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

Preliminary and Incomplete

September 12, 2023

*We are grateful to Mikhail Chernov, Xiang Fang, Neils Gormsen, Lukas Kremens, Yang Liu, Yueran Ma, Ian Martin, Thomas Maurer, Piotr Orlowski, Tarun Ramadorai, Valeri Sokolovski, conference participants at the 2022 Vienna Symposium on Foreign Exchange Markets, 2022 Doctoral Consortium of the Asian Finance Association, 2023 QRFE Asset Pricing Workshop at Durham University, 2023 ICEEE Congress at the University of Cagliari, and 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 Markets

Abstract

This paper challenges the prevailing notion that investors' preferences remain independent of their investment horizon by uncovering a term structure of risk preferences. Theoretically, we extract a utility-free measure of risk preferences without temporal or horizon restrictions. Empirically, we estimate this measure using professional forecasts and expected risk premia derived from FX option prices. Our analysis of G30 currencies from 1996 to 2020 reveals that the fear of high-order risk is more pronounced in the shorter term, indicating a downward-sloping term structure of investor risk preferences. Moreover, we find that this negative slope becomes more pronounced during adverse times but shifts to an upward slope during favorable periods. These insights offer valuable guidance for enriching existing asset pricing models with horizon-dependent risk preferences, shedding new light on the dynamics of risk premia across different time horizons.

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

1 Introduction

The asset pricing literature has documented that large and time-varying risk premia are pervasive across asset classes, including the FX market. It is now well accepted that variation in risk premia helps explain the uncovered interest parity (UIP) failure, starting with Hansen and Hodrick (1980) and Fama (1984), the cross-section in currency excess returns (e.g., Lustig, Roussanov, and Verdelhan, 2011), the performance of global investment strategies such as the carry trade (e.g., Menkhoff, Sarno, Schmeling, and Schrimpf, 2012), or contribute to exchange rate predictability (e.g., Della Corte, Ramadorai, and Sarno, 2016; Kremens and Martin, 2019; Della Corte, Jeanneret, and Patelli, 2021).

A fundamental question is obviously: What drives such risk premia? In the asset pricing literature, investor risk preferences can be time-invariant (e.g., models with long-run risk and recursive preferences) or time-varying (e.g., models with habit preferences). In either case, risk preferences have a flat term structure, that is the preferences of agents are independent of their forecast horizon. This is a strong assumption that ought to be challenged.[1] For example, shall we expect FX investors to perceive the risk of a currency crash similarly if their investment horizon is one month vs. one year? Probably not. How should the term structure of such perceived risk look like? It remains unclear to date. As these questions illustrate, one reason we still have a poor understanding of risk premia in the FX market is because we have limited knowledge of the term structure of risk preferences and, in particular, how this term structure varies over time.

In this paper, we uncover a set of new facts on the term structure of risk preferences. We first show with theory that we can extract a utility-free measure of risk aversion for FX market

[1]Such assumption may also appear at odds with the empirical evidence that the term structure of risk premia varies over the business cycle, as we observe in the equity market for example (Gormsen, 2021).

participants, without imposing any temporal or horizon restrictions. We empirically estimate this measure by comparing expected exchange rate returns from professional forecasters with exchange rate premia computed from option prices, through the lens of no-arbitrage conditions in the FX market using forecasts over different horizons. We can then explore the term structure of risk aversion and shed light on how it varies across economic/financial conditions.

Our main results are as follows. First, we find that investor preferences reflect a strong risk aversion and sensitivity to high-order risk, thus departing from the log utility benchmark considered in the recent predictability literature (e.g., Kremens and Martin, 2019; Della Corte, Jeanneret, and Patelli, 2021). Second, the unconditional term structure of risk preferences is downward-sloping, that is FX risk premia reflect a greater compensation for risk in the short term than in the long term. Third, a conditional analysis reveals that the negative term structure slope strengthens in bad times, but becomes upward-sloping in good times. Hence, risk aversion is greater in the shorter term during bad times, but greater in the longer term during good times. Our findings thus provide novel insights on the term structure of risk preferences, both unconditionally and conditionally.

We propose a conceptual framework to measure FX risk premia that is utility-free with one free parameter that could be motivated as risk preferences, or ambiguity aversion. Specifically, our theoretical framework exploits an equivalent version of the no-arbitrage condition in asset pricing —existence of the so-called *growth optimal portfolio*, which implies the following risk-neutral expression of the expected excess exchange rate return (or risk premium), for currency i:

$$\mathbb{E}_t\left[\frac{S_{i,T}}{S_{i,t}}\right]-\frac{R_{f,t}}{R^i_{f,t}}=\frac{1}{R_{f,t}}\mathbb{C}\mathrm{ov}^*_t\left(R_{g,T},\frac{S_{i,T}}{S_{i,t}}\right). \tag{1}$$

The risk-neutral covariance term $\mathbb{C}\text{ov}_t^*\left(R_{g,T}, \frac{S_{i,T}}{S_{i,t}}\right)$ captures the conditional risk-neutral covariance between the gross return of the growth optimal portfolio ($R_{g,T}$) and the gross exchange rate return ($\frac{S_{i,T}}{S_{i,t}}$) over the horizon $T-t$.

An important aspect of our framework is to choose to use the power of market return to estimate the theoretical object $R_{g,T}$.[2] We provide detailed motivation and examples in appendix to justify this choice. Under our assumption that the return of the optimal growth portfolio is a power function of the market return

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

the risk-neutral covariance measures the comovements between currency return and ϕ-th moment of market return. A larger value of ϕ implies that higher-order market return risk are priced in the FX market, such as co-skewness ($\phi = 2$), co-kurtosis ($\phi = 3$), and so on.[3] Our approach, therefore, links the risk premium with a directly-interpretable measure of risk preferences, given by ϕ, which we call 'risk aversion' for convenience. It is important to stress that our measure of risk preferences is general and, as a result, nests various versions studied in the literature. For example, ϕ can be interpreted as the constant risk aversion of an unconstrained representative investor with CRRA utility whose entire wealth is invested in the market. We show in the appendix that ϕ can be endogenously time varying in a model with heterogeneous agents (Chan and Kogan, 2002), thereby rationalizing a variation in risk preferences over the business cycle. Our approach thus addresses the critique of Bekaert, Engstrom, and Xu (2021) that most studies estimate time-varying risk aversion measures motivated by models that essentially assume a constant risk aversion coefficient and hence

[2]The growth optimal portfolio has been early considered in Kelly (1956), Roll (1973), Fama and MacBeth (1974), Markowitz (1976), Long (1990), and more recently in Alvarez and Jermann (2005), Martin (2012), Martin and Wagner (2019) in the context of the equity market.

[3]Backus, Chernov, and Martin (2011) show that, when departing from the special case of lognormal distributions, investors are more sensitive to high-order cumulants when their risk aversion increases.

are inherently inconsistent.

To estimate ϕ empirically, we try to integrate information from survey based consensus forecasts and option prices. First, following our theory, we compute a currency-level measure of risk premium as the risk-neutral covariance in (1), using 1-month to 24-month options on the S&P 500 and exchange rates. We then identify, in a panel, the value of ϕ such that our theoretically-implied risk premium best matches observable expected excess exchange rate returns, using 1-month-ahead to 24-month-ahead professional forecasts. Our estimation exploits a cross-section of G30 currencies against the USD over the sample spanning the 1996 to 2020 period. Unconditionally, we obtain a risk aversion estimate of $\phi = 4.25$ that is statistically significantly different from one, which also means that FX investors are highly sensitive to high-order risk.

Our approach allows us to explore the term structure of risk preferences, which we obtain by estimating ϕ using options of different maturities and forecasts of different horizons (between one month and two years). We find that risk aversion decreases from 4.5 at the one-month horizon to 1.75 at the three-month horizon, and to 1.6 at the two-year horizon. That is, investors care less about (higher-order) risk as their forecast horizon increases. One explanation is that the risk of a currency crash matters less to investors over a longer horizon, as a currency has more time to recover following a severe depreciation.[4] At the one-month horizon, however, a currency would not have time to recover following a crash, which translates into severe losses. Investors are then more fearful towards such tail risk events when their horizon is shorter, which is expressed by a higher ϕ.

It is noteworthy that such downward-sloping term structure in risk preferences is in line

[4] Examples of recent currency crashes include the severe depreciation in the Australian and Canadian dollar in Fall 2008 (Great Financial Crisis) and in Spring 2020 (Covid-19 crisis), the fall in British pound in Winter 2016 (Brexit), and the fall in the Russian Ruble in February 2022 (Russian-Ukrainian war). Most of these currencies have fully recovered within a few months.

with existing evidence that the price of higher-order risk is concentrated in the very short term. For example, Dew-Becker, Giglio, Le, and Rodriguez (2017) show that, while the spot variance risk premium in the S&P500 index market is large, forward premia are insignificant at maturities in excess of a month or two.[5] Della Corte, Kozhan, and Neuberger (2021) obtain a similar result in the FX market, which indicates that equity and currency investors display similar pattern in terms of aversion to high-order risk. In addition, Lustig, Stathopoulos, and Verdelhan (2019) show that the profitability of the carry trade declines to zero as the maturity of the bonds increases, which implies a term structure of carry trade risk premia that is also downward-sloping.

We then turn to a conditional analysis of the term structure in risk preferences. To do so, we split our sample according to the degree of financial conditions prevailing in the market and estimate ϕ on each subsample separately for different horizons. Specifically, we consider periods of high vs. low level of CBOE equity-option implied volatility index (VIX) and, alternatively, use the level of the option-implied volatility for a basket of G7 currencies (VXY). We show that the term structure of risk risk aversion has a steep negative slope in times of market stress (high volatility), which turns positive during more favourable times (low volatility).[6] Interestingly, it is the level of risk aversion in the short end that largely drives the fluctuations in the term structure slope. The level of ϕ at a one-month horizon switches from a high level (high risk aversion and attention to high-order risk) during adverse times to a negative level (strong risk tolerance and laxity towards high-order risk) during favorable periods, creating the sign change of the term structure slope.

This paper provides new insights into the term structure of risk preferences, both uncon-

[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), Gormsen (2021), among others.

[6] Our finding can help explain Bansal, Miller, Song, and Yaron (2021)'s finding that, in the equity market, the term structure of expected dividend strip returns is downward sloping in bad times and upward sloping in good times.

ditionally and conditionally. Moreover, our approach contributes to the literature in four ways. First, we consider a utility-free environment to extract a measure of risk preferences, which jointly encompasses risk aversion and attention to high-order risk. While the relation between the optimal growth portfolio and the power of market return is consistent with various model classes, as we show in the paper, our approach is not tied to specific model assumptions. In this regard, our framework generalizes the assumptions made in Kremens and Martin (2019) so we have one more degree of freedom. Second, using an empirical representation of identity (1), we can estimate ϕ by simple OLS regressions. The simplicity of this approach is in contrast to existing methods to extract preferences from macroeconomic data and financial asset prices (e.g., Bekaert, Engstrom, and Xu, 2021; Orłowski, Sokolovski, and Sverdrup, 2021). Third, we use observable expected exchange rate returns to measure the left-hand side of (1). While the literature has typically considered past or ex post realized returns, we instead exploit survey data from professional forecasters. Our approach allows us, therefore, to compare the risk premium computed from forward-looking option prices and the consensus-based expected excess return at the daily frequency and for a cross-section of currencies. In particular, given that forecasts and options are available for different horizons, we can explore how ϕ varies in the short term vs. the more distant future.

Abundant anecdotal evidence supports the notion that individuals tend to exhibit a greater aversion to risks that are proximate in time, as opposed to those that are distant.[7] For example when planning for distant or future travel, a person may be inclined to take risks, explore unfamiliar destinations, and engage in thrilling adventures. However, as the departure date approaches, her behavior tends to shift toward greater risk aversion, particularly with re-

[7] The finding that risk aversion decreases with the temporal horizon is also consistent with various field and laboratory experiments. See, for example, Holt and Laury (2002), Coble and Lusk (2010), Abdellaoui, Diecidue, and Öncüler (2011), Eisenbach and Schmalz (2016) for a review on horizon-dependent risk aversion. Building on these insights, Eisenbach and Schmalz (2016) and Andries, Eisenbach, and Schmalz (2019) develop models that account for investors being more risk averse for shorter horizons. The authors show that allowing for horizon-dependent risk aversion helps resolve various asset pricing puzzles.

gard to immediate travel hazards such as extreme sports or potentially dangerous activities. Likewise, when it comes to immediate health and safety concerns, people frequently adopt a more risk-averse stance, exercising caution in engaging in activities that pose a high risk of injury or illness. Conversely, when confronted with long-term health risks associated with lifestyle choices, such as the development of chronic illnesses due to smoking, individuals may display a higher tolerance for risk. Our contribution is to quantify how risk preferences vary with the investors' horizon and to show a strongly downward-sloping term structure.

Regarding the conditional variation in the term structure slope, our paper relates to a growing literature on extracting time-varying preferences from surveys, experiments, or asset prices. Specifically, our result that risk aversion for short horizons varies countercyclically with economic conditions is consistent with evidence observed in other markets. For example, Guiso, Sapienza, and Zingales (2018) show that investors' risk aversion increases after the 2008 crisis, by comparing the risk premium investors would pay to eliminate a simple gamble. Baker and Wurgler (2006) estimate a time-varying measure of sentiment for stock investors. Cohn, Engelmann, Fehr, and Maréchal (2015) show, in a lab experiment, that investors' fear increases as the financial environment becomes riskier. Pflueger, Siriwardane, and Sunderam (2020) extract a measure of perceived risk from stock investors and show that it varies over the business cycle. Finally, Bekaert, Engstrom, and Xu (2021) and Orłowski, Sokolovski, and Sverdrup (2021) use macro data and financial asset prices to extract a measure of time-varying aggregate risk aversion and the stochastic discount factor, respectively. Consistent with these studies, our measure ϕ increases during recessions and periods of heightened uncertainty, suggesting that FX market participants are more risk averse as economic/financial conditions worsen. A fundamental difference between our paper and this literature, however, is that we can provide new valuable insights on how the term structure of risk preferences varies over time.

The remainder of the paper is organized as follows. Section 2 illustrates our conceptual framework of risk premium in the foreign exchange (FX) market.. Section 3 describes the construction of the main variables and presents our framework to estimate the term structure of risk preferences. Section 4 reports and discusses the results. We conclude in Section 5. The Internet Appendix contains technical details and presents additional results not included in the main body of the paper.

2 Theory

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 foreign currency to ease the notation, this strategy's gross return is equivalent to

$$R_T = \frac{S_T}{S_t} R_{f,t}^{€} \,,$$

where S_t is the spot exchange rate at time t defined as units of dollars per unit of euro such that an increase in S_t 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 fundamental asset pricing equation states

$$\mathbb{E}_t[M_T R_T] = 1 \,,$$

and it implies the following expression of currency excess return

$$\mathbb{E}_t\left[\frac{S_T}{S_t}\right]-\frac{R^{\$}_{f,t}}{R^{€}_{f,t}}=\mathbb{C}\text{ov}_t\left(\frac{-M_T}{\mathbb{E}_t[M_T]},\frac{S_T}{S_t}\right), \tag{2}$$

where $\mathbb{E}_t$ and $\mathbb{C}\text{ov}_t$ are the real-world expectation and covariance operators, respectively, conditional on all information available at time t, $R^{\$}_{f,t}$ 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 assets denominated in dollars at time T with $\mathbb{E}_t[M_T]=1/R^{\$}_{f,t}$.[8]

The identity presented in Equation (2) states that FX investors demand a time-varying risk premium that depends on the conditional covariance between the SDF and the gross exchange rate return. Note that if the SDF were constant conditional on information available at time t (or the physical measure is close to the risk-neutral one), the covariance term would disappear and the expected excess return would be zero — it corresponds to the *Uncovered Interest Parity*. Two aspects of Equation (2) are worth stressing. First, the risk-adjustment component is important to understand currency excess returns and, therefore, cannot be empirically neglected (e.g., Fama, 1984; Lustig, Roussanov, and Verdelhan, 2011). Second, the SDF is unobservable ex-ante and likely to change over time, thus making it difficult to determine how investor preferences exactly shape this risk compensation.

2.1 A risk-neutral representation

We specify the general conceptual framework of equation (1) with a free parameter by assuming the gross return on the growth optimal portfolio[9] is a power function of the gross

[8]We do not need to assume complete markets, which means there may exist alternative SDFs that price assets denominated in dollars. However, all these SDFs must agree with M_T on the prices of future dollar payoffs we focus on.

[9]The optimal growth portfolio can also be seen as a levered market portfolio, with $R_{g,T}=R^{\phi}_{m,T}$ and ϕ measuring the degree of leverage (e.g., Long, 1990; Martin, 2017).

return on the market portfolio

$$R_{g,T} = R^{\phi}_{m,T}.$$

Proposition 1. *Starting from the existence of a growth optimal portfolio, i.e.* $M_T R_{g,T} = 1$, *and the assumption* $R_{g,T} = R^{\phi}_{m,T}$, *we can rewrite Equation (2) as follows:*

$$\mathbb{E}_t\left[\frac{S_T}{S_t}\right] - \frac{R^{\$}_{f,t}}{R^{€}_{f,t}} = \underbrace{\mathbb{C}\mathrm{ov}^{\star}_t\left(\frac{R^{\phi}_{m,T}}{\mathbb{E}^{\star}_t[R^{\phi}_{m,T}]}, \frac{S_T}{S_t}\right)}_{\mathrm{ERP}^{\phi}_t}, \tag{3}$$

where $\mathbb{E}^{\star}_t$ *and* $\mathbb{C}ov^{\star}_t$ *are the risk-neutral expectation and covariance operators, respectively, conditional on all information available at time* t, *and* $\mathbb{E}^{\star}_t[R^{\phi}_{m,T}] = R^{\$}_{f,t}$.

Proof. See online appendix B.1, Della Corte, Gao, and Alexandre (2023). □

We hereafter refer to the above risk-neutral covariance term as the expected risk premium or simply ERP^{ϕ}_t.

We could show that this specification is consistent with various theoretical models. For example, in the case of an unconstrained CRRA agent who invests in the market, the solution of the static portfolio choice problem implies that $R_{g,T} = R^{\phi}_{m,T}$ is the growth optimal portfolio return (Martin, 2017), confirming that the return of the optimal growth portfolio is proportional to a power function of the market return. A similar result can be obtained in a setting with ambiguity aversion (Hansen, 2007). In addition, it is easy to show that ϕ can be endogenously time varying in a model with heterogeneous agents, as in Chan and Kogan (2002) and Longstaff and Wang (2012), which will later motivate our conditional analysis of risk preferences. Our online appendix (section A) Della Corte, Gao, and Alexandre (2023) provides details for each of these model classes.

We could derive the following explicit expression for the expected risk premium ERP_t^{ϕ}.

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

$$\mathrm{ERP}_t^{\phi} = \phi\,\mathrm{ERP}_t + \frac{\phi(\phi-1)}{2R_{f,t}^{\$}}\mathbb{C}\mathrm{ov}_t^{\star}\left(\xi_T^{\phi-2}(R_{m,T}-1)^2, \frac{S_T}{S_t}\right), \tag{4}$$

where

$$\mathrm{ERP}_t^{1} = \frac{1}{R_{f,t}^{\$}}\mathbb{C}\mathrm{ov}_t^{\star}\Big(R_{m,T}, \frac{S_T}{S_t}\Big), \tag{5}$$

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

Proof. See online appendix B.2. □

Proposition 2 shows that ERP_t^{ϕ} is equivalent to ϕ times ERP_t plus a non-linear term that increases with ϕ. This non-linear term can be positive or negative depending on the characteristics of the currency (being a safe haven or rather speculative, for example). Fan, Londono, and Xiao (2022) shows the equity tail risk is priced in the currency market.

Note that, in the special case of $\phi = 1$, the expected risk premium corresponds to the *quanto risk premium* of Kremens and Martin (2019). Another simple special case is when $\phi = 2$, and we could simplify the right-hand side of (4) as

$$\mathrm{ERP}_t^{2} = 2\,\mathrm{ERP}_t + \frac{1}{R_{f,t}^{\$}}\mathbb{C}\mathrm{ov}_t^{\star}\left((R_{m,T}-1)^2, \frac{S_T}{S_t}\right), \tag{6}$$

where the first one amounts to twice ERP_t, and the second one captures a non-linear term reflecting how the currency comoves with market variance. The latter 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 squared of the VIX index).

We now provide an interpretation of ϕ as a measure of investors' risk preferences. A key aspect of our framework is that we do not need to impose any utility function. However, it can be convenient to think of the power utility case with lognormal shocks for the interpretation of the results. In this case, the measure of risk preferences, ϕ, corresponds to the level of risk aversion. So the higher the risk aversion, the higher the expected risk premium, as one would expect. However, we find that the risk premium does not linearly increase with risk aversion, as it would be the case with the lognormal model. We indeed allow for any distribution of market returns and exchange rates, which implies that investors also care more about high-order risk when their risk aversion increases (see Backus, Chernov, and Martin, 2011). For convenience, we hereafter use the term 'risk aversion' when referring to ϕ, although it can also encompass attention to high-order risk such as co-skewness ($\phi = 2$), co-kurtosis ($\phi = 3$), and so on. We discuss the role of these high-order moments in online appendix (see section B.3 in Della Corte, Gao, and Alexandre (2023)).

In sum, we have used the properties of the growth optimal portfolio to derive a currency's expected risk premium in function of a utility-free measure of risk preferences, ϕ. The level of ϕ reflects the degree of risk aversion, as well as the attention to high-order risk when one departs from the lognormal case. Our objective in the rest of the paper is to extract a T-dependent estimate of ϕ from the data and to shed new light on the term structure of risk preferences in the FX market.

3 Empirical Methodology

In this section, we present our approach to empirically estimate ϕ, based on the framework developed in Section 2. We then describe the construction of the main variables, using professional forecasts to measure expected currency excess returns and option data to estimate the expected currency risk premium. Finally, we discuss the main contributions of our empirical approach, before presenting the results in Section 4.

3.1 Approaches to empirically estimate ϕ

We can write an empirical representation of Equation (3) as follows

$$\underbrace{\mathbb{E}_t\left[\frac{S_{i,T}}{S_{i,t}}\right] - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}}}_{\text{ERX}_{i,t,T}} = \alpha_\phi + \beta_\phi \underbrace{\mathbb{C}\text{ov}^*_t\left(\frac{R^{\phi}_{m,T}}{\mathbb{E}^*_t[R^{\phi}_{m,T}]}, \frac{S_{i,T}}{S_{i,t}}\right)}_{\text{ERP}^{\phi}_{i,t,T}}, \tag{7}$$

where $\text{ERX}_{i,t,T}$ denotes the expected excess return and $\text{ERP}^{\phi}_{i,t,T}$ is the expected risk premium for a given ϕ, both measured at time t over the horizon $T-t$ for currency i relative to the US dollar.

The expected risk premium we derive in Section 2 implies that $\alpha_\phi = 0$ and $\beta_\phi = 1$ in the empirical representation (7). We can thus exploit this condition in our empirical analysis to estimate the coefficient ϕ. Specifically, we propose to estimate specification (7) by OLS for different values of ϕ, using a panel of currencies, and then select the regression slope coefficient β_ϕ that is closest to one.[10]

[10]We do not use the condition $\alpha_\phi = 0$ because α_ϕ is expected to vary across currencies, capturing trading frictions or transaction costs that typically differ across currency pairs.

Our theory also suggests a complementary approach for inferring ϕ from the data. If the non-linear term in Equation (4) is negligible, we have $\mathrm{ERX}_t^{\phi} \approx \phi\, \mathrm{ERP}_t^{(1)}$ and we can then use a simplified version of the specification (7):

$$\mathrm{ERX}_{i,t,T} = \alpha + \beta\, \mathrm{ERP}_{i,t,T}^{(1)}, \tag{8}$$

where $\mathrm{ERP}_{i,t,T}^{(1)}$ represents the expected risk premium for currency i in the case of $\phi = 1$. In this case, we can infer ϕ from the estimate of β directly. If, however, the non-linear term becomes non-negligible, e.g., in periods of financial market stress, the estimate of β would not provide a good approximation of ϕ. We can easily compare the ϕ estimated obtained using both approaches and, thus, assess the importance of the non-linear term in Equation (4).

We now describe how we construct, for each currency pair, the dependent variable (the expected excess return) and the independent variable (the expected risk premium) used to estimate specifications (7) and (8).

3.2 Measuring expected currency excess returns

For each currency i, we measure expected currency excess return as follows

$$\mathrm{ERX}_{i,t,T} = \frac{\mathbb{E}_t\left[S_{i,T}\right]}{S_{i,t}} - \frac{R_{f,t}^{\$}}{R_{f,t}^{i}}, \tag{9}$$

where $\mathbb{E}_t\left[S_{i,T}\right]$ is the exchange rate forecast at time t over the horizon $T - t$ and $S_{i,t}$ is the spot exchange rate at time t. For the former, we collect monthly forecasts with a horizon of 1, 3, 12, and 24 months from the Foreign Exchange Consensus Forecasts Database, a comprehensive monthly survey compiled by Consensus Economics; we then use a linear extrapolation method to retrieve daily forecasts. For the latter, we use WM/Refinitiv Spot Rates from Datastream. The last term in Equation (9) is the ratio of gross interest rates observed at time t with maturity $T-t$ for currency i ($R^i_{f,t}$) and the US dollar ($R^{\$}_{f,t}$). For the construction of these components, we rely on daily zero-coupon rates bootstrapped from money market rates and interest rate swaps obtained from Bloomberg. Ultimately, we match the maturity of interest rates to that of exchange rate forecasts.

3.3 Measuring the expected currency risk premium

We now describe the construction of the expected risk premium, which is given by

$$\mathrm{ERP}^{\phi}_{i,t,T} = \mathbb{C}\mathrm{ov}^*_t\left(\frac{R^{\phi}_{m,T}}{\mathbb{E}^*_t[R^{\phi}_{m,T}]}, \frac{S_{i,T}}{S_{i,t}}\right). \tag{10}$$

The above risk-neutral covariance is not directly observable from market prices, except for the case of $\phi = 1$, which is discussed in Kremens and Martin (2019). To overcome this challenge, we decompose the risk-neutral covariance into its three distinct components:

$$\mathbb{C}\mathrm{ov}^*_t\left(\frac{R^{\phi}_{m,T}}{\mathbb{E}^*_t[R^{\phi}_{m,T}]}, \frac{S_{i,T}}{S_{i,t}}\right) = \rho^*_{\phi,i,t}\sqrt{\mathrm{var}^*_t\left(\frac{R^{\phi}_{m,T}}{\mathbb{E}^*_t[R^{\phi}_{m,T}]}\right)}\sqrt{\mathrm{var}^*_t\left(\frac{S_{i,T}}{S_{i,t}}\right)}, \tag{11}$$

where $\rho^*_{\phi,i,t}$ captures the risk-neutral correlation between $R^{\phi}_{m,T}$ and $S_{i,T}/S_{i,t}$, while the two var^*_t terms denote the risk-neutral variance of the stock market return (for different levels of ϕ) and that of exchange rate returns, respectively. These quantities are measured at time t over the same horizon $T-t$. While the risk-neutral variances can be constructed using the Breeden-Litzenberger method applied to equity index and currency options (see Appendix C.2), the risk-neutral correlations for different level of ϕ are not directly observable. We instead use the backward-looking realized correlation between $R^{\phi}_{m,T}$ and $S_{i,T}/S_{i,t}$ between times $t-T$ and t as a proxy. For example, if $\rho^*_{\phi,i,t}$ is the 1-year risk-neutral correlation at time t, our proxy is the 1-year realized correlation measured over the previous year.[11]

Equity index option prices are based on the daily implied volatility surface of SPX European options, as provided by OptionMetrics. We use observations from January 1996 to December 2020. We also take the yield curve term structure from OptionMetrics. We revert the implied volatility back to option prices to compute the risk-neutral moments of the S&P 500 index returns with maturities of 1, 3, 12 and 24 months.

FX option prices are converted from implied volatility data collected over-the-counter (OTC) currency options from JP Morgan and Bloomberg. The quoted implied volatilities, in terms of Garman and Kohlhagen (1983), are on baskets of constant maturity plain vanilla options for fixed deltas (δ). From these data, we recover the implied volatility smile ranging from a 10δ put to a 10δ call option. To convert deltas into strike prices and implied volatilities into option prices, we employ exchange rates and zero-yield rates obtained by bootstrapping money market rates and interest rate swap data from Datastream and Bloomberg.

Using the realized correlation and the risk-neutral variances, we can finally compute the daily expected risk premium $\text{ERP}^{\phi}_{i,t,T}$ for each of the 30 currency pairs and different levels of ϕ between 1 and 24 months.

[11] We show in Appendix C.1 that the risk-neutral correlation, implied from the Quanto contracts used in Kremens and Martin (2019), is close to the realized correlation between the market and exchange rate returns. So we can conclude that the realized correlation is a reasonable proxy for the risk-neutral correlation.

3.4 Advantages of our methodology

Our methodology contributes to the literature in four major ways. First, we consider a model-free environment to extract risk preferences. While the relation between the optimal growth portfolio and the power of market return is consistent with various model classes, as discussed in Appendix A, our approach is not tied to specific model assumptions. In this regard, our framework generalizes Kremens and Martin (2019), which builds on investors having log utility. Second, using the two empirical representations provided in Section 3.1, we can estimate ϕ by comparing the simple OLS regressions with $\mathrm{ERP}^{\phi}_{i,t,T}$ when the non-linear term is negligible. The simplicity of this approach is in contrast to existing methods to extract preferences from macroeconomic data and financial asset prices, as in Bekaert, Engstrom, and Xu (2021) or Orłowski, Sokolovski, and Sverdrup (2021) for example. Third, we use observable expected exchange rate return to measure the left-hand side of Equations (7) and (8), i.e., $\mathrm{ERX}_{i,t,T}$. While the literature has typically considered past or ex post realized returns, we instead exploit survey data from professional forecasters. Our approach allows us, therefore, to compare the expected risk premium ($\mathrm{ERP}^{\phi}_{i,t,T}$) computed from forward-looking option prices and the consensus-based expected excess return ($\mathrm{ERX}_{i,t,T}$) at the daily frequency and for a cross-section of currencies. Last but not least, given that forecasts and options are available for different horizons, we can shed light on the term structure of our risk aversion measure ϕ.

4 Empirical Results

In this section, we estimate ϕ using the two approaches presented in Section 3.1, both unconditionally and conditionally, as well as for different horizons. Our analysis allows us to quantify the aversion of FX market participants to high-order risk, explore how such aversion varies over time, and shed light on the term structure of risk preferences.

4.1 Unconditional analysis

We start our empirical analysis by estimating the unconditional level of ϕ. We do so by first pooling together expected excess returns and expected risk premia across all maturities and then running panel regressions based on the following specification:

$$\mathrm{ERX}_{i,t,T} = \alpha_i + \alpha_\ell + \alpha_t + \beta_\phi \, \mathrm{ERP}^{\phi}_{i,t,T} + \varepsilon_{i,t,T}, \tag{12}$$

where $\mathrm{ERX}_{i,t,T}$ is the expected currency excess return observed at time t over the horizon $T-t$ for currency i and $\mathrm{ERP}^{\phi}_{i,t,T}$ is the expected currency risk premium computed at time t for the same currency pair/maturity. We complement our specification with currency fixed effects (α_i) that control for time-invariant differences in exchange rate forecasts, maturity fixed effects (α_ℓ) to control for average variation in forecasts across different horizons, and time fixed effects (α_t) that control for unobservable time-variant global factors driving exchange rate forecasts. We use daily observations and report our results in Table 1 for a large cross-section of 30 currency pairs (Panel A) as well as for a subset of 20 developed and most liquid emerging market currency pairs (Panel B).

In Column (1), we first run the specification (12) for different values of ϕ, using a fine grid between 0 and 6, and then report the value of ϕ associated with an estimate of β_ϕ that is closest to 1. This corresponds to our first approach. We find an estimate of ϕ equal to 4.25 and 4.55 in Panels A and B, respectively. Both estimates are highly statistically significant from $\phi = 0$ (case of risk-neutrality) and $\phi = 1$ (log utility). For completeness, the remaining columns present estimates of β_ϕ for selected values of ϕ ranging between 1 and 6. We can clearly see that β_ϕ is about 1 for values ϕ between 4 and 5, but quite different from 1 outside this range. This analysis provides confidence to the point estimates of ϕ we obtain in Column (1).

Table 1 about here

Alternatively, we can exploit our second approach for estimating ϕ. Recall that if the risk-neutral co-skewness between exchange rate returns and the stock market variance is not too large, we can also infer ϕ from the slope coefficient β_ϕ in the case of $\phi = 1$, which corresponds to regressing $\mathrm{ERX}_{i,t,T}$ on $\mathrm{ERP}^1_{i,t,T}$. Based on this method, Column (2) of Table 1 shows that the estimate of ϕ is 3.87 in Panel A and 4.17 in Panel B. The results imply that this simple alternative approach is a good approximation to the baseline approach, which requires a grid search over hundreds of panel regressions. In addition, the explanatory power (R^2) is quite high and remarkably similar across both approaches (e.g., 18% and 18.2% in Columns 1 and 2, respectively).

Overall, both cases imply a level of ϕ around 4, which means that FX market participants are strongly averse to higher-order risk. Recall from Section **??** that the higher the ϕ, the more investors put weights to (risk-neutral) higher-order moments when computing currency risk premia. Specifically, when $\phi = 4$, FX participants require a positive (negative) currency premium that largely reflects the negative (positive) comovement between the currency and market variance. Such risk preferences imply a significant departure from the log utility case (i.e., $\phi = 1$) considered in the recent exchange rate predictability literature (e.g., Kremens and Martin, 2019; Della Corte, Jeanneret, and Patelli, 2021).

4.2 Time variation in risk preferences

We now investigate how the risk preferences of FX investors vary over time. A recent strand of empirical studies suggests that risk aversion increases in "bad times". For example, surveys indicate that investors are willing to pay a higher risk premium to eliminate a simple gamble after (compared to before) the 2008 crisis (Guiso, Sapienza, and Zingales, 2018). Also, investors' fear appears to increase as financial conditions become riskier, as reported in a lab experiment by Cohn, Engelmann, Fehr, and Maréchal (2015). Investors' perceived risk, as measured by comparing the valuation of stocks with different volatility, decreases as economic conditions improve (Pflueger, Siriwardane, and Sunderam, 2020). Stock investors are also found to have higher risk aversion in times of greater market uncertainty (Bekaert, Engstrom, and Xu, 2021). Surprisingly, however, we thus far have limited knowledge of how the aversion to higher-order risk changes over different market conditions. We thus turn our attention to how ϕ varies over time.

Figure 1 displays the one-year rolling estimate of ϕ, based on the second approach (i.e., using $\text{ERP}^1_{i,t,T}$ as the independent variable in specification (12)). We can see that ϕ varies substantially over time and appears to be elevated in multiple instances of FX market stress, including NBER recessions. In particular, we observe a strong increase in the aversion to higher-order risk during the Russian debt crisis in 1998, during the severe carry trade reversal in 2008, during the European debt crisis in 2011-2012, when the Fed started a monetary tightening phase in 2015, or at the peak of the Covid-19 crisis in 2020, when major commodity currencies have depreciated strongly against the US dollar. We can thus conclude that FX participants' aversion to higher-order risk tends to peak at times currencies experience large movements and are thus more subject to crashes.

Figure 1 about here

4.3 The term structure of risk preferences

A key advantage of our methodology is that we can construct expected excess returns ($\text{ERX}_{i,t,T}$) and the expected risk premium ($\text{ERP}^{\phi}_{i,t,T}$) over different horizons. We are thus able to explore the term structure of risk preferences, as given by estimated values of ϕ for each forecast horizon. For this analysis, we run one panel regression by forecast horizon using the following specification:

$$\text{ERX}_{i,t,T} = \alpha_{i,\tau} + \alpha_{t,\tau} + \beta_{\phi,\tau}\,\text{ERP}^{\phi}_{i,t,T} + \varepsilon_{i,t,T}, \tag{13}$$

where $\alpha_{i,tau}$ and $\alpha_{t,tau}$ are respectively currency and time fixed effects for a given horizon $\tau = T - t$. Table A.4 reports estimates based on four forecast horizons, i.e., 1 month in Panel A, 3 months in Panel B, 1 year in Panel C, and 2 years in Panel D. Two results are worth mentioning at this stage. First, all estimates of ϕ are highly statistically significant. Second, the explanatory power increases significantly when estimating ϕ for a specific maturity, compared to a pooled analysis: the R^2 in Column (1) of Table A.4 ranges between 32.6% and 53.5% vs. about 18% in Table 1. These findings confirm the importance of considering maturity-based estimates of ϕ.

Table A.4 about here

When comparing Panel A and Panel D, we can see that the term structure is unconditionally downward-sloping. In particular, Column (1) indicates that ϕ decreases from 4.80 at the one-month horizon to 1.72 at the two-year horizon. This finding suggests that investors care less about higher-order risk as their forecast horizon increases. One explanation is that the risk of a currency crash becomes less relevant over a longer horizon, as a currency has more time to recover following a severe depreciation. At the one-month horizon, however, a currency would not have time to recover following a crash, which translates into severe losses to FX market participants. Over the short term, investors are then more averse to such tail risk events, which is expressed by a higher level of ϕ.

4.3.1 Robustness analysis

We now summarize a set of robustness analyses that corroborate our previous findings. First, we show that the term structure of ϕ remains qualitatively similar when adding a set of control variables $X_{i,t}$ to our panel specification, i.e., we run panel regressions based on

$$\mathrm{ERX}_{i,t,T} = \alpha_{t,\tau} + \alpha_{i,\tau} + \beta_{\phi,\tau}\, \mathrm{ERP}^{\phi}_{i,t,T} + {\delta_\tau}' X_{i,t} + \varepsilon_{i,t,T}, \tag{14}$$

where $X_{i,t}$ includes the interest rate differential between the US and country i at time t, the dollar basis constructed at time t using the one-month US dollar interest rate, and the year-on-year inflation differential between the US and country i at time t. We also control for the realized covariance of exchange rate changes with the negative reciprocal of the S&P 500 return, observable at time t and computed between times $t-T$ and t, as in Kremens and Martin (2019). We report the estimates in Table 3 and show in Column (1) that ϕ decreases from 4.43 at the one-month horizon to 3.30 at the two-year horizon.

Table 3 about here

In Table 4, we employ expected exchange return (as opposed to currency expected excess returns) as the dependent variable, adding the corresponding interest rate differential as an explanatory variable, in our panel regressions:

$$\mathrm{EFX}_{i,t,T} = \alpha_{t,\tau} + \alpha_{i,\tau} + \beta_{\phi,\tau}\,\mathrm{ERP}^{\phi}_{i,t,T} + \gamma_{\tau}\,\mathrm{IRD}_{i,t,T} + \varepsilon_{i,t,T}, \tag{15}$$

where $\mathrm{EFX}_{i,t,T} = \mathbb{E}_t\left[S_{i,T}\right]/S_{i,t}$ is the expected exchange rate return based on consensus forecasts observed at time t and $\mathrm{IRD}_{i,t,T} = R^{\$}_{f,T}/R^{i}_{f,T} - 1$ is the interest rate differential at time t. We find that estimates of ϕ remain remarkably close to those reported in Table A.4. For example, Column (1) shows that ϕ decreases from 4.95 at the one-month horizon to 1.68 at the two-year horizon. These results confirm that FX participants are more (less) averse to tail risk events when their investment horizon is shorter (longer).

Tables 4 about here

4.3.2 Conditional term structure

We then explore how the term structure of risk preferences varies across economic and financial conditions. To do so, we split our sample according to NBER-dated recessions and expansions, estimate the specification (13) on each subsample separately, and then plot the estimates of ϕ in the top-left Panel of Figure 2. Consistent with the preliminary evidence of Figure 1, we find that ϕ increases during recessions, suggesting that FX market participants are indeed more averse to higher-order risk as economic conditions worsen. More interestingly, the term structure has a steep negative slope in recessions, while it becomes much flatter (and also upward-sloping) during expansions.

We find similar results when we analyze risk preferences across different measures of financial conditions. The top-right Panel, for example, shows the results when we separate the sample by high and low levels of CBOE equity-option implied volatility index (VIX), based on the sample mean. The bottom panels replace the VIX index with the option-implied volatility for a basket of G7 currencies (VXY) and the implied volatility on one-month U.S. Treasury options (MOVE), respectively. In all cases, the term structure of risk preferences is countercyclical with respect to aggregate economic/financial conditions.

FIGURE 2 ABOUT HERE

Overall, our paper contributes to the literature by providing new insights into how the aversion to higher-order risk varies over different horizons as well as into the dynamics of its term structure. In particular, we find that fear of high-order risk is greater in the shorter term during bad times, but becomes greater in the longer term during good times.

5 Concluding Remarks

This paper sheds light on the FX risk premium and the term structure of risk preferences. We first show theoretically that we can extract a utility-free measure of risk preferences for FX market participants. We then estimate this measure by comparing expected exchange rate returns from professional forecasters with exchange rate premia computed from option prices, through the lens of no-arbitrage condition in the FX market. We can then explore how the term structure of risk preferences varies across economic/financial conditions.

The main results are as follows. Investor preferences reflect a strong aversion to high-order risk, thus departing from the log utility considered recently (e.g., Kremens and Martin, 2019; Della Corte, Jeanneret, and Patelli, 2021). Unconditionally, the term structure of risk preferences is downward-sloping, that is FX risk premia provide a greater compensation for high-order risk as the forecast horizon decreases. Conditionally, this negative term structure slope strengthens in bad times, but becomes upward-sloping in good times. Hence, fear of high-order risk is greater in the shorter term during bad times, but greater in the longer term during good times. We therefore provide novel insights on the conditional term structure of risk preferences.

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, 2019, Horizon-dependent risk aversion and the timing and pricing of uncertainty, *FRB of New York Staff Report*.

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, Shane Miller, Dongho Song, and Amir Yaron, 2021, The term structure of equity risk premia, *Journal of Financial Economics* 142, 1209–1228.

Bekaert, Geert, Eric C. Engstrom, and Nancy R. Xu, 2021, The time variation in risk appetite and uncertainty, *Management Science*.

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

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

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.

Della Corte, Pasquale, Can Gao, and Jeanneret Alexandre, 2023, Online appendix for ”title of paper”, *Working Paper.*

Della Corte, Pasquale, Alexandre Jeanneret, and Ella Patelli, 2021, A credit-based theory of the currency risk premium, *Journal of Financial Economics,* Forthcoming.

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.

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.

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).

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

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

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.

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.

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

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.

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.

Orłowski, Piotr, Valeri Sokolovski, and Erik Sverdrup, 2021, Benchmark currency stochastic discount factors, *Working Paper, HEC Montreal*.

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

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

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

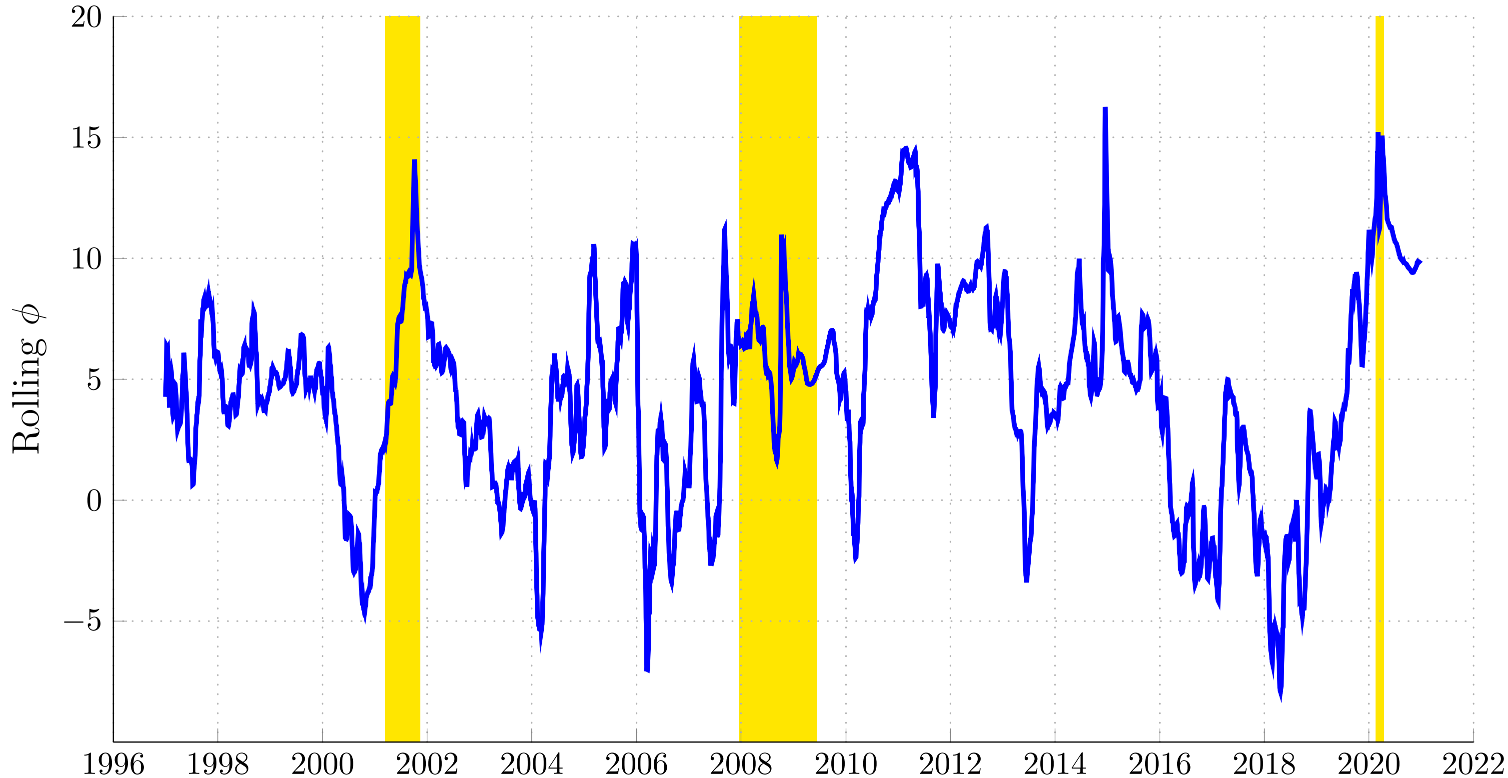

Figure 1. Time-varying measure of risk preferences, ϕ

This figure displays the time variation in the risk preferences of FX participants. The series reflects one-year rolling estimates of ϕ implied from the following specification

$$\mathrm{ERX}_{i,t,T} = \alpha_i + \alpha_\ell + \alpha_t + \beta_1 \, \mathrm{ERP}^1_{i,t,T} + \varepsilon_{i,t,T},$$

where the slope coefficient β_1 is the proxy for ϕ, as described in Section 3.1, while α_i, α_ℓ, and α_t denote currency, maturity, and time fixed effects, respectively. $\mathrm{ERX}_{i,t,T}$ is the expected currency excess return observed at time t over the horizon $T - t$ for currency i, calculated using exchange rate consensus forecasts net of interest rate differentials (see Section 3.2). $\mathrm{ERP}^1_{i,t,T}$ is the expected currency risk premium computed at time t, for the same currency pair/maturity, in the case of $\phi = 1$ based on S&P 500 index and currency options with maturities ranging between one month and two years (see Section 3.3). The sample runs at the daily frequency between January 1996 and December 2020 for a cross-section of 30 currency pairs relative to the US dollar. We use linear extrapolation to retrieve daily forecasts from monthly forecasts. Shaded areas denote NBER-dated recessions. Data are from Bloomberg, FRED, JP Morgan, and OptionMetrics.

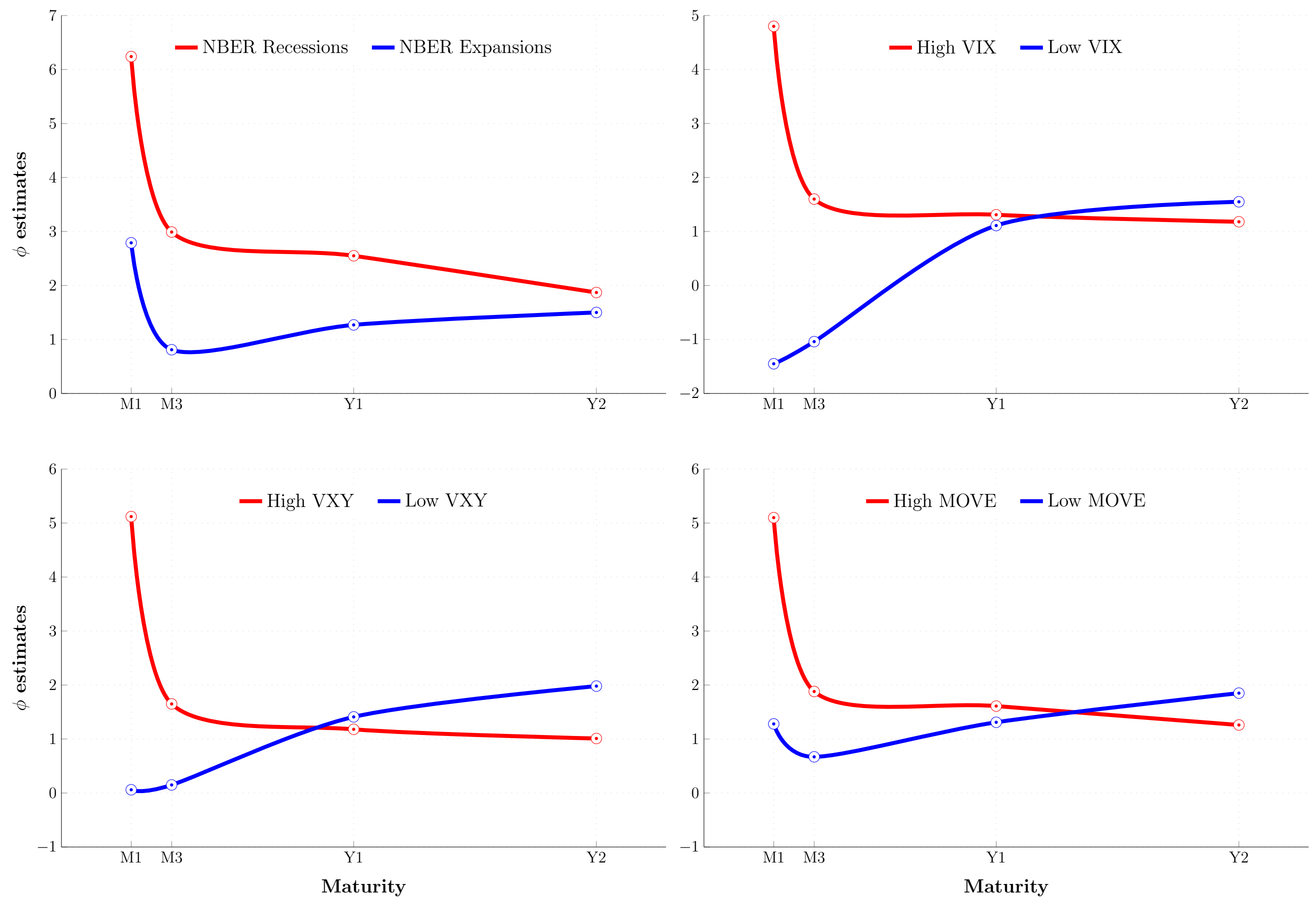

Figure 2. Term structure of risk preferences ϕ by market conditions

This figure displays the conditional term structure of risk preferences, based on panel estimates of ϕ across different market conditions. The estimated specification is Equation (13) for maturities equal to 1 month, 2 months, 1 year, and 2 years. The top-left panel reports results during NBER recession and expansion periods. The top-right panel relates to low and high CBOE equity-option implied volatility index (VIX) periods, the bottom-right panel to low and high option-implied volatility for a basket of G7 currencies (VXY) periods, and the bottom-left panel to low and high implied volatility on one-month U.S. Treasury options (MOVE) periods, all three defined relative to the sample average. The sample runs at the daily frequency between January 1996 and December 2020. Data are from Bloomberg, FRED, JP Morgan, and OptionMetrics.

Table 1. Panel regression estimates of the risk preferences ϕ

This table presents estimates of the risk preferences ϕ from a panel regression of expected currency excess returns on expected currency risk premium. The results are obtained from the following specification

$$\text{ERX}_{i,t,T} = \alpha_i + \alpha_\ell + \alpha_t + \beta_\phi \,\text{ERP}^{\phi}_{i,t,T} + \varepsilon_{i,t,T},$$

where $\text{ERX}_{i,t,T}$ is the expected currency excess return observed at time t over the horizon $T-t$ for currency i and $\text{ERP}^{\phi}_{i,t,T}$ is the expected currency risk premium computed at time t for the same currency pair/maturity. α_i, α_ℓ, and α_t denote currency, maturity, and time fixed effects (FE), respectively. $\text{ERX}_{i,t,T}$ is calculated using exchange rate consensus forecasts net of interest rate differentials, whereas $\text{ERP}^{\phi}_{i,t,T}$ is based on S&P 500 index and currency options for different levels of ϕ and maturities ranging between one month and two years. In Column (1), we first run the above specification for different values of ϕ and then report the value of ϕ associated with $\beta_\phi = 1$. In Column (2), we estimate ϕ from the slope coefficient β_ϕ in the case of $\phi = 1$. The remaining columns present estimates of β_ϕ for selected values of ϕ ranging between 2 and 6. Standard errors, reported in parenthesis, are clustered by currency and time. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. The sample runs at the daily frequency between January 1996 and December 2020 for a cross-section of 30 (20) currency pairs relative to the US dollar in Panel A (B). We use linear extrapolation to retrieve daily forecasts from monthly forecasts. Data are from Bloomberg, FRED, JP Morgan, and OptionMetrics.

	Grid	*Results for different ϕ*					
	$\beta_\phi = 1$	$\phi = 1$	$\phi = 2$	$\phi = 3$	$\phi = 4$	$\phi = 5$	$\phi = 6$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: All Currencies							
ϕ	4.25***						
	(0.85)						
ERP^{ϕ}		3.87***	2.01***	1.38***	1.06***	0.86***	0.73***
		(0.71)	(0.38)	(0.27)	(0.21)	(0.18)	(0.15)
R^2 (%)	18.0	18.2	18.1	18.0	18.0	17.9	17.9
N	604,802	604,802	604,802	604,802	604,802	604,802	604,802
Panel B: Most Liquid Currencies							
ϕ	4.55***						
	(0.96)						
ERP^{ϕ}		4.17***	2.16***	1.48***	1.13***	0.92***	0.77***
		(0.80)	(0.42)	(0.30)	(0.23)	(0.19)	(0.17)
R^2 (%)	19.4	19.7	19.6	19.5	19.5	19.4	19.3
N	436,492	436,492	436,492	436,492	436,492	436,492	436,492
Currency FE	✓	✓	✓	✓	✓	✓	✓
Maturity FE	✓	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓	✓

Table 2. Term structure of risk preferences ϕ

This table presents estimates of the risk preferences ϕ from a panel regression of expected currency excess returns on expected currency risk premium for different horizons. The results are obtained from the following specification

$$\mathrm{ERX}_{i,t,T} = \alpha_t + \alpha_i + \beta_\phi \, \mathrm{ERP}^{\phi}_{i,t,T} + \varepsilon_{i,t,T},$$

where $\mathrm{ERX}_{i,t,T}$ is the expected currency excess return observed at time t over the horizon $T-t$ for currency i and $\mathrm{ERP}^{\phi}_{i,t,T}$ is the expected currency risk premium computed at time t for the same currency pair/maturity. α_i and α_t denote currency and time fixed effects (FE), respectively. $\mathrm{ERX}_{i,t,T}$ is calculated using exchange rate consensus forecasts for a given horizon net of interest rate differentials, whereas $\mathrm{ERP}^{\phi}_{i,t,T}$ is based on S&P 500 index and currency options for different levels of ϕ and a corresponding maturity. In Panel A, we first run the above specification for different values of ϕ and then report the value of ϕ associated with $\beta_\phi = 1$. In Panel B, we estimate ϕ from the slope coefficient β_ϕ in the case of $\phi = 1$. Standard errors, reported in parenthesis, are clustered by currency and time. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. The sample runs at the daily frequency between January 1996 and December 2020 for a cross-section of 30 currency pairs relative to the US dollar. We use linear extrapolation to retrieve daily forecasts from monthly forecasts. Data are from Bloomberg, FRED, JP Morgan, and OptionMetrics.

Panel A: First Approach (Grid search for $\beta_\phi = 1$)				
	Forecast horizon			
	1 month (1)	3 months (2)	12 months (3)	24 months (4)
ϕ	4.54***	1.75***	1.74***	1.56***
	(0.82)	(0.54)	(0.43)	(0.47)
R^2 (%)	32.6	32.7	45.3	53.5
N	153,975	155,714	155,540	139,573
Panel B: Second Approach (Slope when $\phi = 1$)				
ERP^1	4.80***	1.78***	1.74***	1.72***
	(0.96)	(0.57)	(0.43)	(0.53)
R^2 (%)	32.6	32.7	45.3	53.5
N	153,975	155,714	155,540	139,573
Currency FE	✓	✓	✓	✓
Time FE	✓	✓	✓	✓

Table 3. Term structure of risk preferences ϕ – with controls

This table presents panel regression estimates based on the following specification

$$\mathrm{ERX}_{i,t,T} = \alpha_t + \alpha_i + \beta_\phi \, \mathrm{ERP}^{\phi}_{i,t,T} + \delta' X_t + \varepsilon_{i,t,T},$$

where $\mathrm{ERX}_{i,t,T}$ is the expected currency excess return observed at time t for currency i and maturity $T-t$ and $\mathrm{ERP}^{\phi}_{i,t,T}$ is the expected currency risk premium computed at time t for the same currency pair/maturity, with α_t and α_i denoting time and currency fixed effects, respectively. X_t refers to control variables, i.e., the interest rate differential between the US and country i, the year-on-year inflation differential between the US and country i at time t, and the realized covariance of exchange rate changes with the negative reciprocal of the S&P 500 return observable at time t and computed between times $t-T$ and t. $\mathrm{ERX}_{i,t,T}$ is calculated using exchange rate consensus forecasts and interest rate differentials whereas $\mathrm{ERP}^{\phi}_{i,t,T}$ is based on S&P 500 and currency options for different levels of ϕ and maturities ranging between one month and two years. Standard errors, reported in parenthesis, are clustered by currency and time (calendar days) dimension. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. The sample runs at the daily frequency between January 1996 and December 2020 for a cross-section of 30 currency pairs relative to the US dollar. We use linear interpolation to retrieve daily forecasts from monthly forecasts. Data are from Bloomberg, FRED, JP Morgan, and OptionMetrics.

β_ϕ	1 (*grid*)	*different levels of* ϕ (1)	(2)	(3)	(4)	(5)	(6)
Panel A: 1-month Maturity							
ERP	4.43**	5.56***	2.64***	1.65***	1.15***	0.85**	0.65**
	(1.60)	(1.38)	(0.73)	(0.51)	(0.40)	(0.33)	(0.28)
R^2 (%)	32.9	33.1	33.0	33.0	32.9	32.9	32.8
N	153,975	153,975	153,975	153,975	153,975	153,975	153,975
Panel B: 3-month Maturity							
ERP	4.77***	4.97***	2.48***	1.63***	1.21***	0.95***	0.78***
	(1.16)	(0.94)	(0.50)	(0.36)	(0.28)	(0.23)	(0.20)
R^2 (%)	33.0	33.4	33.3	33.2	33.1	33.0	32.9
N	155,714	155,714	155,714	155,714	155,714	155,714	155,714
Panel C: 1-year Maturity							
ERP	2.65***	2.31***	1.26***	0.91***	0.72***	0.61***	0.53***
	(0.83)	(0.68)	(0.39)	(0.29)	(0.24)	(0.21)	(0.19)
R^2 (%)	46.4	46.5	46.5	46.4	46.4	46.4	46.3
N	155,540	155,540	155,540	155,540	155,540	155,540	155,540
Panel D: 2-year Maturity							
ERP	3.30***	2.47***	1.42***	1.07***	0.88***	0.77***	0.69***
	(0.89)	(0.64)	(0.38)	(0.29)	(0.24)	(0.21)	(0.19)
R^2 (%)	56.7	56.8	56.8	56.8	56.7	56.7	56.7
N	139,573	139,573	139,573	139,573	139,573	139,573	139,573
controls	✗	✗	✗	✗	✗	✗	✗
currency fe	✓	✓	✓	✓	✓	✓	✓
time fe	✓	✓	✓	✓	✓	✓	✓

Table 4. Term structure of risk preferences ϕ – using expected FX returns

This table presents panel regression estimates based on the following specification

$$\text{EFX}_{i,t,T} = \alpha_t + \alpha_i + \beta_\phi \, \text{ERP}^{\phi}_{i,t,T} + \gamma \, \text{IRD}_{i,t,T} + \varepsilon_{i,t,T},$$

where $\text{EFX}_{i,t,T}$ is the expected exchange rate return observed at time t over the horizon $T-t$ for currency i, $\text{ERP}^{\phi}_{i,t,T}$ is the expected currency risk premium computed at time t for the same currency pair/maturity, and $\text{IRD}_{i,t,T}$ is the interest rate differential between the US and country i at time t for the same currency pair/maturity. α_i and α_t denote currency and time (calendar date) fixed effects, respectively. $\text{EFX}_{i,t,T}$ is calculated using exchange rate consensus forecasts whereas $\text{ERP}^{\phi}_{i,t,T}$ is based on S&P 500 index and currency options for different levels of ϕ and maturities ranging between one month and two years. Standard errors, reported in parenthesis, are clustered by currency and time (calendar date) dimension. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. The sample runs at the daily frequency between January 1996 and December 2020 for a cross-section of 30 currency pairs relative to the US dollar. We use linear extrapolation to retrieve daily forecasts from monthly forecasts. Data are from Bloomberg, FRED, JP Morgan, and OptionMetrics.

β_ϕ	1 (*grid*)	*different levels of* ϕ (1)	(2)	(3)	(4)	(5)	(6)
Panel A: 1-month Maturity							
ERP	4.95***	4.67***	2.38***	1.61***	1.23***	0.99***	0.83***
	(1.01)	(0.86)	(0.45)	(0.31)	(0.25)	(0.20)	(0.18)
R^2 (%)	33.4	33.6	33.5	33.5	33.4	33.4	33.3
N	153,975	153,975	153,975	153,975	153,975	153,975	153,975
Panel B: 3-month Maturity							
ERP	1.82***	1.77***	0.91***	0.62***	0.48***	0.39***	0.33**
	(0.59)	(0.56)	(0.30)	(0.21)	(0.17)	(0.14)	(0.12)
R^2 (%)	33.1	33.2	33.1	33.1	33.1	33.0	33.0
N	155,714	155,714	155,714	155,714	155,714	155,714	155,714
Panel C: 1-year Maturity							
ERP	1.70***	1.58***	0.87***	0.63***	0.51***	0.44***	0.39***
	(0.46)	(0.42)	(0.24)	(0.18)	(0.14)	(0.13)	(0.11)
R^2 (%)	37.2	37.2	37.2	37.2	37.1	37.1	37.1
N	155,540	155,540	155,540	155,540	155,540	155,540	155,540
Panel D: 2-year Maturity							
ERP	1.68***	1.53***	0.87***	0.66***	0.54***	0.48***	0.43***
	(0.49)	(0.46)	(0.26)	(0.20)	(0.16)	(0.14)	(0.13)
R^2 (%)	35.7	35.7	35.7	35.6	35.6	35.6	35.6
N	139,573	139,573	139,573	139,573	139,573	139,573	139,573
IRD	✗	✗	✗	✗	✗	✗	✗
currency fe	✓	✓	✓	✓	✓	✓	✓
time fe	✓	✓	✓	✓	✓	✓	✓

Internet 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 Understanding the optimal growth portfolio

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.

A.1 Rationalizing the leveraged market return

We start with a simple static portfolio choice problem of an unconstrained CRRA agent who invests in the market.[12] 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 $R_{g,T} = R_{m,T}^{\phi}$ is the growth optimal portfolio return.

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 his modeling mistake, i.e. the distance (relative entropy) between his subjective belief and the rational

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

expectation (through 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 Section A, 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 the ambiguity aversion of the log agent.

Consider now, as an extension, that the representative agent holds only part of her 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 number is *dense* in the set of real number. Also, the random variable has compact support $R_{m,T}/R_{f,t} \in [0, 2(\omega^{-1} - 1)]$.

A.2 A case of time-varying ϕ

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, $\gamma_2 > \gamma_1 \geq 1$.[13] 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 $a =$

[13]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.

$4W_{1,t}/(e_tW_{2,t})^2$ reflects the relative wealth of the two agents. Although agents 1 and 2 are not representative—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 $e_TW_{2,T}^{-2\gamma}$. Integrating across agents, this representative agent's utility function at time T is

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

such that her 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 that in terms powers of W_T between $[0,2\gamma]$.

B Theory

B.1 Proof of Proposition 1

Derivation of Equation (3). Using the property that $M_T R_{g,T} = 1$, we start from the following identity

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

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

and then decompose the above risk-neutral expectation as

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

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

where $\mathbb{E}^{*}_t\left[R_{g,T}\right] = R^{\$}_{f,t}$ follows from the relation between the risk-neutral probability and the SDF valuation and $\mathbb{E}^{*}_t\left[S_T/S_t\right] = R^{\$}_{f,t}/R^{€}_{f,t}$ from the Uncovered Interest Rate Parity (UIP) condition. Finally, by combining and rearranging the above equations, we obtain

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

or equivalently

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

$\square$

B.2 Proposition 2

Proof. We here provide the derivation of Equation (4). 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 used $f(x) = x^{\phi}$, $c = 1$, and $x = R_{m,T}$. This allows us to express $R^{\phi}_{m,T} = 1 + \phi(R_{m,T} - 1) + \frac{\phi(\phi-1)}{2}\xi_T^{\phi-2}(R_{m,T} - 1)^2$. Take 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. $\square$

B.3 Interpretation of ϕ and higher-order moments

An elegant way of exploring how higher-order risk drives the expected risk premium is to expand the risk-neutral covariance in Equation (3) as follows

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

where $r_{m,T} \approx R_{m,T} - 1$ is the continuously compounded stock market return, and $w^{\star}_{\phi,n} = \left(e^{\phi} - 1\right)^{-1}\phi^n/n!$ denote weights that sum up to one. Each weight is a bell-shaped function of n with its maximum value around ϕ. The higher ϕ, the higher the factor $e^{\phi} - 1$ in front of the infinite sum and more weights are shifted to the (risk-neutral) higher-order terms, as illustrated by the left panel of Figure A.2. So the level of ϕ is intrinsically related to an aversion towards higher-order risk.[14]

[14] This is done by expanding the exponential function in the covariance $\mathbb{C}\text{ov}^{\star}_t\left(R^{\phi}_{m,T}, \frac{S_{i,T}}{S_{i,t}}\right) = \mathbb{C}\text{ov}^{\star}_t\left(e^{\phi r_{m,T}}, \frac{S_{i,T}}{S_{i,t}}\right)$ using power series, i.e., $e^x = 1 + x + \frac{1}{2!}x^2 + \ldots$.

FIGURE A.2 ABOUT HERE

Alternatively, we can write the analogous version of (A.7) under the real-world measure by simply replacing $M_T = 1/R_{g,T} = \lambda^{-1} R_{m,T}^{-\phi}$ in (2). We then obtain:

$$\mathbb{E}_t \left[\frac{S_{i,T}}{S_{i,t}} \right] - \frac{R_{f,t}}{R_{f,t}^i} = \frac{1}{R_{f,t}} \left(1 - e^{-\phi}\right) \sum_{n=1}^{\infty} w_{\phi,n} \mathbb{C}\text{ov}_t \left(r_{m,T}^n, \frac{S_{i,T}}{S_{i,t}} - 1 \right), \tag{A.8}$$

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 A.2 shows the change in signs between odd (even) market return moments and how the weights vary with ϕ.

C Ingredients of $\text{ERP}_{i,t}^{\phi}$

C.1 Estimating the risk-neutral correlation $\rho_{\phi,i,t}^*$

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 $\text{var}^*[\cdot]$ and $\mathbb{E}^*[\cdot]$ using options between t and T. We then calculate $\rho_{i,\phi}$ 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, we compare Quanto implied risk-neutral correlation (thanks to authors of Kremens and Martin (2019) for sharing the data) and realised correlation, which is the case when $\phi = 1$, we find those two are similar in terms of variation (see Figure A.1 for an example with EURUSD pair). Second, we find the realised correlations are similar across different value of ϕ. Third, we run similar regressions as Kremens and Martin (2019) by using the $\mathrm{ERP}^1_{i,t,T}$ that we constructed using the realized correlations and match the same sample period.

Figure A.1. The correlation $\rho_{\phi=1,i,t}$ for $i = \mathrm{EURUSD}_t$

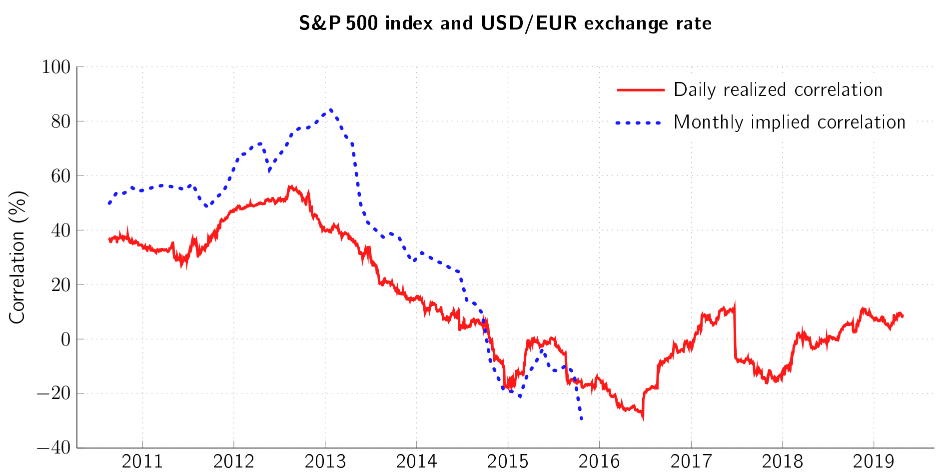

C.2 Risk Neutral Moments

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

Risk Neutral Moments of Equity Index Returns. To compute the risk neutral variance of $\frac{R_{m,T}^{\phi}}{\mathbb{E}_t^*[R_{i,T}^{\phi}]}$, we recall the formula

$$\text{var}_t^*\left(\frac{R_{m,T}^{\phi}}{\mathbb{E}_t^*\left[R_{m,T}^{\phi}\right]}\right) = \frac{\mathbb{E}_t^*\left[R_{m,T}^{2\phi}\right]}{\mathbb{E}_t^*\left[R_{m,T}^{\phi}\right]^2} - 1\,,$$

where we could compute the risk neutral momements using the following formula

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

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 $\Omega_{t,T}(K)$ is the out-of-money 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 $E_tR_{f,t} = F_{t,T}$ in our computation to ignore the dividend payment, i.e. we assume the payment of dividends are not reinvested. The formula simplifies to

$$\mathbb{E}_t^*[R_{m,T}^{\theta}] = R_{f,t}^{\theta} + R_{f,t}\int_0^{F_{t,T}} \frac{\theta(\theta-1)}{E_t^{\theta}}K^{\theta-2}P_{t,T}(K)dK + R_{f,t}\int_{F_{t,T}}^{\infty} \frac{\theta(\theta-1)}{E_t^{\theta}}K^{\theta-2}C_{t,T}(K)dK\,. \tag{A.10}$$

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

$$\text{var}_t^*\left(\frac{S_T}{S_t}\right) = \mathbb{E}_t^*\left(\frac{S_T}{S_t}\right)^2 - \left(\mathbb{E}_t^*\frac{S_T}{S_t}\right)^2\,, \tag{A.11}$$

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

$$\operatorname{var}_t^*\left(\frac{S_T}{S_t}\right) = \frac{2}{B_{t,T}S_t^2}\left(\int_0^{F_{t,T}} P_{t,T}(K)dK + \int_{F_{t,T}}^{\infty} C_{t,T}(K)dK\right), \tag{A.12}$$

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. $B_{t,T}$ 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.12) via trapezoidal integration.

D 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.

D.1 Risk-neutral Correlation

[put comparison to Quanto theory paper]

D.2 Consensus v.s. 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

$$\underbrace{\widetilde{\mathbb{E}}_t\left[\frac{S_T}{S_t}\right]-\frac{R_{f,T}}{R^i_{f,T}}}_{\text{forecast risk premium}}=\underbrace{\mathbb{E}_t\left[\frac{S_T}{S_t}\right]-\frac{R_{f,T}}{R^i_{f,T}}}_{\text{UIP premium}}+\underbrace{\widetilde{\mathbb{E}}_t\left[\frac{S_T}{S_t}\right]-\mathbb{E}_t\left[\frac{S_T}{S_t}\right]}_{\text{none risk premium component}}\ , \tag{A.13}$$

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.13) is directly observable from the data of surveys. Comparing the above to (A.16), we could view the subjective expected risk premium as

$$\widetilde{\mathbb{E}}_t\left[\frac{S_T}{S_t}\right]-\frac{R_{f,T}}{R^i_{f,T}}=\mathrm{ERP}^{(\phi)}_{i,t}+\xi_{i,t}\,, \tag{A.14}$$

where the residual has two parts

$$\xi_{i,t}=\underbrace{-\frac{R_{f,T}\mathbb{C}\mathrm{ov}_t\left(M_T R^{\phi}_{m,T},\frac{S_T}{S_t}\right)}{\mathbb{E}^*_t[R^{\phi}_{m,T}]}}_{\xi^{(1)}_{i,t}}+\underbrace{\widetilde{\mathbb{E}}_t\left[\frac{S_T}{S_t}\right]-\mathbb{E}_t\left[\frac{S_T}{S_t}\right]}_{\xi^{(2)}_{i,t}}\,. \tag{A.15}$$

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.

D.3 Residual Term for $\mathrm{ERP}_{i,T}^{(\phi)}$

A potential misspecification of $R_{g,T}$ would generate an error term in (**??**). 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*

$$\mathbb{E}_t\left[\frac{S_T}{S_t}\right] - \frac{R_{f,t}}{R_{f,t}^i} = \frac{\mathbb{C}ov_t^*\left(X_T, \frac{S_T}{S_t}\right)}{\mathbb{E}_t^*[X_T]} - \frac{R_{f,T}\mathbb{C}ov_t\left(M_T X_T, \frac{S_T}{S_t}\right)}{\mathbb{E}_t^*[X_T]}, \tag{A.16}$$

where M_T is the SDF, $R_{e_i,T}$ is the return of FX and $\mathbb{E}^*[\cdot]$ stands for the risk-neutral measured expectation.

Given the choice of $R_{g,T} = R_{m,T}^{\phi}$, which value of ϕ would minimize the residual term in (A.16) is an empirical question.

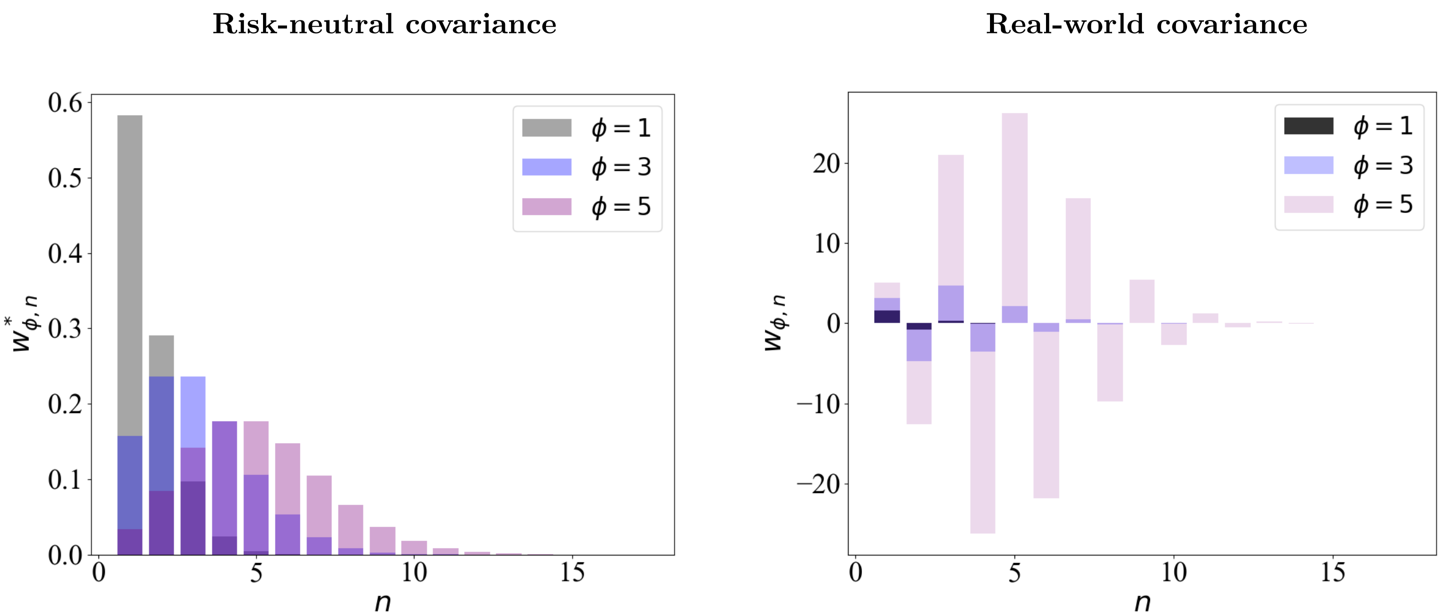

Figure A.2. Higher-order risk and the expected risk premium

This figure displays the contribution of higher-order risk to the expected risk premium for different degrees of risk preferences, ϕ. The left panel expands the risk-neutral covariance in Equation (3) using Equation (A.7), while the right panel is based on Equation (A.8). Panels (A) and (B) plot the risk-neutral weight $w^{\star}_{\phi,n}$ and the real-world weight $w_{\phi,n}$ in function of the market return moment n. The weights are respectively associated with the risk-neutral and real-world covariances between $r^n_{m,T}$, where $r_{m,T}$ is the market return for horizon T, and the exchange rate return $\frac{S_{i,T}}{S_{i,t}}$. The risk-neutral weight takes its maximum value around ϕ. The real-world weight switches signs between the odd and even market return moments because FX investors require a positive risk premium when the covariance between exchange rate returns and odd (even) market return moments is positive (negative).

Table A.1. Panel regression estimates of the risk preferences ϕ – with interpolated forecasts

This table presents panel regression estimates based on the following specification

$$\text{ERX}_{i,t,T} = \alpha_i + \alpha_\ell + \alpha_t + \beta_\phi \, \text{ERP}^{\phi}_{i,t,T} + \varepsilon_{i,t,T},$$

where $\text{ERX}_{i,t,T}$ is the expected currency excess return observed at time t over the horizon $T-t$ for currency i and $\text{ERP}^{\phi}_{i,t,T}$ is the expected currency risk premium computed at time t for the same currency pair/maturity. α_i, α_ℓ, and α_t denote currency, maturity, and time (calendar date) fixed effects, respectively. $\text{ERX}_{i,t,T}$ is calculated using exchange rate consensus forecasts and interest rate differentials whereas $\text{ERP}^{\phi}_{i,t,T}$ is based on S&P 500 index and currency options for different levels of ϕ and maturities ranging between one month and two years. Standard errors, reported in parenthesis, are clustered by currency and time (calendar date) dimension. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. The sample runs at the daily frequency between January 1996 and December 2020 for a cross-section of 30 (20) currency pairs relative to the US dollar in Panel A (B). We use linear interpolation to retrieve daily forecasts from monthly forecasts. Data are from Bloomberg, FRED, JP Morgan, and OptionMetrics.

Panel A: All Currencies							
β_ϕ	1	*different levels of* ϕ					
	(*grid*)	(1)	(2)	(3)	(4)	(5)	(6)
ERP	4.15***	3.71***	1.94***	1.34***	1.03***	0.84***	0.72***
	(0.75)	(0.61)	(0.33)	(0.23)	(0.18)	(0.15)	(0.13)
R^2 (%)	20.1	20.4	20.3	20.2	20.1	20.0	19.9
N	604,897	604,897	604,897	604,897	604,897	604,897	604,897
Panel B: Most Liquid Currencies							
ERP	4.30***	3.89***	2.03***	1.39***	1.07***	0.88***	0.74***
	(0.77)	(0.65)	(0.35)	(0.25)	(0.19)	(0.16)	(0.14)
R^2 (%)	21.1	21.4	21.3	21.2	21.1	21.0	20.9
N	436,568	436,568	436,568	436,568	436,568	436,568	436,568
currency fe	✓	✓	✓	✓	✓	✓	✓
maturity fe	✓	✓	✓	✓	✓	✓	✓
time fe	✓	✓	✓	✓	✓	✓	✓

Table A.2. Expected Excess Returns and Risk Premia

This table presents panel regression estimates based on the following specification

$$\mathrm{ERX}_{i,t,T} = \alpha_t + \alpha_i + \beta_\phi \, \mathrm{ERP}^{\phi}_{i,t,T} + \varepsilon_{i,t,T},$$

where $\mathrm{ERX}_{i,t,T}$ is the expected currency excess return observed at time t over the horizon $T-t$ for currency i and $\mathrm{ERP}^{\phi}_{i,t,T}$ is the expected currency risk premium computed at time t for the same currency pair/maturity. α_i and α_t denote currency and time (calendar date) fixed effects, respectively. $\mathrm{ERX}_{i,t,T}$ is calculated using exchange rate consensus forecasts and interest rate differentials whereas $\mathrm{ERP}^{\phi}_{i,t,T}$ is based on S&P 500 index and currency options for different levels of ϕ and maturities ranging between one month and two years. Standard errors, reported in parenthesis, are clustered by currency and time (calendar date) dimension. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. The sample runs at the daily frequency between January 1996 and December 2020 for a cross-section of 30 currency pairs relative to the US dollar. We use linear interpolation to retrieve daily forecasts from monthly forecasts. Data are from Bloomberg, FRED, JP Morgan, and OptionMetrics.

β_ϕ	1 (*grid*)	*different levels of* ϕ (1)	(2)	(3)	(4)	(5)	(6)
Panel A: 1-month Maturity							
ERP	4.80***	4.46***	2.28***	1.55***	1.19***	0.97***	0.82***
	(0.82)	(0.69)	(0.36)	(0.25)	(0.20)	(0.17)	(0.14)
R^2 (%)	31.1	31.3	31.3	31.2	31.1	31.0	31.0
N	154,013	154,013	154,013	154,013	154,013	154,013	154,013
Panel B: 3-month Maturity							
ERP	2.08***	1.99***	1.03***	0.71***	0.55***	0.45***	0.38***
	(0.54)	(0.50)	(0.27)	(0.19)	(0.15)	(0.13)	(0.11)
R^2 (%)	32.5	32.6	32.5	32.4	32.4	32.3	32.3
N	155,733	155,733	155,733	155,733	155,733	155,733	155,733
Panel C: 1-year Maturity							
ERP	1.80***	1.66***	0.91***	0.67***	0.54***	0.46***	0.41***
	(0.43)	(0.40)	(0.22)	(0.17)	(0.14)	(0.12)	(0.11)
R^2 (%)	48.4	48.5	48.4	48.4	48.4	48.3	48.3
N	155,559	155,559	155,559	155,559	155,559	155,559	155,559
Panel D: 2-year Maturity							
ERP	1.74***	1.58***	0.90***	0.68***	0.56***	0.49***	0.44***
	(0.53)	(0.47)	(0.27)	(0.21)	(0.17)	(0.15)	(0.14)
R^2 (%)	56.5	56.5	56.5	56.5	56.5	56.5	56.5
N	139,592	139,592	139,592	139,592	139,592	139,592	139,592
currency fe	✓	✓	✓	✓	✓	✓	✓
time fe	✓	✓	✓	✓	✓	✓	✓

Table A.3. Expected FX Returns and Risk Premia

This table presents panel regression estimates based on the following specification

$$\mathrm{EFX}_{i,t,T} = \alpha_t + \alpha_i + \beta_\phi \, \mathrm{ERP}^{\phi}_{i,t,T} + \gamma \, \mathrm{IRD}_{i,t,T} + \varepsilon_{i,t,T},$$

where $\mathrm{EFX}_{i,t,T}$ is the expected exchange rate return observed at time t over the horizon $T-t$ for currency i, $\mathrm{ERP}^{\phi}_{i,t,T}$ is the expected currency risk premium computed at time t for the same currency pair/maturity, and $\mathrm{IRD}_{i,t,T}$ is the interest rate differential between the US and country i at time t for the same currency pair/maturity. α_i and α_t denote currency and time (calendar date) fixed effects, respectively. $\mathrm{EFX}_{i,t,T}$ is calculated using exchange rate consensus forecasts whereas $\mathrm{ERP}^{\phi}_{i,t,T}$ is based on S&P 500 index and currency options for different levels of ϕ and maturities ranging between one month and two years. Standard errors, reported in parenthesis, are clustered by currency and time (calendar date) dimension. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. The sample runs at the daily frequency between January 1996 and December 2020 for a cross-section of 30 currency pairs relative to the US dollar. We use linear interpolation to retrieve daily forecasts from monthly forecasts. Data are from Bloomberg, FRED, JP Morgan, and OptionMetrics.

β_ϕ	1	*different levels of ϕ*					
	(*grid*)	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: 1-month Maturity							
ERP	4.95***	4.56***	2.34***	1.59***	1.22***	0.99***	0.84***
	(0.89)	(0.73)	(0.38)	(0.27)	(0.21)	(0.18)	(0.15)
R^2 (%)	32.3	32.6	32.5	32.4	32.4	32.3	32.2
N	154,013	154,013	154,013	154,013	154,013	154,013	154,013
Panel B: 3-month Maturity							
ERP	2.10***	2.01***	1.04***	0.72***	0.56***	0.46***	0.39***
	(0.55)	(0.52)	(0.28)	(0.20)	(0.16)	(0.13)	(0.11)
R^2 (%)	32.8	32.9	32.8	32.7	32.7	32.6	32.6
N	155,733	155,733	155,733	155,733	155,733	155,733	155,733
Panel C: 1-year Maturity							
ERP	1.74***	1.62***	0.89***	0.65***	0.53***	0.45***	0.40***
	(0.46)	(0.42)	(0.24)	(0.18)	(0.15)	(0.13)	(0.11)
R^2 (%)	39.1	39.2	39.1	39.1	39.0	39.0	39.0
N	155,559	155,559	155,559	155,559	155,559	155,559	155,559
Panel D: 2-year Maturity							
ERP	1.70***	1.54***	0.88***	0.66***	0.55***	0.48***	0.44***
	(0.49)	(0.45)	(0.26)	(0.20)	(0.16)	(0.14)	(0.13)
R^2 (%)	38.1	38.1	38.0	38.0	38.0	38.0	38.0
N	139,592	139,592	139,592	139,592	139,592	139,592	139,592
currency fe	✓	✓	✓	✓	✓	✓	✓
time fe	✓	✓	✓	✓	✓	✓	✓

Table A.4. Term structure of risk preferences ϕ – more details

This table presents estimates of the risk preferences ϕ from a panel regression of expected currency excess returns on expected currency risk premium for different horizons. The results are obtained from the following specification

$$\mathrm{ERX}_{i,t,T} = \alpha_t + \alpha_i + \beta_\phi \, \mathrm{ERP}^{\phi}_{i,t,T} + \varepsilon_{i,t,T},$$

where $\mathrm{ERX}_{i,t,T}$ is the expected currency excess return observed at time t over the horizon $T-t$ for currency i and $\mathrm{ERP}^{\phi}_{i,t,T}$ is the expected currency risk premium computed at time t for the same currency pair/maturity. α_i and α_t denote currency and time fixed effects (FE), respectively. $\mathrm{ERX}_{i,t,T}$ is calculated using exchange rate consensus forecasts for a given horizon net of interest rate differentials, whereas $\mathrm{ERP}^{\phi}_{i,t,T}$ is based on S&P 500 index and currency options for different levels of ϕ and a corresponding maturity. Each column presents estimates of β_ϕ for selected values of ϕ ranging between 1 and 6. Standard errors, reported in parenthesis, are clustered by currency and time. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. The sample runs at the daily frequency between January 1996 and December 2020 for a cross-section of 30 currency pairs relative to the US dollar. We use linear extrapolation to retrieve daily forecasts from monthly forecasts. Data are from Bloomberg, FRED, JP Morgan, and OptionMetrics.

	Results for different ϕ					
	$\phi = 1$	$\phi = 2$	$\phi = 3$	$\phi = 4$	$\phi = 5$	$\phi = 6$
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: 1-month Horizon						
ERP^{ϕ}	4.54***	2.31***	1.57***	1.19***	0.96***	0.81***
	(0.82)	(0.43)	(0.30)	(0.24)	(0.20)	(0.17)
R^2 (%)	32.8	32.7	32.7	32.6	32.6	32.5
N	153,975	153,975	153,975	153,975	153,975	153,975
Panel B: 3-month Horizon						
ERP^{ϕ}	1.75***	0.90***	0.61***	0.47***	0.38***	0.32**
	(0.54)	(0.29)	(0.20)	(0.16)	(0.14)	(0.12)
R^2 (%)	32.7	32.7	32.6	32.6	32.6	32.6
N	155,714	155,714	155,714	155,714	155,714	155,714
Panel C: 1-year Horizon						
ERP^{ϕ}	1.62***	0.89***	0.65***	0.53***	0.45***	0.40***
	(0.40)	(0.22)	(0.17)	(0.14)	(0.12)	(0.11)
R^2 (%)	45.3	45.3	45.3	45.2	45.2	45.2
N	155,540	155,540	155,540	155,540	155,540	155,540
Panel D: 2-year Horizon						
ERP^{ϕ}	1.56***	0.89***	0.67***	0.56***	0.49***	0.44***
	(0.47)	(0.27)	(0.21)	(0.17)	(0.15)	(0.14)
R^2 (%)	53.5	53.5	53.5	53.5	53.4	53.4
N	139,573	139,573	139,573	139,573	139,573	139,573
Currency FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓