Option-Implied Currency Risk Premia

Jakub W. Jurek and Zhikai Xu*

Abstract

We use cross-sectional information on the prices of G10 currency options to calibrate a non-Gaussian model of pricing kernel dynamics and construct estimates of conditional currency risk premia. We find that the mean historical returns to short dollar and carry factors (HML_{FX}) are statistically indistinguishable from their option-implied counterparts, which are free from peso problems. Skewness and higher-order moments of the pricing kernel innovations on average account for only 15% of the HML_{FX} risk premium in G10 currencies. These results are consistent with the observation that crash-hedged currency carry trades continue to deliver positive excess returns.

FIRST DRAFT: MARCH 2013

THIS DRAFT: DECEMBER 2013

---

*Jurek: Bendheim Center for Finance, Princeton University and NBER; *jjurek@princeton.edu*. Xu: Department of Operations Research and Financial Engineering, Princeton University; *zhikaixu@princeton.edu*. We thank John Campbell, Peter Carr, Pierre Collin-Dufresne, Valentin Haddad, Ian Martin, Peter Ritchken, and seminar participants at Case Western Reserve University, Dartmouth, Georgetown University, Harvard Business School, Princeton University, University of Rochester, the 2^{nd} EDHEC-Princeton Institutional Asset Management Conference, and the Princeton-Lausanne Quantitative Finance Workshop for providing valuable comments.

The absence of arbitrage links the exchange rate between two currencies and the economies' respective stochastic discount factors. We exploit this link to extract the dynamics of stochastic discount factors from the cross-section of currency options and produce a time series of option-implied currency risk premia. The option-implied currency risk premia are free of peso problems, and allow us to elucidate the importance of global jump risks in the determination of premia for common risk factors in currency returns.

Estimates of currency risk premia are commonly derived by studying the time-series and cross-sectional properties of historical (realized) returns. This paper develops an alternative approach to this question, which does not rely on historical currency returns, but instead uses cross-sectional data on the pricing of exchange rate options. We demonstrate that, even though no single currency option is informative about the expected return of the underlying currency pair, observing a sufficiently broad collection of options on currency cross-rates allows us to deduce the structure of the stochastic discount factors, and therefore the dynamics of currency risk premia.[1] In order to operationalize this methodology, we impose a factor structure on the pricing kernel dynamics, which assumes pricing kernels are driven by a combination of common (global) shocks and idiosyncratic (country-specific) shocks, with individual countries differing in their exposure to global shocks.[2] Cross-sectional differences in global factor loadings generate variation in option-implied exchange rate distributions, which we exploit to infer the structure of the pricing kernel innovations. We show that currency risk premia can be expressed using the cumulant generating functions of these innovations, similar to Backus, et al. (2011) and Martin (2013), generating a two-factor pricing model for currency returns. The global factor in our model corresponds to the HML_{FX} risk factor identified by Lustig, et al. (2011) – and is reflected in the returns to currency carry trade strategies – while the second factor represents compensation demanded by investors for being short their own (local) currency.

The model of pricing kernel dynamics we calibrate to exchange rate option data can be interpreted as a discretized version of a continuous-time model with time-changed Lévy increments (Carr and Wu (2004)). The global and country-specific pricing kernel innovations in our model are comprised of a Gaussian component and a non-Gaussian component, each of which have time-varying distributions. The non-Gaussian component can intuitively be thought of as capturing the joint effects of intra-period stochastic volatility and jumps ("disaster

[1] Our ability to recover risk premia from exchange rate options can be traced back to the fact that option-implied exchange rate distributions reveal the conditional, country-level pricing kernel distributions. It is not an application of the Recovery Theorem (Ross (2013)), and therefore does not rely on its underlying assumptions (e.g. finite-state Markov chains, time-homogenous dynamics, state-variables following bounded diffusions), detailed in Carr and Yu (2012).

[2] The presence of asymmetric global factor loadings can be used to rationalize violations of uncovered interest parity (Backus, et al. (2001)), is consistent evidence from historical currency returns (Lustig, et al. (2011), Hassan and Mano (2013)), as well as, evidence from exchange rate option markets (Bakshi, et al. (2008)). Models with imperfect risk-sharing provide a microfoundation for asymmetric global factor loadings (Verdelhan (2010), Ready, Roussanov, and Ward (2013)).

risks").[3] These features combine to produce a model that not only accommodates non-Gaussian innovations at each point in time, but also flexibly incorporates stochastic variation in second and higher moments, characteristic of currency option data (Carr and Wu (2007), Bakshi, et al. (2008)).[4] The availability of a closed-form expression for the generalized Fourier transform of the log currency return, enables us to efficiently recalibrate the model on each day in our sample (1999:1-2012:6; 3520 days) to match the cross-section of observed G10 exchange rate option prices (up to 45 cross-rate pairs x 5 strikes).[5] Finally, using the calibrated time-series of model parameters, we compute estimates of option-implied currency risk premia for the currency pairs in our panel.

We confront the option-implied currency risk premia with the historical returns to empirical factor mimicking portfolios identified by Lustig, et al. (2011, 2013), and examine the model's ability to forecast currency pair returns in the cross-section and time series. Crucial to note is that – due to limitations on the availability of exchange rate options – our investigation focuses only on the G10 (developed market) currencies, unlike Lustig, et al. (2011, 2013) who examine both developed and emerging market currencies. We find that our model matches the historical returns on common factor mimicking portfolios, indicating a substantial degree of integration in the pricing of currency risks across the spot and option markets. Moreover, the lack of a statistically significant wedge in the historical and option-implied risk premia suggests historical estimates of HML_{FX} risk premia in developed markets are not plagued by peso problems (Burnside, et al. (2011)). Our model is never rejected in cross-sectional (Fama-MacBeth) tests, and achieves adjusted R^2 values up to 50%, which are roughly double those achieved from using interest rate differentials alone. Finally, we decompose the model-implied currency risk premia across innovations to the log pricing kernels, differentiating between: (a) the type of shock (Gaussian vs. non-Gaussian); and, (b) the individual moments of the shock (variance, skewness, etc.). These results provide a structural decomposition of risk premia, which complements reduced-form studies using crash-hedged strategies to examine the effect of tail risks on currency risk premia (Burnside, et al. (2011), and Farhi, et al. (2013), Jurek (2013)). We show that the skewness and higher-order moments of the log pricing kernel innovations, on average, account for only 15% of the model-implied HML_{FX} risk premium. Consequently, the main

[3] Frameworks with (time-varying) disaster risks have been used to match the equity risk premium in consumption models (Martin (2013)), explain aggregate stock market volatility (Wachter (2013)), and as a mechanism for generating violations of uncovered interest rate parity (Farhi and Gabaix (2011). Gabaix (2012) demonstrates the power of this methodology in context of ten classical macro-finance puzzles.

[4] Our specification of pricing kernel dynamics shares the features of the continuous-time model of Bakshi, et al. (2008), which allows both global and country-specific innovations to be non-Gaussian. However, rather than imposing a flat term structure, our model additionally links the dynamics of the interest rates to the underlying state variables. This creates a link between the level of interest rates and risk premia, which is necessary for a consistent, risk-based explanation of violations of uncovered interest rate parity. The specification we propose generalizes the conditionally Gaussian affine model in Lustig, et al. (2011, 2013), and the jump-diffusive setup in Farhi, et al. (2013), which does not allow for country-specific jumps.

[5] The availability of the full cross-section of G10 cross-rates (45 pairs) plays an important role in model identification in a discrete-time setting. The model can also be identified by specifying time-consistent dynamics for the state variables governing the time-change, and using multiple tenors in the calibration, as in Bakshi, et al. (2008).

channel through which non-Gaussian innovations act to determine risk premia in our model – both global and country-specific – is as contributors of pricing kernel variance. For comparison, the contribution of skewness and higher-order moments to *equity* risk premia is on average 35% in models calibrated to historical consumption disasters (Barro (2006), Barro and Ursua (2008), Barro, et al. (2013)), but only 2% in a model calibrated to match the pricing of equity index options (Backus, et al. (2011)).

Our preferred calibration focuses on pricing options on a set of twenty four currency pairs, designed to capture the most liquidly traded exchange rate options, as well as, options on "typical" carry trade currencies. This set includes the nine X/USD pairs, and an additional fifteen pairs formed from currencies which had either the highest or lowest interest rate among G10 countries at some point in time in our sample. The preferred model calibration delivers a root mean squared option pricing error – measured in volatility points – of 1.12, which is in line with typical bid-ask spreads for exchange rate options (Table I, Panel C). Using the calibrated model, we compute pair-level estimates of conditional currency risk premia. We then aggregate these option-implied risk premia using the portfolio weights used in the construction of empirical factor mimicking portfolios, to produce the corresponding model-implied portfolio risk premia. We examine the historical returns to *conditional* factor mimicking portfolios – sorted on the basis of contemporaneous interest rates – and *unconditional* factor mimicking portfolios – sorted on the basis of backward looking averages of interest rates.

We find that option-implied risk premia for the conditional and unconditional HML_{FX} factor mimicking portfolios are 3.55% (3.44%) *per annum*, respectively (Table IV). Neither of these quantities is statistically distinguishable from the corresponding mean realized returns of 4.96% and 3.32% *per annum*; the t-statistics of the differences are equal to 0.43 and -0.16, respectively. We then decompose the model-implied HML_{FX} risk into contributions stemming from Gaussian and non-Gaussian pricing kernel shocks, as well as, across the moments of the pricing kernel innovations. The decomposition reveals that (on average) 58% of the risk premium for the HML_{FX} factor mimicking portfolio is due to non-Gaussian innovations in the global factor. To explore the mechanism through which these innovations contribute to the HML_{FX} risk premium, we decompose the risk-premium across the moments of the global pricing shock. This reveals a striking result. Although non-Gaussian shocks are the dominant drivers of the HML_{FX} risk premium, they exert their influence primarily by contributing variance, rather than skewness (or higher moments), to the global innovation. Specifically, we find that (on average) 85% of the total risk premium is attributed to the variance of the pricing kernel shocks, with only 10% due to skewness of the shocks, and roughly 5% due to the higher moments. This parallels the empirical results reported by Jurek (2013) based on an analysis of returns to crash-hedged G10 portfolios, but contrasts with Farhi, et al. (2013), who argue that "disaster risk accounts for more than a third of the currency risk premia in developed economies." This discrepancy reflects two significant differences in the empirical identification strategies. First, Farhi, et al. (2013) only observe data for X/USD options ($N-1$ option-implied distributions), and are therefore unable to fully identify the pricing kernel dynamics (N kernels) exclusively using exchange rate option data unlike this paper. Instead, they measure the total HML_{FX} risk premium from historical returns, and only use exchange rate options to pin down the disaster risk component. Second, their identification of the disaster risk component leans heavily on the assumption of Gaussian country-specific innovations. This assumption is at odds with evidence in Bakshi, et al. (2008) and the results of our calibration, both of which point to significantly non-Gaussian country-specific innovations (Table A.I). For example, the average skewness values of the country-specific innovations range from -0.82 (SEK) to -3.69 (AUD) in the cross-section; kurtosis values are routinely in excess of ten.

We conduct a similar analysis for the risk premia demanded by investors in each of the G10 countries for being short their own local currency (Table III). We refer to this risk premium as the short reference risk premium, and show that – within our model – it is determined exclusively by the even moments of the country-specific pricing kernel innovations. Empirically, we find that the mean share of the risk premium accounted for by the variance of the country-specific innovations is greater than 98% in all ten G10 currencies. The short reference risk premia range from 0.54% (AUD) to 1.33% (SEK) *per annum*. We find that the historical returns (3.32% *per annum*) to a strategy which is unconditionally short the U.S. dollar against an equal-weighted basket of the G10 currencies, are statistically indistinguishable from the model-implied premium of 1.59% *per annum* (difference t-stat: 0.63). However, the model is unable to explain the high historical returns to the "short dollar carry trade" identified in Lustig, et al. (2013), which goes long (short) the U.S. dollar against a basket of foreign currencies when foreign interest rates are low (high).

To further evaluate the ability of the model to explain realized currency returns we conduct regressions of currency pair excess returns onto the option-implied risk premia (Table V). We find that the single model-implied variable achieves an adjusted R^2 of 30% in cross-sectional (Fama-MacBeth) regressions, which is comparable to that achievable using the interest rate differential (26%). The explanatory power of the model-implied quantities increases even further when the risk premium is disaggregated into its two constituents (47%). The option-implied risk premia are never driven out by interest rate differentials, though they also do not subsume their cross-sectional forecasting power, suggesting that currency options and interest rates carry non-redundant information about currency risk premia. Although the null hypothesis of the model is never rejected in these regressions, the rejections of the null of no predictability are relatively weak. The model fares poorly in pooled panel regressions, mimicking the empirical difficulties of identifying a statistically significant and positive, risk-return tradeoff in equity markets (Campbell (1987), French, et al. (1987), Campbell and Hentschel (1992)).

Our preferred calibration allows for time-varying global factor loadings, ξ_t^i, which are parameterized as a function of the prevailing interest rate differentials, $\xi_t^i = \xi^i - \Psi_t \cdot (r_{t,t+1}^i - r_{t,t+1}^{US})$. This parameterization captures the empirical regularity that portfolios sorted on prevailing interest differentials (conditional factor mimicking portfolios) outperform portfolios sorted on historical averages of interest rate differentials (unconditional factor mimicking portfolios). Indeed, we find a negative within-time-period/across pair association between the loadings and interest rate differentials ($\Psi_t > 0$), and a negative across-time/across-pair relation between the value of ξ^i and the unconditional mean interest rate differential (Table I, Panel A), consistent with a risk-based explanation for currency excess returns. The average estimate of the cross-sectional sensitivity, Ψ_t, is 0.42 (t-stat: 13.09). Although this result is consistent with the preferred ("unrestricted") model specification in Lustig, et al. (2011), the covariation between loadings and exchange rates indicated by exchange rate options is quantitatively weak within *developed* market currencies. As a result, the mean option-implied risk premium for the conditional HML_{FX} replicating portfolio is only 11bps higher than for the corresponding unconditional factor mimicking portfolio in developed market currencies. Equivalently, generating a large wedge between the model-implied conditional and unconditional risk premia by linking global factor loadings to interest rate differentials, induces excessive comovement of option-implied exchange rate moments with interest rate differentials, relative to that observed in the data. Finally, we document that while allowing for time-varying loadings leads to a modest improvement in the model's fit to option prices, it also leads to an overall *decline* in model-implied risk premia. The intuition for this result is that since the models are constrained to match option prices (i.e. exchange rate volatilities), an increase in the global factor loading spread requires some combination of a decline in the global and/or country-specific state variables governing the quantity of risk, such that the level of the risk premium may decline.

The remainder of the paper is organized as follows. Section 1 introduces the model of country-level pricing kernel dynamics and derives key results for currency returns, interest rates, risk premia, and currency option prices. Section 2 describes the data, model calibration, and parameter estimates. Section 3 reports the time series of option-implied currency risk premia, conducts various decompositions, and compares the model risk premia to historical returns on empirical factor mimicking portfolios. Section 4 discuses the effect of time-varying loadings on currency risk premia within our option pricing calibration, and examines the implications of linking time-variation in loadings to interest rate differentials for option-implied moments. Section 5 concludes. Appendix A provides auxiliary results related to the cumulant generating function of the time-changed Lévy increments.

1 Pricing Kernels, Currency Risk Premia, and FX Option Pricing

We develop a parsimonious model of exchange rates based on a specification of pricing kernel dynamics driven by a combination of global and country-specific wealth shocks, and derive its implications for the time-series and cross-section of currency risk premia. Unlike previous papers which have studied these questions using historical (realized) currency returns, we extract estimates of instantaneous currency risk premia from data on foreign exchange options. The calibrated model allows us to: (1) obtain a time-series of option-implied risk premia for the currency carry factor (HML_{FX}; Lustig, et al. (2011)) and country-specific factors (e.g. such as premium for short dollar exposure; Lustig, et al. (2013)); (2) decompose pair- and portfolio-level risk premia across moments of the pricing kernel innovations; and, (3) assess the importance of non-Gaussian (jump) risks in the determination of currency risk premia, complementing earlier empirical work (Burnside, et al. (2011), Farhi, et al. (2013), Jurek (2013)).

Our model begins with a specification of country-level log pricing kernel dynamics of the form:

$$m^i_{t+1} - m^i_t \quad = \quad -\alpha^i_t - \xi^i_t \cdot L^g_{Z_t} - L^i_{Y^i_t} \tag{1}$$

where $L^g_{Z_t}$ and $L^i_{Y^i_t}$ are independent non-Gaussian shocks. The $L^g_{Z_t}$ shock is global, and is common to all countries; the $L^i_{Y^i_t}$ shocks are country-specific, and are cross-sectionally independent. This factor representation corresponds to a discretized version of the continuous-time dynamics in the time-changed Lévy model of Bakshi, et al. (2008), and shares the spirit of the affine models of Backus, et al. (2001), and Lustig, et al. (2011, 2013). The distribution of the shocks is time-varying, and depends on the underlying model state variables, Z_t and Y^i_t, which correspond to the time-change parameters of a dynamic model (Carr and Wu (2004)). Intuitively, the values of the state-variables can be thought of as determining the periodic variance of the pricing kernel shocks (i.e. the prices of risk). We leave the dynamics of the state variables unspecified, and recover them period-by-period from cross-sectional data on foreign exchange options.[6]

Cross-sectional differences in pricing kernel dynamics are driven by a combination of differences in: (a) the level of α^i_t; (b) countries' loadings on global shocks, ξ^i_t; (c) the time-varying level of the country-specific state-variables, Y^i_t; and, (d) parameters determining the higher-moments of the country-specific innovations, $L^i_{Y^i_t}$. The drifts, α^i_t, have no effect on currency risk premia, and can be substituted out using the yields of bonds maturing at $t+1$ for the purpose of option pricing. The importance of allowing for asymmetric country-level loadings on

[6] A supplementary online Technical Appendix derives the continuous-time analog of our model. The continuous-time model requires: (a) specifying the dynamics of the model state variables, Z_t and Y^i_t; and, (b) imposing the requirement that global factor loadings remain constant.

global innovations, ξ_t^i, for rationalizing violations of uncovered interest parity was first highlighted by Backus, et al. (2001).[7] Since then, this feature has been incorporated in the context of option pricing (Bakshi, Carr, and Wu (2008)), and – in the identification of a factor structure in currency returns (Lustig, et al. (2011, 2013)). Following Lustig, et al. (2011), we allow the global factor loadings to be time-varying, and explore the implications of this feature for the determination of currency risk premia.

1.1 Pricing kernel innovations

The dynamics of the pricing kernels and exchange rates in our model are driven by time-changed Lévy increments, such that each shock can be thought of as a combination of a Gaussian (diffusive) shock and non-Gaussian (jump) shock.[8] Moreover, the distributions of these two components are time-varying through their dependence on the state-variables, Z_t, and, Y_t^i, controlling the time-change. These features combine to produce a model that not only accommodates non-Gaussian innovations at each point in time, but also flexibly incorporates stochastic variation in second and higher moments, characteristic of currency option data (Carr and Wu (2007)).

In our model, *both* the global, $L^g_{Z_t}$, and country-specific increments, $L^i_{Y_t^i}$, are allowed to be non-Gaussian. This mimics the features of the continuous-time model of Bakshi, et al. (2008), but is different from Farhi, et al. (2013), who allow for global jump shocks, but fix the country-specific component to be Gaussian. The latter assumption is restrictive since it forces the model to match any non-normalities present in exchange rate option data using a combination of global factor loadings and the parameters governing the global jump distribution. Consistent with Bakshi, et al. (2008) we find that allowing for non-Gaussianity in country-specific innovations plays an important role in matching the data.

To gain intuition about the pricing kernel innovations, it is useful to first consider their non-time-changed equivalents (i.e. a model with i.i.d innovations). We denote the non-time-changed pricing kernel shocks corresponding to the time interval between t and $t+1$, L^g_t and L^i_t, to differentiate them from their time-changed counterparts, $L^g_{Z_t}$ and $L^i_{Y_t^i}$. Without loss of generality we normalize the non-time-changed innovations to have unit variance and decompose each of them ($j \in \{g, i\}$) into the sum of two components:

$$L_t^j \;\; = \;\; W^j_{(1-\eta_t^j)} + X^j_{\eta_t^j} \qquad (2)$$

[7] Backus, Telmer and Foresi (2001) show that in order to account for the anomaly in an affine model, one has to either allow for state variables to have asymmetric effects on state prices in different currencies or abandon the requirement that interest rates be strictly positive.

[8] Since we are in a discrete-time setting the distinction between diffusive and jump shocks is somewhat semantic. More generally, in the continuous-time version of the model the diffusive shock can also be non-Gaussian, if the instantaneous innovations to the state-variable governing the time change and the innovations to the pricing kernels are correlated (Bakshi, et al. (2008)). Consequently, the non-Gaussian component in our discrete-time innovation is designed to simultaneously capture the effects of stochastic variation in the state-variable within the modeled time interval, and the effect of jumps.

where $W^j_{(1-\eta^j_t)}$ is a Gaussian innovation with variance $\left(1-\eta^j_t\right)$ and $X^j_{\eta^j_t}$ is a non-Gaussian innovation with variance η^j_t, and $\eta^j_t \in [0,1]$. With this normalization, the parameter, η^j_t, is interpretable as the time-varying share of innovation variance due to jumps. The jump component has a CGMY distribution, introduced by Carr, et al. (2002), which is characterized by a quartet of potentially time-varying parameters, $\{C,\, G,\, M,\, Y\}^j_t$:

$$\mu^j_t[dx] \quad = \quad \begin{cases} C^j_t \cdot e^{G^j_t \cdot x} \cdot |x|^{-Y^j_t - 1} \cdot dx & x \leq 0 \\ C^j_t \cdot e^{-M^j_t \cdot x} \cdot x^{-Y^j_t - 1} \cdot dx & x > 0 \end{cases} \tag{3}$$

The C^j_t parameter is a scaling factor, which is set such that the variance of the jump component is η^j_t; G^j_t and M^j_t determine the exponential dampening of the distribution for negative and positive jumps, respectively. The Y^j_t parameter can be interpreted as measuring the degree of similarity between the jump process and a Brownian motion. The CGMY process nests compound Poisson jumps ($-1 \leq Y < 0$), infinite-activity jumps with finite variation ($0 \leq Y < 1$), as well as, infinite-activity jumps with infinite variation ($1 \leq Y < 2$). Specifically, we assume that: (1) global jumps are one-sided, allowing only for positive shocks to marginal utility ($M^g_t = \infty$); and, (2) country-specific jumps are two-sided, capturing both positive and negative idiosyncratic shocks. Finally, we use the (time-change) state-variables, Z_t and Y^i_t, to set the periodic volatility of the Lévy increments. Appendix A discusses our specification in more detail.

Similar to Martin (2013), we rely on the cumulant generating functions (CGF) of the pricing kernel innovations, $k\,[u]$, to express quantities of interest such as yields, currency risk premia and and option prices.[9] The cumulant generating function for the non-time-changed Lévy increments, L^j_t, are reported in Appendix A. To derive the CGF for the corresponding *time-changed* increments we rely on Theorem 1 in Carr and Wu (2004). In our setup, the time-change is controlled by *pre-determined* state-variables, Z_t and Y^i_t, which allow the model to have non-identically distributed innovations over time. Unlike in a more typical stochastic volatility model, the time-change variables affect not only volatility, but also the higher order moments of the pricing kernel, enabling the model to better match the empirical features of foreign exchange option data. Theorem 1 of Carr and Wu (2004) states that for a generic time change, $\mathcal{T}$, the cumulant generating function of the time-changed Lévy process, $L_{\mathcal{T}}$, is given by $k_{\mathcal{T}}\left[k_L\left[u\right]\right]$, where $k_L[u]$ is the cumulant generating function of the non-time-changed

[9]Recall that the cumulant generating function of a random variable, ϵ_{t+1}, is defined as follows:

$$k_\epsilon[u] = \ln E_t\left[\exp\left(u \cdot \epsilon_{t+1}\right)\right] = \sum_{j=1}^{\infty} \frac{\kappa_j \cdot u^j}{j!}$$

where κ_j, are the cumulants of the random variable, which can be computed by taking the j-th derivative of $k_\epsilon[u]$ and evaluating the resulting expression at zero. The cumulant generating function of the sum of two independent random variables is equal to the sum of their cumulant generating functions.

process and $k_{\mathcal{T}}[u]$ is the cumulant generating function of the time-change. In our case, the time-change variables are fixed within the measurement interval (i.e. they follow a degenerate stochastic process with zero drift and volatility), such that:

$$k_{L^g_{Z_t}}[u] = k_{Z_t}\left[k_{L^g_t}[u]\right] = k_{L^g_t}[u]\cdot Z_t \tag{4a}$$

$$k_{L^i_{Y^i_t}}[u] = k_{Y^i_t}\left[k_{L^i_t}[u]\right] = k_{L^i_t}[u]\cdot Y^i_t \tag{4b}$$

Unless specifically noted with superscripts, cumulant generating functions are computed under the historical (objective) measure, $\mathbb{P}$. Using these results, along with the cumulant generating functions of the Gaussian and non-Gaussian components (Appendix A), the cumulant generating function of the time-changed Lévy increments, $k_{L^j_{S_t}}[u]$ – for the empirically relevant case when $Y \neq \{0,1\}$ – is given by:

$$\begin{aligned} k_{L^j_{S_t}}[u] &= k_{W^j_{\left(1-\eta^j_t\right)\cdot S_t}}[u] + k_{X^j_{\eta^j_t\cdot S_t}}[u] \\ &= \left(1-\eta^j_t\right)\cdot\frac{u^2}{2}\cdot S_t + \eta^j_t\cdot\frac{\left((M-u)^Y - M^Y + (G+u)^Y - G^Y\right)}{Y\cdot(Y-1)\cdot\left(M^{Y-2}+G^{Y-2}\right)}\cdot S_t \end{aligned} \tag{5}$$

To fix intuition about the analytical results, we will occasionally specialize to the case of a conditionally Gaussian model, in which case the cumulant generating function of the increments can be obtained by setting, $\eta^j_t = 0$.

1.2 Term structure of interest rates

Since the stochastic discount factors must price the risk-free claims in their respective economies, we have $E_t\left[M^i_{t+1}\right] = M^i_t\cdot\exp\left(-y^i_{t,t+1}\cdot\tau\right)$. Using the CGFs of the Lévy increments we can express the yield on a one-period bond in country i as:

$$\begin{aligned} y^i_{t,t+1} &= \alpha^i_t - k_{L^g_{Z_t}}[-\xi^i_t] - k_{L^i_{Y^i_t}}[-1] \\ &= \alpha^i_t - k_{L^g_t}[-\xi^i_t]\cdot Z_t - k_{L^i_t}[-1]\cdot Y^i_t \end{aligned} \tag{6}$$

In our model, yields across countries share an exposure to the common global factor, Z_t, with each country having its own – potentially time-varying – loading, ξ^i_t. If the loadings are constant and the state-variables additionally follow affine diffusions under the relevant measures, e.g. as in Bakshi, et al. (2008), the term-structure of interest rates will obey a two-factor affine model. In the cross-section, this expression links interest rates to underlying risk exposures, ξ^i_t, establishing an important channel for a risk-based explanation of violations of uncovered interest parity. For example, if the global innovation is Gaussian ($\eta_t^g = 0$), the cumulant generating function of the global increment, $k_{L_t^g}[-\xi_t^i] = \frac{1}{2} \cdot \left(\xi_t^i\right)^2$, such that countries with higher (lower) global factor loadings tend to have lower (higher) interest rates, all else equal. More generally, if the first derivative of the cumulant generating function of the global increment evaluated at $-\xi_t^i$ is negative, $k^{'}_{L_t^g}[-\xi_t^i] < 0$, and the second derivative is positive, $k^{''}_{L_t^g}[-\xi_t^i] < 0$, interest rates will be inversely related to global factor loadings. A sufficient condition for this relationship to hold is that $G > \xi_t^i$ and $Y < 1$, which we show is satisfied under our preferred model calibration. In the time-series, the expression places a restriction on the dynamics of interest rates (yields) as a function of the global, Z_t, and country-specific, Y_t^i, state variables, and the parameters of the cumulant generating functions.

Finally, the above expression highlights that term-structure data can be used to extract information about currency dynamics, which is an approach taken up by Brennan and Xia (2006), Ang and Chen (2010) and Sarno, et al. (2012). While term structure dynamics reflect all moments of the underlying pricing kernels, this information is actually more than sufficient for characterizing the dynamics of risk premia, since these depend only on moments two and higher. Our approach exploits this feature to extract currency risk premia exclusively from exchange rate option data.

1.3 Exchange rates and currency risk premia

In the absence of restrictions on the trade of financial assets, no arbitrage requires that the pricing kernel in country I price all assets denominated in currency I, as well as, the currency-I returns of assets denominated in all foreign currencies, J. This links the dynamics of the exchange rate, S_{t+1}^{ji} – specified as the currency I price of currency J – and the realizations of the respective pricing kernels, M_{t+1}^i and M_{t+1}^j, through:

$$E_t^{\mathbb{P}}\left[\frac{M_{t+1}^i}{M_t^i} \cdot \frac{S_{t+1}^{ji}}{S_t^{ji}} \cdot x_{t+1}^j\right] = E_t^{\mathbb{P}}\left[\frac{M_{t+1}^j}{M_t^j} \cdot x_{t+1}^j\right] \tag{7}$$

$$E_t^{\mathbb{P}}\left[\frac{M_{t+1}^i}{M_t^i} \cdot x_{t+1}^i\right] = E_t^{\mathbb{P}}\left[\frac{M_{t+1}^j}{M_t^j} \cdot \left(\frac{S_{t+1}^{ji}}{S_t^{ji}}\right)^{-1} \cdot x_{t+1}^i\right] \tag{8}$$

where x_{t+1}^i and x_{t+1}^j are payoffs of the traded assets denominated in currencies I and J, respectively. If we could construct Arrow-Debreu securities corresponding to each $(N+1)$-tuple of realizations, $\{L_{Z_t}^g, \{L_{Y_t^k}^k\}_{k=1}^N\}$, the market would be complete and each of the above no arbitrage restrictions would have to hold state-by-state, such that the exchange rate would be given by the ratio of the corresponding pricing kernels (Fama (1984), Dumas

(1992), Backus, et al. (2001), Brandt, et al. (2006), Bakshi, et al. (2008)):[10]

$$\frac{S_{t+1}^{ji}}{S_t^{ji}} = \frac{M_{t+1}^{j}}{M_t^{j}} \cdot \left(\frac{M_{t+1}^{i}}{M_t^{i}}\right)^{-1} \tag{9}$$

The set of tradeable claims in our model economy includes risk-free bonds, exchange rates, and exchange rate options, such that markets are incomplete. To see this, note that although we observe exchange rate options on all $\frac{N\cdot(N-1)}{2}$ pairs, only $N-1$ are non-redundant, preventing us from inferring the realization of the N underlying stochastic discount factors. When markets are incomplete, the restriction, (9), provides a sufficient, but not necessary, condition for the exchange rate to satisfy the pair of no arbitrage conditions. For example, any exchange rate process formed by taking the ratio of a pair of candidate, strictly positive, stochastic discount factors is admissible. In these circumstances, exchange rates are non-redundant assets, such that trading in exchange rate options can be viewed as helping complete markets.

We interpret the pricing kernel dynamics, (1), as a representation of candidate stochastic discount factors in incomplete markets. Correspondingly, we form a candidate exchange rate process using the ratio of the proposed pricing kernels, and recover its dynamics by matching the prices of exchange rate options. Given our pricing kernel parametrization the log currency return can be written as:

$$\begin{aligned} s_{t+1}^{ji} - s_t^{ji} &= \left(m_{t+1}^{j} - m_t^{j}\right) - \left(m_{t+1}^{i} - m_t^{i}\right) \\ &= -\alpha_t^j + \alpha_t^i - \left(\xi_t^j - \xi_t^i\right) \cdot L_{Z_t}^{g} - L_{Y_t^j}^{j} + L_{Y_t^i}^{i} \end{aligned} \tag{10}$$

Our primary interest is the characterization of the model-implied risk premium for a zero-investment position which is long currency J, and is financed in the investor's home currency I. Risk premia on generic long-short positions, which do not involve the investor's home currency I, can then be simply formed as the difference between the risk premia on long positions in the two foreign currencies. Without loss of generality, if we take the perspective of an investor in country I, the excess return from investing in a currency J is given by:

$$\begin{aligned} E_t^{\mathbb{P}}\left[\exp\left(y_{t,t+1}^{j}\right) \cdot \frac{S_{t+1}^{ji}}{S_t^{ji}} - \exp\left(y_{t,t+1}^{i}\right)\right] &= \exp\left(y_{t,t+1}^{j} + \ln E_t^{\mathbb{P}}\left[\frac{S_{t+1}^{ji}}{S_t^{ji}}\right]\right) - \exp\left(y_{t,t+1}^{i}\right) \\ &\approx \ln E_t^{\mathbb{P}}\left[\frac{S_{t+1}^{ji}}{S_t^{ji}}\right] + y_{t,t+1}^{j} - y_{t,t+1}^{i} \end{aligned} \tag{11}$$

[10]Graveline and Burnside (2013) critique economic inference about real exchange rate determination and risk-sharing based on models in which the exchange rate is represented as as a ratio of pricing kernels ("asset-market view"). In particular, they emphasize the importance of modeling trade frictions, the span of tradable assets, and differences in preferences and consumption bundles. By contrast, our interest is in building an arbitrage free model of exchange rates, rather than studying the economic model giving rise to them.

In our empirical analysis, we study the relation between the measured currency excess returns (left hand side), and the model-implied risk premia (right hand side) obtained by calibrating the model to exchange rate options. An attractive feature of focusing on simple excess returns is that both sides of the equation can be aggregated using portfolio weights to obtain model risk premia for empirical factor mimicking portfolios (e.g. for the HML_{FX} and "dollar carry" factor mimicking portfolios studied by Lustig, et al. (2011, 2013)), which is not the case for *log* excess returns.[11] We verify *ex post* that the use of the approximation results in quantitatively negligible errors.

Following Bakshi, et al. (2008), we define the currency risk premium for currency pair J/I as the difference between the log expected return on the currency adjusted for the interest rate differential:

$$\lambda_t^{ji} \equiv \ln E_t^{\mathbb{P}}\left[\frac{S_{t+1}^{ji}}{S_t^{ji}}\right] + \left(y_{t,t+1}^j - y_{t,t+1}^i\right) \tag{12}$$

Substituting in the expressions for the risk-free rates, and using the cumulant generating function to re-write the conditional expectation we obtain:

$$\begin{aligned} \lambda_t^{ji} &= \left(k_{L^g_{Z_t}}[\xi_t^i - \xi_t^j] + k_{L^g_{Z_t}}[-\xi_t^i] - k_{L^g_{Z_t}}[-\xi_t^j]\right) + \left(k_{L^i_{Y^i_t}}[1] + k_{L^i_{Y^i_t}}[-1]\right) \\ &= \left(k_{L^g_t}[\xi_t^i - \xi_t^j] + k_{L^g_t}[-\xi_t^i] - k_{L^g_t}[-\xi_t^j]\right) \cdot Z_t + \left(k_{L^i_t}[1] + k_{L^i_t}[-1]\right) \cdot Y_t^i \end{aligned} \tag{13}$$

This expression illustrates that the expected excess return on an individual currency pair is comprised of two components. The first component – controlled by the global state variable, Z_t – represents compensation for exposure to the global "slope" factor (HML_{FX}). The second component – controlled by the country-specific state variable, Y_t^i – represents the compensation demanded by an investor in country i for being short his own reference (local) currency:[12]

$$\lambda_{HML,t}^{ji} = \left(k_{L^g_t}[\xi_t^i - \xi_t^j] + k_{L^g_t}[-\xi_t^i] - k_{L^g_t}[-\xi_t^j]\right) \cdot Z_t \tag{14}$$

$$\lambda_{refFX,t}^{i} = \left(k_{L^i_t}[1] + k_{L^i_t}[-1]\right) \cdot Y_t^i \tag{15}$$

[11] An alternative measure of the currency risk premium used by Backus, et al. (2001) and Lustig, et al. (2013) is the mean *log* excess return, $E_t^{\mathbb{P}}\left[s_{t+1}^{ji} - s_t^{ji}\right] + \left(y_{t,t+1}^j - y_{t,t+1}^i\right)$. Although the log risk premium can be computed without reliance on an approximation, the disadvantage of this measure is that it cannot be aggregated linearly using portfolio weights to produce model-implied estimates of portfolio risk premia. The log risk premium can be expressed as a series expansion in terms of the cumulants of the time-changed Lévy increments as:

$$E_t^{\mathbb{P}}\left[s_{t+1}^{ji} - s_t^{ji}\right] + \left(y_{t,t+1}^j - y_{t,t+1}^i\right) = \sum_{k=2}^{\infty}(-1)^k \cdot \frac{\left(\left(\xi_t^i\right)^k - \left(\xi_t^j\right)^k\right) \cdot \kappa^k_{L^g_{Z_t}} + \kappa^k_{L^i_{Y^i_t}} - \kappa^k_{L^j_{Y^j_t}}}{k!}$$

[12] The online Technical Appendix demonstrates that in a model with parametric state-variable dynamics (e.g. as in Bakshi, et al. (2008)), there will generally be a full term-structure of option-implied currency risk premia. Lustig, et al. (2013) study the corresponding term-structure of *realized* risk premia.

To facilitate interpretation, consider the case where both the global and country-specific innovations are Gaussian ($\eta_t^g = \eta_t^j = \eta_t^i$), such that the cumulant generating function of the increments is given by, $k_{L_t}[u] = \frac{u^2}{2}$. In this case, the two risk premia are equal to:

$$\lambda_{HML,t}^{ji} = \xi_t^i \cdot \left(\xi_t^i - \xi_t^j\right) \cdot Z_t \qquad \lambda_{refFX,t}^i = Y_t^i \tag{16}$$

Suppose that – as in a typical currency carry trade – the foreign currency J has an interest rate higher than the home country I. Since loadings are inversely linked to interest rate levels, (6), the first risk premium is positive and reflects the compensation demanded by an investor in country I for exposure to a risky asset (i.e. the exchange rate), whose loading on the global factors is given by $\left(\xi_t^i - \xi_t^j\right)$. The magnitude of the risk premium is time-varying and given by the level of the global state-variable, Z_t. Consequently, the risk premium for exposure to the global risk factor, $L_{Z_t}^g$, is interpretable as exposure to the HML_{FX} ("slope") factor identified by Lustig, et al. (2011). The investor in country I also requires compensation for being short exposure to his local (reference) currency shocks, $L_{Y_t^i}^i$; the magnitude of this risk premium is controlled by the local state variable, Y_t^i. By combining the risk premia on currency pairs J/I and K/I (e.g. AUD/USD and JPY/USD) we find that the risk premium demanded by an agent in country I for a generic currency pair, J/K, not involving his home currency are: $\lambda_{HML,t}^{jk} = \xi_t^i \cdot \left(\xi_t^k - \xi_t^j\right) \cdot Z_t$ and $\lambda_{refFX,t}^i = 0$.

Finally, to gain more intuition about the two risk premium components, we explore two decompositions. The first decomposition expresses each risk premium in terms of the moments (variance, skewness, etc.) of the corresponding pricing kernel shocks ($L_{Z_t}^g$ and $L_{Y_t^i}^i$). These decomposition shares the flavor of the analysis in Backus, et al. (2011), who study disaster risk models by relating the pricing kernel entropy to the cumulants of log consumption growth.[13] The second decomposition expresses each risk premium in terms of the contributions attributable to the Gaussian ($W_{Z_t}^g$ and $W_{Y_t^i}^i$) and non-Gaussian ($X_{Z_t}^g$ and $X_{Y_t^i}^i$) components of the pricing kernel innovations.

1.3.1 HML_{FX} risk premium

In order to express the two risk premia as a function of the moments of the underlying pricing kernel innovations, we take advantage of the fact that cumulant generating functions are linked to the cumulants of the corresponding random variable through an infinite series expansion. Using this expansion, the contribution to the

[13]The entropy of the pricing kernel is defined as, $L(M_{t+1}^i) = lnE_t\left[M_t^i\right] - E_t\left[\ln M_{t+1}^i\right]$, and provides an upper bound on the log excess return of any asset in economy I. For details see Appendix B in Backus, et al. (2011).

excess return of currency pair J/I from exposure to the global factor (HML_{FX}) is given by:

$$\lambda_{HML,t}^{ji} = \sum_{n=2}^{\infty} \frac{\left(\xi_t^i - \xi_t^j\right)^n + \left(-\xi_t^i\right)^n - \left(-\xi_t^j\right)^n}{n!} \cdot \kappa_{L_t^g}^n \cdot Z_t \tag{17}$$

where $\kappa_{L_t^g}^n$ denotes the n-th cumulant of the global increment, L_t^g. If we approximate this expression with terms up to order three – capturing the premium for bearing variance and skewness – we obtain:

$$\lambda_{HML,t}^{ji} \approx \left(\xi_t^i \cdot \left(\xi_t^i - \xi_t^j\right) \cdot \mathcal{V}_t^g \cdot \left(1 - \frac{1}{2} \cdot \xi_t^j \cdot \mathcal{S}_t^g \cdot \sqrt{\mathcal{V}_t^g}\right)\right) \cdot Z_t + O(\xi^4) \tag{18}$$

where $\mathcal{V}_t^g = \kappa_{L_t^g}^2$ and $\mathcal{S}_t^g = \kappa_{L_t^g}^3 \cdot \left(\kappa_{L_t^g}^2\right)^{-\frac{3}{2}}$, are the variance and skewness of the Lévy increments of the global factor, L_t^g. In particular, given our increment normalization, we have $\mathcal{V}_t^g = 1$. The compensation for a pair's HML_{FX} risk is high whenever: (a) the loading differential is large and positive ($\xi_t^i - \xi_t^j \gg 0$); (b) the price of risk is high (i.e. volatility of global pricing kernel shocks, Z_t, is high); and, (c) the skewness of the pricing kernel innovations is negative ($\mathcal{S}_t^g < 0$). Although the fourth order terms can be shown to be positive, we find that they are empirically negligible.

Risk-based explanations of violations of uncovered interest rate parity require a link between interest rate differentials and currency risk premia. To illustrate this feature in our model consider a similar, third-order expansion of the interest rate differential:

$$\begin{aligned} y_{t,t+1}^j - y_{t,t+1}^i &= \alpha_t^j - \alpha_t^i + \left(k_{L_t^g}[-\xi_t^i] - k_{L_t^g}[-\xi_t^j]\right) \cdot Z_t - k_{L_t^j}[-1] \cdot Y_t^j + k_{L_t^i}[-1] \cdot Y_t^i \\ &\approx \alpha_t^j - \alpha_t^i + \left(\left(\xi_t^j - \xi_t^i\right) \cdot \kappa_{L_t^g}^1 + \frac{\left(\xi_t^i\right)^2 - \left(\xi_t^j\right)^2}{2} \cdot \kappa_{L_t^g}^2 + \frac{\left(\xi_t^j\right)^3 - \left(\xi_t^i\right)^3}{6} \cdot \kappa_{L_t^g}^3\right) \cdot Z_t \\ &- k_{L_t^j}[-1] \cdot Y_t^j + k_{L_t^i}[-1] \cdot Y_t^i + O(\xi^4) \end{aligned} \tag{19}$$

Whenever the global jump innovations have negative mean ($\kappa_{L_t^g}^1 < 0$) and are negatively skewed ($\kappa_{L_t^g}^3 < 0$), the interest rate differential will increase with the pair's loading differential, $\xi_t^i - \xi_t^j$. These pairs are also predicted to have higher currency excess returns within our model, consistent with empirical evidence on violations of uncovered interest parity.

An alternative methodology for decomposing the HML_{FX} component of the currency pair risk premium is to compute the contribution from Gaussian and non-Gaussian components of the global pricing kernel innovation, $L_{Z_t}^g$. This decomposition formalizes the intuition underlying the papers by Burnside, et al. (2011), Farhi, et al. (2013), and Jurek (2013), which study the contribution of tail risk premia to the excess returns of carry trade portfolios by comparing the realized returns to unhedged and crash-hedged portfolios. Specifically, since the diffusive and jump components of the Lévy increments are independent at each point in time, the cumulant generating functions appearing in the risk premium formula, (14), can be expressed as sums of the cumulant generating functions of the time-changed diffusive and jump shocks. This allows us to isolate the component of the pair's HML_{FX} risk premium due to jump shocks, and represent it as a share, ϕ^{ji}_{HML}, of the total risk premium:

$$\phi^{ji}_{HML,t} \quad = \quad \frac{k_{X^g_t}\left[\xi^i_t - \xi^j_t\right] + k_{X^g_t}\left[-\xi^i_t\right] - k_{X^g_t}\left[-\xi^j_t\right]}{k_{L^g_t}\left[\xi^i_t - \xi^j_t\right] + k_{L^g_t}\left[-\xi^i_t\right] - k_{L^g_t}\left[-\xi^j_t\right]} \cdot \eta^g_t \tag{20}$$

1.3.2 Short reference risk premium

We perform the same moment and shock-type decompositions on the premium demanded by an investor in country i for being short his local, or reference, currency. Applying the series expansion to the cumulant generating function of the country-specific shock, $L^i_{Y^i_t}$, governing the short reference risk premium, (14), we obtain:

$$\lambda^i_{refFX,t} \quad = \quad \sum_{k=1}^{\infty} \frac{1+(-1)^k}{k!} \cdot \kappa^k_{L^i_t} \cdot Y^i_t = \sum_{n=1}^{\infty} \frac{1}{(2\cdot n)!} \cdot \kappa^{2\cdot n}_{L^i_t} \cdot Y^i_t \tag{21}$$

A stark feature of the reference currency premium is that it does *not* depend on the odd moments of the country-specific shock. In other words, economic rationalizations of a reference currency premium, e.g. a short dollar premium, cannot be based on one-sided events, such as crashes or sporadic flights to quality. Finally, we compute the fraction, $\phi^i_{refFX,t}$, of the short reference risk premium driven by the non-Gaussian (jump) component of the pricing kernel innovation.

$$\phi^i_{refFX,t} \quad = \quad \frac{k_{X^i_t}[1] + k_{X^i_t}[-1]}{k_{L^i_t}[1] + k_{L^i_t}[-1]} \cdot \eta^i_t \tag{22}$$

Since our model does not ascribe a special role to any particular reference currency, we compute short reference risk premia for all currencies in our sample. Empirically, Lustig, et al. (2013) find evidence of large compensation for short dollar exposure. We revisit these results in the model calibration and results section (Section 3).

1.4 FX option pricing

Given the choice of our model parametrization, option pricing can be efficiently accomplished using standard Fourier transform option pricing methods described in Carr and Madan (1999). The central input to this pricing methodology is the cumulant generating function (or characteristic function) of the log exchange rate, $s_t^{ji} = \log S_t^{ji}$, computed under the pricing measure. In general, if interest rates are time varying over the life of the option, it is convenient to work under the risk-forward measure, $\mathbb{F}_\tau^i$, whose numeraire is the τ-period zero coupon bond in country i. In our discretized model, we assume the interest rate is fixed over the life of the option at the level given by (6), such that the risk-forward and risk-neutral, $\mathbb{Q}^i$, measures coincide. Although we write the subsequent formulas under the risk-forward measure, these can be interpreted as corresponding to risk-neutral quantities in the model. Finally, it is important to emphasize, that the pricing measure depends on the home country of the investor, since investors in different countries have distinct pricing kernels.

The risk-forward measure for an investor from country I, associated with a zero-coupon bond maturing at time $t+1$, is defined as follows:

$$\frac{d\mathbb{F}^i}{d\mathbb{P}} = \frac{M_{t+1}^i}{M_t^i} \cdot \exp\left(y_{t,t+1}^i\right) \tag{23}$$

where $\mathbb{P}$ denotes the physical (historical) measure. The virtue of working under the risk-forward measure is that we can price claims denominated in units of currency I (e.g. bonds or exchange rate options on the currency pairs J/I) as the product of their expected payoff, $x(s_{t+1}^{ji})$, under the $\mathbb{F}^i$ measure multiplied by the value of the one-period zero-coupon bond in country I. To see this, recall that any payoff satisfies the following pricing equation: $V_t^i = E_t^{\mathbb{P}}\left[\frac{M_{t+1}^i}{M_t^i} \cdot x(s_{t+1}^{ji})\right]$. Dividing both sides by the value of the zero-coupon maturing at $t+1$, and recognizing the measure change we obtain:

$$V_t^i = Z_{t,t+1}^i \cdot E_t^{\mathbb{P}}\left[\frac{M_{t+1}^i}{M_t^i} \cdot \exp\left(y_{t,t+1}^i\right) \cdot x(s_{t+1}^{ji})\right] = Z_{t,t+1}^i \cdot E_t^{\mathbb{F}^i}\left[x(s_{t+1}^{ji})\right] \tag{24}$$

where $Z_{t,t+1}^i = \exp\left(-y_{t,t+1}^i\right)$ is the value of the numeraire, zero-coupon bond. In order to apply this valuation approach to exchange rate options, we rely on Fourier pricing methods and characterize the distribution of the exchange rate at $t+1$ using its cumulant generating function under $\mathbb{F}^i$.

To derive the cumulant generating functions of the global and country-specific Lévy increments under the risk-forward measure, $\mathbb{F}^i$, we proceed directly from the definitions of the cumulant generating function and the measure change:

$$
\begin{aligned}
k^{\mathbb{F}^i}_{L^g_{Z_t}}[u] &= \ln E^{\mathbb{F}^i}_t\left[\exp\left(u \cdot L^g_{Z_t}\right)\right] = \ln E^{\mathbb{P}}_t\left[\exp\left(y^i_{t,t+1} + \left(m^i_{t+1} - m^i_t\right) + u \cdot L^g_{Z_t}\right)\right] \\
&= \left(k_{L^g_t}\left[u - \xi^i_t\right] - k_{L^g_t}\left[-\xi^i_t\right]\right) \cdot Z_t \qquad (25a) \\
k^{\mathbb{F}^i}_{L^i_{Y^i_t}}[u] &= \left(k_{L^i_t}[u-1] - k_{L^i_t}[-1]\right) \cdot Y^i_t \qquad (25b) \\
k^{\mathbb{F}^i}_{L^j_{Y^j_t}}[u] &= k_{L^j_t}[u] \cdot Y^j_t \qquad (25c)
\end{aligned}
$$

Contrasting these expressions with the corresponding values under the objective measure, $\mathbb{P}$, we see that the change of measure results in: (1) a change in the distribution of the global factor dependent on country i's loading on the global factor, ξ^i_t; (2) a change in the distribution of the local (reference) shock, $L^i_{Y^i_t}$; and, (3) leaves the distribution of foreign, country-specific shocks unchanged. To obtain the cumulant generating function for the exchange rate at time $t+1$ under the risk forward measure, $\mathbb{F}^i$, we substitute (10), into the definition of the CGF to obtain:

$$
k^{\mathbb{F}^i}_{s^{ji}_t}[u] = \left(s^{ji}_t - \alpha^j_t + \alpha^i_t\right) \cdot u + k^{\mathbb{F}_i}_{L^g_t}\left[\left(\xi^i_t - \xi^j_t\right) \cdot u\right] \cdot Z_t + k^{\mathbb{F}_i}_{L^i_t}[u] \cdot Y^i_t + k^{\mathbb{F}_i}_{L^j_t}[-u] \cdot Y^j_t \qquad (26)
$$

where each of the risk-forward CGFs can be evaluated using the formulas above. In order to substitute out the α^i_t and α^j_t terms, we take advantage of the fact that: $k^{\mathbb{F}^i}_{s^{ji}_t}[1] = s^{ji}_t + y^i_{t,t+1} - y^j_{t,t+1}$.[14] Finally, the exponential of the cumulant generating function, (26), evaluated at $i \cdot u$, $\exp\left(k^{\mathbb{F}^i}_{s^{ji}_t}[i \cdot u]\right)$, corresponds to the generalized Fourier transform of the log currency return. From there, the Fourier transform can be inverted numerically to produce prices for calls and puts, as in Carr and Madan (1999).

1.4.1 Option-implied moments

An attractive feature of our model, which we exploit in the empirical calibration, is that the *cumulants* of the distribution are *linear* functions of the state variables, $\{Z_t,\, Y^i_t,\, Y^j_t\}$. For example, the second and third

[14] To obtain this result consider pricing an investment in currency J from the perspective of an investor in country i:

$$
\begin{aligned}
1 = E_t\left[\frac{M^i_{t+1}}{M^i_t} \cdot \frac{S^{ji}_{t+1}}{S^{ji}_t} \cdot \exp\left(y^j_{t,t+1}\right)\right] &= E_t\left[\frac{M^i_{t+1}}{M^i_t} \cdot \exp\left(y^i_{t,t+1}\right) \cdot \frac{S^{ji}_{t+1}}{S^{ji}_t} \cdot \exp\left(y^j_{t,t+1} - y^i_{t,t+1}\right)\right] \\
&= E^{\mathbb{F}^i}_t\left[\frac{S^{ji}_{t+1}}{S^{ji}_t} \cdot \exp\left(y^j_{t,t+1} - y^i_{t,t+1}\right)\right] = \exp\left(k^{\mathbb{F}^i}_{s^{ji}_t}[1] + y^j_{t,t+1} - y^i_{t,t+1}\right)
\end{aligned}
$$

cumulants under the risk-forward measure, $\mathbb{F}^i$ are given by:

$$\kappa^{2,\mathbb{F}^i}_{s^{ji}_t} = \left(\xi^i_t - \xi^j_t\right)^2 \cdot \left(k_{L^g_t}\right)''\left[-\xi^i_t\right] \cdot Z_t + \left(k_{L^i_t}\right)''\left[-1\right] \cdot Y^i_t + \left(k_{L^j_t}\right)''\left[0\right] \cdot Y^j_t \tag{27a}$$

$$\kappa^{3,\mathbb{F}^i}_{s^{ji}_t} = \left(\xi^i_t - \xi^j_t\right)^3 \cdot \left(k_{L^g_t}\right)'''\left[-\xi^i_t\right] \cdot Z_t + \left(k_{L^i_t}\right)'''\left[-1\right] \cdot Y^i_t - \left(k_{L^j_t}\right)'''\left[0\right] \cdot Y^j_t \tag{27b}$$

Given a set of model parameters at time t, the state variables can be recovered using a cross-sectional regression of cumulants measured from contemporaneous exchange rate option data onto the model-implied coefficients. Although claims on cumulants are not traded, they can be readily reconstructed from option prices. Specifically, using the insights from Breeden and Litzenberger (1978) and Bakshi, et al. (2003), we compute option-implied swap rates for variance ($\hat{\mathcal{V}}_{s^{ji}_t}$) and skewness ($\hat{\mathcal{S}}_{s^{ji}_t}$). The corresponding empirical estimates of the risk-forward cumulants can then be recovered from the definitions linking moments and cumulants, $\hat{\kappa}^{2,\mathbb{F}^i}_{s^{ji}_t} = \hat{\mathcal{V}}_{s^{ji}_t}$, and $\hat{\kappa}^{3,\mathbb{F}^i}_{s^{ji}_t} = \hat{\mathcal{S}}_{s^{ji}_t} \cdot \left(\hat{\mathcal{V}}_{s^{ji}_t}\right)^{\frac{3}{2}}$. With the prices of these claims in hand, the values of the state variables can be inferred directly by means of a cross-sectional linear regression, sidestepping more complicated filtering procedures.

Finally, these expressions highlight that the moments of the option-implied exchange rate distribution – and, more generally, option prices – depend not only on the loading *differentials*, but also on the *levels* of the loadings, ξ^i_t. The dependence on the level of the loadings is crucial for model identification, since otherwise option-implied moments would only allow us to pin down loading differentials, which are insufficient for recovering risk premia, e.g. (14). For example, if the global and country-specific innovations are Gaussian, the higher order cumulants ($j \geq 3$) are equal to zero, and the second cumulant is only a function of the loading differential, precluding identification of currency risk premia. Similarly, the $\mathbb{P}$-measure moments only depend on the loading differentials.[15]

2 Data and Model Calibration

We calibrate the model of pricing kernel dynamics described in Section 1 using panel data on G10 currency exchange rate options, and the time-series of G10 exchange rates and one-month LIBOR rates. A novel feature of our approach is that it provides cross-sectional estimates of instantaneous currency risk premia, without relying on the time-series of realized (historical) returns. The option-implied estimates of risk premia are free from peso problems (Burnside, et al. (2011)), and complement the evidence on the factor structure in currency returns documented by Lustig, et al. (2011, 2013). We then use the model to decompose the HML_{FX} risk premium

[15] The cumulants under the historical measure, $\mathbb{P}$, are given by the same expressions, but with the derivatives of the consecutive cumulant generating functions evaluated at zero, rather than $\{-\xi^i, -1, 0\}$.

and the premium for short exposure to various G10 currencies into: (a) diffusive and jump shocks; and (b) contributions from various factor moments (e.g. variance, skewness, etc.).

2.1 Data

The key dataset used in the analysis includes price data on foreign exchange options covering the full cross-section of 45 G10 cross-pairs, spanning the period from January 1999 to June 2012 ($T = 3520$ days).[16] The dataset provides daily price quotes in the form of implied volatilities for European options at constant maturities and five strikes, and was obtained via J.P. Morgan DataQuery. FX option prices are quoted in terms of their Garman-Kohlhagen (1983) implied volatilities, which correspond to Black-Scholes (1973) implied volatilities adjusted for the fact that both currencies pay a continuous "dividend" given by their respective interest rates. We focus attention on constant-maturity one-month exchange rate options. For each day and currency pair, we have quotes for five options at fixed levels of option delta (10δ puts, 25δ puts, 50δ options, 25δ calls, and 10δ calls), which correspond to strikes below and above the prevailing forward price. In standard FX option nomenclature an option with a delta of δ is typically referred to as a $|100 \cdot \delta|$ option; we adopt this convention throughout. The specifics of foreign exchange option conventions are further described in Wystup (2006), Carr and Wu (2007), and Jurek (2013). In general, an option on pair J/I gives its owner the right to buy (sell) currency J at option expiration at an exchange rate corresponding to the strike price, which is expressed as the currency J price of one unit of currency I. The remaining data we use are one-month Eurocurrency (LIBOR) rates and daily exchange rates for the nine G10 currencies versus the U.S. dollar obtained from Reuters via Datastream.

2.2 Calibration: Intuition

The final goal of our model calibration procedure is to produce a time series of instantaneous currency risk premia. Since these risk premia are expressed in terms of the cumulant generating function of the pricing kernel innovations, (13), the calibration procedure can be understood in terms of pinning down the cumulants in the series expansion of the respective CGFs. Even more precisely, the identification of currency risk premia only requires knowledge of cumulants of order two or higher, as can be seen in (17) and (21). Our calibration procedure recovers these cumulants from information in the option-implied exchange rate distribution.

To present the intuition behind our empirical identification strategy it is useful to begin by specializing to a world in which all pricing kernel innovations are Gaussian ($\eta_t^j = 0$), such that there are no CGMY jump

[16]The G10 currency set is comprised of the Australian dollar (AUD), Canadian dollar (CAD), Swiss franc (CHF), Euro (EUR), U.K. pound (GBP), Japanese yen (JPY), Norwegian kronor (NOK), New Zealand dollar (NZD), Swedish krone (SEK), and the U.S. dollar (USD). There are a total of 45 possible cross-pairs.

parameters to estimate. In this case, exchange rate distributions are Gaussian, and the only cumulant of order two or higher – either in the distribution of the pricing kernel innovations, or in the exchange rate distribution – which is non-zero, is the second cumulant. In particular, the second cumulant (or, variance) of the option-implied distribution is immediately revealed by the implied volatility of any one of the five, quoted options. The quotes on the remaining four options are redundant. The identification of the model parameters proceeds on the basis of equation ($27a$), which links the variance of the option-implied distribution under the risk-forward measure with the cumulants of the pricing kernel innovations. Given we have data on all exchange rate cross-pairs, we have a total of $\frac{N\cdot(N-1)}{2}$ observation equations for the option-implied variances. The model parameters to be pinned down on each day are: (1) the N country loadings, ξ_t^i; (2) the variance of the global innovation, Z_t; and, (3) the variances of the N country-specific innovations, Y_t^i. Since we can normalize one of the loadings to one, this leaves us with $2\cdot N$ parameters. In particular, the G10 currency set has ten currencies, yielding twenty parameters and forty five observation equations. Consequently, the parameters of the Gaussian model are fully-identifiable using our cross-sectional information on *each* day individually.

Our ability to identify the Gaussian model owes to the availability of cross-rate options. By contrast, consider the identification strategy proposed by Farhi, et al. (2013), which *only* uses information on X/USD options (i.e. options on the exchange rate of foreign currencies against the U.S. dollar). This approach yields a total of nine $N-1$ observation equations for the twenty model parameters, such that their model is not fully identified even in the Gaussian setting. This identification problem does not go away when distributions are non-Gaussian, but rather highlights that the variance of the pricing kernels – which plays a first-order role in determining the risk premium contributions of the Gaussian and non-Gaussian shocks – remains unidentified in their setup. Importantly, the non-identification problem can be resolved by specifying time-consistent dynamics for the state-variables and simultaneously calibrating to X/USD options with multiple tenors, e.g. as in Bakshi, et al. (2008).

How does the identification intuition change if we go to the more general setting with arbitrary option-implied distributions? To focus attention on the conceptual issues of model identification, rather than the empirical implementation, suppose that one observes options at a continuum of strike values. Consequently, one could use the methodology of Breeden and Litzenberger (1976) to extract the option-implied distributions for any exchange rate for which one has option data. In this case, the goal at each point in time is to identify the $N+1$ factor *distributions* (global + N country-specific innovations) on the basis of the exchange rate distributions extracted from the option data. This is clearly not possible using information solely on the $N-1$ X/USD exchange rate options. Again, our ability to identify the model relies on the availability of cross-rate options, which yield a total of $\frac{N\cdot(N-1)}{2}$ option-implied distributions to match. From here, the identification exercise relies on the specific parametric specification adopted for the time-changed Lévy increments.

We are faced with matching all the moments of the option-implied exchange rate distributions, by varying the global factor loadings, state-variables, and parameters describing the pricing kernel innovations (or, equivalently, the cumulants of the pricing kernel innovations). Again, the goal is to identify all pricing kernel cumulants of order two or higher, since only these determine currency risk premia. Observing cross-rate options yields a total of $\frac{N\cdot(N-1)}{2}$ independent observation equations per exchange rate cumulant used in model identification. Contrast this with the number of parameters that would need to be estimated for our pricing kernel parametrization based on time-changed Lévy increments. We continue to have $2 \cdot N$ country loadings and state-variables, but we now also have the CGMY parameters to estimate for the global and country-specific jump components. Our parametrization only allows for one-sided global jumps, such that we have three global jump parameters, $\{C, G, Y\}_t^g$, and N quartets for the country-specific components, $\{C, G, M, Y\}_t^i$. This would result in a total of $6 \cdot N + 3$ parameters to be identified on each day. With a full set of cross-rate currency options for ten countries, the model could – at least in principle – be identified using data on only two option-implied moments (e.g. variance and skewness). Of course, practical implementation faces both the numerical challenges of high dimensional optimization, and the lack of availability of a dense set of options. The next section describes the simplifications made to facilitate estimation in light of these issues.

2.3 Calibration: Implementation

To facilitate model calibration we make two simplifying assumptions. First, we assume that on any given day, t, the country-specific jump parameters, $\{G, M, Y\}_t^i$, are identical across all countries, and only allow the jump share, η_t^i (equivalently, C_t^i) to be country-specific. This reduces the number of country-specific jump parameters from $4 \cdot N$ to $N + 3$. Second, we parameterize the global factor loadings to be a function of each country's interest rate differential relative to the U.S. one-month rate, $\xi_t^i = \xi^i - \Psi_t \cdot \left(r_{t,t+1}^i - r_{t,t+1}^{US}\right)$, as in Lustig, et al. (2011). This allows us to parsimoniously explore the importance of co-variation of loadings and interest rates for the pricing of exchange rate options, and its effect on the determination of currency risk premia in the time series and cross section.

Our model calibration proceeds in two stages. The first stage assumes that global factor loadings and the CGMY jump parameters are time-invariant. It then uses repeated regressions of option-implied cumulants onto the model-implied loadings, $(27a) - (27b)$, to identify the time series of state variables and estimates of the global factor loadings, ξ^i, by minimizing the sum of squared regression errors over the full sample. The outputs of the first stage are then used to initialize the second stage of the estimation, which steps through the dataset day-by-day and perturbs the first stage jump parameters, state variables, and the interest rate sensitivity of global factor loadings (Ψ_t) to minimize option pricing errors on each day individually.[17] Effectively, the two stage calibration allows for the jump parameters, state variables, and global factor loadings to change over time. The variation of the jump parameters and state variables is left unconstrained, while variation in the factor loadings, ξ_t^i, is related to interest rate differentials and is controlled by the estimate of Ψ_t.

To identify the first-stage parameters, we match the empirical cumulants computed from the prices of variance and skewness swaps. The procedure for computing moment swaps is based on Bakshi, et al. (2003), and requires interpolating and extrapolating the observed implied volatility functions. Following Jurek (2013), we interpolate the implied volatility function within the observed range using the vanna-volga method of Castagna and Mercurio (2007), and append flat tails beyond the observed strike rage (below the 10δ puts and above the 10δ calls). Because the procedure requires data augmentation, we only rely on it to provide starting values for the second-stage of the calibration. The first stage is essentially an application of non-linear least squares to the full sample of data, where we are minimizing:[18]

$$\min_{Z_t, \{Y_t^i\}_{i=1}^{10}, \{\xi^i\}_{i=1}^{9}, \{\eta^i\}_{i=1}^{10}, G, M, Y, \eta^g, G^g, Y^g} \sum_{t=1}^{T} \sum_{i=1}^{N} \left(\left(\sqrt{m_{i,t}^2} - \sqrt{\hat{m}_{i,t}^2} \right)^2 + \omega \cdot \left(m_{i,t}^3 - \hat{m}_{i,t}^3 \right)^2 \right) \tag{28}$$

where $\hat{m}_{i,t}^k$ ($m_{i,t}^k$) is the empirical (model) value of the option-implied moment k for cross-rate i at time t. The first summation is over the dates in our sample, while the second – is over the number of pairs. For each observation we match the volatility and skewness of the exchange rate, setting the weight ω equal to 0.01, to account for the fact that the average option-implied skewness is an order of magnitude larger than the average implied volatility. We sample data every five days ($T = 744$ weeks; 1999:1-2012:6), and use two different sets of currency cross-pairs. The first set (HLX + X/USD) includes 24 currency pairs and is comprised of: (a) all X/USD currency pairs (9 pairs); and, (2) cross-pairs formed on the basis of currencies which had the highest or lowest interest rates in the G10 set at any point in our sample (15 pairs).[19] The virtue of this subset is that it captures the pricing of the liquid X/USD options, and options on "typical" carry trade currency pairs. The second set uses all 45 G10 cross-rates.

The second stage refines the first stage estimates by minimizing the option pricing errors for the specified

[17]The loading intercepts, ξ^i, are left unchanged in the second stage, which enables us to nest calibrations with and without time-varying global factor loadings by imposing, $\Psi_t = 0$, in the second stage.

[18]Longstaff and Rajan (2008) rely on this methodology to calibrate their collateralized debt obligation pricing model, and provide additional references of the use of this model calibration approach. As they do, we utilize a direct search algorithm, which does not rely on the gradient or Hessian of the objective function.

[19]The set of unique currencies which – at some point in the sample – had the highest one-month LIBOR rates includes: AUD, GBP, NOK, and NZD. The corresponding set of unique currencies with the lowest interest rates includes: CHF and JPY. This yields a total of 15 cross-pairs for use in estimation.

option set by comparing the observed and model option-implied volatility. Specifically, we step through each day in the sample, and perturb the first stage estimates to best match the cross-section of the option-implied volatilities on that day. Formally, the second stage of the calibration minimizes:

$$\min_{Z_t, \{Y_t^i\}_{i=1}^{10}, \Psi_t, \{\eta_t^i\}_{i=1}^{10}, G_t, M_t, Y_t, \eta_t^g, G_t^g, Y_t^g} \sum_{i=1}^{N} \sum_{j=1}^{5} \left(\sigma_{j,t}^i - \hat{\sigma}_{j,t}^i\right)^2 \tag{29}$$

where $\hat{\sigma}_{j,\,t}^i$ ($\sigma_{j,\,t}^i$) is the empirical (model) value of the option-implied volatility for currency pair i at strike j at time t. The first summation is over the set of currency pairs used in the estimation, and the second summation is over the five strikes quoted for each currency pair. We repeat this procedure daily ($T = 3520$ days; 1999:1-2012:6). The net effect of the second stage is to allow for: (a) adjustments to the level of the state variables on that day; (b) time-variation in the global and country-specific jump parameters; and, (c) variation in the global factor loading sensitivity, Ψ_t. Depending on the option set under consideration, this procedure involves pricing 120 (225) individual options on each day in the sample using the Fourier inversion method described in Section 1. Crucially, the second stage of the calibration does *not* rely on data augmentation, and uses exclusively the quoted option prices.[20]

In our baseline specification, we focus on pricing the cross section of high-low cross-pairs and X/USD options (24 pairs). We allow the global factor loadings to be time-varying and depend on the prevailing interest rate differentials ($\Psi_t \neq 0$), and place no constraints on the type of jump process (Y_t) used to explain the option prices. We examine the robustness of our model calibration results with respect to: (1) the choice of the option set used to calibrate the model (all G10 cross pairs vs. high-low cross pairs + X/USD options); and, (2) the choice of fixed ($\Psi_t = 0$) vs. time-varying global factor loadings.

2.4 Model parameters

The calibrated model parameters are reported in Table I. Panel A reports the values of global factor loadings, ξ^i, along with estimates of their standard errors (in parentheses), obtained from the first stage of the calibration. Throughout our implementation the U.S. global factor loading is normalized to one. Panel B reports summary statistics for the time series of model parameters obtained in the second stage of the calibration, which minimizes option pricing errors day-by-day. We report the time series mean and volatility of the sensitivity of global factor loadings to interest rate differentials (Ψ_t), and the parameters governing the distribution of the global pricing

[20]Since inverting model option prices to obtain estimates of option-implied volatility is computationally costly, we implement an approximate version of this minimization criterion. Specifically, we minimize the (squared) difference between the model and market option price, scaled by the value of Black-Scholes vega at that strike.

kernel innovation. The latter include the share of variance attributable to the non-Gaussian component of the innovation, η_t^g, and estimates of the parameters of the global CGMY jump component, G_t (dampening coefficient) and Y_t (power coefficient). To facilitate interpretation we also report the mean instantaneous skewness and kurtosis of the time-changed global innovation, $L_{Z_t}^g$, induced by variation in the jump parameters and the global state variable, Z_t. Panel C reports the quality of the model's fit to the data in the form of the root mean squared fitting error for implied volatilities by strike, and in total, measured in volatility points.

We consider a total of four specifications. Specifications I and II are based on the set of high/low interest rate cross pairs, combined with the nine X/USD pairs (24 pairs); Specifications III and IV are calibrated using data on all G10 cross pairs (45 pairs). Specifications II and IV are constrained versions of Specifications I and III, respectively, which impose the restriction that the global factor loadings remain constant in the time series ($\Psi_t = 0$). Based on a combination of parsimony and quality of fit to the option data, Specification I is our preferred model formulation.

2.4.1 Global factor loadings, ξ_t^i

The global loading parameters (Table I; Panel A) range from a low of 0.79 (AUD) to a high of 1.27 (JPY). These values are broadly consistent with previously reported estimates from models allowing for asymmetries in exposures to global risks. For example, using three years of data on a single currency triangle (JPY/USD, GBP/USD, GBP/JPY), Bakshi, et al. (2008), find loadings of 1.53 (JPY) and 1.01 (GBP). Our estimate of the loading for the British Pound is slightly lower and ranges from 0.93 to 0.98. These estimates can also be contrasted with the parameters used in the simulation framework in Lustig, at al. (2011), which was calibrated to match the historical risk and return properties of currency returns, risk-free rates, and inflation. Under their "restricted" model with constant loadings, which is most comparable to our first stage output, the calibration requires loadings to range from $\sqrt{\frac{\delta}{\delta^*}} = 0.81$ (high interest rate currencies) to $\sqrt{\frac{\bar{\delta}}{\delta^*}} = 1.16$ (low interest rate currencies), when measured as a fraction of the loading of the reference currency (δ^*).[21] Our estimates of global factor loadings based on option data fall into a comparable, but slightly broader range.

The second stage of the calibration relaxes the assumption that loadings are constant (Specifications I and III), allowing them to depend on the contemporaneous interest rate differential, $\xi_t^i = \xi^i - \Psi_t \cdot \left(r_{t,t+1}^i - r_{t,t+1}^{US}\right)$. We find strong evidence that Ψ_t is positive on average, indicating that currencies with higher prevailing interest rates tend to have lower global factor loadings. The estimate of the mean value of Ψ_t is 0.42 (t-stat: 13.09) for

[21] Their "restricted" specification fixes the sensitivity of the global loading to the country-specific state variable, Y_t^i, at zero by setting $\kappa = 0$. Our estimates of the global loading coefficient, ξ^i, map into $\sqrt{\delta^i}$ under their notation, with the added normalization that we fix the loading of the U.S. dollar at unity.

the option set including the high-low cross rates and X/USD currency pairs, and increases to 2.37 (t-stat: 62.17) as the option set is expanded to include the full cross-section of 45 G10 cross rates (Table I, Panel B). Aside from this conditional variation in loadings, we observe a strong negative unconditional cross-sectional relationship between the ξ^i and the mean interest rate differential relative to the U.S. (Table I, Panel A; Figure 2, top left panel). Both of these dimensions of variation are consistent with the preferred ("unrestricted") specification in Lustig, et al. (2011). We plot the time series of global factor loadings from Specification I for two high interest rate currencies (AUD, NOK) and two low interest rate currencies (CHF, JPY) in the top panel of Figure 1.

Figure 2 summarizes the unconditional, cross-sectional relation between the global factor loading differentials and various data quantities. The top two panels consider the relation between the loading differentials and interest rate differentials (left panel), as well as, the realized currency excess return (right panel). Recall that neither of these two quantities was used as a target in the model calibration, which is designed to match FX option prices. The top left panel indicates a strong negative relation between the mean loading differential and the mean one-month interest rate differential. We find that this unconditional relation is driven primarily by the negative relation between the estimates of ξ^i and the interest rate differentials, rather than the estimate of Ψ_t. This is consistent with the link between loadings and interest rate differentials implied by our model of pricing kernel dynamics, (19). Similarly, we observe a negative relation between the global factor loadings recovered from exchange rate options and the mean realized currency pair excess returns. Likewise, this is consistent with the model-predicted relation between currency risk premia and loading differentials, (13), suggesting that exchange rate options carry information that is useful for explaining currency excess returns. The two bottom panels examine the link between the calibrated global factor loadings and option-implied exchange rate moments (volatility and skewness), both of which are used in the model calibration. The model roughly predicts that: (a) option-implied volatility is a function of the absolute loading differential $(27a)$; and, (b) option-implied skewness is an increasing function of the loading differential $(27b)$, whenever the skewness of the global factor, $L^g_{Z_t}$ is negative. We find empirical support for both of these cross-sectional relations.

2.4.2 Global factor innovations, $L^g_{Z_t}$

Panel B of Table I reports the distributional properties of the global factor innovations, $L^g_{Z_t}$ – the fraction of variance attributable to non-Gaussian shocks (η^g_t), and the parameters of the one-sided CGMY jump distribution (G^g_t, Y^g_t). All four specifications indicate that jump risks play a dominant role in describing global innovations, accounting for at least 50% of the total shock variance. In each of the specifications, the skewness of the global Lévy innovation, L^g, is stochastic and depends on the level of the global state variable, Z_t, and the jump parameters. The mean skewness of the global factor innovation is slightly below -1. These values are in line with the realized return skewness of currency carry trade portfolios reported in Jurek (2013), which represent empirical factor mimicking portfolios for the HML_{FX} (global) factor in G10 currencies.

The bottom panel of Figure 1 plots the time series of the state variable, Z_t, which controls the variance of the global factor innovations and the HML_{FX} factor risk premium, obtained from Specification I. The figure shows that the dynamics of the global state variable share the features of the first principal component extracted from the corresponding panel of exchange rate variance swap rates. In a supplementary table (Table A.I) we report the parameters of square root diffusions fitted to the time series of the global and country-specific state variables, $\{Z_t, Y_t^i\}$. We find that the global state variable is relatively persistent, exhibiting a half-life of 76 business days. This feature is common to all of four specifications considered in Table I, and confirms that treating the level of Z_t as fixed over the course of a month – as is implicit in our discrete time model – is a benign simplification.

2.5 FX option pricing

Panel C of Table I reports the results of the second stage option pricing. Specifically, the table reports the root mean squared option pricing errors (in volatility points) by strike, and across all all strikes jointly. For each exchange rate pair included in the calibration set, we price options on each day (1999:1-2012:6; $T = 3520$ days) for each of the five quoted strikes (10δ put to 10δ call). Our model RMSE stands at roughly 1.1 volatility points for the combined HLX (high-low cross pairs) and X/USD option set; and at 1.6 volatility points – for the full panel of 45 G10 cross-rate options. These errors are comparable to typical bid-ask spreads in FX option markets (Jurek (2013)), suggesting that the model is doing a reasonable job of matching the data. Interestingly, although the option data confirm the presence of time-varying loadings, their effect on prices appears to be economically small. Figure 3 illustrates the quality of the fit under Specification I by plotting the mean actual option-implied volatilities (blue) and their fitted counterparts (dashed red) for cross pairs formed by combining two high interest rate currencies (AUD, NOK) with two low interest rate currencies (CHF, JPY), as well as, pairs involving each of these currencies against the U.S. dollar. We additionally plot a typical bid-ask spread (dashed blue lines) equal to 0.1 times the mean quoted implied volatility at each strike.

Since our goal is to produce option-implied estimates of currency risk premia, we want the calibrated model to match the underlying option prices as well as possible. For this reason, we focus our subsequent analysis on the output based on Specification I, which produces the lowest pricing errors. However, to guard against the risk of overfitting, we have also verified the robustness of our results to using the full G10 cross-rate set. We find that our conclusions regarding currency risk premia are qualitatively unaffected by the choice of the option set.

3 Option-Implied Currency Risk Premia

The calibrated option pricing model allows us to: (a) decompose currency risk premia across Gaussian and non-Gaussian pricing kernel innovations; (b) evaluate the mechanism through which jumps contribute to equilibrium currency risk premia, by decomposing risk premia across the moments of the pricing kernel innovations (variance, skewness, etc.); and, (c) compare the option-implied risk premia with the historical returns to individual currency pairs and empirical factor mimicking portfolios.

The direct output of the calibrated option model are the time series of currency risk premia for individual currency pairs, (13). These can be compared to realized currency pair returns, or aggregated using portfolio weights for comparison with realized currency portfolio returns. Throughout this analysis we take the perspective of a U.S. dollar investor, to match the convention in previous papers. Note that the identity of the reference currency plays an important role in the determination of risk premia. First, it determines the level of the global (HML_{FX}) factor risk premium, since pricing kernels in different counties have different exposures to the global shocks. Second, investors demand a positive premium for short exposure to their country-specific pricing kernel shocks, giving rise to a local currency risk premium.

3.1 Model risk premium decomposition

To compute the HML_{FX} factor risk premium within the model, we construct a hypothetical factor mimicking portfolio on the basis of the calibrated time series of the global factor loadings, ξ^i_t. At each point in time, we sort the G10 currencies – excluding the U.S. dollar – into long and short portfolios on the basis of their prevailing loadings, and weight the currencies within each portfolio by the absolute deviation of their loading from the average loadings of currencies with ranks five and six. The resulting portfolio is dollar-neutral and loads exclusively on the global factor innovations. This construction mimics the construction of spread-weighted, dollar-neutral currency carry trade portfolios in Jurek (2013). However, there the portfolio weights are computed on the basis of the prevailing one-month LIBOR rates, which act as an empirical proxy for global factor loadings. The monthly return correlation for the portfolios sorted on option-implied loadings, ξ^i_t, and one-month interest rates, $y^i_{t,t+1}$, is equal to 0.92. During the time period in which they overlap (Jan. 1999 - Mar. 2010), the two portfolios are also highly correlated with the developed-market HML_{FX} factor return series reported by Lustig, et al. (2011) with correlations of 0.90 and 0.95, respectively, explaining 80-90% of the factor's time series variation. The mean return of the HML_{FX} factor in developed (all) currencies during this period is 4.86% (7.26%) *per annum*, and their monthly return correlation is 0.61. The mean return on the G10 currency portfolio we form using loadings (interest rates) is 4.46% (5.49%) *per annum*, respectively.

Using the time series of portfolio weights, we aggregate the pair level global factor risk premia, (14), to produce an option-implied HML_{FX} risk premium. Table II reports the mean level of this risk premium, the global factor loadings of the long and short portfolios and two risk premium decompositions. To demonstrate the robustness of our results to the specifics of the calibration exercise, we repeat this procedure for all four specifications considered in Table I. A supplementary table confirms that the results are robust to: forming portfolios on the basis of the prevailing one-month interest rates, rather than global factor loadings (Table A.II), and equal-weighting currencies within the long and short portfolios (Table A.III).

We defer the discussion of the *level* of the model HML_{FX} risk premium until the next section, where we compare the model-implied risk premia with the returns realized by empirical factor mimicking portfolios formed on the basis of interest rate differentials. Abstracting from the level, we instead focus on: (a) the *share* of the risk premium due to the Gaussian and non-Gaussian components of the global factor innovation, $L^g_{Z_t}$; and, (b) the *share* of the risk premium due to various moments of the pricing kernel innovation. Across all four specifications, we find that on average the non-Gaussian (jump) component of the global factor innovation accounts for approximately 55% of the HML_{FX} risk premium, with the balance accruing to the Gaussian shocks. However, we observe considerable variation in these shares, with the maximum contribution of each shock exceeding 90% in the time series. These values roughly match the mean level of the global shock variance due to the non-Gaussian innovation, η^g_t, reported in Table I.

The evidence uniformly points to an important role for non-Gaussian innovations as determinants of the risks and returns to the HML_{FX} factor mimicking portfolios. To examine this in detail we decompose the model HML_{FX} risk premium across the moments of the global factor innovation in the pricing kernel (variance, skewness, and higher-order moments). Again we report *shares* rather than levels. On average, we find that 85% of the HML_{FX} risk premium is due to the variance of the innovation, roughly 10% is due to the skewness of the innovation, and the balance is due to higher order moments. We find that this decomposition is robust across all four specifications. In the time series, the share of the risk premium due to variance ranges from 28-100%, while skewness accounts from 0% to 26%, across the four specifications. Consequently, the primary mechanism through which non-Gaussian innovations contribute to the global factor risk premium is by contributing variance, rather than higher-order moments. For comparison, the contribution of skewness and higher-order moments to *equity* risk premia is on average 35% in models calibrated to historical consumption disasters (Barro (2006), Barro and Ursua (2008), Barro, et al. (2013)), but only 2% in a model calibrated to match the pricing of equity index options (Backus, et al. (2011)).[22]

[22]A supplementary online Technical Appendix explores the mechanism through which disasters contribute to the determination of equity risk premia. The expressions for equity risk premia are obtained from Martin (2013), and the parameters values describing the

Table III implements the same decompositions for the short reference risk premium. The short reference risk premium can be computed on the basis of the calibrated model parameters directly from (15), and does not require the formation of factor mimicking portfolios. We report the results of decompositions for investors based in each of the G10 countries in Table III, using the model parameters obtained under our preferred specification (Specification I). The results based on the other specifications are qualitatively identical, and are omitted to conserve space. The mean levels of the short reference risk premium range from 0.48% (JPY) to 1.33% (SEK) *per annum*, and are largely determined by the mean level of volatility in the country-specific factors, Y_t^i (Table A.I). There is also considerable time-series variation in the short reference risk premium within each currency. For example, the premium U.S. investors demand for being short the U.S. dollar ranges from a low of 0.20% *per annum*, to a high of 3.17% *per annum*.

Panel B of Table III reports the share of the short reference risk premium for each currency due to Gaussian and non-Gaussian risks, and a decomposition across the moments of the country-specific pricing kernel innovations, $L^i_{Y^i_t}$. The shock-type decomposition reveals significant cross-sectional and time-series variation in the share of the risk-premium due to non-Gaussian (jump) risks. The mean share attributable to non-Gaussian risks ranges from 20-60% across the G10 currency set, but reaches over 80% at some point in the sample for nine out of ten currencies. This highlights the importance of allowing for a non-Gaussian component in modeling the country-specific pricing kernel innovations. This feature is common to our model and Bakshi, et al. (2008), but is absent from Lustig, et al. (2011) and Farhi, et al. (2013), who force country-specific shocks to be conditionally Gaussian. Finally, we decompose the short reference currency risk premium across the moments of the country-specific pricing kernel shocks. We find that the variance of the shocks accounts for at least 98% of the mean risk premium for each country, with higher order moments playing a very minor share. The share attributable to higher moments never exceeds 15% at any point in time for the G10 currencies. Specifically, recall that the contribution of skewness to the short reference risk premium is zero, since the short reference risk premium is determined solely by the even moments of the pricing kernel innovation, (21). We conclude that while non-Gaussian country-specific innovations play an important role in matching the features of exchange rate options across strikes and in the time series, they exert their influence on short reference risk premia through their contribution to the variance of the pricing kernel, rather than its higher (even) moments.

consumption growth process in macro models and implied by equity index options are from Backus, et al. (2011).

3.2 Comparison with empirical factor mimicking portfolios

The previous section focused on the *share* of currency risk premia due to different shock types and moments. In this section, we compute the model-implied *level* of risk premia corresponding to empirical factor mimicking portfolios. Specifically, we use the actual portfolio composition used in the construction of historical returns to currency carry and short dollar portfolios, and use them to aggregate the model-implied, pair-level risk premia. This allows us to contrast the realized (historical) performance of these trades with our model required rate of return.

The basic unit of observation in the empirical analysis of currency risk premia is a currency excess return, capturing the net return to a zero-investment portfolio which borrows one unit of currency I, at interest rate $y^i_{t,t+1}$, to lend at short-term rate $y^j_{t,t+1}$ in market J. The short-term interest rates (yields) are expressed in annualized terms. At time t, the one unit of borrowed currency I buys $\frac{1}{S^{ji}_t}$ units of currency J, such that S^{ji}_t has the interpretation of the currency I price of one unit of currency J. Finally, at time $t+1$ the trade is unwound and the proceeds converted back to currency I, generating an excess return of:

$$r^{ji}_{t+1} = \exp\left(y^j_{t,t+1} \cdot \tau\right) \cdot \frac{S^{ji}_{t+1}}{S^{ji}_t} - \exp\left(y^i_{t,t+1} \cdot \tau\right) \tag{30}$$

Unless otherwise noted, we take the perspective of a U.S. dollar investor, and report U.S. dollar denominated returns. Consequently, if I is not the investor's home currency, the above return needs to be converted to the home currency, H, by multiplying it by $\frac{S^{ih}_{t+1}}{S^{ih}_t}$. Pair-level currency excess returns are then aggregated using portfolio weights to produce the portfolio excess returns.

Following Lustig, et al. (2011, 2013), we report results for two types of historical factor mimicking portfolios: conditional and unconditional. In *unconditional* factor mimicking portfolios currencies are sorted on the basis of backward-looking averages of interest rates, or no historical information at all (i.e. the portfolio composition is fixed). By contrast, the composition of the *conditional* factor mimicking portfolio is determined on the basis of the prevailing interest rates. We repeat the analysis for empirical factor mimicking portfolios which: (a) exploit violations of UIP by borrowing funds in low interest rate currencies to purchase high interest rate currencies (i.e. carry trades); and, (b) short the U.S. dollar. We focus attention on G10 currencies and the period from January 1999 to June 2012, matching the span of our exchange rate option data. We compute a time series of buy-and-hold returns, rebalancing positions at month ends ($N = 162$ months). The top two panels of Figure 4 plot the cumulative returns to the four factor mimicking portfolios. The bottom panels plot two estimates of the instantaneous portfolio return volatility. The first measure is based on in-sample estimates of an EGARCH(1,

1) model. The second is based on an option-implied variance-covariance matrix computed on the basis of contemporaneous observations of variance swap rates for the full panel of G10 cross-rates. Both models confirm the presence of stochastic volatility in all four factor mimicking portfolios, consistent with the structure of our model.

3.2.1 Slope (carry) factor

To construct conditional and unconditional factor mimicking portfolios for the G10 HML_{FX} (slope) factor we use a dollar-neutral, spread-weighted portfolio of currency carry trades. Each of the portfolios is long (short) the currencies with the highest (lowest) one-month LIBOR interest rates. We compute spread-weighted portfolio returns by assigning portfolio weights on the basis of the absolute distance of country i's interest rate from the average of the interest rates in countries with ranks five and six. The spread-weighting procedure is similar in spirit to forming portfolios of currencies based on interest rate sorts, and computing a long-short return between the extremal portfolios, but is more pragmatic given the small cross-section. Importantly, we impose the restriction that carry trade portfolios be dollar-neutral by constraining the sum of the nine remaining country weights to equal zero. This allows us to cleanly separate the pricing of the short dollar and the HML_{FX} factors.

Panel A of Table IV reports summary statistics for the historical returns of the conditional and unconditional HML_{FX} factor mimicking portfolios formed using developed market (G10) currencies. Consistent with Lustig, et al. (2011) we find that the conditional factor mimicking portfolio delivers a mean currency excess return of 4.96% *per annum*, relative to a mean excess return of 3.32% *per annum* for the unconditional portfolio.[23] Roughly half of the 1.63% (t-stat: 1.99) return differential is accounted for by the difference in the carry (interest rate differential) earned by the two portfolios. Specifically, by sorting and weighting currencies using the prevailing – rather than historical – interest rates, the conditional factor mimicking portfolio earns 0.88% (t-stat: 27.95) more *per annum* in carry than the unconditional portfolio. In fact, among G10 currencies the carry accounts for almost the entirety of the currency excess return, indicating that raw exchange rates roughly follow a random walk. Panel A further illustrates that excess returns to carry trades are generally non-normal and exhibit high, negative skewness. Even after adjusting for the effects of stochastic volatility, the standardized monthly return innovations are non-Gaussian, consistent with the presence of jumps. The Jarque-Bera test rejects the null of Gaussianity both for the returns and standardized (log) returns based on the in-sample output of a EGARCH(1, 1) model.[24]

[23] Using an expanded set of 35 currencies (15 developed, 20 emerging), Lustig, et al. (2011) find that the conditional HML_{FX} portfolio earns 4.8% *per annum*. After transaction costs the unconditional carry premium accounts for roughly half of the total conditional carry trade premium. See also Hassan and Mano (2013).

[24] Chernov, Graveline, and Zviadadze (2012) use a combination of historical returns and option data to estimates stochastic volatility

We compute the model-implied risk premia for the empirical factor replicating portfolio on each rebalancing date (monthly) using Specification I. We find a mean model-implied risk premium for the conditional (unconditional) portfolio of 3.55% (3.44%) *per annum*. We further decompose the required portfolio risk premium into compensation for exposure to the global risk factor (HML_{FX}), and for compensation for short U.S. dollar exposure. Since the empirical factor replicating portfolios were constructed to be dollar-neutral the value of this risk premium component is zero at all points in time. We find that: (a) the model-implied risk premium for the conditional and unconditional portfolios is statistically indistinguishable; and, (b) the empirical portfolios realized historical mean returns, which were indistinguishable from the corresponding model-implied premia. The first result indicates that, although we find evidence of time-variation with the global factor loadings (Table I), it is insufficiently correlated with the variation in the global price of risk to drive a meaningful wedge between the risk premia on the conditional and unconditional factor portfolios. We discuss the impact of time-varying global factor loadings on risk premia in more detail in Section 4. The second result indicates that foreign exchange spot and option markets consistently price currency risks, and suggests that estimates of currency risk premia based on historical returns are not significantly upward biased due to peso problems, as argued by Burnside, et al. (2011).

Figure 5 decomposes the model-implied HML_{FX} risk premium for the conditional empirical factor replicating portfolio based on Specification I of the calibration. The top panel plots the time series of the total risk premium and the contribution from exposure to the global risk factor; the middle panel – decomposes the risk premium into contributions from Gaussian and non-Gaussian (CGMY) global factor innovations; and, the bottom panel – decomposes the risk premium across moments of the global factor innovation. At the portfolio level, we find that the option-implied risk premium for the empirical HML_{FX} factor mimicking portfolio is dominated by compensation for jump risks, which account for 58% of the total portfolio risk premium. To evaluate the channel through which non-Gaussian innovations contribute to the determination of currency risk premia, we exploit the series expansion of the cumulant generating function for the global innovation to decompose the risk premia across the moments of the underlying shocks, (17) and (21). We decompose the model-implied currency risk premia into contributions from variance, skewness, and higher moments, and plot their shares in the bottom panel of Figure 5.

In contrast to the typical "rare disasters" intuition, we find that global jumps exert their effect on currency risk premia through their contribution to the total variance of the global factor innovation, rather than its skewness (or higher moments). For example, roughly 85% of the 3.55% annualized model-implied risk premium for the

---

jump-diffusion models characterizing currency returns. Brunnermeier, Nagel and Pedersen (2008) argue that realized skewness is related to rapid unwinds of carry trade positions, precipitated by shocks to funding liquidity. Plantin and Shin (2008) provide a game-theoretic motivation of how strategic complementarities, which lead to crowding in carry trades, can generate currency crashes.

conditional HML_{FX} replicating portfolio is accounted for by compensation for variance. The skewness and higher-moments of the global factor innovation, $L^g_{Z_t}$, contribute 9.5% and 5.3% of the annualized risk premium (Figure 5, bottom panel). These findings are in line with the empirical analysis of returns to crash-hedged currency carry trades, which indicates that tail risks account for 8-10% of the historical excess returns earned by dollar-neutral, spread-weighted portfolios of G10 currency carry trades (see Panel B of Table III in Jurek (2013)).

3.2.2 Short U.S. dollar factor

Lustig, et al. (2013) report evidence that investors earn high excess returns for being short the U.S. dollar in bad times, when aggregate risk premia are high. We construct two versions of this short dollar strategy. An *unconditional* strategy in which the investor is short the U.S. dollar currency against an equally-weighted basket of foreign currencies independent of the level of interest rates; and a *conditional* strategy in which the investor is long (short) an equally-weighted basket of foreign currencies whenever the U.S. dollar short rate is below (above) the average G10 short rate. Again, we compute the historical realized returns to these strategies, and report them alongside the model-implied risk premia in Panel B of Table IV.

The table confirms that the conditional portfolio earns a statistically significant excess return (4.93% *per annum*; t-stat: 2.05), consistent with evidence on the "short dollar carry trade" reported in Lustig, et al. (2013). However, we do not find evidence of a statistically distinguishable return differential between the conditional and unconditional variants of this trade. Although the historical returns are only modestly negatively skewed and kurtotic, once we adjust for stochastic volatility, the Jarque-Bera test rejects normality for the standardized (log) return innovations (Z-scores).

Unsurprisingly, we find a *lower* model-implied risk premium for the conditional trade (0.76% *per annum*; t-stat: 4.38) than for the unconditional trade (1.69% *per annum*; t-stat: 13.46). Within our model, the dollar investor is compensated with a positive risk premium only when he is *short* the U.S. dollar. Consequently, the conditional trade, which is long the U.S. dollar when U.S. interest rates are high relative to the other countries in the G10 set, occasionally earns a negative model risk premium. Our model is consistent with the realized risk premium on the unconditional short dollar portfolio, but is unable to match the returns of the "short dollar carry trade."

The top panel of Figure 6 plots the time series of the model-implied risk premium for the unconditional short dollar portfolio. Since this portfolio is not orthogonal to the HML_{FX} factor, a portion of the realized excess return reflects compensation for exposure to the global (slope) factor. Specifically, our model attributes 0.37% (t-stat: 3.10) and 0.94% (t-stat: 9.54) of the model-implied portfolio risk premium to exposure to the HML_{FX}

factor for the conditional and unconditional portfolios, respectively. To facilitate exposition, the middle and bottom panels examine only the component of the model-implied portfolio risk premium attributable to the short dollar exposure. These panels decompose this risk premium into contributions from Gaussian and non-Gaussian innovations (middle panel), and the moments of the pricing kernel innovations (bottom panel). In contrast to global shocks, we find that a U.S. investor's premium for being short the U.S. dollar is dominated by Gaussian shocks. This is consistent with evidence that, on average, only 30% of the total variance of the country-specific innovation is due to the non-Gaussian innovation (η_t^{US}).[25] As our theoretical results indicate, the risk premium for short exposure to one's reference (local) currency is determined exclusively by the even moments of the country-specific pricing kernel innovation, (21). We find that nearly the entirety (99%) of the short U.S. dollar risk premium is determined by the variance of the country-specific pricing kernel innovations.

3.3 Forecasting currency returns

Finally, we examine the ability of option-implied currency risk premia to forecast one-month ahead currency returns. We consider two sets of predictive regressions (Table V): (a) repeated cross-sectional regressions (Fama-MacBeth); and (b) pooled panel regressions. In both instances the dependent variable is the simple excess return, (30), and the regressions are based on the high-low cross pairs and X/USD currency pairs (24 pairs). The cross-sectional regressions are repeated every 21 days ($N = 167$); the panel regressions are also run with non-overlapping 21-day returns. Panel regressions include pair fixed effects, and report standard errors adjusted for cross-sectional correlation and time series auto- and cross-correlations using the methodology from Thompson (2011).

The cross-sectional regressions point to a positive and marginally significant relation between the option-implied risk premium and the subsequent realized excess return (regression (1)). The hypothesis that the model correctly prices the cross-section of currency returns (i.e. the intercept is zero and the slope coefficient is one) has a p-value of 68.9%. Simultaneously, the p-value of the alternative that the model has no explanatory power (i.e. all the coefficients with the exception of the intercept are zero) is 11.0%. The adjusted R^2 from the cross-sectional regression is 30%, and is slightly higher than the 26% adjusted R^2 from an *ad hoc* predictive regression using the interest rate differential (regression (3)). To further evaluate the ability of the model to explain currency returns, we regress the realized pair-level excess return onto the model-implied HML_{FX} and short reference components of the risk premium (regression (2)). The regression adjusted R^2 rises to 47%, and neither of the coefficients is statistically distinguishable from one. The p-value of the null hypothesis of the model is

[25]We relegate estimates of the jump variance shares, η_t^i, and parameters of square root processes fitted to country-specific state variable time series, Y_i^t, to a supplementary appendix (Table A.I).

45.2%; and, that of the alternative of no explanatory power is 12%. The fit of the regression improves even further when the model-implied risk premia are combined with the interest rate differential (regressions (4) and (5)). The regression adjusted R^2 peaks at 57% for the specification combining the disaggregated model risk premium with the interest rate differential, even though the interest rate differential itself ceases to be statistically significant. Overall, the cross-sectional regressions support the theoretical predictions of the model, albeit with modest degrees of statistical significance.

Figure 7 contrasts the ability of option-implied risk premia to explain the cross-section of currency excess returns, relative to the random walk model of exchange rate dynamics. The left (right) panel plots the relation between the unconditional mean of the realized currency excess return, and the unconditional mean of the model-implied risk premium (interest rate differential). We find that our model of risk premia – which was calibrated exclusively to exchange rate options – achieves a 77.0% explanatory, adjusted R^2. By contrast, the random walk model of exchange rates achieves an R^2 of 57.2%.

The second part of Table V examines the predictive ability of the model-implied risk premia in the context of panel regressions. Despite the ability of the model to match the cross-sectional properties of currency risk premia well, the time series dimension remains a challenge. The slope coefficients on the model risk premia are negative, though statistically insignificant, in all specifications. Formal tests strongly reject the theoretical prediction of zero currency-pair fixed effects and unit coefficients on the model risk premia. This echoes the features of equity return data, where the identification of a positive risk-return tradeoff in the time series has remained elusive.[26]

4 Discussion

Empirical evidence indicates that currency carry trades formed on the basis of *prevailing* interest rate differentials (conditional carry trades), outperform portfolios formed on the basis of historical *average* interest rate differentials (unconditional carry trades). Using a calibrated affine pricing kernel model, Lustig, et al. (2011) show this finding can be rationalized if the global factor loadings, ξ_t^i, are functions of the country-specific state variable, Y_t^i.[27] With this feature, the instantaneous interest rate, (6) – which is affected by Y_t^i, – carries additional information about the conditional global factor loading, and therefore currency risk premia. However, their calibration was not designed to match exchange rate option data, and therefore places comparatively mild

[26]French, et al. (1987) and Campbell and Hentschel (1992) find a positive albeit insignificant relation between conditional variance and the conditional expected return in equities. Campbell (1987) and Nelson (1991) find a significantly *negative* relation. Ghysels, et al. (2005) find a positive and significant risk-return tradeoff using a mixed data sampling volatility estimator.

[27]Our model can be roughly mapped into the preferred ("unrestricted") specification in Lustig, et al. (2011) by: (a) shutting down jump risks in the global and diffusive Lévy increments by setting $\eta_t^g = \eta_t^i = 0$; and, (b) parameterizing the global factor loading as, $\tilde{\xi}_t^i = \xi^i \cdot \sqrt{1 + \frac{\kappa_i}{(\xi^i)^2} \cdot \frac{Y_t^i}{Z_t}}$.

constraints on the structure of the pricing kernel volatility. We show that, once these constraints are imposed, increasing global factor loading differentials may act to *decrease* currency risk premia.

To understand the implications of time-varying loadings for currency risk premia, in the context of a calibration designed to match exchange rate option prices, it will be useful to specialize to a setting in which all pricing kernel innovations are (conditionally) Gaussian. Furthermore, suppose the model fits the data perfectly. Since the model is Gaussian, the variance of the (log) currency return, $(27a)$, is simply, $\sigma^2_{s_t^{ji}} = \left(\xi_t^j - \xi_t^i\right)^2 \cdot Z_t + Y_t^i + Y_t^j$, and the objective and risk-neutral (i.e. option-implied) variances will be identical. In this context, the expression for the HML_{FX} component of the risk premium for currency pair J/I is given by:

$$\lambda^{ji}_{HML,t} \;=\; \xi_t^i \cdot \left(\xi_t^i - \xi_t^j\right) \cdot Z_t = \xi_t^i \cdot \frac{\sigma^2_{s_t^{ji}} - \left(Y_t^i + Y_t^j\right)}{\left|\xi_t^i - \xi_t^j\right|} \tag{31}$$

where the second equality follows by substituting out the level of the global state variable using the implied variance of the option. As the loading differential widens, the currency risk premium will *decline*, unless there is a sufficiently large offsetting decline in the sum of the country-specific variances. Put differently, an increase in the global factor loading differential, increases the model-implied exchange rate variance, which has to be offset by a decline in one of the model state variables, $\{Z_t,\, Y_t^i,\, Y_t^j\}$, in order to continue matching the data $\sigma^2_{s_t^{ji}}$. If this happens through a decline in Z_t, currency risk premia may fall. We find empirical support for this effect.

Two of the specifications we consider (Specifications I and III) allow global factor loadings to depend on the prevailing interest rate differential, $\xi_t^i = \xi^i - \Psi_t \cdot \left(r^i_{t,t+1} - r^{US}_{t,t+1}\right)$. This parsimonious parametrization fixes the U.S. loading ($\xi_t^{US} = \xi^{US} = 1$), implicitly links the loadings to the model state variables, and matches the specification used by Lustig, et al. (2011) in empirical tests of conditional factor pricing models. Strikingly, we consistently find that specifications in which global factor loadings are time-varying deliver *lower* average estimates of the HML_{FX} risk premium, than the corresponding specifications in which loadings are fixed ($\Psi_t = 0$). This can be seen in Table II, and also in a supplementary table (Table A.IV) which computes the model-implied risk premia for empirical factor replicating portfolios under Specification II (fixed loadings). The model-implied risk premium for the conditional (unconditional) HML_{FX} replicating portfolios rises relative to the values reported in Table IV to 3.81% (3.73%) *per annum*, respectively. Consequently, we find that although allowing for time-varying loadings leads to a modest improvement in the model's fit to option prices, it also leads to an overall decline in model-implied risk premia.

To explore the link between time-varying loadings and currency risk premia in greater depth we conduct two experiments. First, we ask what are the implications of increasing Ψ_t to match the empirical point estimates of risk premia earned by conditional currency carry trades in G10 currencies (Table IV) on option prices. Specifically, we contrast the link induced by high values of Ψ_t between interest rate differentials and the model-implied exchange rate volatility and skewness, with its counterpart in the historical data. Second, we empirically examine whether the idiosyncratic factors driving FX option prices, Y_t^i, are also the drivers of instantaneous interest rates. This relation is a feature of our model, (6), but need not hold in the data. We find that interest rate variation is, in fact, largely orthogonal to factors extracted from foreign exchange option data. Consequently, we examine an alternative parametrization of global factor loadings, $\xi_t^i = \xi^i + \tilde{\Psi}_t \cdot \left(Y_t^i - Y_t^{US}\right)$, which links them directly to the country-specific state variables, Y_t^i.

4.1 Matching risk premia with Ψ_t

The basic design of our empirical investigation has been to match data on exchange rate options, and derive implications for currency risk premia. However, we could also proceed in the reverse direction, and ask what are the model's implications for exchange rate option prices, once the model has been calibrated to match empirically observed currency risk premia. To illustrate this complementary analysis, we begin with our preferred model specification and perturb the calibrated values of Ψ_t by a constant Ψ^* at each point in the time series, until the mean model-implied risk premium for the empirical HML_{FX} factor replicating portfolio *exactly* matches its historical mean excess return of 4.96% *per annum* (Table IV). We leave the calibrated dynamics of the state variables, Z_t and Y_t^i, unaltered. Effectively, we are forcing the model to match the observed excess return through the time-varying global factor loadings, while allowing it to arbitrarily misprice the cross-section of options. We find that this requires increasing the cross-sectional dependence of the global factor loadings on interest rate differential by $\Psi^* = 2.1$, in our sample.

Using the perturbed model ($\Psi_t + \Psi^*$), we examine the implications of time-varying global factor loadings for option price dynamics. As can be clearly seen in, $(27a)$ and $(27b)$, our model links the variance and skewness of the model-implied exchange rate distribution with the loading differential. Consequently, an implication of increasing the dependence of global factor loadings on interest rate differentials ($\Psi^* > 0$) is that option-implied moments will co-move more strongly with interest rate differentials. We examine this effect by running pooled panel regressions of option-implied moments – in the data, in our preferred calibrated specification, and in the perturbed specification – onto the pairwise interest rate differentials.[28]

[28]The computation of option-implied moments in the data requires the interpolation and extrapolation of option-implied volatilities. To ensure the robustness of our results we report the regression for two different computation schemes. The first, interpolates implied volatilities in the observed range and then appends flat tails. The second, extrapolates the observed implied volatilities to the 1δ strike, and then appends flat tails. Extrapolated volatilities in the $(1\delta_p,\, 10\delta_p)$ range are based on the $(10\delta_p,\, 25\delta_p,\, ATM)$ option triplet, and – in the $(10\delta_c,\, 1\delta_c)$ range – on the $(ATM,\, 25\delta_c,\, 10\delta_c)$. Implied volatilities below $1\delta_p$ and above $1\delta_c$ are set equal to their values at those thresholds, unless otherwise noted.

Panel A of Table VI reports the results of panel regressions of option-implied volatility onto the *absolute* interest rate differential – as suggested by $(27a)$ – and the corresponding regression of option-implied skewness onto the raw interest rate differential. We find that the relation between option-implied volatility and the absolute interest rate differential is statistically significant and negative, with a coefficient estimate of -0.45 (t-stat: -3.28). This result is insensitive to the details of the computation of the option-implied variance. Our calibrated model (Specification I) reproduces the relation exactly, with a coefficient estimate of -0.46 (t-stat: -3.06). In contrast, once loadings have been perturbed by Ψ^* – in order to match the historical point estimate of the mean return on the empirical factor mimicking portfolio for the HML_{FX} factor – the relation becomes statistically insignificant, with a point estimate of -0.03 (t-stat: -0.20).

The intuition from skewness regressions is similar in spirit. In the data, the relation between the option-implied skewness and the interest rate differential is weakly negative, with a point estimate between -0.89 (t-stat: -1.66) and -1.21 (t-stat: -1.81). Unsurprisingly, the value of the coefficient here is more sensitive to the details of the implied volatility extrapolation scheme used to compute the value of the option-implied skewness. Our preferred calibration produces a noticeably stronger link, with a point estimate of -6.66 (t-stat: -3.35). This link is strengthened even further once the model is forced to match historical risk premia ($\Psi^* > 0$), raising the point estimate to -12.41 (t-stat: -6.20). Once again, we find that inducing strong covariation between loadings and interest rate differentials, pushes the relation between model-implied exchange rate moments and interest rate differentials further away from that observed in the data.

4.2 Are currency option and bond markets integrated?

Within the model, the dynamics of FX option prices and interest rates, (6), are driven by the global and country-specific state variables. In Panel B of Table VI, we examine the existence of this relationship by regressing the change in the one-month interest rate onto terms governing the model-implied change, $\Delta\left(-k_t^g\left[-\xi_t^i\right]\cdot Z_t\right)$ and $\Delta\left(-k_t^i\left[-1\right]\cdot Y_t^i\right)$. Notice that while this regression is the closest proxy to our model, it suffers from omitted variable bias due to the absence of an empirical proxy for α_t^i. We find that the country-level regressions consistently point to the absence of a link between interest rates and the state variables extracted from foreign exchange options. This observation is consistent with the results of the cross-sectional forecasting regressions (Table V), which indicate that option-implied risk premia and interest rate differentials carry non-redundant information about currency excess returns. The p-values of the joint test that the coefficients on the model-implied variables are equal to one strongly rejects the null of integration between currency option and bond markets. This negative result leaves open the possibility that linking global factor loadings to the Y_t^i state variables directly, rather than interest rate differential, may generate the desired increase in risk premia for models with time-varying loadings. As mentioned earlier, this type of parametrization is explicitly built into the "unrestricted" model of Lustig, et al. (2011).

We repeat the calibration of our preferred model specification, allowing global factor loadings to be tied to local state variable dynamics via $\xi_t^i = \xi^i - \tilde{\Psi}_t \cdot \left(Y_t^i - Y_t^{US}\right)$. Similar to the parametrization based on interest rate differentials, this expression normalizes the U.S. loading to equal ξ^{US} at each point in time, which is set equal to one. Repeating the calibration of our preferred specification, we find that the mean value of $\tilde{\Psi}_t$ in the time series is -0.37 (t-stat: -1.84). The RMSE of the option fitted error is 1.05, which is slightly better than under Specification I. The mean model-implied risk premium for the conditional (unconditional) HML_{FX} replicating portfolio stands at 3.82% (3.76%). Again, we find that: (a) the historical mean return of the empirical portfolios are statistically indistinguishable from the corresponding model-implied quantities; and, (b) the variation in global factor loadings only contributes an additional 6bps to the model-implied risk premium for the conditional HML_{FX} factor replicating portfolio, relative to its unconditional counterpart. These results parallel the baseline results of the paper, and are relegated to the supplementary data appendix (Table A.V).

5 Conclusion

This paper develops a new methodology for computing conditional currency risk premia on the basis of observations of contemporaneous prices of G10 exchange rate options. Conceptually, our methodology requires no arbitrage, such that exchange rates can be interpreted as ratios of pricing kernels, and a factor structure in the pricing kernel dynamics. The factor structure imposes that each country's marginal utility is driven by a combination of common (global) and an idiosyncratic (country-specific) innovations, both of which are modeled using time-changed Lévy increments. Cross-sectional differences in the global factor loadings generate differences in interest rates, currency risk premia, and option-implied exchange rate distributions. We exploit this variation to calibrate our model to exchange rate options, and construct a time-series of option-implied risk premia at the currency pair level. Crucially, our methodology does *not* rely on historical realized returns.

We show that option-implied risk premia for conditional and unconditional empirical factor replicating portfolios for the HML_{FX} (Lustig, et al. (2011)), implemented in developed markets currencies, are statistically indistinguishable from the realized returns of these portfolios. These results indicate that historical estimates of currency returns in developed economies are unlikely to suffer from peso problems, as argued by Burnside, et al. (2011). Furthermore, the structure of our model allows us to decompose the model risk premia across the types of pricing kernel innovations (Gaussian and non-Gaussian), as well as, their moments (variance, skewness, etc.). We find an important role for non-Gaussian global and country-specific innovations in describing the cross-section of option-implied exchange rate distributions. For example, the non-Gaussian innovation in the global factor – which can be interpreted as a catch-all for the effects of stochastic volatility and jumps – accounts for nearly 60% of the model's 3.55% annual risk premium on the HML_{FX} factor. Interestingly, the mechanism through which non-Gaussian innovations contribute to the determination of currency risk premia is primarily as a source of variance in the pricing kernel, rather than higher-order moments. This contrasts with insights from rare disaster models of the equity risk premium calibrated to match historical consumption disasters (Barro and Ursua (2008)). In general, we find that the skewness and higher-order moments of the pricing kernel, on average account for only 15% of the HML_{FX} risk premium in developed markets (G10 currencies). These results are consistent with the observation that crash-hedged currency carry trades continue to deliver positive excess returns (Jurek (2013)). Finally, we demonstrate that the model is never rejected by cross-sectional tests, and that option-implied risk premia forecast 47% of the cross-sectional variation in realized returns, exceeding the forecasting power of interest rate differentials.

Our model calibration confirms the presence of time-varying global factor loadings, although their economic effect on the quality of the fit to option prices is modest (i.e. within plausible estimates of bid-ask spreads). Specifically, we find a negative within-time-period (conditional) association between the loadings and interest rate differentials, and a negative (unconditional) relation with the mean historical interest rate differential. The quantitative variation in the option-implied global factor loadings produces a small spread between the model risk premia for conditional and unconditional HML_{FX} factor mimicking portfolios constructed using developed market (G10) currencies. More generally, in a model with asymmetries in global factor loadings and non-Gaussian innovations, linking loadings to interest rate differentials induces strong comovement between option-implied moments and interest rates. Although we find evidence of this form of comovement in the exchange rate option data, it is relatively weak.

A Cumulant Generating Functions of Lévy Increments

In our model, the non-time-changed Lévy increments are given by L_t^g and L_t^i, and correspond to shocks occurring over the interval from t to $t+1$. Their variance is normalized to be equal to one, and is subsequently set through the time change state variables, Z_t and Y_t^i. The Lévy increments ($j \in \{g, i\}$) can be decomposed into the sum of two independent components:

$$L_t^j \quad = \quad W^j_{\left(1-\eta_t^j\right)} + X^j_{\eta_t^j} \tag{A.1}$$

where $W^j_{\left(1-\eta_t^j\right)}$ is a Gaussian innovation with variance $\left(1-\eta_t^j\right)$ and $X^j_{\eta_t^j}$ is a non-Gaussian innovation with variance η_t^j, and $\eta_t^j \in [0,1]$. With this normalization, the parameter, η_t^j, is interpretable as the time-varying share of innovation variance due to jumps. Note that the random variable $W^j_{1-\eta_t^j}$ is equivalent to the random variable $\sqrt{1-\eta_t^j} \cdot W_1^j$, since each variable has zero mean and equal variance. However, the random variable $X^j_{\eta_t^j}$ is *not* equivalent to $\sqrt{\eta_t^j} \cdot X_1^j$. This can be seen by computing the cumulants of these random variables. Given the independence of the Gaussian and non-Gaussian innovations in the non-time-changed Lévy increment, its cumulant generating function is given by the sum of the cumulant generating functions of the two innovations:

$$k_{L_t^j}[u] \quad = \quad k_{W^j_{\left(1-\eta_t^j\right)}}[u] + k_{X^j_{\eta_t^j}}[u] \tag{A.2}$$

In our application, the $X^j_{\eta_t^j}$ shock follows a CGMY distribution, with variance normalized to equal η^j.

The cumulant generating function of the Gaussian innovation is given by:

$$k_{W^j_{\left(1-\eta_t^j\right)}}[u] \quad = \quad \frac{1-\eta_t^j}{2} \cdot u^2 \tag{A.3}$$

and the corresponding cumulant generating function of the CGMY random variable is given by:

$$k_{X_{\eta_t^j}}[u] \quad = \quad \begin{cases} C \cdot \Gamma[-Y] \cdot \Big((M-u)^Y - M^Y + (G+u)^Y - G^Y \Big) \cdot \tau & Y \neq \{0,\, 1\} \\ -C \cdot \left(\ln\left(1 - \frac{u}{M}\right) + \ln\left(1 + \frac{u}{G}\right) \right) \cdot \tau & Y = 0 \\ C \cdot \left((M-u) \cdot \left(\ln\left(1 - \frac{u}{M}\right) + (G+u) \cdot \ln\left(1 + \frac{u}{G}\right) \right) \right) \cdot \tau & Y = 1 \end{cases} \tag{A.4}$$

where $\Gamma[\cdot]$ is the gamma function. In general, the parameters of the CGMY process, $\{C,\, G,\, M,\, Y\}_t^j$, are allowed to be time-varying our model, and we drop subscripts and superscripts only for parsimony. Since we have normalized the increments such that the non-Gaussian component contributes a fraction η_t^j of the periodic variance, we use the constraint that the second cumulant (variance), κ_X^2, equal to η_t^j to pin down the value of C. For example, consider the generic case when $Y \neq \{0,\, 1\}$. In this case, the second and third cumulant of the CGMY shock are given by:

$$\kappa_X^2 = C \cdot \Gamma[-Y] \cdot Y \cdot (Y-1) \cdot \left(M^{Y-2} + G^{Y-2}\right) = \eta_t^j \tag{A.5a}$$

$$\kappa_X^3 = C \cdot \Gamma[-Y] \cdot Y \cdot (Y-1) \cdot (Y-2) \cdot \left(M^{Y-3} + G^{Y-3}\right) \tag{A.5b}$$

To derive the CGF for the corresponding *time-changed* increments we rely on Theorem 1 in Carr and Wu (2004). In our setup, the time-change is controlled by *pre-determined* state variable, $S_t \in \{Z_t,\, Y_t^i\}$, which allows the model to have non-identically distributed innovations over time. Theorem 1 of Carr and Wu (2004) states that for a generic time change, $\mathcal{T}$, the cumulant generating function of the time-changed Lévy process, $L_{\mathcal{T}}$, is given by $k_{\mathcal{T}}\left[k_L\left[u\right]\right]$, where $k_L[u]$ is the cumulant generating function of the non-time-changed process and $k_{\mathcal{T}}[u]$ is the cumulant generating function of the time-change. In our case, the time-change variables are fixed within the measurement interval (i.e. they follow a degenerate stochastic process with zero drift and volatility), such that:

$$k_{L^j_{S_t}}[u] = k_{S_t}\left[k_{L^j_t}\left[u\right]\right] = k_{L^j_t}\left[u\right] \cdot S_t \tag{A.6}$$

Consequently, the second and third cumulants of the time-changed increments are given by:

$$\kappa^2_{L^j_{S_t}} = \left(\left(1-\eta_t^j\right) + \kappa_X^2\right) \cdot S_t = S_t \tag{A.7a}$$

$$\kappa^3_{L^j_{S_t}} = \left(0 + \kappa_X^3\right) \cdot S_t \tag{A.7b}$$

Unlike in a more traditional stochastic volatility, *both* the variance and skewness of the total time-changed Lévy increment are time-varying and governed by the state variable, S_t. The variance is given by: $\kappa^2_{L^j_{S_t}} = S_t$, and the skewness is given by: $\kappa^3_{L^j_{S_t}} \cdot \left(\kappa^2_{L^j_{S_t}}\right)^{-\frac{3}{2}} = \frac{\kappa_X^3}{\sqrt{S_t}}$.

Putting these results together, the cumulant generating function for the time-changed Lévy increments, $k_{L^j_{S_t}}[u]$ – for the empirically relevant case when $Y \neq \{0, 1\}$ – is given by:

$$\begin{aligned} k_{L^j_{S_t}}[u] &= k_{W^j_{\left(1-\eta_t^j\right)\cdot S_t}}[u] + k_{X^j_{\eta_t^j \cdot S_t}}[u] \\ &= \left(1-\eta_t^j\right) \cdot \frac{u^2}{2} \cdot S_t + \eta_t^j \cdot \frac{\left((M-u)^Y - M^Y + (G+u)^Y - G^Y\right)}{Y \cdot (Y-1) \cdot \left(M^{Y-2} + G^{Y-2}\right)} \cdot S_t \end{aligned} \tag{A.8}$$

References

[1] Ang, Andrew and Joseph S. Chen, 2010, Yield Curve Predictors of Foreign Exchange Returns, Columbia University, *working paper*.

[2] Backus, David K., Silverio Foresi and Chris Telmer, 2001, Affine Term Structure Models and the Forward Premium Anomaly, *Journal of Fianance*, 56(1), p. 279-304.

[3] Backus, David K., Mikhail Chernov, and Ian Martin, 2011, Disasters Implied by Equity Index Options, *Journal of Fianance*, 66, p. 1969-2012.

[4] Bakshi, Gurdip, Carr, Peter, and Liuren Wu, 2008, Stochastic Risk Premiums, Stochastic Skewness in Currency Options, and Stochastic Discount Factors in International Economies, *Journal of Financial Economics*, 87, p. 132-156.

[5] Bakshi, Gurdip, Nikunj Kapadia and Dilip Madan, 2003, Stock Return Characteristics, Skew Laws and the Differential Pricing of Individual Equity Options, *Review of Financial Studies*, 16(1), p. 101-143.

[6] Barro, Robert, 2006, Rare Disasters and Asset Markets in the Twentieth Century, *Quarterly Journal of Economics*, 121, p. 823-866.

[7] Barro, Robert J., Emi Nakamura, Jon Steinsson, and Jose F. Ursua, 2013, Crises and Recoveries in an Empirical Model of Consumption Disasters, *American Economic Journal: Macroeconomics*, 5, p. 35-74.

[8] Barro, Robert J., and Jose F. Ursua, 2008, Macroeconomic Crises Since 1870, *Brookings Papers on Economic Activity*, 39, p. 255350.

[9] Black, Fischer and Myron Scholes, 1973, The Pricing of Options and Corporate Liabilities, *Journal of Political Economy*, 81, p. 617-654.

[10] Brandt, Michael, Cochrane, John, Santa-Clara, Pedro, 2006. International Risk Sharing Is Better Than You Think, or exchange rates are too smooth, *Journal of Monetary Economics*, 53, p. 671698.

[11] Breeden, Douglas and Robert Litzenberger, 1978, Prices of state-contingent claims implicit in option prices, *Journal of Business*, p. 621-651.

[12] Brennan, Michael J., and Yihong Xia, 2006, International Capital Markets and Foreign Exchange Risk, *Review of Financial Studies*, 19(3), p. 753-795.

[13] Brunnermeier, Markus, Stefan Nagel and Lasse Pedersen, 2009, Currency Crashes and the Carry Trade, *NBER Macroeconomics Annual 2008*, 23(1), pp. 77-100.

[14] Burnside, A. Craig, Martin Eichenbaum, Isaac Kleshchelski and Sergio Rebelo, 2011, Do Peso Problems Explain the Returns to the Carry Trade?, *Review of Financial Studies*, 24(3), p. 853-891.

[15] Burnside, A. Craig, and Jeremy J. Graveline, 2013, Exchange Rate Determination, Risk Sharing, and the Asset Market View, *working paper.*

[16] Caballero, Ricardo J., and Joseph Doyle, 2012, Carry Trade and Systemic Risk: Why Are FX Options So Cheap?, NBER Working Paper No. 18644.

[17] Campbell, John Y., 1987, Stock Returns and the Term Structure, *Journal of Financial Economics*, 18(2), 373399.

[18] Campbell, John Y., Hentschel, Ludger, 1992, No News is Good News: An asymmetric model of changing volatility in stock returns, *Journal of Financial Economics*, 31(3), p. 281318.

[19] Carr, Peter, and Dilip Madan, 1999, Option Pricing and the Fast Fourier Transform, *Journal of Computational Finance*, 2(4), p. 61-73.

[20] Carr, Peter, Helyette Geman, Dilip Madan, and Marc Yor, 2002, The Fine Structure of Asset Returns: An Empirical Investigation, *Journal of Business*, 75(2), 305-332.

[21] Carr, Peter and Liuren Wu, 2004, Time-Changed Lévy Processes and Option Pricing, *Journal of Financial Economics*, 71(1), p. 113-141.

[22] Carr, Peter and Liuren Wu, 2007, Stochastic Skew for Currency Options, *Journal of Financial Economics*, 86(1), p. 213-247.

[23] Carr, Peter and Jiming Yu, 2012, Risk, Return, and Ross Recovery, *The Journalf of Derivatives*, 20(1), p. 38-59.

[24] Castagna, Antonio and Fabio Mercurio 2007, Consistent Pricing of FX Options, *Risk*, p. 106-111.

[25] Chernov, Mikhail, Jeremy Graveline, and I. Zviadadze, 2012, Crash Risk in Currency Returns, London School of Economics, *working paper*.

[26] Dumas, Bernard, 1992, Dynamic Equilibrium and the Real Exchange Rate in a Spatially Separated World, *Review of Financial Studies*, 5, p. 153180.

[27] Engel, Charles, 2013, Exchange Rates and Interest Parity, NBER Working Paper No. 19336.

[28] Fama, Eugene F., Forward and Spot Exchange Rates, 1984, *Journal of Monetary Economics*, 14, p. 319-338.

[29] Farhi, Emmanuel and Xavier Gabaix, 2011, Rare Disasters and Exchange Rates, NBER Working Paper #13805.

[30] Farhi, Emmanuel, Samuel P. Fraiberger, Xavier Gabaix, Romain Ranciere and Adrien Verhelhan, 2013, Crash Risk in Currency Markets, NBER Working Paper #15062.

[31] French, Kenneth R., Schwert, G. William, Stambaugh, Robert F., 1987, Expected Stock Returns and Volatility, *Journal of Financial Economics*, 19(1), p. 329.

[32] Gabaix, Xavier, 2012, Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-finance, *The Quarterly Journal of Economics*, 127, p. 645-700.

[33] Garman, Mark B. and Steven W. Kohlhagen, 1983, Foreign Currency Option Values, *Journal of International Money and Finance*, 2, p. 231-237.

[34] Ghysels, Eric, Pedro Santa-Clara, and Rossen Valkanov, 2005, There is a Risk-Return Trade-off After All, *Journal of Financial Economics*, 76(3), p. 509-548.

[35] Hassan, Tarek, and Rui Mano, 2013, Forward and Spot Exchange Rates in a Multi-Currency World, University of Chicago, *working paper*.

[36] Jurek, Jakub W., 2013, Crash-Neutral Currency Carry Trades, *Journal of Financial Economics*, forthcoming.

[37] Longstaff, Francis and Arvind Rajan, 2008, An Empirical Analysis of the Pricing of Collateralized Debt Obligations, *Journal of Finance*, 63, 509-563.

[38] Lustig, Hanno, Nickolai Roussanov and Adrien Verdelhan, 2011, Common Risk Factors in Currency Markets, *Review of Financial Studies*, 24(11), p. 3731-3777.

[39] Lustig, Hanno, Nickolai Roussanov and Adrien Verdelhan, 2013, Countercyclical Currency Risk Premia, *Journal of Financial Economics*, forthcoming.

[40] Lustig, Hanno, Andreas Stathopoulos, and Adren Veredelhan, 2013, The Term Structure of Currency Risk Premia, NBER Working Paper No. 19623.

[41] Martin, Ian, 2013, Consumption-Based Asset Pricing with Higher Cumulants, *Review of Economic Studies*, forthcoming.

[42] Nelson, Daniel B., 1991, Conditional Heteroskedasticity in Asset Returns: A new approach, *Econometrica*, 59(2), p. 347370.

[43] Plantin, Guillaume, and Hyun S. Shin, 2011, Carry Trades, Monetary Policy and Speculative Dynamics, working paper.

[44] Ready, Robert, Nikolai Roussanov, and Colin Ward, 2013, Commodity Trade and the Carry Trade: A Tale of Two Countries, NBER Working Paper #19371.

[45] Ross, Stephen, 2013, The Recovery Theorem, *Journal of Finance*, forthcoming.

[46] Rietz, Thomas A., 1988, The Equity Risk Premium: A Solution, *Journal of Monetary Economics*, 22, p. 117-131.

[47] Sarno, L., P. Schneider, and C. Wagner, 2012, Properties of Foreign Exchange Risk Premia, *Journal of Financial Economics*, 105, p. 279-310.

[48] Thompson, Samuel, 2011, Simple Formulas for Standard Errors that Cluster by Both Firm and Time, *Journal of Financial Economics*, 99(1), p. 1-10.

[49] Verdelhan, Adrien, 2010, A Habit-Based Explanation of the Exchange Rate Risk Premium, *Journal of Finance*, 65(1), p. 123-145.

[50] Wachter, Jessica A., 2013, Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility?, *Journal of Finance*, 68(3), p. 987-1035.

[51] Wystup, Uwe, *FX Options and Structured Products*, Chichester: John Wiley and Sons, 2006.

Technical Appendix to "Option-Implied Currency Risk Premia"

Jakub W. Jurek and Zhikai Xu*

1 A Continuous-Time Model of Exchange Rate Dynamics

We develop a parsimonious model of exchange rates based on a continuous-time specification of pricing kernel dynamics driven by a combination of global and country-specific wealth shocks, and derive its implications for the time-series and cross-section of currency risk premia. This specification extends the model in Bakshi, et al. (2008), to allow for time-varying interest rates and therefore a non-trivial term structure of interest rates. This dimension is particularly important in the study of currency risk premia, which have empirically been linked to interest rates (Lustig, et al. (2011)). This section derives the key expressions of our paper in a continuous-time setting, and highlights the costs and benefits of discretizing the continuous-time model to obtain the specification used in the paper.

1.1 Pricing kernel dynamics

The dynamic model begins with a specification of country-level log pricing kernel dynamics of the form:

$$m_t^i \quad = \quad -\alpha_t^i \cdot t - \xi^i \cdot L_{\Pi_t}^g - L_{\Lambda_t^i}^i \tag{1}$$

where $L_{\Pi_t}^g$ and $L_{\Lambda_t^i}^i$ are independent, time-changed Lévy processes, corresponding to the global and country-specific factors, respectively. Each of these processes can be interpreted as being comprised of a combination of independent diffusive and jump components. The ξ^i reflect each country's loading on the global factor innovations; cross-sectional differences in these loadings play a central role in determining exchange rate dynamics and risk-premia. Note, unlike in the model presented in the paper, the global factor loadings are assumed to be constant through time.

Following Bakshi, et al. (2008), the global factor time-change, Π_t, is defined as function of the the instantaneous stochastic variance process, Z_t:,

$$\Pi_t \quad = \quad \int_0^t Z_s ds \tag{2}$$

$$dZ_t \quad = \quad \kappa_Z \cdot (\theta_Z - Z_t) \cdot dt + \omega_Z \cdot \sqrt{Z_t} \cdot dW_t^Z \tag{3}$$

*Jurek: Bendheim Center for Finance, Princeton University and NBER; *jjurek@princeton.edu*. Xu: Department of Operations Research and Financial Engineering, Princeton University; *zhikaixu@princeton.edu*.

where W_t^Z is a Brownian motion. Similarly, the country-specific time change, Λ_t^i, is modeled as:

$$\Lambda_t^i = \int_0^t Y_s^i ds \tag{4}$$

$$dY_t^i = \kappa_Y^i \cdot (\theta_Y^i - Y_t^i) \cdot dt + \omega_Y^i \cdot \sqrt{Y_t^i} \cdot dW_t^{Y,i} \tag{5}$$

where $(W_t^Z)_{t\geq 0}$ (resp. $(W_t^{Y,i})_{t\geq 0}$) is independent from $(L_t^i)_{t\geq 0}$ and $(W_t^{Y,i})_{t\geq 0}$ (resp. $(L_t^g)_{t\geq 0}$ and $(W_t^Z)_{t\geq 0}$), but can be correlated with the diffusive part of $(L_t^g)_{t\geq 0}$ (resp. $(L_t^i)_{t\geq 0}$). We fix the correlation between the diffusive part of $(L_t^g)_{t\geq 0}$ (resp. $(L_t^i)_{t\geq 0}$) and $(W_t^Z)_{t\geq 0}$ (resp. $(W_t^{Y,i})_{t\geq 0}$) at ρ_Z (resp. ρ_Y^i). Allowing the instantaneous innovations to the pricing kernels and the time change state variables to be correlated, induces non-normalities into the pricing kernel distribution, over and above, those induced by the jump component of the Lévy increments.

In order to obtain closed form expressions within the continuous time-framework we make two key modeling assumptions: (a) the global factor loadings are constant, $\xi_t^i = \xi^i$; and, (b) the dynamics of the time-change variables are specified parametrically. These two assumptions contrast with the discretized model used in the paper. Since the calibration of the discrete-time model is carried out period-by-period using a single tenor of exchange rate options (1-month maturity), we are able to leave the dynamics of the global factor loadings and state variables unrestricted. Allowing for variation in loadings is potentially important, since it offers a mechanism for reconciling the differential performance of conditional and unconditional currency carry trades (Lustig, et al. (2011)). Avoiding parametric restrictions on the state variable dynamics mitigates the risk of model misspecification, but prevents us from being able to simultaneously calibrate the model to the full surface of exchange rate option implied volatilities.

1.2 Cumulant generating functions and the risk-forward measure

The key expression in our model are given in terms of the cumulant generating function (CGF) of the Leǘy increments of $L^g_{\Pi_t}$ and $L^i_{\Lambda^i_t}$. Recall that the cumulant generating function of a random variable, ϵ, is defined as: $k_\epsilon[u] = \ln E_t^{\mathbb{P}}\left[\exp\left(u \cdot \epsilon\right)\right]$. Since the CGF is closely linked to the characteristic function, this formulation also facilitates the pricing of foreign exchange options via generalized Fourier transform methods. The CGFs of the Leǘy increments over an interval of time τ are:

$$k^z_{L^g,\tau}[u] = \ln E_t^{\mathbb{P}}\left[e^{u\cdot(L^g_{\Pi_{t+\tau}} - L^g_{\Pi_t})}\right] = \ln E_t^{\mathbb{P}}\left[e^{u\cdot(L^g_{\Pi_{t+\tau}} - L^g_{\Pi_t})} \mid Z_t = z\right] \tag{6}$$

$$k^y_{L^i,\tau}[u] = \ln E_t^{\mathbb{P}}\left[e^{u\cdot(L^i_{\Lambda^i_{t+\tau}} - L^i_{\Lambda^i_t})}\right] = \ln E^{\mathbb{P}}\left[e^{u\cdot(L^i_{\Lambda^i_{t+\tau}} - L^i_{\Lambda^i_t})} \mid Y_t^i = y\right] \tag{7}$$

For example, the time t yield on a τ-period zero-coupon bond can be expressed using the cumulant generating functions of the Lévy increments by noting that the pricing kernel, M_t^i, has to price risk-free claims denominated in currency I:

$$\begin{aligned} y^i_{t,t+\tau} &= -\frac{1}{\tau} \cdot \ln E_t^{\mathbb{P}}\left[\frac{M^i_{t+\tau}}{M^i_t}\right] = -\frac{1}{\tau} \cdot \ln E^{\mathbb{P}}\left[\frac{M^i_{t+\tau}}{M^i_t} \mid Z_t = z, Y_t^i = y\right] \\ &= \alpha_t^i - \frac{1}{\tau} \cdot k^z_{L^g,\tau}[-\xi^i] - \frac{1}{\tau} \cdot k^y_{L^i,\tau}[-1]. \end{aligned} \tag{8}$$

As we will subsequently show, given the parametrization of the state variable dynamics, the cumulant generating functions, $k^z_{L^g,\tau}[u]$ and $k^y_{L^i,\tau}[u]$, are affine in the prevailing level of the state variables Z_t and Y^i_t. Consequently, the risk-free term structures in all of the countries fall into the $A_2(2)$ family (Dai and Singleton (2002)), and share exposure to the global state variable, Z_t. By contrast, the continuous-time model in Bakshi, et al. (2008) features flat term structures.

Since interest rates are stochastic, it will be convenient to derive option pricing formulas under the risk-forward measure. Recall that under the risk-forward measure, $\mathbb{F}^i_\tau$ – with the τ-period bond denominated in currency I as its numeraire – the price of any risky claim whose payoff is denominated in currency I (e.g. options on exchange rate $S^{ji}_{t+\tau}$), can be priced as a product of the time t value of the numeraire bond, $\exp\left(y^i_{t,t+\tau}\cdot\tau\right)$, and the risk-forward expectation of the risky payoff. Formally, this risk-forward measure, $\mathbb{F}^i_\tau$, is defined through the following change of measure:

$$\frac{d\mathbb{F}^i_\tau}{d\mathbb{P}} = \frac{M^i_{t+\tau}}{M^i_t}\cdot e^{y^i_{t,t+\tau}\cdot\tau} \tag{9}$$

The cumulant generating functions of the Lević increments of $L^g_{\Pi_t}$ and $L^i_{\Lambda^i_t}$, under the $\mathbb{F}^i_\tau$-measure, can be related to their $\mathbb{P}$-measure counterparts as follows:

- For the global factor:

$$\begin{aligned} k^{z,\mathbb{F}^i_\tau}_{L^g,\tau}[u] &= \ln E^{\mathbb{F}^i_\tau}_t\left[e^{u\cdot L^g_{\Pi_{t+\tau}}}\mid Z_t=z\right] = \ln E^{\mathbb{P}}_t\left[e^{y^i_{t,t+\tau}\tau+m^i_{t+\tau}+u\cdot L^g_{\Pi_{t+\tau}}}\mid Z_t=z\right] \\ &= \ln E^{\mathbb{P}}_t\left[e^{-k^z_{L^g,t}[-\xi^i]-\xi^i\cdot L^g_{\Pi_{t+\tau}}+u\cdot L^g_{\Pi_{t+\tau}}}\mid Z_t=z\right] \\ &= k^z_{L^g,\tau}[u-\xi^i]-k^z_{L^g,\tau}[-\xi^i]. \end{aligned} \tag{10}$$

- For country I's country-specific factor:

$$\begin{aligned} k^{y,\mathbb{F}^i_\tau}_{L^i,\tau}[u] &= \ln E^{\mathbb{F}^i_\tau}_t\left[e^{u\cdot L^i_{\Lambda^i_{t+\tau}}}\mid Y^i_t=y\right] = \ln E^{\mathbb{P}}_t\left[e^{y^i_{t,t+\tau}\cdot\tau+m^i_{t+\tau}+u\cdot L^i_{\Lambda^i_{t+\tau}}}\mid Y^i_t=y\right] \\ &= \ln E^{\mathbb{P}}\left[e^{-k^y_{L^i,\tau}[-1]-L^i_{\Lambda^i_{t+\tau}}+u\cdot L^i_{\Lambda^i_{t+\tau}}}\mid Y^i_t=y\right] \\ &= k^y_{L^i,\tau}[u-1]-k^y_{L^i,\tau}[-1]. \end{aligned} \tag{11}$$

- By independence, all the other country-specific factors ($\jmath\neq i$) are unaffected by the measure change.

What remains to be computed are the cumulant generating functions, $k^z_{L^g,\tau}[u]$ and $k^y_{L^i,\tau}[u]$. This computation is made possible via Theorem 1 in Carr and Wu (2004), which states that for any time-changed Lévy process L_{T_τ}, if the CGF of the process L_t is $k_L[u]=\frac{1}{t}\ln E[e^{u' L_t}]$, then the CGF of the process L_{T_τ} is given by $k^u_{T,t}[k_L[u]]$ where $k^u_{T,\tau}$ denotes the CGF of the time-change process $(T_t)_{t\geq 0}$ under the measure $\mathbb{M}^u$ defined by:

$$\frac{d\mathbb{M}^u}{d\mathbb{P}} = \exp\left(u\cdot L_{T_\tau}-T_\tau\cdot k_L[u]\right). \tag{12}$$

Applying this result to our case,

$$k^z_{L^g,\tau}[u]=k^{z,\mathbb{M}^u_Z}_{T,\tau}[k_{L^g}[u]] \quad \text{and} \quad k^y_{L^i,\tau}[u]=k^{y,\mathbb{M}^u_{Y^i}}_{T,\tau}[k_{L^i}[u]] \tag{13}$$

where $k^{z,\mathbb{M}^u_Z}_{T,\tau}$ (resp. $k^{y,\mathbb{M}^u_{Y^i}}_{T,\tau}$) denotes the cumulant generating of Π_τ (resp. Λ^i_τ) under measure $\mathbb{M}^u_Z$ (resp. $\mathbb{M}^u_{Y^i}$) defined by:

$$\frac{d\mathbb{M}^u_Z}{d\mathbb{P}} = \exp\left(u \cdot L^g_{\Pi_\tau} - \Pi_\tau \cdot k_{L^g}[u]\right) \tag{14}$$

$$\frac{d\mathbb{M}^u_{Y^i}}{d\mathbb{P}} = \exp\left(u \cdot L^i_{\Lambda^i_\tau} - \Lambda^i_\tau \cdot k_{L^i}[u]\right). \tag{15}$$

By Girsanov's Theorem, under $\mathbb{M}^u_Z$, the $(Z_t)_{t\geq 0}$ process remains a CIR process, with parameters $\kappa_Z(u) = \kappa_Z - \rho_Z\omega_Z u$, $\theta_Z(u) = \frac{\kappa_Z\theta_Z}{\kappa_Z(u)}$, and $\omega = \omega_Z$. Similarly, under $\mathbb{M}^u_{Y^i}$, $(Y^i_t)_{t\geq 0}$ is still a CIR process, with parameters $\kappa^i_Y(u) = \kappa^i_Y - \rho^i_Y\omega^i_Y u$, $\theta^i_Y(u) = \frac{\kappa^i_Y\theta^i_Y}{\kappa^i_Y(u)}$, and $\omega = \omega^i_Y$.

Finally, since the time-changes are defined as integrals of the underlying CIR state variables, Z_t and Y^i_t, the cumulant generating functions, $k^{z,\mathbb{M}^u_Z}_{T,\tau}$ and $k^{y,\mathbb{M}^u_{Y^i}}_{T,\tau}$, have closed-form expressions (Bakshi, et al. (2008)):

$$k^x_{T,\tau}[u] = -b_\tau(u)x - c_\tau(u) \tag{16}$$

with

$$b_\tau(u) = \frac{-2u(1-e^{-\eta(u)\tau})}{2\eta(u) - (\eta(u)-\kappa)(1-e^{-\eta(u)\tau})} \tag{17a}$$

$$c_\tau(u) = \frac{\kappa\theta}{\omega^2}\left(2\ln\left(1 - \frac{\eta(u)-\kappa}{2\eta(u)}(1-e^{-\eta(u)\tau})\right) + (\eta(u)-\kappa)\tau\right) \tag{17b}$$

$$\eta(u) = \sqrt{\kappa^2 - 2\omega^2 u} \tag{17c}$$

Consequently, the cumulant generating function of the time-changed Lévy innovations is also available in closed-form, so long as $(L^g_t)_{t\geq 0}$ and $(L^i_t)_{t\geq 0}$ are chosen to be Lévy processes, which themselves have closed-from cumulant generating functions.

1.3 The term structure of option-implied currency risk premia

The exchange rate dynamics are determined by the ratio of pricing kernels, such that the log exchange rate between country J and I (the price of currency J in units of currency I), at time $t+\tau$, $s^{ji}_{t+\tau}$, is given by:

$$\begin{aligned} s^{ji}_{t+\tau} - s^{ji}_t &= (m^j_{t+\tau} - m^j_t) - (m^i_{t+\tau} - m^i_t) \\ &= \left(\xi^i - \xi^j\right) \cdot (L^g_{\Pi_{t+\tau}} - L^g_{\Pi_t}) + (L^i_{\Lambda^i_{t+\tau}} - L^i_{\Lambda^i_t}) - (L^j_{\Lambda^j_{t+\tau}} - L^j_{\Lambda^j_t}) \end{aligned} \tag{18}$$

Given this expression and the above results, the cumulant generating function for the log exchange rate, $s^{ji}_{t+\tau}$, under the risk-forward measure, $\mathbb{F}^i_\tau$, is:

$$k^{\mathbb{F}^i_\tau}_{s^{ji},\tau}[u] = k^{Z_t}_{L^g,\tau}\left[(\xi^i - \xi^j)\cdot u - \xi^i\right] - k^{Z_t}_{L^g,\tau}\left[-\xi^i\right] + k^{Y^i_t}_{L^i,\tau}[u-1] - k^{Y^i_t}_{L^i,\tau}[-1] + k^{Y^j_t}_{L^j,\tau}[-u] \tag{19}$$

The knowledge of the cumulant generating function of log exchange rate under forward measure is sufficient to compute option prices using the generalized Fourier transform method (e.g. Carr and Madan (1999)), since the characteristic function of the log exchange rate at time $t+\tau$ can be recovered from: $\phi_{s^{ji}_{t,\tau}}(u) = \exp\left(k^{\mathbb{F}^i_\tau}_{s^{ji},\tau}[i\cdot u]\right)$. As in the discrete-time setup, this allows us to fit the model exclusively to data on exchange rate options and compute option-implied currency risk premia for individual currency pairs.

Following Bakshi, et al. (2008), we define the currency risk premium for currency pair J/I as the difference between the log expected return on the currency adjusted for the interest rate differential:

$$\lambda^{ji}_{t,\tau} \equiv \ln E^{\mathbb{P}}_t \left[\frac{S^{ji}_{t+\tau}}{S^{ji}_t} \right] + \left(y^j_{t,t+\tau} - y^i_{t,t+\tau} \right) \tag{20}$$

Substituting in the $\mathbb{P}$-measure cumulant generating function of the log exchange rate between currencies J/I at horizon τ, $s^{ji}_{t,\tau}$, we obtain:

$$\lambda^{ji}_{t,\tau} = \left(k^{Z_t}_{L^g,\tau}[\xi^i_t - \xi^j_t] + k^{Z_t}_{L^g,\tau}[-\xi^i_t] - k^{Z_t}_{L^g,\tau}[-\xi^j_t] \right) + \left(k^{Y^i_t}_{L^i,\tau}[1] + k^{Y^i_t}_{L^i,\tau}[-1] \right) \tag{21}$$

Notice that, unlike in the discrete-time model, the continuous-time framework allows us to compute an entire *term-structure* of conditional currency risk premia. This represents the direct benefit of parametrically specifying the dynamics of the state-variables.

1.4 Discretizing the model

In bringing the continuous-time model to the data, we make the following modifications. First, we assume that the time change state variables, Z_t and Y^i_t, remain constant within the measurement interval from t to $t+\tau$ (one-month), and are only allowed to change across measurement intervals. This allows us to sidestep estimating the parameters of the time-change, $\{\kappa,\, \theta,\, \omega\}$, for the global and country-specific state variables (11 triplets). However, this modification mechanically imposes that the diffusive component of the Lévy innovations is conditionally Gaussian, and the only source of non-Gaussianity in the pricing kernel innovations are the jump components of the Lévy increments. Since this need not be the case in continuous time, we interpret the non-Gaussian component in our discrete-time innovation as simultaneously capturing the effects of stochastic variation in the state-variable within the modeled time interval, and the effect of jumps.

Second, rather than impose CIR dynamics on the state variables and estimate the model using the full panel of option data (as in Bakshi, et al. (2008)), we sequentially recalibrate the model to match a cross-section of option prices on (up to) 45 G10 currency pairs. This effectively leaves the structure of the state-variable process unconstrained, and allows for non-time-homogenous model dynamics, such that the distributions of the shocks, $\{L^g,\, L^i\}$, and factors loadings, ξ^i, can vary over time.

2 What Drives Equity Risk Premia in Rare Disaster Models?

We contrast our finding that the skewness and higher moments of the log pricing kernels play a modest role in the determination of the HML_{FX} risk premium in G10 currencies, with the corresponding results for U.S. equity risk premia. The derivations here follow the notation of Martin (2013) and rely on parameter values from Backus, et al. (2011). Specifically, we construct the contribution of the higher moments of the log pricing kernel to *equity* risk premia under: (a) macro-model parameterizations based on historical evidence on consumption disasters (Barro (2006), Barro and Ursua (2008) and Barro, et al. (2013)); and, (b) evidence obtained from equity index options (Backus, et al. (2011)).

2.1 The equity risk premium

Following Martin (2013), consider a hypothetical claim which pays a periodic dividend equal to C_t^λ, with C_t equal to aggregate consumption. When $\lambda = 0$, this asset corresponds to a risk-free perpetuity; when $\lambda = 1$, the asset corresponding to a claim on aggregate consumption, and when $\lambda > 1$, the asset represents a levered consumption claim. He defines the risk premium as the difference between the log of the expected return on the risky asset and the log riskless rate, paralleling our definition in the context of currencies. When the representative agent has CRRA preferences, the risk premium is given by:

$$\begin{aligned} rp &= \ln E_t\left[\frac{P_{t+1}(\lambda)}{P_t(\lambda)}\right] - r_f \\ &= k\left[\lambda\right] + k\left[-\gamma\right] - k\left[\lambda - \gamma\right] \end{aligned} \tag{22}$$

where $k[u]$ is the cumulant generating function of the log consumption growth process, $g_{t+1} = \ln C_{t+1} - \ln C_t$, and γ is the relative risk aversion coefficient of the representative agent. If the log consumption growth, g_{t+1}, is driven by a Lévy increment with finite variation (e.g. as in the jump-diffusion models of Barro (2006), Barro and Ursua (2008), Backus, et al. (2011) and Barro, et al. (2013)), the cumulant generating function takes the form:

$$k[u] = \mu u + \frac{1}{2}\sigma^2 u^2 + \int_{\mathbb{R}} (e^{ux} - 1) d\nu(x) \tag{23}$$

where ν denotes the Lévy measure of g_{t+1}, and μ is the mean of the log consumption growth process.

Using an infinite series expression for the cumulant generating function we have:

$$rp = \sum_{n=1}^{\infty} \frac{(\lambda^n + (-\gamma)^n - (\lambda-\gamma)^n) \cdot \kappa^n}{n!} \tag{24}$$

where $\kappa^n = \left.\frac{\partial^n k[u]}{\partial u^n}\right|_{u=0}$ is the n-th cumulant of the log consumption growth rate process. Expanding the risk premium to third order, and exploiting the relations between the cumulants of the innovations with variance and skewness, we have:

$$rp \approx \kappa^2 \cdot \lambda \cdot \gamma + \frac{1}{2} \cdot \kappa^3 \cdot \lambda \cdot \gamma \cdot (\lambda - \gamma) = \lambda \cdot \gamma \cdot \mathcal{V} + \frac{1}{2} \cdot \mathcal{S} \cdot \mathcal{V}^{\frac{3}{2}} \cdot \lambda \cdot \gamma \cdot (\lambda - \gamma)$$

where $\mathcal{V}$ and $\mathcal{S}$, denotes the variance and skewness of the log consumption growth process, respectively. Similar to our model for currencies, the two terms reflect the contributions of variance and skewness to the equity risk premium. The contribution of the higher-order terms can be trivially recovered as the difference between the risk premium given by the closed-form expression and the approximate formula.

2.2 Moment decompositions

To compare the results we obtained for the HML_{FX} factor in currencies with the results in the equity literature, we compute the contribution of variance, skewness and higher-moments to the equity risk premium under the calibrations reported in Table II in Backus, et al. (2011). Specifically, we focus on parametrizations (2) and (4), which correspond to a macro model calibrated to historical data (*Poisson Cons Gr*), and a model calibrated to match evidence on equity index options (*Implied Cons Gr*). These two sets of parameters are chosen to generate the same equity risk premium – defined as the expectation of the difference in log returns, $E_t\left[r_{t+1}(\lambda) - r_f\right]$ –

but produce different values for *rp* due to the Jensen effect. The parameters of the jump component for the macro model are derived from studies of international macroeconomic data by Barro (2006), Barro and Ursua (2008), and Barro, et al. (2013). In each of the models, the consumption growth jump distribution is modeled as Poisson mixture of normals.

To facilitate comparisons with our decompositions for currency risk premia, we first report the decomposition for $rp = \ln E_t\left[R_{t+1}(\lambda)\right] - r_f$, following the risk premium definition in Martin (2013). The table below displays the contributions of various moments in levels, and as a fraction of the total model risk premium:

	Poisson Cons Gr	Implied Cons Gr
$\ln E_t\left[R_{t+1}(\lambda)\right] - r_f$	4.95%	5.56%
Variance	3.24%	5.44%
Skewness	0.06%	0.10%
Residual	1.65%	0.02%
Variance [share]	65.5%	97.9%
Skewness [share]	1.2%	1.7%
Residual [share]	33.3%	0.3%

The same decomposition for the equity risk premium, defined as the expectation of the log equity return over the log risk-free rate, following Backus, et al. (2011), yields:

	Poisson Cons Gr	Implied Cons Gr
$E_t\left[r_{t+1}(\lambda)\right] - r_f$	4.00%	3.99%
Variance	1.65%	3.85%
Skewness	1.10%	0.12%
Residual	1.25%	0.02%
Variance [share]	41.3%	96.5%
Skewness [share]	27.5%	3.0%
Residual [share]	31.3%	0.5%

Decompositions of the equity risk premium into contributions from the variance of the log pricing kernel and the higher-order ($j \geq 3$) odd and even cumulants, yield quantitatively and qualitatively similar results. For example, the decomposition of rp under the *Poisson Cons Gr* parametrization yields: 3.24% (variance), 0.09% (higher-order odd cumulants), and 1.62% (higher-order even cumulants). The decomposition under the equity-index option implied parametrization (*Implied Cons Gr*) is unchanged.

Table I
Calibrated Model Parameters

This table reports summary statistics for the calibrated parameters of the pricing kernel factor model. Results are reported for four specifications which either: (a) match the prices of the high/low interest rate cross pairs and X/USD options (HLX + X/USD option set; 24 pairs) or the full panel of G10 cross rates (X/Y option set; 45 pairs); and, (b) allow the global factors loadings, ξ_i, to be time-varying or not. All calibrations are performed using daily option quotes from January 1999 to June 2012 ($T = 3520$ days) at five individual option strikes (10δ put, 25δ put, at-the-money, 25δ call, 10δ call). In specifications with time-varying global factor loadings, the loadings are parameterized on the basis of the one-month interest rate differential, $\xi_t^i = \xi^i - \Psi_t \cdot (r_{t,t+1}^i - r_{t,t+1}^{US})$. The slope coefficient, Ψ_t, is time-varying, but common to all currencies in the cross-section. Panel A reports the values of the country-specific loadings, ξ^i, obtained from the first stage of the calibration, with standard errors in parentheses. We also report the mean one-month LIBOR interest rate differential for each country relative to the U.S.. Panel B summarizes the cross-sectional dependence of the global factor loadings on interest rates differentials, Ψ_t, and the characteristics of the global factor innovation, $L_{Z_t}^g$. We report the share of variance in the global innovations due to jumps, η_t^g, and estimates of the parameters of the global CGMY jump component, G_t^g (dampening coefficient) and Y_t^g (power coefficient). Finally, we compute the skewness (Skewness$_t^g$) and kurtosis (Kurtosis$_t^g$) of the global factor innovation induced by variation in the global factor, Z_t. We obtain each of these quantities day-by-day by minimizing option pricing errors for the target option set, and report their time series means and volatilities. Panel C reports the root mean squared option pricing error measured in volatility points. We compute root mean squared errors (RMSE) for each pair in the target option set, and report their mean pooled: (a) across all strikes and currency pairs; (b) across all pairs with a given option strike.

Specification	Fixed ξ_t^i	Option Set
I	No	HLX and X/USD pairs
II	Yes ($\Psi_t = 0$)	HLX and X/USD pairs
III	No	All pairs (X/Y)
IV	Yes ($\Psi_t = 0$)	All pairs (X/Y)

Panel A: Global Factor Loadings, ξ^i

Specification	AUD	CAD	CHF	EUR	GBP	JPY	NOK	NZD	SEK	USD
I and II	0.79	1.00	0.96	1.00	0.93	1.17	0.92	0.81	1.00	1.00
	(0.0101)	(0.0000)	(0.0033)	(0.0001)	(0.0044)	(0.0072)	(0.0054)	(0.0084)	(0.0001)	-
III and IV	0.84	0.96	1.02	1.00	0.98	1.27	1.00	0.87	0.99	1.00
	(0.0046)	(0.0044)	(0.0050)	(0.0049)	(0.0047)	(0.0062)	(0.0050)	(0.0049)	(0.0051)	-
$\overline{r_{t,t+1}^i - r_{t,t+1}^{US}}$ [%]	2.46	0.20	-1.58	-0.15	1.06	-2.54	1.53	2.78	0.07	0.00

Panel B: Global Factor Summary Statistics

Specification		Ψ_t	η_t^g	G_t^g	Y_t^g	Skewness$_t^g$	Kurtosis$_t^g$
I	Mean	0.42	0.51	7.91	-0.75	-1.10	6.03
	Volatility	1.89	0.19	5.60	0.16	1.35	4.04
II	Mean	0.00	0.53	7.13	-0.74	-1.16	6.21
	Volatility	-	0.18	3.80	0.17	1.29	4.06
III	Mean	2.37	0.53	9.98	-0.74	-1.25	5.72
	Volatility	2.26	0.20	6.83	0.15	1.73	4.20
IV	Mean	0.00	0.53	7.43	-0.75	-1.12	5.30
	Volatility	-	0.19	3.16	0.14	1.53	3.46

Panel C: Option Pricing RMSE (Volatility Points)

Specification	All	$10\delta_p$	$25\delta_p$	50δ	$25\delta_c$	$10\delta_c$
I	1.12	1.04	1.27	1.27	1.16	0.85
II	1.17	1.09	1.34	1.33	1.22	0.89
III	1.55	1.28	1.77	1.85	1.67	1.19
IV	1.60	1.30	1.83	1.92	1.74	1.22

Table II
Model HML_{FX} Factor Risk Premium Decomposition

This table uses the calibrated pricing kernel model to compute and decompose the HML_{FX} factor risk premium from the perspective of a U.S. dollar investor. We compute the model-implied HML_{FX} factor risk premium by constructing a hypothetical, dollar-neutral factor mimicking portfolio on the basis of the calibrated time series of the global factor loadings, ξ_t^i. At each point in time, we sort the G10 currencies – excluding the U.S. dollar – into long and short portfolios on the basis of their prevailing loadings, and weight the currencies within each portfolio on the basis of the absolute deviation of their loading from the average loadings of currencies with ranks five and six. The model risk premia are computed daily and span from January 1999 to June 2012 ($T = 3520$ days). We report the global factor loadings of the high and low interest rate portfolios (ξ^H and ξ^L), as well as, the mean portfolio risk premium (λ^{HML}; *% per annum*). We then decompose the portfolio risk premium into contributions from the Gaussian and non-Gaussian components of the global factor innovation, $L^g_{Z_t}$. Finally, we evaluate the mechanism through which jumps contribute to equilibrium currency risk premia by decomposing risk premia across the moments of the pricing kernel innovations (variance, skewness, etc.). For each quantity, we report its time-series mean, volatility, and range. Results are reported for each of the four specifications reported in Table I.

		Loadings		Risk premium	By Shock Type		By Moment		
Specification		ξ^H	ξ^L	λ^{HML}	Gaussian	Non-Gaussian	Variance	Skewness	Other
I	Mean	0.81	1.12	5.03	42.07	57.93	85.30	9.57	5.13
	Volatility	0.05	0.05	6.28	18.18	18.18	10.79	5.17	6.12
	Min.	0.46	0.99	0.22	0.01	0.18	38.03	0.01	0.00
	Max.	0.88	1.49	74.87	99.82	99.99	99.99	26.14	37.24
II	Mean	0.83	1.13	5.31	40.12	59.88	85.05	10.26	4.69
	Volatility	0.00	0.00	6.65	17.11	17.11	9.70	4.80	5.35
	Min.	0.83	1.13	0.25	0.01	5.09	36.04	0.08	0.00
	Max.	0.83	1.13	79.70	94.91	99.99	99.92	26.08	39.37
III	Mean	0.80	1.25	4.46	41.82	58.18	86.87	9.17	3.96
	Volatility	0.07	0.06	4.99	20.36	20.36	11.01	5.68	5.96
	Min.	0.48	1.13	0.10	0.00	0.16	21.62	0.02	0.00
	Max.	0.92	1.68	57.92	99.84	100.00	99.97	26.88	59.98
IV	Mean	0.87	1.23	5.44	40.95	59.05	84.98	10.65	4.38
	Volatility	0.00	0.00	5.96	18.76	18.76	9.90	5.27	5.11
	Min.	0.87	1.23	0.17	0.00	5.97	28.47	0.71	0.06
	Max.	0.87	1.23	75.25	94.03	100.00	99.21	26.98	48.20

Table III
Model Short Reference Currency Factor Risk Premium Decomposition

This table uses the calibrated pricing kernel model (Specification I) to compute and decompose the risk premium demanded by investors in each of the G10 countries for being short their local (reference) currency. The model risk premia are computed daily and span from January 1999 to June 2012 ($T = 3520$ days). Panel A reports the mean, volatility, and range of the model-implied short reference risk premium for each country (*% per annum*). Panel B decomposes the risk premium into: (a) contributions from the Gaussian and non-Gaussian components of the country-specific innovation, $L^i_{Y^i_t}$; (b) contributions from variance and higher moments of the kernel innovations. The risk premium decompositions report the mean share of the total risk premium due to each component and its volatility. Since the decompositions are exhaustive, we only report the range for one of the two components.

Panel A: Short Reference Risk Premia [%]

		AUD	CAD	CHF	EUR	GBP	JPY	NOK	NZD	SEK	USD
Total (λ^i_{refFX})	Mean	0.54	0.54	0.69	0.92	0.63	0.48	0.76	0.79	1.33	0.74
	Volatility	0.45	0.81	0.50	0.75	0.51	0.51	0.61	0.38	1.11	0.41
	Min.	0.00	0.00	0.15	0.12	0.13	0.00	0.16	0.05	0.33	0.20
	Max.	5.28	7.40	4.31	5.96	4.12	3.49	5.26	3.74	8.86	3.17

Panel B: Risk Premium Decompositions [share, %]

Component		AUD	CAD	CHF	EUR	GBP	JPY	NOK	NZD	SEK	USD
Gaussian (Share [%])	Mean	39.37	52.98	39.77	65.70	62.74	56.97	58.70	61.18	75.57	64.86
	Volatility	19.68	30.68	15.97	17.30	15.89	33.01	19.06	18.26	12.11	18.26
Non-Gaussian (Share [%])	Mean	60.63	47.02	60.23	34.30	37.26	43.03	41.30	38.82	24.43	35.14
	Volatility	19.68	30.68	15.97	17.30	15.89	33.01	19.06	18.26	12.11	18.26
	Min.	0.19	0.02	1.54	0.18	0.11	0.01	0.33	0.12	0.08	0.02
	Max.	100.00	100.00	99.99	97.01	88.71	100.00	92.69	99.96	62.40	87.62
Variance (Share [%])	Mean	98.05	98.57	98.10	98.92	98.83	98.30	98.69	98.66	99.24	98.75
	Volatility	2.31	1.94	2.17	1.35	1.45	2.71	1.63	1.77	1.00	1.68
Other (Share [%])	Mean	1.95	1.43	1.90	1.08	1.17	1.70	1.31	1.34	0.76	1.25
	Volatility	2.31	1.94	2.17	1.35	1.45	2.71	1.63	1.77	1.00	1.68
	Min.	0.00	0.00	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	Max.	10.97	11.03	10.34	8.51	8.68	11.04	9.08	10.87	6.17	8.83

Table IV
Option-Implied Currency Risk Premia: Empirical factor mimicking portfolio returns

The table compares the historical (realized) returns to factor mimicking portfolios, with the corresponding model-implied risk premia. The data span the period from January 1999 to June 2012 ($N = 162$ months). Realized returns are computed on the basis of monthly buy-and-hold returns. The conditional HML_{FX} factor mimicking portfolios is a dollar-neutral portfolio formed by sorting currencies into long and short portfolios on the basis of their prevailing one-month LIBOR rates. Within each portfolio currencies are spread weighted, on the basis of the distance of their respective interest rates to the mean of the interest rates in countries with ranks five and six. The unconditional HML_{FX} portfolios is formed on the basis of the historical interest rate differentials computed using an expanding window starting in January 1990. The conditional short USD factor mimicking goes long (short) the U.S. dollar when the prevailing U.S. dollar interest rate is above (below) the prevailing average of the nine other G10 interest rates. The unconditional short USD factor mimicking portfolio is long an equal-weighted basket of G10 currencies against the U.S. dollar. Model risk premia are computed on the basis of Specification I. We report the moments of various quantities of interest, the p-values of the Jarque-Bera test for the realized portfolio returns, and comparisons of the mean realized and model-implied risk premia (t-stats in square brackets). Portfolio ξ^L and ξ^S report the time-series mean of the global factor loading for the long and short leg of the portfolio.

Panel A: HML_{FX} Factor Mimicking Portfolio

		Conditional	Unconditional	Difference
Realized	Mean	4.96	3.32	1.63
		[1.92]	[1.32]	[1.99]
	Volatility	9.51	9.26	
	Skewness	-1.07	-0.89	
	Kurtosis	7.03	6.90	
	JB (returns)	0.00	0.00	
	JB (Z-scores)	0.00	0.00	
	Carry	4.57	3.69	0.88
	t-stat	[60.81]	[51.65]	[27.95]
Model	Total ($\bar{\lambda}$)	3.55	3.44	0.11
		[8.72]	[8.37]	[2.63]
	HML (global, ($\bar{\lambda}^{HML}$)	3.55	3.44	0.11
		[8.72]	[8.37]	[2.63]
	Short USD ($\bar{\lambda}^{USD}$)	0.00	0.00	0.00
		-	-	-
	Portfolio ξ^L (long)	0.83	0.87	
	Portfolio ξ^S (short)	1.06	1.07	
Difference	Mean	1.41	-0.11	1.53
	t-stat	[0.53]	[-0.04]	[1.86]

Panel B: Short USD Factor Mimicking Portfolio

		Conditional	Unconditional	Difference
Realized	Mean	4.93	3.12	1.81
		[2.05]	[1.29]	[0.78]
	Volatility	8.84	8.90	
	Skewness	-0.45	-0.17	
	Kurtosis	3.96	3.72	
	JB (returns)	0.01	0.09	
	JB (Z-scores)	0.00	0.00	
Model	Total ($\bar{\lambda}$)	0.76	1.69	-0.93
		[4.38]	[13.46]	[-7.61]
	HML (global, ($\bar{\lambda}^{HML}$)	0.37	0.94	-0.57
		[3.10]	[9.53]	[-6.32]
	Short USD ($\bar{\lambda}^{USD}$)	0.38	0.74	-0.36
		[3.19]	[22.73]	[-8.93]
	Portfolio ξ^L (long)	0.96	0.94	
	Portfolio ξ^S (short)	0.98	1.00	
Difference	Mean	4.18	1.44	2.74
	t-stat	[1.73]	[0.59]	[1.17]

Table V
Option-Implied Currency Risk Premia: Individual currency pairs

The table reports the results of a regression analysis of realized carry trade returns for individual currency pairs on option-implied currency risk premia and interest rate differentials. The returns are measured using 21-day non-overlapping intervals, and the explanatory variables are assumed to be measured as of the last day preceding that interval. We report two sets of results. The first is based on repeated cross-sectional regressions (Fama-MacBeth); the second, is based on a pooled panel regression. The reported coefficients are time series averages and standard errors are computed on the basis of the time series of estimates. $Adj.\ R^2$ reports the average cross-sectional adjusted R^2. Standard errors of coefficient estimates are reported in parentheses ($N = 167$). The second set of regressions is based on pooled panel regressions with currency-pair fixed effects ($N = 24 \cdot 167 = 4008$). For panel regressions, standard errors are adjusted for cross-sectional correlation and time series auto- and cross-correlations using the methodology from Thompson (2011) with three lags. The panel regression is run with country-fixed effects; however, the adjusted R^2 is reported net of the explanatory power of the fixed effects. Finally, we report the p-value for two hypothesis tests. The first hypothesis asserts that the model is correctly specified, such that the regression intercept (respectively, fixed effects) is zero and the slope coefficients on the model-implied risk premia are one ($\mathcal{H}_0$). When additional regressors are included, the model predicts a zero coefficient on those variables; if no model-implied variables are included in the regression, we do not report the result of the test. The second hypothesis ($\mathcal{H}_1$) is that the included variables have no explanatory power (i.e. all the coefficients with the exception of the intercept are zero).

	Cross-sectional					Panel				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Intercept	0.00	-0.00	0.00	-0.00	-0.00	FE	FE	FE	FE	FE
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)					
λ_t^{ji}	2.09			0.73		-0.22			-0.21	
	(1.31)			(1.44)		(0.72)			(0.74)	
$\lambda_t^{ji,hml}$		2.50			0.75		-0.15			-0.11
		(1.42)			(1.46)		(0.72)			(0.75)
$\lambda_t^{ji,refFX}$		6.01			6.74		-5.25			-6.61
		(4.03)			(4.53)		(4.81)			(4.68)
$r_t^j - r_t^i$			0.14	0.11	0.08			0.06	0.05	0.07
			(0.05)	(0.06)	(0.06)			(0.07)	(0.08)	(0.08)
Adj. R^2	0.30	0.47	0.26	0.43	0.57	0.00	0.00	0.00	0.00	0.00
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	-	-	-	-	-
N	167	167	167	167	167	4008	4008	4008	4008	4,008
$\mathcal{H}_0$ p-value [%]	68.90	45.22	-	27.76	32.35	0.00	0.00	-	0.00	0.00
$\mathcal{H}_1$ p-value [%]	11.01	11.88	0.51	6.20	17.07	75.47	53.86	44.01	68.28	46.25

Table VI
Time-Varying Loadings, Interest Rates, and Option-Implied Moments

Panel A of this table compares the model-implied dynamics of risk-neutral volatility and skewness implied by Specification I with those observed in the data. In particular, we focus on the link between the interest rate differentials driving the global factor loadings and the option-implied moments in the data, and in the model. In order to compute option implied moments in the data we interpolate the implied volatility functions and append flat tails at 10δ, or extrapolate to 1δ using the vanna-volga method of Castagna and Mercurio (2007) before appending flat tails. We report results for both computation schemes to ensure robustness. We then compute the model-implied moments under Specification I, which allows the global factor loadings to be functions of interest rate differentials. Finally, we repeat the computation in a model where Ψ_t has been perturbed by a constant Ψ^*, chosen such that the model-implied portfolio risk premium for the conditional HML_{FX} factor replicating portfolio is equal to the point estimate of the mean excess return realized in the sample. We then run pooled panel regressions of the data and model-implied moments onto interest rate differentials ($N = 24 \cdot 167 = 4008$). The panel regressions include currency pair fixed effects, and adjust standard errors for cross-sectional correlation and time series auto- and cross-correlations using the methodology from Thompson (2011). Panel B of this table examines the existence of model-predicted links between interest rates, and the state-variables extracted from FX options. We report the results from regressions of changes in one-month LIBOR rates in G10 countries onto changes in the model-implied quantities, -$k_{L^g}[-\xi_t^i] \cdot Z_t$ and -$k^i_{L^i}[-1] \cdot Y_t^i$, appearing in the expression for yields. The global factor loadings, ξ_t^i, and the option-implied state variables are obtained from our preferred model calibration (Specification I). We report regression coefficients, standard errors (in parentheses), and the regression adjusted R^2; for panel regressions, the adjusted R^2 is reported net of the explanatory power of the fixed effects. Finally, we report the p-value for two hypothesis tests. The first hypothesis asserts that the model is correctly specified, such that the slope coefficients on the model-implied quantities are jointly equal to one ($\mathcal{H}_0$). The second hypothesis ($\mathcal{H}_1$) is that the slope coefficients are jointly equal to zero. The data span the period from January 1999 to June 2012 ($N = 166$).

Panel A: Option-Implied Moments

	Option-Implied Volatility				Option-Implied Skewness			
	Data		Model	Model	Data		Model	Model
	(10δ)	(1δ)	(Ψ_t)	($\Psi_t + \Psi^*$)	(10δ)	(1δ)	(Ψ_t)	($\Psi_t + \Psi^*$)
Intercept	FE	FE	FE	FE	FE	FE	FE	FE
$r^j_{t,t+1} - r^i_{t,t+1}$	-0.45	-0.46	-0.46	-0.03	-0.89	-1.21	-6.66	-12.41
	(0.14)	(0.14)	(0.15)	(0.17)	(0.53)	(0.67)	(1.99)	(2.00)
Adj. R^2	0.02	0.02	0.02	0.00	0.00	0.00	0.01	0.04
N	4008	4008	4008	4008	4008	4008	4008	4008

Panel B: Interest rate changes, Δr_t^i

	AUD	CAD	CHF	EUR	GBP	JPY	NOK	NZD	SEK	USD
Intercept	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
$-\Delta(k_{L_t^g}[-\xi_t^i] \cdot Z_t)$	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
$-\Delta(k^i_{L_t^i}[-1] \cdot Y_t^i)$	0.04	-0.01	0.01	-0.01	-0.03	-0.01	0.05	0.03	0.01	0.00
	(0.02)	(0.02)	(0.02)	(0.02)	(0.05)	(0.01)	(0.05)	(0.03)	(0.02)	(0.04)
Adj. R^2	0.00	-0.01	0.00	-0.01	0.00	0.01	0.00	0.00	-0.01	0.00
N	166	166	166	166	166	166	166	166	166	166
$\mathcal{H}_0$ p-value [%]	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
$\mathcal{H}_1$ p-value [%]	31.71	82.37	46.95	81.35	30.97	14.82	26.77	40.29	69.10	44.73

Table A.I
State-variable Dynamics and CGMY Jump Structures in the Preferred Model Specification

This table reports detailed information on the state variable dynamics, Z_t and Y_t^i, and calibrated CGMY jump structures under Specification I. We report the parameters of square-root diffusions fitted to the time series of the global, Z_t, and country-specific state variables, Y_t^i, corresponding to each of the G10 currencies (coefficient standard errors in parentheses). The second section of the table reports the time-series means (first row) and volatilities (second row) of the calibrated CGMY parameters, as well as, the fraction of the innovation variance due to the non-Gaussian innovation, η_t^j. Our model allows for one-sided global factor innovations ($M^g \to \infty$), and two-sided country-specific innovations. The jump structures in the countries are assumed to share the same parameters, $\{G,\, M,\, Y\}_t^i$, which we report on the line corresponding to AUD, and omit them elsewhere. Finally, we compute the skewness (Skewness$_t^j$) and kurtosis (Kurtosis$_t^j$) of the global, $L_{Z_t}^g$, and country-specific factor innovation, $L_{Y_t^i}^i$, induced by variation in the corresponding state variable.

State variable	κ^j	θ^j	ω^j	η_t^j	G_t^j	M_t^j	Y_t^j	Skewness$_t^j$	Kurtosis$_t^j$
Global (Z_t)	2.3061	0.1416	0.6083	0.51	7.91	∞	-0.75	-1.10	6.03
	(0.3384)	(0.0007)	(0.0003)	0.19	5.59	-	0.16	1.35	4.05
AUD	2.8655	0.0051	0.0972	0.59	7.75	78.83	-0.73	-3.69	15.22
	(0.4317)	(0.0000)	(0.0000)	0.20	3.39	13.48	0.12	1.86	4.03
CAD	1.3574	0.0053	0.1235	0.38				-2.06	11.92
	(0.2379)	(0.0000)	(0.0000)	0.38				2.79	4.93
CHF	2.1419	0.0064	0.0886	0.59				-3.50	15.17
	(0.2926)	(0.0000)	(0.0000)	0.16				1.64	4.42
EUR	1.7826	0.0087	0.1135	0.27				-1.70	11.04
	(0.2500)	(0.0000)	(0.0000)	0.26				1.70	4.86
GBP	1.3111	0.0060	0.0691	0.36				-2.34	12.72
	(0.1840)	(0.0000)	(0.0000)	0.16				1.46	4.66
JPY	4.4177	0.0041	0.1982	0.31				-1.42	10.65
	(0.7917)	(0.0000)	(0.0000)	0.40				2.87	4.52
NOK	1.6399	0.0071	0.0766	0.40				-2.42	12.23
	(0.2350)	(0.0000)	(0.0000)	0.19				1.66	4.70
NZD	3.3117	0.0077	0.0928	0.37				-2.04	11.30
	(0.5183)	(0.0000)	(0.0000)	0.19				1.42	4.51
SEK	1.7331	0.0127	0.1185	0.15				-0.82	8.49
	(0.2456)	(0.0000)	(0.0000)	0.21				1.00	3.73
USD	2.2455	0.0073	0.0772	0.34				-1.97	10.92
	(0.3370)	(0.0000)	(0.0000)	0.18				1.41	4.59

Table A.II

Model HML_{FX} Factor Risk Premium Decomposition: Portfolios formed using one-month interest rates

This table uses the calibrated pricing kernel model to compute and decompose the HML_{FX} factor risk premium from the perspective of a U.S. dollar investor. We compute the model-implied HML_{FX} factor risk premium by constructing a hypothetical, dollar-neutral factor mimicking portfolio on the basis of the one-month interest rates, $y^i_{t,t+1}$. At each point in time, we sort the G10 currencies – excluding the U.S. dollar – into long and short portfolios on the basis of their yields, and weight the currencies within each portfolio on the basis of the absolute deviation of their loading from the average yields of currencies with ranks five and six. The model risk premia are computed daily and span from January 1999 to June 2012 ($T = 3520$ days). We report the global factor loadings of the high and low interest rate portfolios (ξ^H and ξ^L), as well as, the mean portfolio risk premium (λ^{HML}; % *per annum*). We then decompose the portfolio risk premium into contributions from the Gaussian and non-Gaussian components of the global factor innovation, $L^g_{Z_t}$. Finally, we evaluate the mechanism through which jumps contribute to equilibrium currency risk premia by decomposing risk premia across the moments of the pricing kernel innovations (variance, skewness, etc.). For each quantity, we report its time-series mean, volatility, and range. Results are reported for each of the four specifications reported in Table I.

		Loadings		Risk premium	By Shock Type		By Moment		
Specification		ξ^H	ξ^L	λ^{HML}	Gaussian	Non-Gaussian	Variance	Skewness	Other
I	Mean	0.83	1.06	3.46	42.01	57.99	85.22	9.48	5.30
	Volatility	0.07	0.04	4.53	17.21	17.21	10.19	4.95	5.63
	Min.	0.47	0.95	0.21	0.79	2.44	45.70	0.35	0.07
	Max.	0.96	1.46	50.26	97.56	99.21	99.54	25.18	29.27
II	Mean	0.85	1.06	3.74	40.22	59.78	85.25	10.09	4.65
	Volatility	0.03	0.02	4.73	16.29	16.29	8.78	4.52	4.58
	Min.	0.82	1.03	0.22	2.01	10.30	48.36	0.71	0.06
	Max.	0.93	1.12	54.01	89.70	97.99	99.23	24.93	27.12
III	Mean	0.84	1.17	2.89	41.68	58.32	86.72	9.16	4.12
	Volatility	0.09	0.07	3.44	19.36	19.36	10.40	5.46	5.35
	Min.	0.49	1.06	0.09	0.04	0.92	33.09	0.18	0.02
	Max.	1.01	1.64	39.12	99.08	99.96	99.76	26.44	44.13
IV	Mean	0.91	1.12	3.14	40.96	59.04	85.09	10.53	4.38
	Volatility	0.03	0.02	3.46	17.83	17.83	9.05	4.95	4.44
	Min.	0.87	1.08	0.13	1.03	13.91	40.22	1.90	0.16
	Max.	0.99	1.21	40.57	86.09	98.97	97.72	26.41	34.78

Table A.III
Model HML_{FX} Factor Risk Premium Decomposition: Equal-weighted portfolios

This table uses the calibrated pricing kernel model to compute and decompose the HML_{FX} factor risk premium from the perspective of a U.S. dollar investor. We compute the model-implied HML_{FX} factor risk premium by constructing a hypothetical, dollar-neutral factor mimicking portfolio on the basis of the calibrated time series of the global factor loadings, ξ_t^i. At each point in time, we sort the G10 currencies – excluding the U.S. dollar – into long and short portfolios on the basis of their prevailing loadings, and equal-weight the currencies within each portfolio. The model risk premia are computed daily and span from January 1999 to June 2012 ($T = 3520$ days). We report the global factor loadings of the high and low interest rate portfolios (ξ^H and ξ^L), as well as, the mean portfolio risk premium (λ^{HML}; *% per annum*). We then decompose the portfolio risk premium into contributions from the Gaussian and non-Gaussian components of the global factor innovation, $L^g_{Z_t}$. Finally, we evaluate the mechanism through which jumps contribute to equilibrium currency risk premia by decomposing risk premia across the moments of the pricing kernel innovations (variance, skewness, etc.). For each quantity, we report its time-series mean, volatility, and range. Results are reported for each of the four specifications reported in Table A.I.

		Loadings		Risk premium	By Shock Type		By Moment		
Specification		ξ^H	ξ^L	λ^{HML}	Gaussian	Non-Gaussian	Variance	Skewness	Other
I	Mean	0.87	1.04	2.66	42.28	57.72	85.78	9.32	4.90
	Volatility	0.04	0.03	3.16	18.18	18.18	10.54	5.13	5.88
	Min.	0.62	0.96	0.11	0.01	0.18	39.08	0.01	0.00
	Max.	0.92	1.34	37.48	99.82	99.99	99.99	26.00	36.10
II	Mean	0.88	1.04	2.82	40.41	59.59	85.68	9.94	4.38
	Volatility	0.00	0.00	3.53	17.17	17.17	9.33	4.73	5.01
	Min.	0.88	1.04	0.13	0.01	5.08	37.90	0.08	0.00
	Max.	0.88	1.04	42.18	94.92	99.99	99.92	25.89	37.24
III	Mean	0.88	1.09	1.95	42.12	57.88	87.64	8.79	3.57
	Volatility	0.05	0.05	2.20	20.38	20.38	10.28	5.50	5.31
	Min.	0.66	1.00	0.04	0.00	0.16	25.49	0.02	0.00
	Max.	0.95	1.45	25.60	99.84	100.00	99.97	26.59	54.47
IV	Mean	0.93	1.07	2.12	41.26	58.74	85.70	10.24	4.06
	Volatility	0.00	0.00	2.32	18.83	18.83	9.50	5.16	4.76
	Min.	0.93	1.07	0.07	0.00	5.84	30.50	0.67	0.06
	Max.	0.93	1.07	29.20	94.16	100.00	99.25	26.67	45.74

Table A.IV
Option-Implied Currency Risk Premia: Empirical factor mimicking portfolio returns
(Specification II; constant global factor loadings)

The table compares the historical (realized) returns to factor mimicking portfolios, with the corresponding model-implied risk premia. The data span the period from January 1999 to June 2012 ($N = 162$ months). Realized returns are computed on the basis of monthly buy-and-hold returns. The conditional HML_{FX} factor mimicking portfolios is a dollar-neutral portfolio formed by sorting currencies into long and short portfolios on the basis of their prevailing one-month LIBOR rates. Within each portfolio currencies are spread weighted, on the basis of the distance of their respective interest rates to the mean of the interest rates in countries with ranks five and six. The unconditional HML_{FX} portfolios is formed on the basis of the historical interest rate differentials computed using an expanding window starting in January 1990. The conditional short USD factor mimicking goes long (short) the U.S. dollar when the prevailing U.S. dollar interest rate is above (below) the prevailing average of the nine other G10 interest rates. The unconditional short USD factor mimicking portfolio is long an equal-weighted basket of G10 currencies against the U.S. dollar. Model risk premia are computed on the basis of Specification II. We report the moments of various quantities of interest, the p-values of the Jarque-Bera test for the realized portfolio returns, and comparisons of the mean realized and model-implied risk premia (t-stats in square brackets).

Panel A: HML_{FX} Factor Mimicking Portfolio

		Conditional	Unconditional	Difference
Realized	Mean	4.96	3.32	1.63
		[1.92]	[1.32]	[1.99]
	Volatility	9.51	9.26	
	Skewness	-1.07	-0.89	
	Kurtosis	7.03	6.90	
	JB (returns)	0.00	0.00	
	JB (Z-scores)	0.00	0.00	
	Carry	4.57	3.69	0.88
	t-stat	[60.81]	[51.65]	[27.95]
Model	Total ($\bar{\lambda}$)	3.81	3.73	0.08
		[8.97]	[8.54]	[1.74]
	HML (global, ($\bar{\lambda}^{HML}$)	3.81	3.73	0.08
		[8.97]	[8.54]	[1.74]
	Short USD ($\bar{\lambda}^{USD}$)	0.00	0.00	0.00
		-	-	-
	Portfolio ξ^H	0.85	0.88	
	Portfolio ξ^L	1.06	1.08	
Difference	Mean	1.14	-0.41	1.55
	t-stat	[0.43]	[-0.16]	[1.89]

Panel B: Short USD Factor Mimicking Portfolio

		Conditional	Unconditional	Difference
Realized	Mean	4.93	3.12	1.81
		[2.05]	[1.29]	[0.78]
	Volatility	8.84	8.90	8.56
	Skewness	-0.45	-0.17	-0.62
	Kurtosis	3.96	3.72	8.07
	JB (returns)	0.01	0.09	0.00
	JB (Z-scores)	0.00	0.00	0.00
Model	Total ($\bar{\lambda}$)	0.84	1.59	-0.74
		[5.35]	[13.59]	[-8.30]
	HML (global, ($\bar{\lambda}^{HML}$)	0.46	0.81	-0.35
		[4.47]	[9.21]	[-6.81]
	Short USD ($\bar{\lambda}^{USD}$)	0.39	0.78	-0.39
		[3.76]	[22.96]	[-8.88]
	Portfolio ξ^H	0.97	0.95	
	Portfolio ξ^L	0.98	1.00	
Difference	Mean	4.09	1.54	2.56
		[1.70]	[0.63]	[1.09]

Table A.V
Calibrated Model Parameters: Linking ξ_t^i to Y_t^i

This table reports summary statistics for the calibrated parameters of the pricing kernel factor model. Results are reported for four specifications which either: (a) match the prices of the high/low interest rate cross pairs and X/USD options (HLX + X/USD option set; 24 pairs) or the full panel of G10 cross rates (X/Y option set; 45 pairs); and, (b) allow the global factors loadings, ξ_i, to be time-varying or not. All calibrations are performed using daily option quotes from January 1999 to June 2012 ($T = 3520$ days) at five individual option strikes (10δ put, 25δ put, at-the-money, 25δ call, 10δ call). In specifications with time-varying global factor loadings, the loadings are parameterized on the basis of the one-month interest rate differential, $\xi_t^i = \xi^i - \tilde{\Psi}_t \cdot (Y_t^i - Y_t^{US})$. The slope coefficient, $\tilde{\Psi}_t$, is time-varying, but common to all currencies in the cross-section. Panel A reports the values of the country-specific loadings, ξ^i, obtained from the first stage of the calibration, with standard errors in parentheses. We also report the mean one-month LIBOR interest rate differential for each country relative to the U.S.. Panel B summarizes the cross-sectional dependence of the global factor loadings on interest rates differentials, Ψ_t, and the characteristics of the global factor innovation, $L_{Z_t}^g$. We report the share of variance in the global innovations due to jumps, η_t^g, and estimates of the parameters of the global CGMY jump component, G_t^g (dampening coefficient) and Y_t^g (power coefficient). Finally, we compute the skewness (Skewness$_t^g$) and kurtosis (Kurtosis$_t^g$) of the global factor innovation induced by variation in the global factor, Z_t. We obtain each of these quantities day-by-day by minimizing option pricing errors for the target option set, and report their time series means and volatilities. Panel C reports the root mean squared option pricing error measured in volatility points. We compute root mean squared errors (RMSE) for each pair in the target option set, and report their mean pooled: (a) across all strikes and currency pairs; (b) across all pairs with a given option strike.

Specification	ξ_t^i Parametrization	Option Set
I	$\xi^i - \Psi_t \cdot (r_{t,t+1}^i - r_{t,t+1}^{US})$	HLX and X/USD pairs
I'	$\xi^i - \tilde{\Psi}_t \cdot (Y_t^i - Y_t^{US})$	HLX and X/USD pairs

Panel A: Global Factor Loadings, ξ^i

Specification	AUD	CAD	CHF	EUR	GBP	JPY	NOK	NZD	SEK	USD
I and I'	0.79	1.00	0.96	1.00	0.93	1.17	0.92	0.81	1.00	1.00
	(0.0101)	(0.0000)	(0.0033)	(0.0001)	(0.0044)	(0.0072)	(0.0054)	(0.0084)	(0.0001)	-
$\overline{r_{t,t+1}^i - r_{t,t+1}^{US}}$ [%]	2.46	0.20	-1.58	-0.15	1.06	-2.54	1.53	2.78	0.07	0.00

Panel B: Global Factor Summary Statistics

Specification		Ψ_t	η_t^g	G_t^g	Y_t^g	Skewness$_t^g$	Kurtosis$_t^g$
I	Mean	0.42	0.51	7.91	-0.75	-1.10	6.03
	Volatility	1.89	0.19	5.60	0.16	1.35	4.04
I'	Mean	-0.37	0.54	7.60	-0.74	-1.13	6.19
	Volatility	11.94	0.20	4.86	0.17	1.35	4.20

Panel C: Option Pricing RMSE (Volatility Points)

Specification	All	$10\delta_p$	$25\delta_p$	50δ	$25\delta_c$	$10\delta_c$
I	1.12	1.04	1.27	1.27	1.16	0.85
I'	1.05	1.01	1.17	1.19	1.05	0.83

Panel D: HML_{FX} Factor Mimicking Portfolio

		Conditional	Unconditional	Difference
Realized	Mean	4.96	3.32	1.63
		[1.92]	[1.32]	[1.99]
Model I'	Total ($\bar{\lambda}$)	3.82	3.76	0.06
		[8.37]	[8.03]	[1.19]
Difference	Mean	1.14	-0.43	1.57
	t-stat	[0.42]	[-0.17]	[1.90]

Figure 1. Global Factor Loadings and State Variable Dynamics. The top panel plots the time series of global factor loadings, ξ_t^i, for two high interest rate currencies (AUD, NOK), two low interest rate currencies (CHF, JPY), and the U.S. dollar. The bottom panel plots the dynamics of the global state variable, Z_t, from the preferred model specification. Overlaid is the time-series of the first principal component extracted from the time series of option-implied variance swap rates for the 24 G10 currency pairs used in the calibration. This set includes: (a) all X/USD currency pairs (9 pairs); and, (2) cross-pairs formed on the basis of currencies which had the highest or lowest interest rates in the G10 set at any point in our sample (15 pairs). The principal component is re-scaled to have the same mean and volatility as the global state variable. The plots report data spanning the period 1999:1-2012:6, and are based on the output of Specification I.

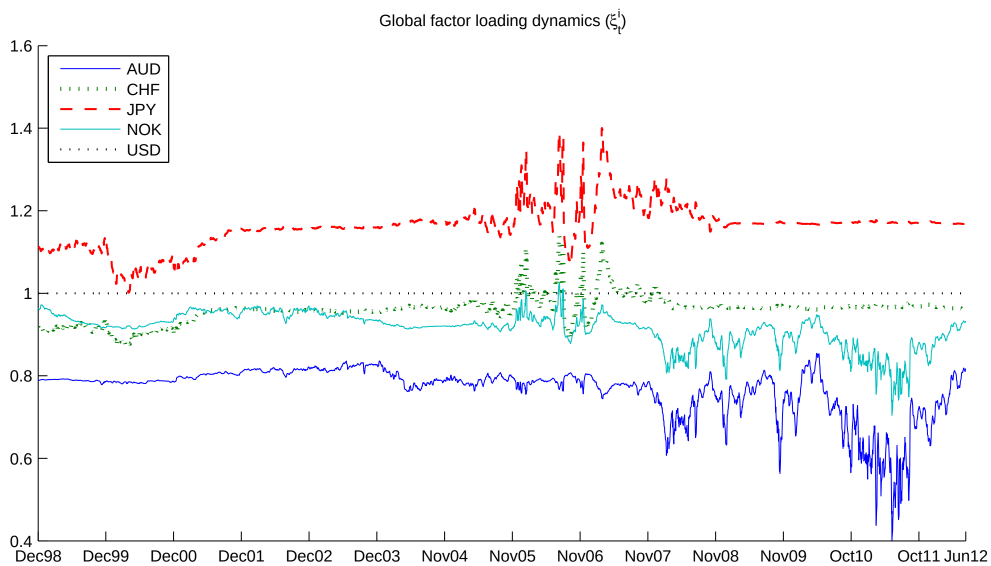

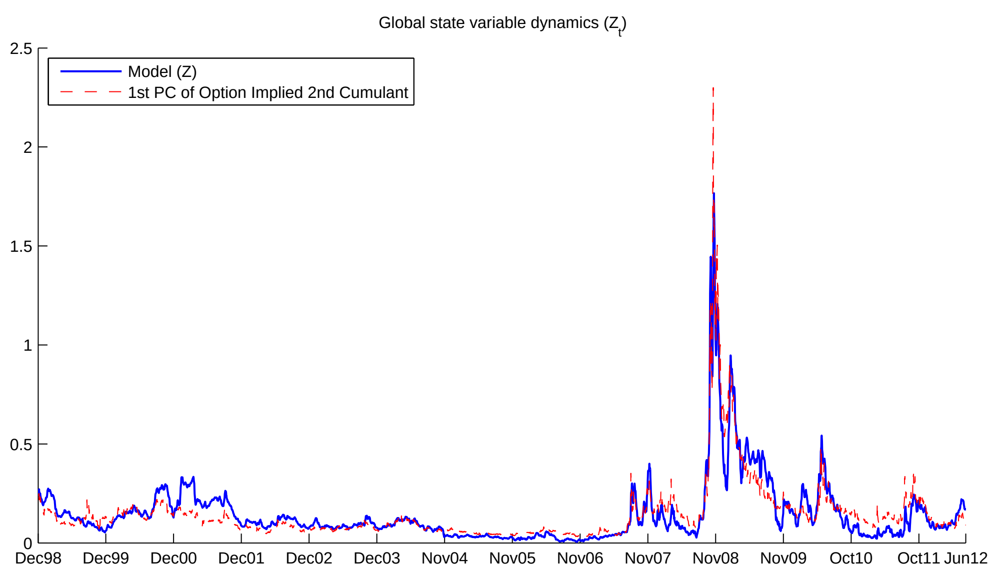

Figure 2. Cross-Sectional Relation between Calibrated Model Loading Differentials and Data Quantities. This figure plots the unconditional pair-level relation between the time-series mean of the calibrated global factor loading differential and the time-series mean of the corresponding one-month interest rate differential (top left), realized currency excess return (top right), option-implied volatility (bottom left), and the option-implied skewness (bottom right). The option-implied moments are computed on the basis of implied volatility functions which have been interpolated using the method of Castagna and Mercurio (2007), and extrapolated beyond the lowest and highest quoted strikes (10δ) by appending flat tails. The results are plotted for the 24 G10 currency pairs used to calibrate the model (Specification I). This set includes: (a) all X/USD currency pairs (9 pairs); and, (2) cross-pairs formed on the basis of currencies which had the highest or lowest interest rates in the G10 set at any point in our sample (15 pairs). X/USD pairs are additionally denoted with red dots. The plots report data spanning the period 1999:1-2012:6.

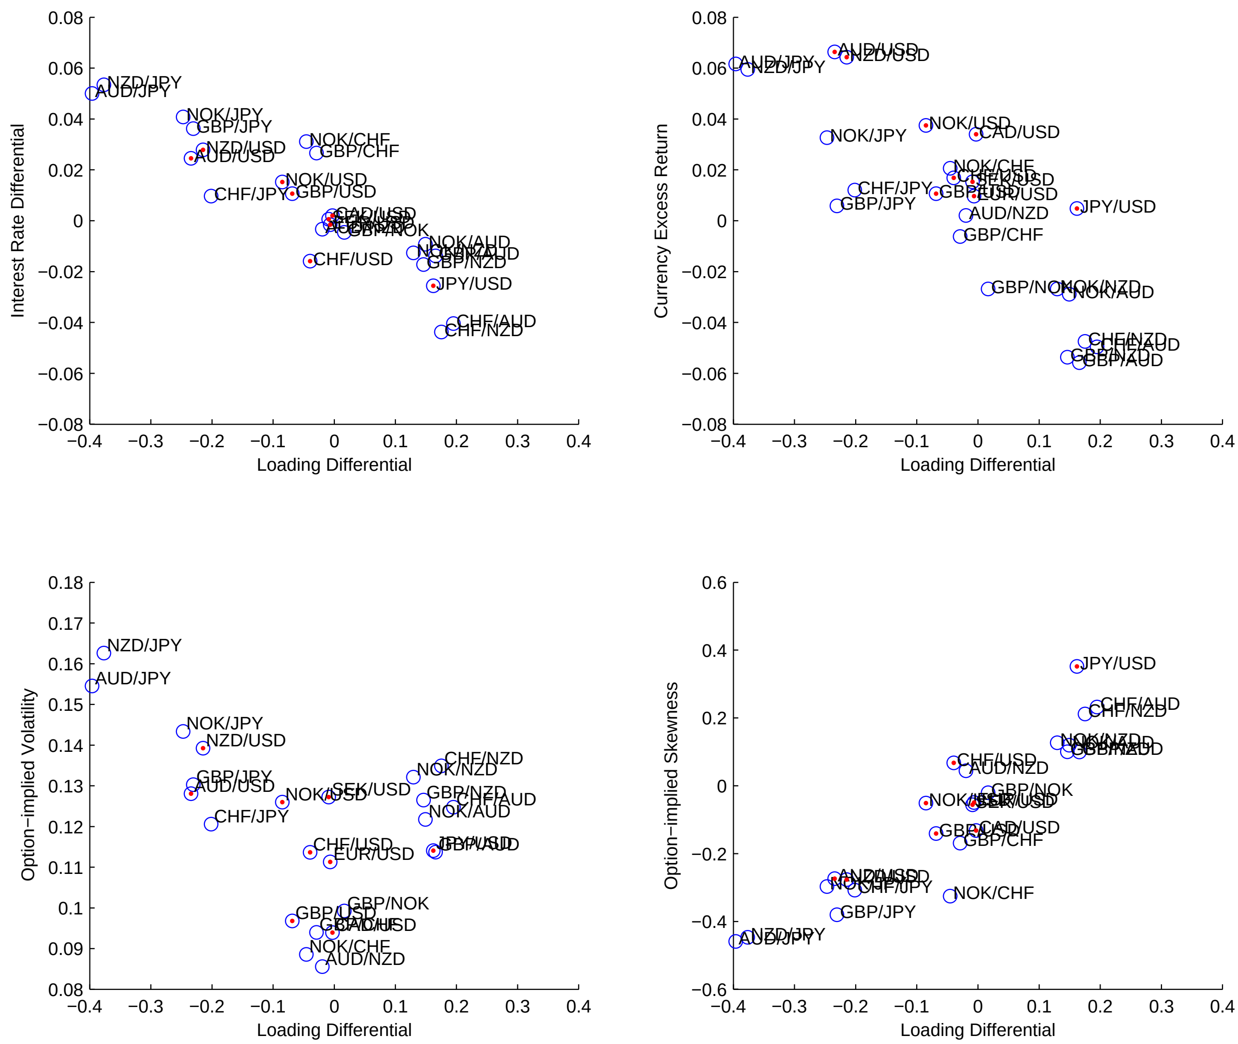

Figure 3. Fitted and Actual Option-Implied Volatilities. This figure plots the mean actual option-implied volatilities (blue) and their fitted counterparts (dashed red) for cross pairs formed by combining two high interest rate currencies (AUD, NOK) with two low interest rate currencies (CHF, JPY), as well as, pairs involving each of these currencies against the U.S. dollar. These pairs represent a subset of the 24 pairs used in the model calibration (Specification I) on which these plots are based. All values are annualized and reported in units of percent *per annum*. To illustrate the economic quality of the fit, the actual mean implied volatility is plotted with a typical bid-ask spread (dashed blue lines), equal to 0.1 times the implied volatility at that strike; the mean at-the-money volatility is reported in the title of each subplot. For each currency pair we plot option-implied volatilities at the five quoted strikes (10δ put, 25δ put, at-the-money, 25δ call, 10δ call). The plot reports time-series means computed using daily data for the period 1999:1-2012:6.

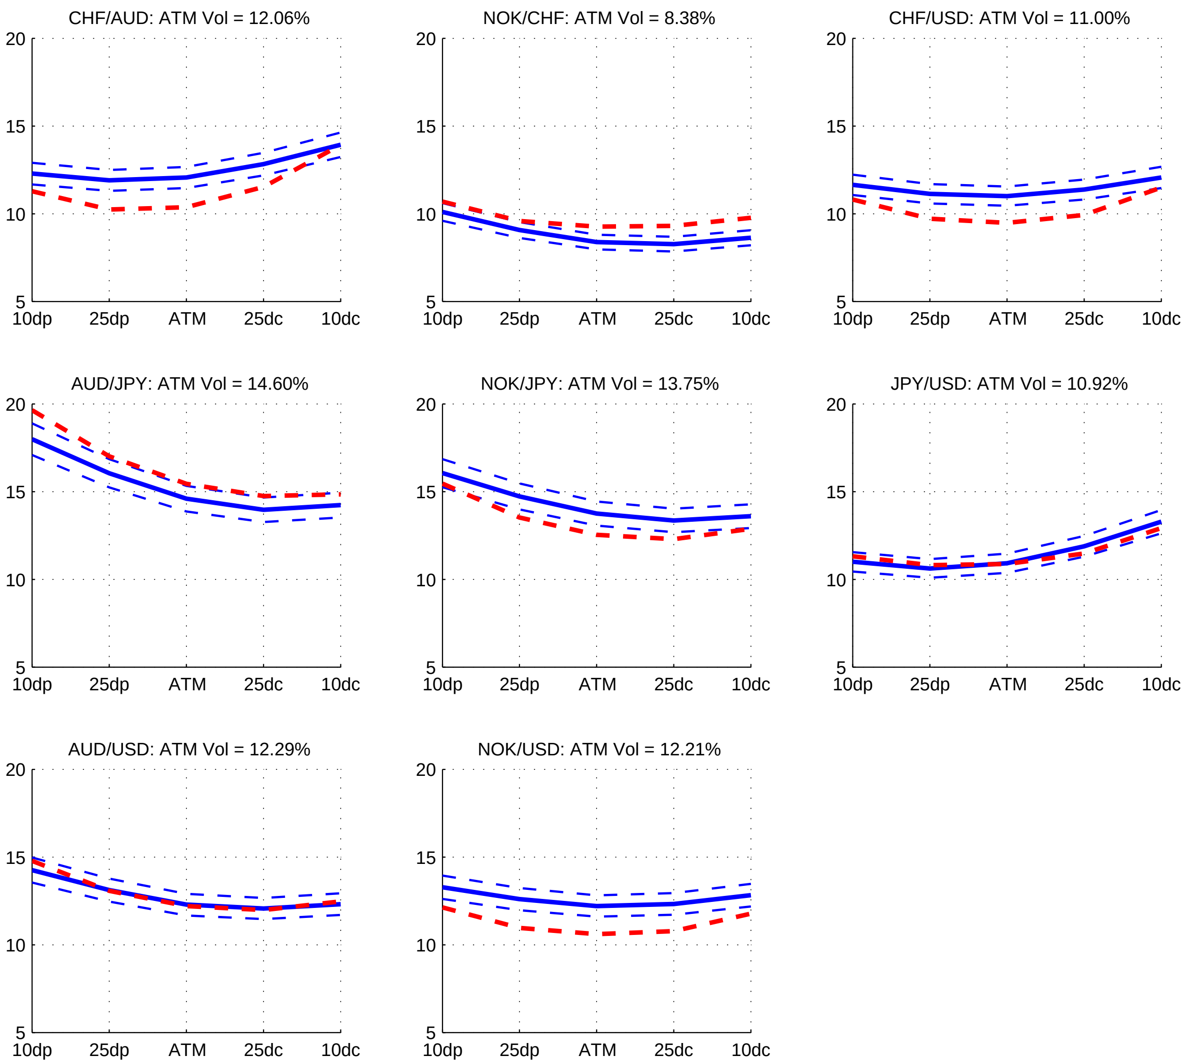

Figure 4. Returns to Empirical Factor Mimicking Portfolios. This figure plots the returns to empirical factor mimicking portfolios for the HML_{FX} (left panels) and short dollar (right panels) factors, constructed using G10 currencies, over the period from January 1999 to June 2012. We report results for conditional and unconditional factor mimicking portfolios. The conditional factor mimicking portfolios are formed by sorting currencies on the basis of the prevailing one-month interest rates; unconditional factor mimicking portfolios sort currencies on the basis of a backward looking average of one-month interest rates. The HML_{FX} factor replicating portfolios are spread-weighted, and dollar-neutral; the short dollar factor replicating portfolios are equal-weighted. The top two panels plot the cumulative returns from investing \$1 in each of the factor replicating portfolios. The bottom panels plot estimates of the instantaneous portfolio return volatility based on an EGARCH(1, 1) model, and – on the basis of a variance-covariance matrix imputed from the contemporaneous cross-section of option-implied variance swap rates for G10 cross-rates ('FX Option-Implied').

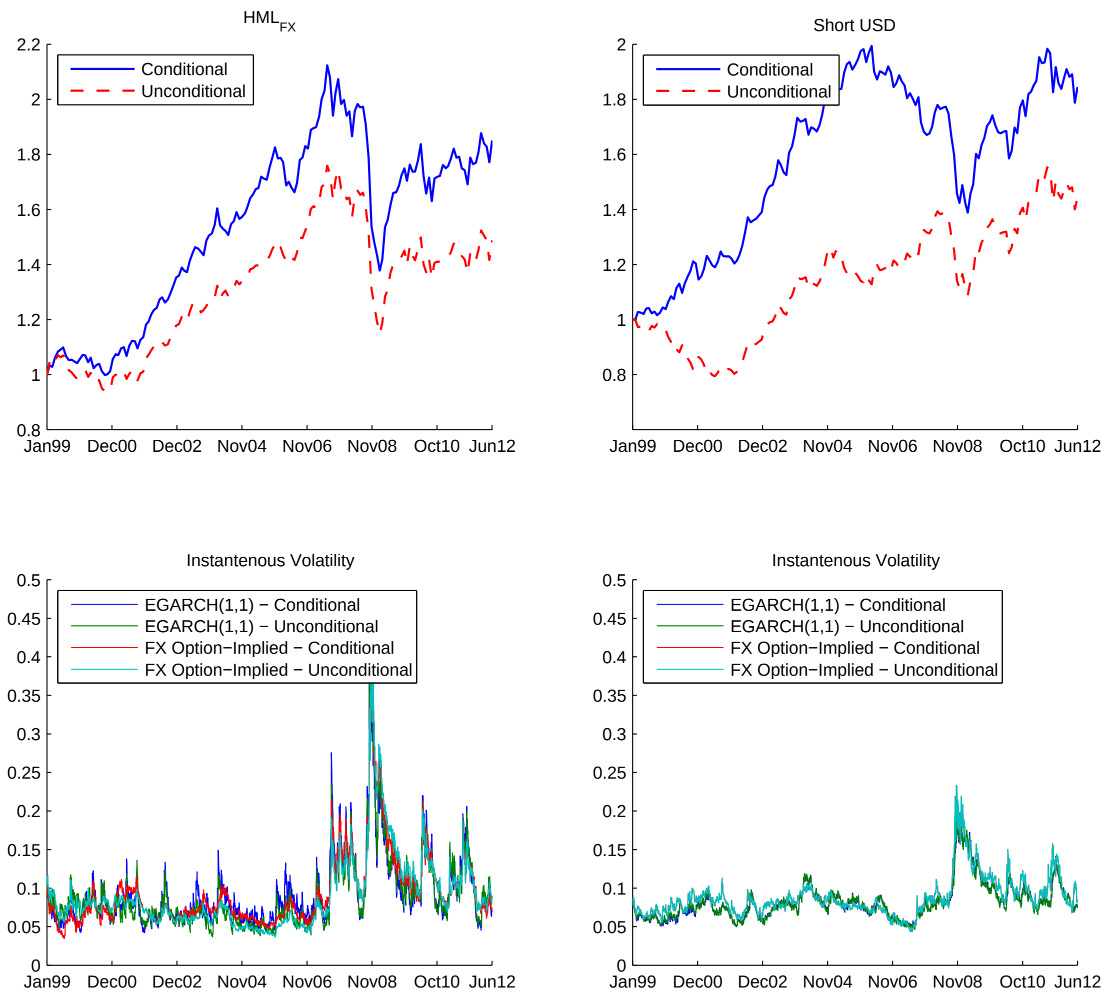

Figure 5. Option-Implied Currency Risk Premia: Conditional carry trade portfolio. This figure illustrates the model-implied currency risk premium decomposition for the conditional G10 currency carry trade portfolio, which is spread-weighted and dollar-neutral. The top panel plots the time series of the total, model-implied risk premium, and the contribution from exposure to the global currency risk factor (HML_{FX}). Since the empirical factor mimicking portfolio is constructed to be dollar-neutral, the entirety of the model risk premium reflects compensation for exposure to the global factor. The middle panel decomposes the model-implied portfolio risk premium into compensation for Gaussian and non-Gaussian innovations. The panel plots the share of the risk premium due to each component; the mean share due to each component is reported in the title of the respective subplots. The bottom panel decomposes the portfolio risk premium into contributions from the variance, skewness, and higher moments (other) of the global pricing kernel shock, $L^g_{Z_t}$. The plots report data spanning the period 1999:1-2012:6, and are based on the output of Specification I.

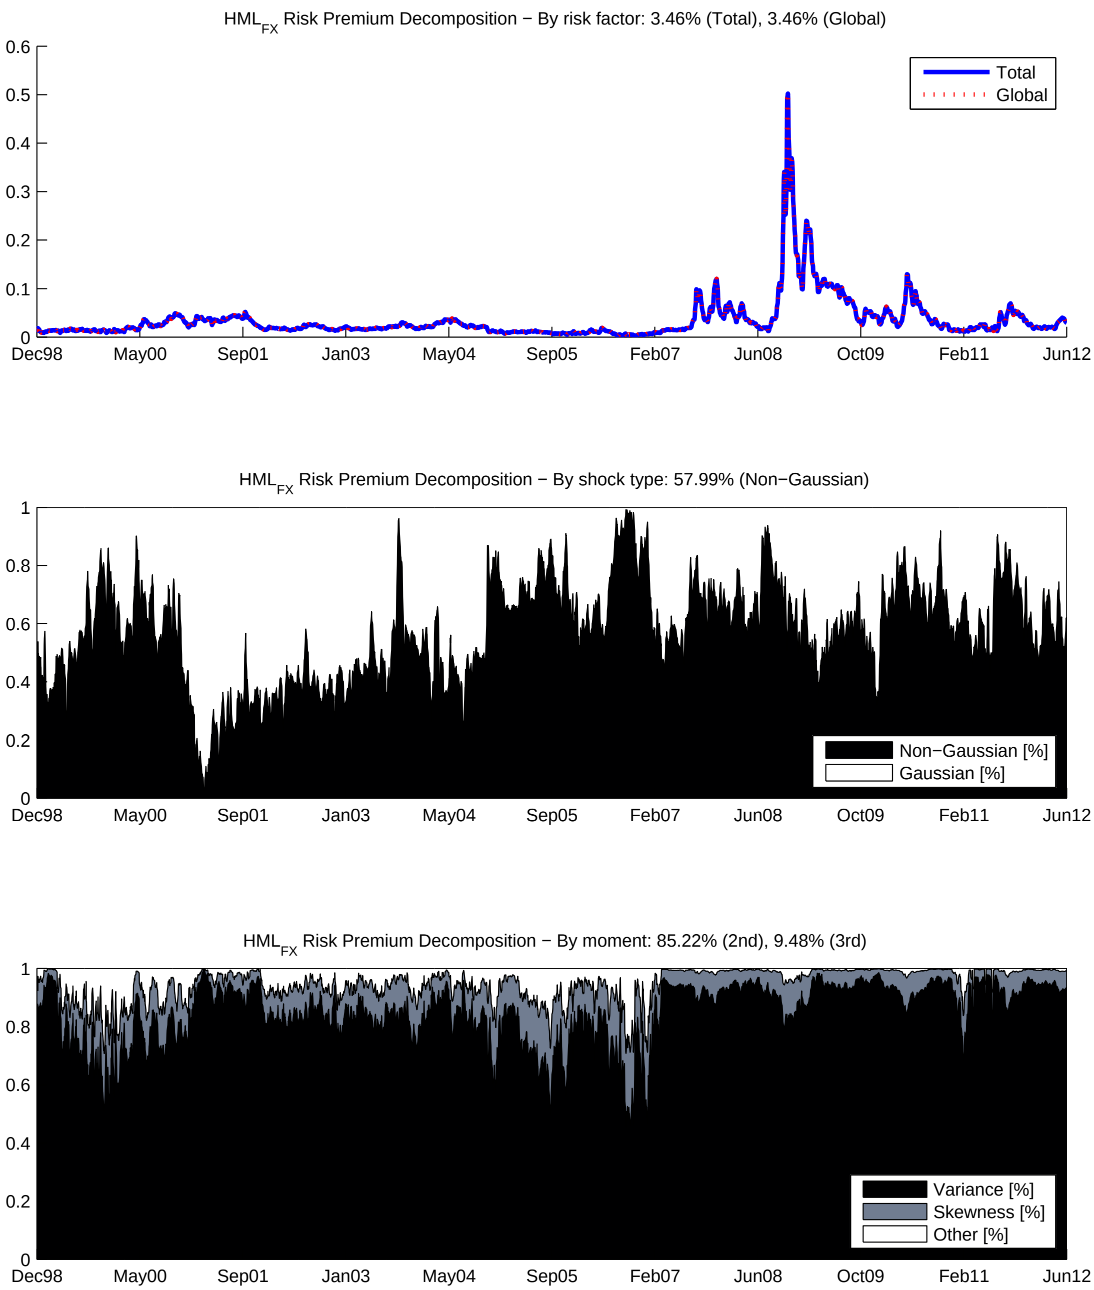

Figure 6. Option-Implied Currency Risk Premia: Unconditional short U.S. dollar portfolio. This figure illustrates the model-implied currency risk premium decomposition for the unconditional short dollar portfolio, which is short the U.S. dollar against an equal-weighted portfolio of G10 currencies. The top panel plots the time series of the total, model-implied risk premium, and the contribution from exposure to the global currency risk factor (HML_{FX}). The middle and bottom panels decompose the component of the portfolio risk premium not attributable to exposure to the global factor. The middle panel decomposes the risk premium into compensation for Gaussian and non-Gaussian innovations. The panel plots the share of the risk premium due to each component; the mean share due to each component is reported in the title of the respective subplots. The bottom panel decomposes the short dollar risk premium into contributions from the variance, skewness, and higher moments of the country-specific U.S. pricing kernel innovation, $L^{US}_{Y^{US}_t}$. The plots report data spanning the period 1999:1-2012:6, and are based on the output of Specification I.

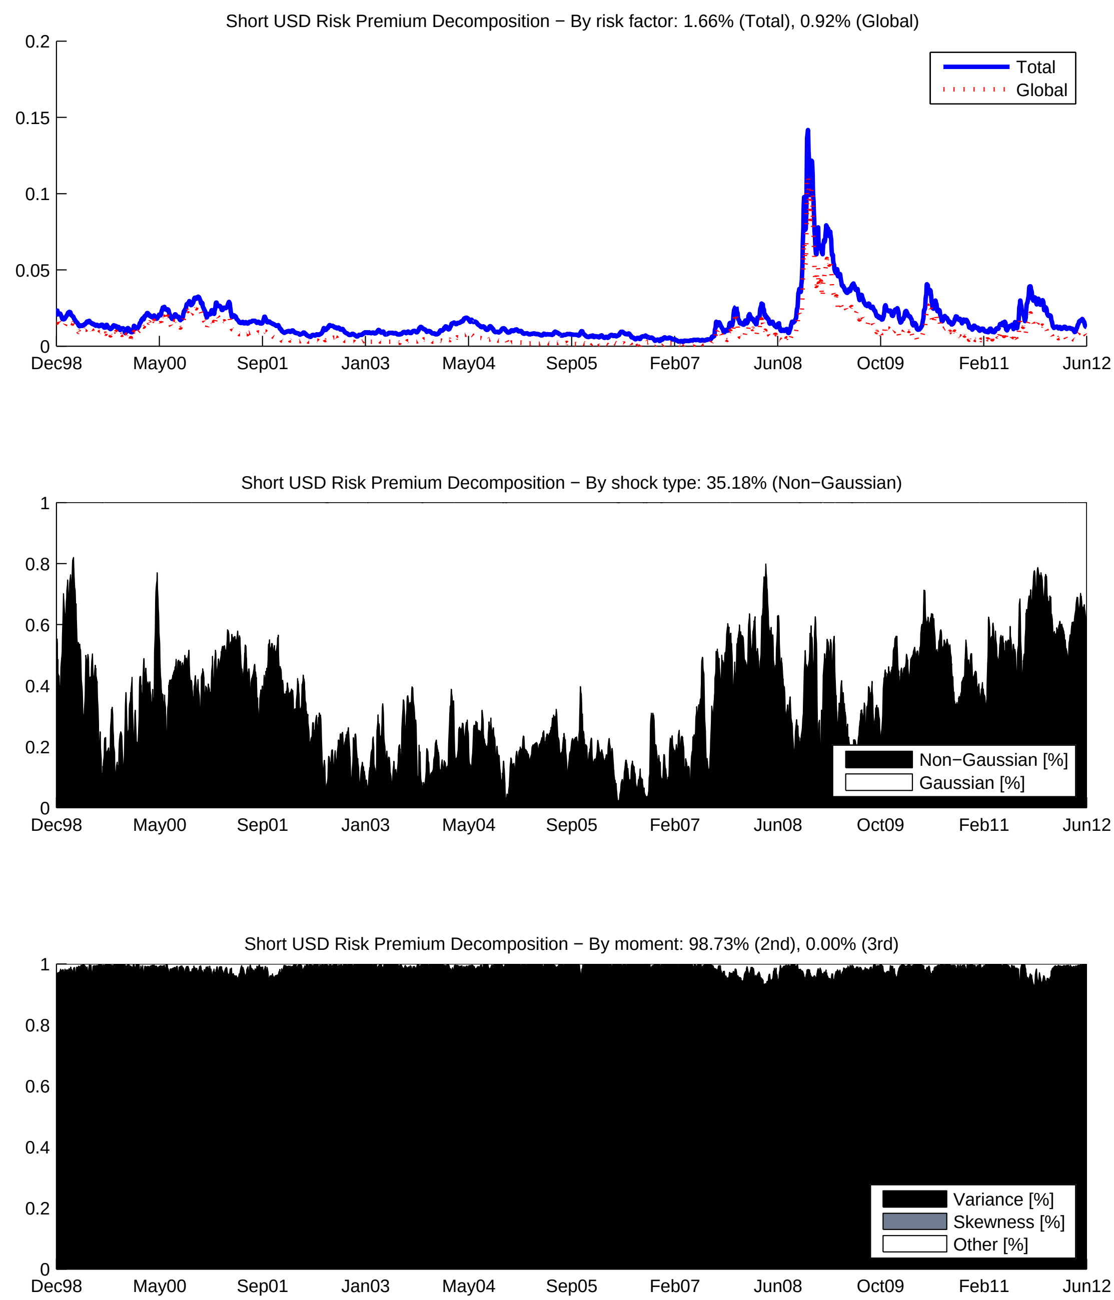

Figure 7. The Cross-Section of Model and Realized Currency Risk Premia. The left panel plots the cross-sectional relation between the mean realized currency pair excess return, and the mean option-implied currency pair risk premium obtained from the calibrated model. The right panel plots the cross-sectional relation between the mean realized currency pair excess return and the mean one-month interest rate differential. Each subplot is equipped with a 45-degree line with an intercept of zero. In the left panel, this line corresponds to the hypothesis that the calibrated, option-implied risk premia are an unbiased predictor of currency excess returns. In the right panel, this line corresponds to a random walk model of exchange rate dynamics. The results are plotted for the 24 G10 currency pairs used to calibrate the model (Specification I). This set includes: (a) all X/USD currency pairs (9 pairs); and, (2) cross-pairs formed on the basis of currencies which had the highest or lowest interest rates in the G10 set at any point in our sample (15 pairs). X/USD pairs are additionally denoted with red dots. The plots report data spanning the period 1999:1-2012:6.