Expected Currency Returns*

Pasquale Della Corte **Can Gao** **Alexandre Jeanneret**

March 2, 2022

Preliminary and Incomplete

*We are grateful to Neils Gormsen, Yueran Ma, Tarun Ramadorai, and seminar participants at Imperial College London for comments and suggestions. All remaining errors are our own. Pasquale Della Corte is with Imperial College Business School, Imperial College London and the Centre for Economic Policy Research (CEPR). Email: p.dellacorte@imperial.ac.uk. Can Gao is with Leibniz Institute for Financial Research SAFE, Frankfurt. Email: gao@safe-frankfurt.de. Alexandre Jeanneret is with the School of Banking and Finance, UNSW Business School. Email: a.jeanneret@unsw.edu.au.

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 are 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 analyze the FX risk premium, both theoretically and empirically, and 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 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 different economic/financial conditions.

[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).

Our main results are as follows. First, 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). Second, the unconditional 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. Third, a conditional analysis reveals that the 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. Our findings thus provide novel insights on the conditional term structure of risk preferences.

Our theoretical approach exploits the no-arbitrage condition based on the existence of the so-called *growth optimal portfolio*.[2] We show that the expected excess exchange rate return (or risk premium) can be rewritten as the following identity:

$$\mathbb{E}_t\left[\frac{e_{i,T}}{e_{i,t}}\right] - \frac{R_{f,T}}{R^i_{f,T}} = \frac{1}{R_{f,T}}\operatorname{cov}^*_t\left(R_{g,T}, \frac{e_{i,T}}{e_{i,t}}\right), \tag{1}$$

where $\operatorname{cov}^*_t\left(R_{g,T}, \frac{e_{i,T}}{e_{i,t}}\right)$ captures the conditional risk-neutral covariance between the gross return of the growth optimal portfolio ($R_{g,T}$) and the gross exchange rate return ($\frac{e_{i,T}}{e_{i,t}}$) over the horizon $T-t$. The return of the optimal growth portfolio is a power function of the market return $R_{mkt,T}$, such that $R_{g,T} = \lambda R^{\phi}_{mkt,T}$, where λ is a level parameter and $\phi > 0$ is a coefficient capturing the degree of risk preferences. Our approach therefore links the risk premium with a directly-interpretable measure of risk preferences, given by ϕ, whereby a large value implies that FX investors are more "averse" against higher-order risk, such as co-skewness ($\phi = 3$), co-kurtosis ($\phi = 4$), and so on.

Our measure of risk preferences is general and nests various interesting versions studied in the literature. For example, ϕ can be interpreted as the risk aversion of an unconstrained representative investor with CRRA utility whose entire wealth is invested in the market. We also show that ϕ can be endogenously time varying in a model with heterogeneous agents (Chan and Kogan, 2002). Our approach thus addresses the critique of Bekaert, Engstrom, and Xu (2021) that most studies estimate time-varying risk aversion measures motivated by

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

models that essentially assume a constant risk aversion coefficient and hence are inherently inconsistent.

To empirically measure the FX risk premium, we compute the risk-neutral covariance in (1) using options on the S&P 500 and on exchange rates for different horizons (between 1 month and 24 months). We then identify the value of ϕ such that our theoretically-implied risk premium best matches observable expected excess exchange rate returns, using survey data from professional forecasters. We exploit a cross-section of G30 currencies against the USD over the sample spanning the 1996.01.01 to 2020.12.31 period. Unconditionally, we obtain an estimate of $\phi = 4.5$, which suggests that FX investors are substantially averse 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). We find that ϕ decreases from 4.5 at the one-month horizon to 1.4 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 (e.g., the Australian dollar in October 2008 or the Russian Ruble in February 2022) matters less to investors 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. Investors are then more fearful towards such tail risk events when their horizon is shorter, which is expressed by a higher ϕ and thus a downward-sloping term structure.

We then turn to a conditional analysis of this term structure. To do so, we split our sample according to NBER-dated recessions and expansions and estimate ϕ on each subsample separately for different horizons. We show that the term structure of risk preferences has a steep negative slope in recessions, which turns positive during expansions. We find similar results when we analyze risk preferences across different measures of financial conditions. For example, we separate the sample by high and low levels of CBOE equity-option implied volatility index (VIX), based on the sample mean. Alternatively, we split the sample according to the level of the option-implied volatility for a basket of G7 currencies (VXY) or using the implied volatility on one-month U.S. Treasury options (MOVE). In all cases, the term structure of risk preferences is countercyclical with respect to aggregate economic/financial conditions.

In sum, our approach contributes to the literature in three ways. First, we consider a utility-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 we show in the paper, 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 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). 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. Last but not least, given that forecasts and options are available for different horizons, we can shed light on the conditional term structure of ϕ. Our paper is the first to provide insights on how risk preferences vary over different horizons.

This paper relates to a growing literature on extracting time-varying preferences from surveys, experiments, or asset prices. Guiso, Sapienza, and Zingales (2018) find 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) extract a measure of time-varying aggregate risk aversion from macro data and financial asset prices. Consistent with these studies, we find that ϕ increases during recessions and periods of heightened uncertainty, suggesting that FX market participants are indeed more averse to higher-order risk as economic/financial conditions worsen. A fundamental difference between our paper and this literature, however, is that we are able to extract and provide insights on the term structure of risk preferences.

The remainder of the paper is organized as follows. Section 2 shows theoretically how the growth optimal portfolio helps identify the role of risk preferences for expected exchange rate excess returns. 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

In this section, we derive the expected excess return for exchange rates and shed new light on the role of investor preferences in driving risk premia in the FX market. Our approach extends the theory of Kremens and Martin (2019) by exploiting the growth optimal portfolio, which allows to consider general risk preferences in a model-free environment.

2.1 Environment

Consider a currency strategy that converts a dollar into foreign currency at time t, lends at the foreign riskless rate between time t and T, and then exchanges the proceeds in foreign currency for dollars at time T. According to the fundamental equation of asset pricing, the expected excess exchange rate return is given by

$$\mathbb{E}_t\left[\frac{e_{i,T}}{e_{i,t}}\right] - \frac{R_{f,T}}{R^i_{f,T}} = -R_{f,t}\,\mathrm{cov}_t\left(M_T, \frac{e_{i,T}}{e_{i,t}}\right), \tag{2}$$

where $\mathbb{E}_t$ is the expectation operator (under the physical measure) conditional on the information available at time t, $e_{i,t}$ is the spot exchange rate defined as units of dollars per foreign currency i such that an increase in $e_{i,t}$ reflects an appreciation of the foreign currency, $R^i_{f,T}$ ($R_{f,T}$) denotes the gross risk-free rate in foreign country (US) from time t to T (known at t), and M_T is a stochastic discount factor (SDF) that prices assets denominated in dollars.

Under the risk-neutral expectation $\mathbb{E}^*_t$, the covariance term disappears in (2) and the expected excess exchange rate return is zero. This condition fails to hold empirically and the non-zero excess return can be interpreted as compensation for time-varying risk (e.g., Fama, 1984; Lustig, Roussanov, and Verdelhan, 2011). This is equivalent to saying that investors demand a risk adjustment component captured by the covariance between the SDF and the gross exchange rate return. Because the SDF is unobservable ex-ante and likely to change over time, it remains challenging to determine how investor preferences exactly shape this risk compensation. The objective of this paper is to overcome this challenge and to shed new light on risk preferences in the FX market.

Our approach relies on a risk-neutral representation, building on the insights of Martin (2017), Kremens and Martin (2019), Martin and Wagner (2019) and Della Corte, Jeanneret, and Patelli (2021). To do so, we exploit a property of the no-arbitrage condition, largely overlooked in the recent FX literature, which is the existence of a so-called *growth optimal portfolio*. The growth optimal portfolio is particularly useful in our case because its gross return, denoted by $R_{g,T}$, satisfies $M_T R_{g,T} = 1$, every state of the world.[3] In this case, the expected excess exchange rate return can be rewritten as the following identity:

$$\mathbb{E}_t\left[\frac{e_{i,T}}{e_{i,t}}\right] - \frac{R_{f,T}}{R^i_{f,T}} = \frac{1}{R_{f,T}} \operatorname{cov}^*_t\left(R_{g,T}, \frac{e_{i,T}}{e_{i,t}}\right), \tag{3}$$

where $\operatorname{cov}^*_t\left(R_{g,T}, \frac{e_{i,T}}{e_{i,t}}\right)$ captures the conditional risk-neutral covariance between the gross return of the growth optimal portfolio ($R_{g,T}$) and the gross exchange rate return ($\frac{e_{i,T}}{e_{i,t}}$) over the horizon $T-t$.

Derivation of (3). Using the property that $M_T R_{g,T} = 1$, one can expand the expectation of the gross exchange rate return as follows:

$$\mathbb{E}_t\left[\frac{e_{i,T}}{e_{i,t}}\right] = \mathbb{E}_t\left[M_T R_{g,T}\frac{e_{i,T}}{e_{i,t}}\right] = \frac{1}{R_{f,T}}\,\mathbb{E}^*_t\left[R_{g,T}\frac{e_{i,T}}{e_{i,t}}\right] \tag{4}$$

[3] The growth optimal portfolio reaches the entropy bound of the SDF in Bansal and Lehmann (1997) and Alvarez and Jermann (2005), where

$$\mathbb{E}_t[\log R_{g,T}] \leq -\mathbb{E}_t[\log M_T].$$

The growth optimal portfolio can be viewed as a benchmark portfolio that is analogous to the minimum second moment portfolio from the Hansen-Jagannathan bound. For early work using this portfolio, see e.g. Kelly Jr (1956), Roll (1973), Fama and MacBeth (1974), Markowitz (1976), Long Jr (1990), and Ross (1999).

and then decompose the above risk-neutral expectation:

$$\mathbb{E}_t^*\left[R_{g,T}\frac{e_{i,T}}{e_{i,t}}\right]=\underbrace{\mathbb{E}_t^*\left[R_{g,T}\right]}_{=R_{f,T}}\underbrace{\mathbb{E}_t^*\left[\frac{e_{i,T}}{e_{i,t}}\right]}_{R_{f,T}/R_{f,T}^i}+\operatorname{cov}_t^*\left(R_{g,T},\frac{e_{i,T}}{e_{i,t}}\right)\,, \tag{5}$$

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 $\mathbb{E}_t^*\left[\frac{e_{i,T}}{e_{i,t}}\right]=R_{f,T}/R_{f,T}^i$ from the Uncovered Interest Rate Parity (UIP) condition. By combining the two equations and rearranging, we obtain (3). □

2.2 Optimal growth portfolio and investor preferences

We now derive the optimal growth portfolio return and show how risk preferences enter the expected excess exchange rate return.

We assume that the return of the optimal growth portfolio $R_{g,T}$ is a power function of the market return $R_{mkt,T}$, such that $R_{g,T}=\lambda R_{mkt,T}^{\phi}$, where λ is a level parameter and $\phi>0$, possibly time varying, is a coefficient capturing the degree of risk preferences. We find that such specification arises endogenously in various existing asset pricing models, as we will discuss in Section 2.2.1. We can then rewrite the right-hand side of (3) as follows:

$$\frac{1}{R_{f,T}}\operatorname{cov}_t^*\left(R_{g,T},\frac{e_{i,T}}{e_{i,t}}\right)=\frac{1}{\lambda\,\mathbb{E}_t^*[R_{mkt,T}^{\phi}]}\operatorname{cov}_t^*\left(\lambda R_{mkt,T}^{\phi},\frac{e_{i,T}}{e_{i,t}}\right) \tag{6}$$

$$=\operatorname{cov}_t^*\left(\frac{R_{mkt,T}^{\phi}}{\mathbb{E}_t^*[R_{mkt,T}^{\phi}]},\frac{e_{i,T}}{e_{i,t}}\right)\,, \tag{7}$$

using $\mathbb{E}_t^*[R_{g,T}]=R_{f,T}$, which implies that the constant λ cancels out and thus vanishes from the expected excess exchange rate return.

This optimal growth portfolio is particularly relevant to forward our understanding of the expected excess exchange rate return. To see that, we can expand the risk-neutral covariance

Figure 1. Relative weights of risk-neutral covariances

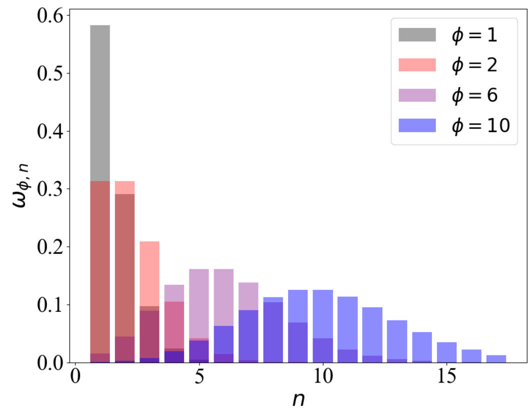

(7) to rewrite the identity (3) as follows:[4]

$$\mathbb{E}_t\left[\frac{e_{i,T}}{e_{i,t}}\right] - \frac{R_{f,T}}{R^i_{f,T}} = \frac{\left(e^{\phi}-1\right)}{R_{f,T}}\sum_{n=1}^{\infty} w_{\phi,n}\,\mathrm{cov}^*_t\left(r^n_{mkt,T}, \frac{e_{i,T}}{e_{i,t}} - 1\right), \tag{8}$$

where $r_{mkt,T} = R_{mkt,T} - 1$ denotes the market net return. The weight $w_{\phi,n} = \left(e^{\phi}-1\right)^{-1}\frac{\phi^n}{n!}$ sums to one and is a bell-shaped function in terms of n with its maximum value around ϕ, as illustrated in Figure 1. The higher the value of ϕ, the higher the factor $e^{\phi} - 1$ in front of the infinite sum and the more weights are shifted to the (risk-neutral) higher-order terms. This decomposition shows that the level of expected excess exchange rate return is thus increasing in ϕ. Note that the value of ϕ is intrinsically related to aversion to higher-order risk, such as co-skewness ($\phi = 3$), co-kurtosis ($\phi = 4$), and so on. So ϕ has a direct economic interpretation in terms of risk preferences.

2.2.1 Rationalizing the return of the optimal growth portfolio

We now show that our specification, where the return of the optimal growth portfolio $R_{g,T}$ is proportional to a power function of the market return, $R_{g,T} \propto R^{\phi}_{mkt,T}$, is consistent with

[4] This is done by expanding the exponential function in the covariance $R^{-1}_{f,T}\,\mathrm{cov}^*_t\left(R^{\phi}_{mkt,T}, \frac{e_{i,T}}{e_{i,t}}\right) = R^{-1}_{f,T}\,\mathrm{cov}^*_t\left(e^{\phi r_{mkt,T}}, \frac{e_{i,T}}{e_{i,t}}\right)$ using power series, i.e. $e^x = 1 + x + \frac{1}{2!}x^2 + \ldots$.

different asset pricing models. To do so, we consider different economic environments and, for each case, derive the return of the optimal growth portfolio.

We start with a simple static portfolio choice problem of an unconstrained CRRA agent who invests in the market. Her 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 $\lambda = \mathbb{E}_t\left[R_{i,T}\left(\sum_i w_i^* R_{i,T}\right)^{-\phi}\right]$, $\forall i$, where λ is the Lagrangian multiplier, and w_i^* is the optimal weight of an asset with gross return $R_{i,T}$ in the representative agent's portfolio. Note that the quantity $\left(\sum_i w_i^* R_{i,T}\right)^{-\phi} = R_{mkt,T}^{-\phi}$ is proportional to the SDF, which proves that $R_{g,T} = \lambda R_{mkt,T}^{\phi}$ is the growth optimal portfolio return. A similar result can be obtained in a setting with ambiguity aversion (see Appendix A.1).

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 $R_{g,T} = \lambda\left(\omega R_{mkt,T} + (1-\omega)R_{f,T}\right)^{\phi}$. Under some reasonable conditions,[5] the above binomial function can be expanded as a Maclaurin series

$$R_{g,T} = \lambda(1-\omega)^{\phi} R_{f,T}^{\phi}\left(1 + \phi\frac{\omega}{1-\omega}\frac{R_{mkt,T}}{R_{f,T}} + \frac{\phi(\phi-1)}{2!}\left(\frac{\omega}{1-\omega}\frac{R_{mkt,T}}{R_{f,T}}\right)^2 + \ldots\right),$$

which is essentially a sum of integer powers of the market return with constant coefficients. Since the value of the coefficient $\frac{\phi(\phi-1)(\phi-2)\ldots(\phi-n)}{(n+1)!}$ converges to zero when $n \to \infty$, the value of ϕ indicates the maximum power of the market return that is relevant for $R_{g,T}$. That is, a higher ϕ is associated with a greater aversion to higher-order risk.

[5] 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 $R_{mkt,T}/R_{f,T} \in [0, 2(\omega^{-1}-1)]$.

2.2.2 An example with 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).

The following example builds on Longstaff and Wang (2012), except that each agent faces a portfolio choice problem. 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, $\gamma_2 > \gamma_1 \geq 1$. 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 $e_T W_{2,T} = \frac{2}{a}\left(\sqrt{1+aW_T} - 1\right)$, where the constant $a = 4W_{1,t}/(e_t W_{2,t})^2$ reflects the relative wealth of the two agents. Although agents 1 and 2 are not representative—neither invests only in the market—they have the same beliefs and SDF as a representative agent. Such representative agent has a wealth W_T (measured in currency 1) and marginal utility $v'(W_T)$ that is proportional to $e_T W_{2,T}^{-2\gamma}$. Integrating across agents, this representative agent's utility function 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.

3 Empirical Methodology

We present our approaches to estimate ϕ, based on the framework developed in Section 2, and then describe the construction of the main variables.

3.1 Two approaches to estimate ϕ

We start by proposing two approaches to estimate ϕ in the data, which we will use to shed light on risk preferences in the FX market. Note that, combining (3) and (7), the expected excess exchange rate return is given by

$$\mathbb{E}_t \left[\frac{e_{i,T}}{e_{i,t}} \right] - \frac{R_{f,T}}{R^i_{f,T}} = \frac{1}{R_{f,T}} \operatorname{cov}^*_t \left(R_{g,T}, \frac{e_{i,T}}{e_{i,t}} \right) = \operatorname{cov}^*_t \left(\frac{R^{\phi}_{mkt,T}}{\mathbb{E}^*_t [R^{\phi}_{mkt,T}]}, \frac{e_{i,T}}{e_{i,t}} \right). \tag{9}$$

Main Approach. We can write an empirical representation of (9) as follows:

$$\underbrace{\mathbb{E}_t \left[\frac{e_{i,T}}{e_{i,t}} \right] - \frac{R_{f,T}}{R^i_{f,T}}}_{\mathrm{ERX}_{i,t}} = \alpha_\phi + \beta_\phi \underbrace{\operatorname{cov}^*_t \left(\frac{R^{\phi}_{mkt,T}}{\mathbb{E}^*_t [R^{\phi}_{mkt,T}]}, \frac{e_{i,T}}{e_{i,t}} \right)}_{\mathrm{ERP}^{(\phi)}_{i,t}}, \tag{10}$$

where $\mathrm{ERX}_{i,t}$ reflects observable expected excess exchange rate return and $\mathrm{ERP}^{(\phi)}_{i,t}$ is the *expected risk premium* for a given ϕ. An accurate specification of the expected risk premium implies $\alpha_\phi = 0$ and $\beta_\phi = 1$ when regressing $\mathrm{ERX}_{i,t}$ on $\mathrm{ERP}^{(\phi)}_{i,t}$. This is the condition we are going to consider in this approach to pin down the coefficient of ϕ. Specifically, each $\mathrm{ERP}_{i,t}^{(\phi)}$ implies a different slope coefficient β_ϕ and we can thus select the level of ϕ such that β_ϕ is closest to one.

A Linear Approximation. To provide a comparison of our $\mathrm{ERP}_{i,t}^{(\phi)}$ to the Quanto risk premium defined in Kremens and Martin (2019), we derive an *exact* decomposition of the risk-neutral covariance term from (10),[6]

$$\underbrace{\frac{1}{R_{f,T}}\mathrm{cov}_t^*\left(R_{mkt,T}^{\phi},\frac{e_{i,T}}{e_{i,t}}\right)}_{\mathrm{ERP}_{i,t}^{(\phi)}}=\phi\underbrace{\frac{1}{R_{f,T}}\mathrm{cov}_t^*\left(R_{mkt,T},\frac{e_{i,T}}{e_{i,t}}\right)}_{\mathrm{ERP}_{i,t}^{(1)}}$$
$$+\frac{\phi(\phi-1)}{2}\frac{1}{R_{f,T}}\mathrm{cov}_t^*\left(\xi_T^{\phi-2}(R_{mkt,T}-R_{f,T})^2,\frac{e_{i,T}}{e_{i,t}}\right),\tag{11}$$

where ξ_T is a random number that lies between $R_{f,T}$ and $R_{mkt,T}$. Note that $\mathrm{ERP}_{i,t}^{(1)}$ on the right hand-side is equivalent to the $\mathrm{QRP}_{i,t}$, i.e. the Quanto risk premium defined in Kremens and Martin (2019).

If, on average, the exchange rate's risk-neutral co-skewness with market's volatility is not too large, one could expect the non-linear term in (11) to be small and negligible. When that is true, a simplified empirical representation of (15) should work well:

$$\mathrm{ERX}_{i,t}=\alpha+\beta\,\mathrm{ERP}_{i,t}^{(1)},\tag{12}$$

where α and β can be estimated by jointly regressing $\mathrm{ERX}_{i,t}$ on $\mathrm{ERP}_{i,t}^{(1)}$. In the absence of the non-linear term in (11), we should have the coefficients in (12) to satisfy $\alpha=0$ and $\beta=\phi$.

However, conditionally, when downside risk spikes in the international financial market, the non-linear term might be important and non-negligible. As an illustrating example, we could get rid of the random variable ξ_T in the non-linear term when $\phi=2$ and derive a simpler

[6] The decomposition is exact due to the mean value theorem. Recall that $f(x)=f(c)+f'(c)(x-c)+\frac{1}{2}f"(\xi)(x-c)^2$ for $\xi\in[c,x]$. Here we used $f(x)=x^\phi$, $c=R_{f,T}$, and $x=R_{mkt,T}$.

representation of (11):

$$\mathrm{ERP}_{i,t}^{(2)} = 2\underbrace{\mathrm{ERP}_{i,t}^{(1)}}_{\mathrm{QRP}_{i,t}} + \frac{1}{R_{f,T}}\,\mathrm{cov}^*\left(\left(R_{mkt,T} - \mathbb{E}_t^*[R_{mkt,T}]\right)^2, \frac{e_{i,T}}{e_{i,t}}\right), \tag{13}$$

where the second covariance term can be synthetically priced as a Quanto contract of $\mathrm{SVIX}_t^2 = \mathbb{E}_t^*\left[(R_{mkt,T} - \mathbb{E}_t^*[R_{mkt,T}])^2\right]$, i.e., the 'simple variance swap' introduced in Martin (2011). Empirically, as shown in Martin (2017), the SVIX_t^2 tracks VIX_t^2 closely. So we can think of this second covariance term as the co-skewness between exchange rate returns and the level of VIX_t^2.[7] In a nutshell, this analysis shows that (simple) variance risk becomes priced by FX investors when $\phi > 1$.

3.2 Construction of the main variables

We now describe how we compute, for each currency pair, the expected excess exchange rate return and the expected risk premium. We exploit a cross-section of the G30 currencies against the USD over the sample spanning the 1996.01.01 to 2020.12.31 period.

3.2.1 Expected excess exchange rate returns

We use exchange rate forecasts to compute expected excess exchange rate returns as follows:

$$\mathrm{ERX}_{i,t,T} = \frac{\mathbb{E}_t\left[e_{i,T}\right]}{e_{i,t}} - \frac{R_{f,T}}{R_{f,T}^i}, \tag{14}$$

where $\mathbb{E}_t\left[e_{i,T}\right]$ is the mean exchange rate forecast for currency i, observed at time t, and for horizon $T-t$. The last term in (14) is the traditional UIP forecast that we approximate with the interest rate differential between USD and the considered currency. For the construction of this component, we rely on daily zero-coupon rates bootstrapped from money market rates and interest rate swaps obtained from Bloomberg. We match the maturity with that of the forecasts. The monthly exchange rate forecast data are from the Foreign Exchange

[7] A caveat is that, to have $\mathrm{SVIX}_t^2 \approx \mathrm{VIX}_t^2$, market return has to be log-normally distributed.

Consensus Forecasts, with forecast horizons of 1, 3, 6, 12 and 24 months. We interpolate monthly data to obtain daily observations.

3.2.2 The expected risk premium

We here describe how we construct an empirical measure of the expected risk premium, which is given by:

$$\mathrm{ERP}_{i,t}^{(\phi)} = \mathrm{cov}_t^* \left(\frac{R_{mkt,T}^{\phi}}{\mathbb{E}_t^*[R_{mkt,T}^{\phi}]}, \frac{e_{i,T}}{e_{i,t}} \right). \tag{15}$$

The above risk-neutral covariance is not directly observable from market prices, as Quanto contracts on a power of the market index are not typically traded (except for $\phi = 1$ as discussed in Kremens and Martin (2019)). To overcome this challenge, we can decompose the risk-neutral covariance into its three distinct components:

$$\mathrm{cov}_t^* \left(\frac{R_{mkt,T}^{\phi}}{\mathbb{E}_t^*[R_{mkt,T}^{\phi}]}, \frac{e_{i,T}}{e_{i,t}} \right) = \rho_{\phi,i,T}^* \sqrt{\mathrm{var}_t^* \left(\frac{R_{mkt,T}^{\phi}}{\mathbb{E}_t^*\left[R_{mkt,T}^{\phi}\right]} \right)} \sqrt{\mathrm{var}_t^* \left(\frac{e_{i,T}}{e_{i,t}} \right)}, \tag{16}$$

where $\rho_{\phi,i,T}^*$ captures the risk-neutral correlation between $R_{mkt,T}^{\phi}$ and $e_{i,T}/e_{i,t}$, while the var_t^* operator denotes the risk-neutral variance.

Armed with this decomposition, we can compute the risk-neutral covariance for each level of ϕ. The risk-neutral correlation is not directly observable — so we use the realized correlation between $R_{mkt,T}^{\phi}$ and $e_{i,T}/e_{i,t}$.[8] We use the return of the S&P 500 over the horizon T as a proxy for $R_{mkt,T}$ and the gross return on spot exchange rates from Bloomberg for $e_{i,T}/e_{i,t}$.

The two risk-neutral variances, however, can be obtained from option prices, by computing risk-neutral moments of the market index returns and exchange rate returns with the Breeden-Litzenberger (or Carr-Madan) method. We provide details on the construction of

[8] We show in Appendix B.1 that the risk-neutral correlation, implied from the Quanto contracts in Kremens and Martin (2019), is close to the realized correlation between the market and exchange rate returns. So the realized correlation is a reasonable proxy of the risk-neutral correlation.

these risk-neutral variances in Appendix B.2. Essentially, we extract the risk-neutral variance of the powered market return and of the exchange rate returns using equity index options and FX options, respectively, which we describe below.

Equity index option prices are based on the daily implied volatility surface of SPX European options, as provided by OptionMetrics. We use observations from 1996.01.01 to 2020.12.31. 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, 6, 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 then compute the daily expected risk premium $\mathrm{ERP}_{i,t,T}^{(\phi)}$ for each of the 30 currency pairs, each level of ϕ ranging between 1 and 10, and each maturity between 1 and 24 months.

4 Empirical Results

We estimate ϕ, using the two approaches presented in Section 3.1, and explore the conditional term structure of risk preferences. We first highlight the contributions of our approach and then discuss the results.

4.1 Contributions

Our methodology contributes to the literature in three 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 Section 2.2.1, 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 $\mathrm{ERP}^{\phi}_{i,t}$ 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 Bekaert, Engstrom, and Xu (2021). Third, we use observable expected exchange rate return to measure the left-hand side of (10) and (12), i.e., $\mathrm{ERX}_{i,t}$. 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 ($\mathrm{ERP}^{(\phi)}_{i,t}$) computed from forward-looking option prices and the consensus based expected excess return ($\mathrm{ERX}_{i,t}$) 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 ϕ.

4.2 Unconditional analysis

We estimate the unconditional level of ϕ by running panel regressions at the daily frequency based on the following specification:

$$\mathrm{ERX}_{i,t,T} = \alpha_i + \alpha_t + \beta_\phi \, \mathrm{ERP}^{(\phi)}_{i,t,T} + \varepsilon_{\phi,i,t,T}, \tag{17}$$

where α_i and α_t are currency and time (calendar date) fixed effects, respectively. We estimate the above specification by forecast horizon $T-t$ and for different values of ϕ. Table 1 reports estimates for different 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. In each panel, we report the results for a different value of ϕ, ranging between 1 and 7.

Table 1 about here

Panel A shows that the slope coefficient β_ϕ is close to one when ϕ is between 4 and 5, as suggested by our identity (10). Consistent with this finding, the slope coefficient β_1, which corresponds to regressing $\mathrm{ERX}_{i,t,T}$ on $\mathrm{ERP}^{(1)}_{i,t,T}$ is equal to 4.497 and is highly statistical significant. In this second approach, the value of β_1 can be directly interpreted as the level of ϕ when the non-linear approximation term in (11) is negligible. The results imply that the second approach is a good approximation of the baseline approach, as both cases imply a level of ϕ lying between 4 and 5. This is a significant departure from the case of log utility (Kremens and Martin, 2019), i.e., $\phi = 1$. Our estimate suggests that investors are thus substantially averse to higher-order risk.

We then explore the term structure of risk preferences, as given by the estimates of ϕ over different forecast horizons. When moving from Panel A to Panel D, we find that the term structure is downward sloping unconditionally. In particular, Column (1) indicates that ϕ decreases from 4.497 at the one-month horizon to 1.439 at the two-year horizon. These findings suggest that investors care less about higher-order risk as their forecast horizon increases. A possible explanation is that the risk of a currency crashes 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 a shorter horizon, investors are then more averse to such tail risk events, which is expressed by a higher ϕ.

4.3 Additional analysis

We now summarize a set of additional results that further corroborate our core findings. First, we show that the dynamics of ϕ at different horizons remain qualitatively similar when adding a set of control variables $X_{i,t}$ to our panel specifications, i.e., we run panel regressions based on

$$\mathrm{ERX}_{i,t,T} = \alpha_t + \alpha_i + \beta_\phi \, \mathrm{ERP}^{(\phi)}_{i,t,T} + \delta' X_{i,t} + \varepsilon_{\phi,i,t,T}, \tag{18}$$

where $X_{i,t}$ includes the year-on-year inflation differential between the US and country i at time t, 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), and the dollar basis constructed at time t using the one-month US dollar interest rate and its synthetic replication based on local currency interest rates and spot/forward exchange rates. We report the estimates in Table 2 and show in Column (1) that ϕ decreases from 5.085 at the one-month horizon to 2.371 at the two-year horizon, thus confirming that investors are more (less) averse to tail risk events when their investment horizon is shorter (longer).

Table 2 about here

In Table 3, moreover, we employ expected exchange return (as opposed to currency expected excess returns) as the dependent variable. At the same time, we add the corresponding interest rate differential as explanatory variable to our panel regressions

$$\mathrm{EFX}_{i,t,T} = \alpha_t + \alpha_i + \beta_\phi \, \mathrm{ERP}^{(\phi)}_{i,t,T} + \gamma \, \mathrm{IRD}_{i,t,T} + \varepsilon_{\phi,i,t,T}, \tag{19}$$

where $\mathrm{EFX}_{i,t,T} = \mathbb{E}_t \left[e_{i,T} \right] / e_{i,t}$ 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 1. In particular, Column (1) shows that ϕ decreases from 4.605 at the one-month horizon to 1.451 at the two-year horizon. In Table 4, we further add the control variables in X_t but results remain virtually identical to those reported in Table 2.

Tables 3 and 4 about here

4.4 Conditional term structure of risk preferences

We then investigate how risk preferences vary over time. Several empirical studies suggest that risk aversion increases in "bad times". For example, 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). It is thus important to analyze how risk preferences of FX investors, as measured by ϕ, vary in good vs. bad times.

To do so, we split our sample according to NBER-dated recessions and expansions, estimate the specification in Equation (17) on each subsample separately, and then plot the estimates of ϕ in the top-left Panel of Figure 2. Consistent with the existing evidence, we find that ϕ tends to increase during recessions, suggesting that FX market participants are indeed more averse to higher-order risk as economic conditions worsen.

FIGURE 2 ABOUT HERE

We then turn to a conditional analysis of this term structure. The top-left Panel of Figure 2 shows that the term structure has a steep negative slope in recessions, while it becomes much flatter 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. Our paper is the first to provide insights on how risk preferences vary over different horizons as well as on the conditional term structure.

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

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

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

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

Bansal, Ravi, and Bruce N Lehmann, 1997, Growth-optimal portfolio restrictions on asset pricing models, *Macroeconomic dynamics* 1, 333–354.

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

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.

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

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

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, *The Journal of Finance* 29, 857–885.

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

Gormsen, Niels Joachim, 2021, Time variation of the equity term structure, *The 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).

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 Jr, JL, 1956, A new interpretation of the information rate, bell systems tech, *J* 35, 917–926.

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

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

Longstaff, Francis A, and Jiang Wang, 2012, Asset pricing and the credit market, *The 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.

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

Martin, Ian, 2011, Simple variance swaps, Discussion paper, National Bureau of Economic Research.

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

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

Martin, Ian WR, and Christian Wagner, 2019, What is the expected return on a stock?, *The 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.

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

Politis, Dimitris N., and Joseph P. Romano, 1994, The stationary bootstrap, *Journal of the American Statistical Association* 89, 1301–1313.

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

Ross, Stephen A, 1999, Adding risks: Samuelson's fallacy of large numbers revisited, *Journal of Financial and Quantitative Analysis* 34, 323–339.

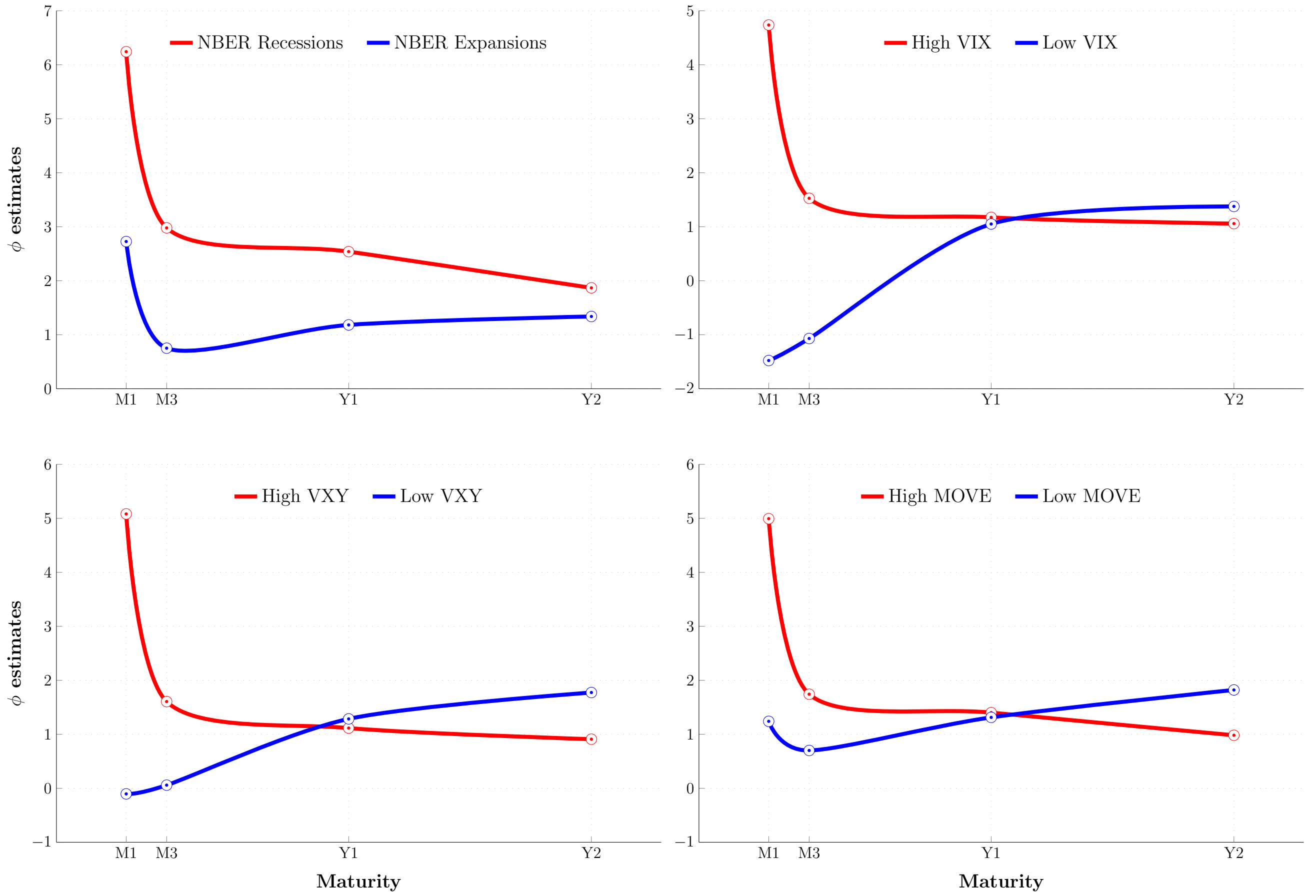

Figure 2. Market Conditions and Estimates of ϕ

This figure displays panel estimates of ϕ across different market conditions. The top-left panel refers to NBER recession and NBER expansion periods, the top-right panel to high VIX and low VIX periods (relative to its sample average), the bottom-right panel to high VXY and VYX periods (defined relative to its sample average), and the bottom-left panel to high and low MOVE periods (defined relative to its 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. Expected Excess Returns and Risk Premia

This table presents panel regression estimates based on the following specification

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

where $\text{ERX}_{i,t,T}$ is the expected currency excess return observed at time t for currency i and maturity $T-t$ and $\text{ERP}^{(\phi)}_{i,t,T}$ 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. $\text{ERX}_{i,t,T}$ is calculated using exchange rate consensus forecasts and interest rate differentials whereas $\text{ERP}^{(\phi)}_{i,t,T}$ is based on S&P 500 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. Data are from Bloomberg, FRED, JP Morgan, and OptionMetrics.

Panel A: 1-month Maturity							
ϕ	1	2	3	4	5	6	7
ERP	4.497***	2.289***	1.550***	1.177***	0.950***	0.798***	0.687***
	(0.828)	(0.435)	(0.305)	(0.239)	(0.199)	(0.172)	(0.153)
R^2	0.326	0.326	0.325	0.325	0.324	0.324	0.324
N	154,876	154,876	154,876	154,876	154,876	154,876	154,876
Panel B: 3-month Maturity							
ERP	1.700***	0.871***	0.593***	0.453**	0.368**	0.311**	0.270**
	(0.553)	(0.295)	(0.209)	(0.166)	(0.139)	(0.122)	(0.109)
R^2	0.325	0.325	0.325	0.324	0.324	0.324	0.324
N	156,635	156,635	156,635	156,635	156,635	156,635	156,635
Panel C: 1-year Maturity							
ERP	1.539***	0.848***	0.617***	0.500***	0.428***	0.378***	0.342***
	(0.435)	(0.245)	(0.182)	(0.150)	(0.130)	(0.117)	(0.107)
R^2	0.448	0.448	0.447	0.447	0.447	0.447	0.446
N	156,548	156,548	156,548	156,548	156,548	156,548	156,548
Panel D: 2-year Maturity							
ERP	1.439**	0.821**	0.615**	0.510**	0.445**	0.402**	0.371**
	(0.538)	(0.310)	(0.234)	(0.195)	(0.171)	(0.155)	(0.143)
R^2	0.526	0.526	0.526	0.525	0.525	0.525	0.525
N	140,704	140,704	140,704	140,704	140,704	140,704	140,704
currency fe	✓	✓	✓	✓	✓	✓	✓
time fe	✓	✓	✓	✓	✓	✓	✓

Table 2. Expected Excess Returns and Risk Premia – With Controls

This table presents panel regression estimates based on the following specification

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

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 year-on-year inflation differential between the US and country i at time t, 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, and the dollar basis constructed at time t using the one-month US dollar interest rate and its synthetic replication based on local currency interest rates and spot/forward exchange rates. $\mathrm{ERX}_{i,t,T}$ 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. Data are from Bloomberg, FRED, JP Morgan, and OptionMetrics.

Panel A: 1-month Maturity							
ϕ	1	2	3	4	5	6	7
ERP	5.085***	2.386***	1.473***	1.010***	0.730**	0.543**	0.409*
	(1.223)	(0.641)	(0.445)	(0.346)	(0.286)	(0.245)	(0.215)
R^2	0.330	0.330	0.329	0.329	0.329	0.328	0.328
N	152,627	152,627	152,627	152,627	152,627	152,627	152,627
Panel B: 3-month Maturity							
ERP	4.948***	2.463***	1.621***	1.196***	0.940***	0.770***	0.648***
	(0.941)	(0.505)	(0.358)	(0.283)	(0.237)	(0.205)	(0.182)
R^2	0.335	0.334	0.332	0.331	0.331	0.330	0.329
N	153,827	153,827	153,827	153,827	153,827	153,827	153,827
Panel C: 1-year Maturity							
ERP	2.235***	1.216***	0.872***	0.694***	0.584**	0.506**	0.448**
	(0.698)	(0.400)	(0.301)	(0.249)	(0.217)	(0.194)	(0.178)
R^2	0.455	0.455	0.454	0.454	0.454	0.453	0.453
N	153,740	153,740	153,740	153,740	153,740	153,740	153,740
Panel D: 2-year Maturity							
ERP	2.371***	1.358***	1.016***	0.837***	0.727***	0.654***	0.601***
	(0.690)	(0.404)	(0.309)	(0.259)	(0.227)	(0.206)	(0.192)
R^2	0.551	0.551	0.551	0.551	0.551	0.551	0.551
N	140,570	140,570	140,570	140,570	140,570	140,570	140,570
controls	✓	✓	✓	✓	✓	✓	✓
currency fe	✓	✓	✓	✓	✓	✓	✓
time fe	✓	✓	✓	✓	✓	✓	✓

Table 3. Expected Exchange Rate Returns and Risk Premia

This table presents panel regression estimates based on the following specification

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

where $\text{EFX}_{i,t,T}$ is the expected exchange rate return observed at time t for currency i and maturity $T-t$, $\text{ERP}^{(\phi)}_{i,t,T}$ 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, with α_t and α_i denoting time and currency 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 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. Data are from Bloomberg, FRED, JP Morgan, and OptionMetrics.

Panel A: 1-month Maturity							
ϕ	1	2	3	4	5	6	7
ERP	4.605***	2.346***	1.589***	1.208***	0.976***	0.820***	0.707***
	(0.862)	(0.453)	(0.317)	(0.248)	(0.207)	(0.179)	(0.159)
R^2	0.334	0.334	0.333	0.333	0.332	0.332	0.332
N	154,876	154,876	154,876	154,876	154,876	154,876	154,876
Panel B: 3-month Maturity							
ERP	1.715***	0.880***	0.599***	0.458**	0.372**	0.315**	0.273**
	(0.565)	(0.301)	(0.213)	(0.169)	(0.142)	(0.124)	(0.111)
R^2	0.329	0.329	0.328	0.328	0.328	0.328	0.327
N	156,635	156,635	156,635	156,635	156,635	156,635	
Panel C: 1-year Maturity							
ERP	1.516***	0.834***	0.606***	0.491***	0.420***	0.372***	0.336***
	(0.456)	(0.256)	(0.191)	(0.157)	(0.137)	(0.123)	(0.113)
R^2	0.368	0.368	0.368	0.367	0.367	0.367	0.367
N	156,548	156,548	156,548	156,548	156,548	156,548	
Panel D: 2-year Maturity							
ERP	1.451***	0.829***	0.622***	0.516***	0.451***	0.408***	0.377***
	(0.512)	(0.294)	(0.222)	(0.185)	(0.162)	(0.146)	(0.136)
R^2	0.351	0.351	0.350	0.350	0.350	0.350	0.350
N	140,704	140,704	140,704	140,704	140,704	140,704	140,704
IRD	✓	✓	✓	✓	✓	✓	✓
currency fe	✓	✓	✓	✓	✓	✓	✓
time fe	✓	✓	✓	✓	✓	✓	✓

Table 4. Expected Exchange Rate Returns and Risk Premia – With Controls

This table presents panel regression estimates based on the following specification

$$\mathrm{EFX}_{i,t,T} = \alpha_t + \alpha_i + \beta_\phi \, \mathrm{ERP}^{(\phi)}_{i,t,T} + \gamma \, \mathrm{IRD}_{i,t,T} + \delta' X_t + \varepsilon_{\phi,i,t,T},$$

where $\mathrm{EFX}_{i,t,T}$ is the expected exchange rate return observed at time t for currency i and maturity $T-t$, $\mathrm{ERP}^{(\phi)}_{i,t,T}$ 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, with α_t and α_i denoting time and currency fixed effects, respectively. X_t refers to control variables, i.e., the year-on-year inflation differential between the US and country i at time t, 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, and the dollar basis constructed at time t using the one-month US dollar interest rate and its synthetic replication based on local currency interest rates and spot/forward exchange rates. $\mathrm{EFX}_{i,t,T}$ is calculated using exchange rate consensus forecasts whereas $\mathrm{ERP}^{(\phi)}_{i,t,T}$ is based on S&P 500 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. Data are from Bloomberg, FRED, JP Morgan, and OptionMetrics.

Panel A: 1-month Maturity							
ϕ	1	2	3	4	5	6	7
ERP	5.099***	2.393***	1.477***	1.013***	0.732**	0.545**	0.411*
	(1.225)	(0.642)	(0.446)	(0.347)	(0.287)	(0.246)	(0.216)
R^2	0.337	0.337	0.336	0.336	0.336	0.335	0.335
N	152,627	152,627	152,627	152,627	152,627	152,627	152,627
Panel B: 3-month Maturity							
ERP	4.938***	2.458***	1.618***	1.193***	0.938***	0.768***	0.647***
	(0.938)	(0.503)	(0.356)	(0.282)	(0.236)	(0.204)	(0.182)
R^2	0.341	0.340	0.339	0.338	0.337	0.337	0.336
N	153,827	153,827	153,827	153,827	153,827	153,827	153,827
Panel C: 1-year Maturity							
ERP	2.212***	1.203***	0.862***	0.686**	0.577**	0.500**	0.442**
	(0.712)	(0.408)	(0.307)	(0.254)	(0.221)	(0.198)	(0.181)
R^2	0.388	0.387	0.387	0.386	0.386	0.385	0.385
N	153,740	153,740	153,740	153,740	153,740	153,740	153,740
Panel D: 2-year Maturity							
ERP	2.439***	1.398***	1.048***	0.865***	0.753***	0.677***	0.624***
	(0.670)	(0.391)	(0.299)	(0.249)	(0.218)	(0.198)	(0.184)
R^2	0.373	0.373	0.373	0.372	0.372	0.372	0.372
N	140,570	140,570	140,570	140,570	140,570	140,570	140,570
IRD & *controls*	✓	✓	✓	✓	✓	✓	✓
currency fe	✓	✓	✓	✓	✓	✓	✓
time fe	✓	✓	✓	✓	✓	✓	✓

Internet appendix to

“Expected Currency Returns”

(not for publication)

Abstract

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

A Optimal growth portfolio

A.1 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) his subjective belief from the rational 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's Ely lecture note (section 6.2) or 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), $\xi_T = R_{mkt,T}^{-\frac{1}{\theta}} / \mathbb{E}_t \left[R_{mkt,T}^{-\frac{1}{\theta}}\right]$. 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 2.2.1, the growth optimal portfolio is $R_{g,T} = \lambda R_{mkt,T}^{\theta^{-1}+1}$. 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.

B Ingredients of $\mathrm{ERP}_{i,t}^{(\phi)}$

B.1 Estimating the risk-neutral correlation $\rho^*_{\phi,i,t}$

We measure $\rho_{\phi,i,t}$ using a backward-looking sample correlation between daily $\mathcal{R}_t^\gamma$ and $\mathcal{E}_t$ over a window that matches the maturity of the options. Suppose that on day t, for example, we compute $\mathrm{var}^*[\cdot]$ and $\mathbb{E}^*[\cdot]$ using options between t and T. We then calculate $\rho_{i,\gamma}$ using daily $\mathcal{R}_t^\gamma$ and $\mathcal{E}_t$ between t and $t-T$.

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 $\mathrm{ERP}_{i,t}^{(1)}$ that we constructed using the realized correlations and match the same sample period.

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

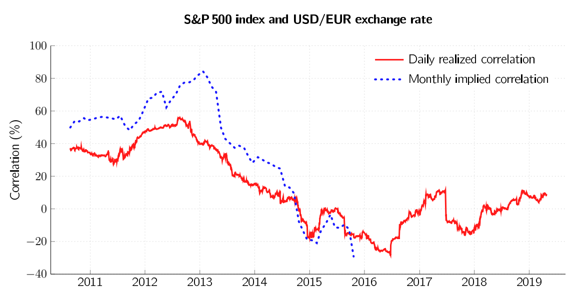

B.2 Risk Neutral Moments

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 $\frac{R^{\phi}_{mkt,T}}{\mathbb{E}^*_t[R^{\phi}_{i,T}]}$, we recall the formula

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

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

$$\mathbb{E}^*_t[R^{\theta}_{mkt,T}] = R^{\theta}_{f,t} + R_{f,t}\int_0^{F_{t,T}} \frac{\theta(\theta-1)}{S^{\theta}_t}(S_tR_{f,t} - F_{t,T} + K)^{\theta-2}\Omega_{t,T}(K)dK\,, \tag{A.1}$$

where $R_{f,t}$ is the risk free rate from yield curve with maturity T, K is the strike of option, S_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 $S_tR_{f,t} = F_{t,T}$ in our computation to ignore the dividend payment, i.e. we assume the payment of dividends are not reinvested. The formula simplifies to

$$\mathbb{E}^*_t[R^{\theta}_{mkt,T}] = R^{\theta}_{f,t} + R_{f,t}\int_0^{F_{t,T}} \frac{\theta(\theta-1)}{S^{\theta}_t}K^{\theta-2}\Omega_{t,T}(K)dK\,. \tag{A.2}$$

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

$$\text{var}^*_t\left(\frac{e_{i,T}}{e_{i,t}}\right) = \mathbb{E}^*_t\left(\frac{e_{i,T}}{e_{i,t}}\right)^2 - \left(\mathbb{E}^*_t\,\frac{e_{i,T}}{e_{i,t}}\right)^2\,, \tag{A.3}$$

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

$$\text{var}^*_t\left(\frac{e_{i,T}}{e_{i,t}}\right) = \frac{2}{B_{t,T}e^2_{i,t}}\left(\int_0^{F_{t,T}} P_{t,T}(K)dK + \int_{F_{t,T}}^{\infty} C_{t,T}(K)dK\right), \tag{A.4}$$

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.4) via trapezoidal integration.