
30 years	1.00	0.02	0.02	0.00	0.20
		Panel C: 33 Currencies			
5 years	0.98	0.04	0.04	−0.02	0.31
10 years	0.99	0.03	0.03	−0.01	0.28
30 years	1.00	0.02	0.02	0.00	0.27