Crash Risk in Currency Returns*

Mikhail Chernov
UCLA and CEPR†

Jeremy Graveline
University of Minnesota‡

Irina Zviadadze
Stockholm School of Economics§

First Draft: July 2010
This Revision: October 13, 2013

Abstract

We quantify crash risk in currency returns. To accomplish this task, we develop and estimate an empirical model of exchange rate dynamics using daily data for four currencies relative to the US dollar: the Australian dollar, the British pound, the Swiss franc, and the Japanese yen. The model includes (i) normal shocks with stochastic variance, (ii) jumps up and down in the exchange rate, and (iii) jumps in the variance. We identify these components using data on exchange rates and at-the-money implied variances. We find that crash risk is time-varying. The probability of an upward (downward) jump in the exchange rate, associated with depreciation (appreciation) of the US dollar, is increasing in the domestic (foreign) interest rate. The probability of a jump in variance is increasing in the variance but is not related to interest rates. Many of the jumps in exchange rates are associated with macroeconomic and political news, but jumps in variance are not. On average, jumps account for 25% of total currency risk (and can be as high as 40%), as measured by the entropy of exchange rate changes, over horizons of one to three months. The dollar carry index, which is based on 21 exchange rates, retains these features. A simple calibration analysis using option-implied smiles suggests that jump risk is priced.

---

*We are grateful to many people who gave us advice, including David Backus, Magnus Dahlquist, Itamar Drechsler, Bjorn Eraker, Peter Feldhutter, Diego Garcia, Francisco Gomes, Michael Johannes, Igor Makarov, Alessandro Palandri, Guillaume Plantin, Nick Roussanov, Lucio Sarno, Paul Schneider, Ivan Shaliastovich, Elvira Sojli, Vladimir Sokolov, Zhaogang Song, Stijn Van Nieuwerburgh, Adrien Verdelhan, Christian Wagner, Amir Yaron, Stan Zin, as well as participants in seminars at, and conferences sponsored by, Chicago, EDHEC, the Federal Reserve, Goldman Sachs, Imperial College, LBS, LSE, Montreal IFSID and Bank of Canada, Moscow HSE, the New York Fed, NYU, SED, SoFiE, the Swedish House of Finance, Tinbergen Institute, UBC, UCLA, UNC, Wharton. Zviadadze acknowledges support by the AXA research fund.

†Anderson School of Management, 110 Westwood Plaza, Los Angeles, CA 90095, USA, Email: mikhail.chernov@anderson.ucla.edu, Web: sites.google.com/site/mbchernov

‡Department of Finance, 3-122 CSOM, 321 19th Ave S, Minneapolis, MN, USA, Email: jeremy@umn.edu, Web: umn.edu/~jeremy

§Department of Finance, Drottninggatan 98, 111 60 Stockholm, Sweden, Email: Irina.Zviadadze@hhs.se, Web: sites.google.com/site/irinazviadadzessite

1 Introduction

The time variation and high magnitude of returns to currency speculation have attracted a lot of recent attention. Much of the literature has focused on measuring risk premiums, or expected excess returns, in this market (e.g., Lustig and Verdelhan, 2007). However, expected returns alone do not tell the whole story. Investors also care about the risks that they must bear to earn these returns. Therefore, the distribution of risks is an important ingredient in understanding currency premiums.

A number of papers have suggested that investors in currency markets require high returns on average as a compensation for crash risk. The merit of this explanation hinges on the magnitude and probability of large moves in currency markets. Our paper is the first empirical study that systematically quantifies crashes, documents when they occurred historically and what their determinants were.

We develop and estimate a model that incorporates both crash risk and regular, or normal, risk. The model can accommodate information from both exchange rates and option-implied volatilities. Our objective is to use the model to establish whether a large move in the exchange rate takes place because of a crash, or because the conditional variance of a normal shock is high. Implied volatilities are particularly helpful for this task because they provide information about the market's view on conditional variance.

We establish the relative importance of three key modelling elements. First, it is well-documented that currency returns are heteroscedastic (e.g., Baillie and Bollerslev, 1989; Engel and Hamilton, 1990; Engle, Ito, and Lin, 1990; Jorion, 1988; Harvey and Huang, 1991). Casual observation of time-series variation in option-implied exchange rate volatility also confirms this point. We capture this feature of the data with a standard stochastic volatility component in our model.

Second, there is also strong empirical evidence that daily changes in exchange rates are not conditionally Gaussian (as would approximately be the case in a model with only stochastic volatility). To account for this feature of the data, our model includes jump risks in exchange rates.[1] We allow the probability of these jumps to be time-varying, in order to capture the variation in conditional skewness that has been previously documented (e.g., Bakshi, Carr, and Wu, 2008; Brunnermeier, Nagel, and Pedersen, 2008; Carr and Wu, 2007; Johnson, 2002).

Third, changes in the at-the-money implied volatility of a typical exchange rate exhibit unconditional skewness of 1 and kurtosis of 10 or more. To accommodate this property, our model allows for jumps in the variance of Gaussian shocks to exchange rates. The

[1] In this paper, we use "crash" and "jump" interchangeably. Because we are only studying asset prices, not the economy, we cannot say anything about economic disasters. Despite this clarification we think that our findings could be useful for researchers interested in modelling the economy, see, e.g., Backus, Chernov, and Martin (2011) for an example of how asset prices could be used to infer behaviour of consumption.

importance of such jumps for modelling equity returns has been emphasized in Broadie, Chernov, and Johannes (2007); Duffie, Pan, and Singleton (2000); Eraker, Johannes, and Polson (2003), among others. To our knowledge, our paper is the first to investigate the role of jumps in the volatility of exchange rates.

A jump in an exchange rate is qualitatively different from a jump in its variance. Almost by definition, large jumps are rare events. Therefore, when there is a direct jump in the exchange rate, one doesn't necessarily expect there to be many subsequent jumps in the near future. By contrast, when there is a jump in the variance of the Gaussian shock to an exchange rate, one expects there to be many large subsequent moves in the exchange rate because of the higher level of variance. We use our model and empirical analysis to determine whether these qualitative distinctions lead to materially quantitative differences.

We use daily joint data for exchange rates and implied volatilities from 1986 to 2010 (the options data start in 1994) on four spot exchange rates: Australian dollar, Swiss franc, British pound, and Japanese yen. Crashes are rare and volatility is persistent, so it is important to use a long time span of data. Short samples are likely to either over- or under-represent jumps and periods of high or low volatility leading to biased estimates of the required probabilities.

We rely on Bayesian MCMC to estimate candidate models. One of the key advantages of this approach is that it provides estimates of the conditional distribution of currency returns, as well as estimates of the realized shocks. This feature allows us to link large shocks, or jumps, to important macro-finance events and thereby illuminate the potential economic channels that are responsible for crash risk in currencies.

Our statistical tests strongly favor both jumps in exchange rates and in their variances. This conclusion is similar to the one in the equity literature. However, the similarity ends there. In contrast to equity models that rely on one Poisson (counting) process controlling the arrival rate of all jumps, we find three such processes in FX. The three types of jumps arise via different mechanisms. The arrival rate of a jump in the variance of currency returns is positively related to the variance itself. Thus, this component belongs to the class of self-exciting processes. The probability of a jump up in the exchange rate, which corresponds to a depreciation of the US dollar, is positively related to the domestic (US) interest rate. The probability of a jump down, which corresponds to an appreciation of the US dollar, is positively related to the foreign interest rate.

Although jumps in currencies and in variance are alternative channels for large currency returns, we find that economically they are quite distinct. We can connect most of the jumps in FX to important macro or political announcements. In contrast, jumps in variance cluster at the moments of high uncertainty in markets, which are captured by comments on current events, political speculation and overall anxiety about upcoming events.

We use entropy (a generalized measure of variance) of changes in an exchange rate to measure the amount of risk associated with currency positions and to decompose this risk into the contributions from different sources of shocks (Alvarez and Jermann, 2005; Backus, Chernov, and Martin, 2011). Appropriately scaled entropy is equal to the variance of an exchange rate return if it is normally distributed, but otherwise includes high-order cumulants. Therefore, entropy is a convenient measure that captures both normal and tail risk in one number. We find that, depending on the currency, the time-series average of the joint contribution of the three types of jumps can be as high as 25% of the total risk and on individual days this contribution can be up to 40%. Jumps in variance contribute about a third to the average contribution and can be as high as 15% of the total risk on individual days. Also, the contribution of jumps in variance to the total risk increases with investment horizon.

To test whether diversification eliminates or reduces the importance of these jumps, we estimate the dynamics of the 21-currency dollar index studied by Lustig, Roussanov, and Verdelhan (2013). While diversification reduces jump frequency, it also reduces the volatility of the normal risks. As a result, the relative importance of jumps stays unchanged. The dollar carry trade switches position of the dollar index from long to short depending on the difference between the US interest rate and the index of foreign interest rates. We find that this trade increases the number of instances of jumps in currencies that go against the speculator. Jumps in variance are not affected by the switching position.

Given the large contribution of jumps to the overall risk, it is natural to ask whether the jump risk is priced. The full answer to this question requires an explicit model of the pricing kernel and the use of assets, such as out-of-the-money options, that are particularly sensitive to jumps for estimation. While such analysis is outside of the scope of this paper, we carry out a limited option valuation exercise. We select representative implied volatility smiles and attempt to fit them by allowing premiums for normal risks only. We find that the resulting theoretical implied volatilities cannot generate the asymmetries present in the data.

Related Literature

We limit our discussion of related literature to papers that highlight the importance of jumps for understanding the properties of exchange rate returns. One exception is the work of Brandt and Santa-Clara (2002) and Graveline (2006). These papers are early antecedents of our paper in terms of methods and research questions. These authors also estimate a time-series model of exchange rates using the time-series of FX and implied variance. However, they do not allow for jumps.

Our paper is related to recent empirical papers that investigate whether high currency returns can be explained as compensation for jump, or crash, risk. Brunnermeier, Nagel, and Pedersen (2008) provide evidence consistent with the hypothesis that large exchange rate moves are related to funding constraints of speculators enagaged in carry trades. In particular, they relate the sign and magnitude of skewness of various exchange rates relative to the USD to those of the respective interest rate differentials. Jurek (2009) analyzes the returns on carry trade portfolios in which the exposure to currency crashes is hedged with options. He concludes that exposure to currency crashes account for 15% to 35% of the excess returns on unhedged carry trade portfolios. Burnside, Eichenbaum, Kleshchelski, and Rebelo (2011) investigate whether carry trade returns reflect a "peso problem" (i.e., a low probability event that did not occur in the sample). They use carry returns hedged with options to argue that any such peso event must be a modest negative return on the carry trade combined with an extremely large value of the stochastic discount factor (i.e., the marginal utility of a representative investor must be very high in the, as yet, unobserved peso state). Jordà and Taylor (2012) propose to manage the risk of carry positions by conditioning on macro information instead of options, but the resulting strategy still yields a very high Sharpe ratio. The common thread in these papers is that they provide indirect evidence on the magnitude of jump risk. Our paper aims to complement this previous work with a formal statistical model and analysis.

Our paper is also related to the option pricing literature, which has focused on modeling the risk-adjusted (risk-neutral) distribution of exchange rates. By construction, these papers do not consider risk premiums. However, the shock structures under the risk-adjusted and actual (true) distributions are usually modelled to be similar. In this respect, this work is complimentary to our analysis. Bates (1996) considers option prices on the Deutsche Mark and is the earliest paper that argues for the inclusion of jumps in currencies. He considers a single normally distributed jump in FX with a constant probability. Carr and Wu (2007) distinguish jumps up and down in FX and also allow for time-varying jump probabilities controlled by unobservable states. Bakshi, Carr, and Wu (2008) extend the Carr-Wu model to a triangle of currencies (GBP, JPY, and USD) and estimate it using 2.25 years of data on exchange rates and option prices. Our analysis provides additional economic intuition, as time variation in jump probabilities are driven by observable interest rates. None of these papers consider jumps in variance or estimate jump times and sizes. Farhi, Fraiberger, Gabaix, Ranciere, and Verdelhan (2009) do not have explicit time-varying states, but allow risk-adjusted parameters to change every period in a nonparametric fashion.

There is also an important literature that attempts to explain the behaviour of exchange rates in macro-founded equilibrium.[2] Our paper is silent about the prices of risk (with the exception of the limited option calibration exercise), but it has implications for how to best model the fundamentals in an equilibrium setting. Gourio, Siemer, and Verdelhan (2013), Guo (2007) and Ready, Roussanov, and Ward (2013) propose production-based models, where productivity is allowed to experience a disastrous decline. Farhi and Gabaix (2008) consider a pure exchange economy and a similar assumption of a disaster in consumption. Disasters are modelled as jumps down, and all papers, with the exception of Ready, Roussanov, and Ward (2013), allow unobservable time-varying processes to drive disaster probabilities. Exchange rates inherit these properties. Our results suggest that it may also

[2]Examples include, but not limited to Bekaert (1996); Backus, Gavazzoni, Telmer, and Zin (2010); Bansal and Shaliastovich (2012); Colacito (2009); Colacito and Croce (2013).

be important to allow for jumps in the volatility of these processes and for the process driving probability of jumps in consumption to be related to interest rates in equilibrium.

Our results also speak to the frictions-based equilibrium model of Plantin and Shin (2011). These authors focus on endogenously generated dynamics of a carry trade. A carry trade gets started in a high-liquidity environment, such as accommodative monetary policy. It is self-enforcing because of the speculators' belief that others will join the trade. The trade crashes when the speculators hit funding constraints. As a result, extended periods of slow appreciations of a high interest rate currency are randomly interrupted by endogenous crashes. Because our analysis is implemented at the daily frequency, we are able to capture, in reduced form, related phenomena.

2 Preliminaries

This section motivates our analysis and highlights properties of the data that our model is designed to capture.

2.1 Excess Returns

In this paper we always treat USD (i.e., U.S. dollars) as the domestic currency. Let r_t be the continuously-compounded USD interest rate, $\tilde{r}_t$ be the analogous interest rate in the foreign currency (e.g., GBP), and S_t be the exchange rate expressed as units of domestic currency per unit of foreign currency. Then one U.S. dollar buys $1/S_t$ units of foreign currency at time t, which grows at the foreign risk-free rate to $1/S_t \cdot \exp(\tilde{r}_t)$ units at time $t+1$, and can be exchanged for $S_{t+1}/S_t \cdot \exp(\tilde{r}_t)$ U.S. dollars. Therefore, the log excess dollar-denominated return to investing in the foreign currency (i.e., the return over and above investing at the USD risk-free rate) is

$$y_{t+1} = s_{t+1} - s_t + \tilde{r}_t - r_t\,,$$

where $s_t = \ln S_t$.

Figures 1 - 4 display the time series of log excess returns, y_{t+1} (panel (a)), and implied volatilities (panel (b)) for the currencies we consider in this paper. We have selected four currencies - Australian dollar (AUD), Swiss franc (CHF), British pound (GBP), and Japanese yen (JPY) based on the availability of daily data, and cross-sectional and time-series variation in the interest rate differential. We use one-month LIBOR to proxy for interest rates. Using one-month rather than overnight rates implicitly assumes a flat term structure at the very short end of the LIBOR curve and allows us to abstract from potential high-frequency idiosyncratic effects associated with fixed-income markets. Because we denominate exchange rates by the U.S. dollar, movements up correspond to a depreciation in the USD.

2.2 Properties of Excess Returns

In Table 1 we provide summary statistics of daily log excess returns and changes in the one-month at-the-money implied volatility. Means are close to zero at daily frequency. Therefore, these summary statistics inform us primarily about the properties of shocks.

All currencies have volatility of about 10% per year. There is evidence of substantial kurtosis (AUD and JPY are the most notable in this regard), which is suggestive of non-normalities. Skewness of all currencies is mild. It turns out that this is a manifestation of time-varying and sign-switching conditional skewness. We produce a rough estimate of conditional skewness by computing a six-month rolling window. The time-series of these estimates are displayed in panels (a) of Figures 1 - 4. Depending on the currency, conditional skewness ranges from -2 to 2. Thus, excess returns are not only fat-tailed, but also asymmetric with the degree of asymmetry changing over time.

The implied volatility is itself quite variable at about 60% per year (the number in the table multiplied by $\sqrt{252}$) and highly non-normal with skewness and kurtosis much higher than those of the currency returns themselves. The implied volatility from the short-dated options should be very close to the true volatility of exchange rates (which is unobservable) and therefore its properties provide insight into the features that a realistic model of variance must require.

As a reference, we report the same summary statistics for the S&P 500 whose risks have been thoroughly studied in the literature. The index returns are more volatile and exhibit much stronger departures from normality as compared to currencies. In particular, negative unconditional skewness is evident (in fact, a measure of conditional skewness rarely becomes positive). In contrast, changes in VIX, a cousin of implied variance, display weaker non-normalities than currencies. These statistics suggest that a model of currency risks could be substantively different from that of equity risks even though one clearly has to use similar building blocks.

2.3 Risks and Expected Excess Returns

We can generically represent excess returns as:

$$y_{t+1} = E_t(y_{t+1}) + \text{shocks}. \tag{2.1}$$

Most research has focused on conditional expected excess returns $E_t(y_{t+1})$. For example, if currencies do not carry a risk premium, then uncovered interest rate parity (UIP) holds and $E_t(y_{t+1}) = 0$. However, Bilson (1981), Fama (1984), and Tryon (1979) establish that the regression

$$s_{t+1} - s_t = a_1 + a_2(f_t - s_t) + \text{shocks}, \tag{2.2}$$

where f_t is is the log of the one-month forward exchange rate, typically yields estimates of a_2 of approximately -2. If covered interest rate parity (i.e., no-arbitrage) holds, then the log forward exchange rate is given by $f_t = s_t + r_t - \tilde{r}_t$, therefore this result is equivalent to:

$$y_{t+1} = a_1 + (a_2 - 1)(r_t - \tilde{r}_t) + \text{shocks}, \tag{2.3}$$

with a slope coefficient of about -3. Subsequent research has extended the specification of risk premiums $E_t(y_{t+1})$ (e.g., Beber, Breedon, and Buraschi, 2010; Bekaert and Hodrick, 1992; Lustig, Roussanov, and Verdelhan, 2011; Menkhoff, Sarno, Schmeling, and Schrimpf, 2012, among others).

Asset pricing theory relates expected excess returns to compensation for bearing risks, that is, it relates "$E_t\left(y_{t+1}\right)$" to the "shocks" in equation (2.1). In the language of pricing kernels, expected excess returns are determined by the covariation of currency risks with a pricing kernel. We don't empirically test any specific asset pricing theories (pricing kernels) in this paper, but a thorough analysis of the shocks is a necessary ingredient for full testing of any dynamic asset pricing model. To illustrate this point, we provide an example of two theories that can lead to identical expected excess returns despite the different shock structures (Appendix A.1). In this situation, the distribution of shocks is the only element that can distinguish one theory from the other.

To measure shocks, we need to model conditional means as well. We use a simple specification that encompasses the UIP regressions result by allowing for linear dependence on the domestic and foreign interest rates, and includes variance of FX returns as an extra variable.[3] Because we are working with daily returns, the magnitude of the drift term is much smaller than the higher order moments and so any omitted variables that might affect expected returns are not likely to introduce much bias in our results. Verdelhan (2011) provides direct evidence supporting this conjecture. As such, to avoid overfitting, we did not include any other variables in the drift of the exchange rate.[4]

While our focus is on careful modelling of currency risks themselves, our conclusions should have implications for modelling of economic channels leading to the observed risk premiums. As highlighted by our examples in Appendix A.1, successful equilibrium models should be able to replicate not only the measured risk premiums, but the distribution of currency shocks as well. To this end, our model can be used to construct portfolios that isolate jump risks and serve as inputs to traditional factor models that examine the pricing of these risks. Moreover, our extensive analysis of the shocks to currency returns provides useful guidance for specifying shocks to fundamentals in equilibrium models.

[3] This addition can be supported in various theoretical settings (Bacchetta and van Wincoop, 2006; Brennan and Xia, 2006). Empirical work with such a term includes Bekaert and Hodrick (1993), Bekaert (1995), Brandt and Santa-Clara (2002), Domowitz and Hakkio (1985), Graveline (2006), Lustig, Roussanov, and Verdelhan (2011), and Menkhoff, Sarno, Schmeling, and Schrimpf (2012).

[4] A recent literature suggests that inference about the conditional mean of excess returns can be improved by considering portfolios of currencies (e.g., Barroso and Santa-Clara, 2011; Lustig, Roussanov, and Verdelhan, 2011; Lustig and Verdelhan, 2007; Menkhoff, Sarno, Schmeling, and Schrimpf, 2012).

3 Empirical Model

We start by presenting our empirical model in Section 3.1. Section 3.2 discusses how we arrived at the assumed functional forms.

3.1 Currency Dynamics

In this paper we model each exchange rate in isolation from others. A large fraction of currency analysis, such as UIP regressions or equilibrium modelling is conducted on a currency-by-currency basis. This approach is able to identify the normal and non-normal shocks, and how they should be modelled. However, we cannot say which fraction of shocks can be explained by common variation in the exchange rates, and which fraction is country-pair specific. The important question of modelling the joint distribution of currency risks goes hand-in-hand with modelling of the pricing kernel and we leave this investigation for future research.[5]

We model log excess FX returns as

$$y_{t+1} \equiv (s_{t+1} - s_t) - (r_t - \tilde{r}_t) = \mu_t + v_t^{1/2} w_{t+1}^s + z_{t+1}^u - z_{t+1}^d \,, \tag{3.1}$$

where w_{t+1}^s is a standard Gaussian shock (i.e., zero mean and unit variance), z_{t+1}^u is a jump up (i.e., depreciation of USD) and the negative of z_{t+1}^d is a jump down (i.e., appreciation of USD). The conditional spot variance is v_t and the jump intensities of z_{t+1}^u and z_{t+1}^d are h_t^u and h_t^d respectively.[6] The discussion of μ_t is postponed until we have further described these three shocks.

The conditional spot variance v_t is assumed to follow a mean-reverting "square-root" process,

$$v_{t+1} = (1-\nu)v + \nu v_t + \sigma_v v_t^{1/2} w_{t+1}^v + z_{t+1}^v \,, \tag{3.2}$$

which itself can jump with intensity h_t^v.[7] The Gaussian shocks to excess returns w^s and to conditional spot variance w^v have a correlation coefficient $\mathrm{corr}\,(w^s, w^v) = \rho$. Finally, to

[5] Lustig, Roussanov, and Verdelhan (2011); Sarno, Schneider, and Wagner (2012) perform such modelling allowing normal shocks only. Bakshi, Carr, and Wu (2008) model a triangle of currencies (GBP, JPY, and USD) allowing for jumps in FX.

[6] This specification can be viewed as a discrete-time model or as a Euler discretization of a continuous-time model (see, e.g., Platen and Rebolledo, 1985). In any case, a discrete-time model is required at the estimation stage, which is why we omit explicit continuous-time formulation. Formally, all shocks, even the Gaussian variables, are jumps in discrete time. We model small jumps via Gaussian shocks and large jumps via the compound Poisson process. We distinguish the small and the large jumps by imposing the respective priors at the estimation stage. We apply the term jump to the large component only for the ease of referral.

[7] In continuous time, the Feller condition $\sigma_v^2 < 2v(1-\nu)$ ensures that the variance stays positive if there are no jumps. A formal modelling of this process in discrete time is achieved via a Poisson mixture of Gamma distributions (e.g., Gourieroux and Jasiak, 2006; Le, Singleton, and Dai, 2010). We use a direct discretization of the continuous-time counterpart so that the model parameters can be easily interpreted. We ensure that the variance stays positive at the estimation stage by a careful design of the simulation strategy.

ensure positivity of the variance when jumps are present, we only allow for upward jumps so that z^v_{t+1} has non-negative support.

The jump arrival rate is controlled by a Poisson distribution. The assumed jump intensities imply that the number of jumps takes non-negative integer values j with probabilities

$$\mathrm{Prob}(j^k_{t+1} = j) = e^{-h^k_t}(h^k_t)^j/j!\,, \quad k = u, d, v\,. \tag{3.3}$$

We allow all of the jump intensities to depend on the domestic and foreign interest rates, as well as on the conditional spot variance,

$$h^k_t = h^k_0 + h^k_r r_t + \tilde{h}^k_r \tilde{r}_t + h^k_v v_t\,, \quad k = u, d, v\,. \tag{3.4}$$

For a given number of jumps per period, the magnitude of a jump size is assumed to be random with a Gamma distribution,

$$z^k_t | j \sim \mathcal{G}\mathrm{amma}(j, \theta_k)\,, \quad k = u, d, v\,. \tag{3.5}$$

Intuitively, because we consider daily data, a Bernoulli distribution is a very good approximation to our model as it is reasonable to assume no more than one jump per day. Then, the probability of a jump is $1 - e^{-h^k_t} \approx h^k_t$ and the distribution of the jump size is exponential with mean parameter θ_k.[8]

We complement our data on exchange rate rates with variances implied from option prices. In this respect we follow the rich options literature that highlights the importance of using information in options for model estimation (e.g., see Aït-Sahalia and Kimmel, 2007; Brandt and Santa-Clara, 2002; Chernov and Ghysels, 2000; Jones, 2003; Pan, 2002; Pastorello, Renault, and Touzi, 2000). Many authors use implied variance in empirical work by interpreting it as a very accurate approximation of the risk-adjusted expectation of the average future variance realized over an option's lifetime. This is certainly true for models with stochastic volatility only. The one-for-one relationship between implied variance and risk-adjusted expected variance may break down in the presence of jumps. For example, Chernov (2007) has to assume that the risk-adjusted mean of jumps in FX is equal to zero to retain the simple relationship. Importance of careful accounting for jumps is manifested more clearly in the literature on model-free implied variance, such as VIX for S&P 500, where analytic expressions are feasible. Martin (2013) shows that, in the presence of jumps, VIX is equal to risk-adjusted expected variance plus additional terms reflecting the higher order risk-adjusted cumulants of returns.

We treat the Black-Scholes implied variance of a short-term (one-month) at-the-money option, IV_t, as a noisy and biased observation of the conditional spot variance v_t. Such a view allows us to avoid the aforementioned difficulties in explicit connection between

[8] Our choice of the variance jump size distribution is frequently used when modelling variance to ensure its positivity as discussed above. The model of variance is also capable of generating quite rapid variance declines after jumps. A jump leads to a large deviation from the long-run mean v, and mean-reversion controlled by parameter ν ensures that the variance is pulled back.

implied variance and risk-adjusted expected future variance. The cost of such approach is our inability to estimate risk-adjusted parameters of the model. Specifically,

$$IV_t = \alpha_{iv} + \beta_{iv} v_t + \sigma_{iv} v_t \sqrt{\lambda_t}\, \varepsilon_t \,, \tag{3.6}$$

where IV_t is expressed in daily terms, ε_t is $\mathcal{N}(0,1)$ and λ_t is $\mathcal{IG}(\nu/2, \nu/2)$, so the product $\sqrt{\lambda_t}\varepsilon_t$ is $t_\nu-$distributed (Cheung, 2008; Jacquier, Polson, and Rossi, 2004).[9] We have considered a version of (3.6) with non-zero loadings on r_t and $\tilde{r}_t$, but this specification did not find empirical support.[10]

The model implies that expected log excess return is equal to

$$E_t\left[y_{t+1}\right] = \mu_t + \underbrace{h_t^u \theta_u}_{E_t\left[z_{t+1}^u\right]} - \underbrace{h_t^d \theta_d}_{E_t\left[z_{t+1}^d\right]} \,. \tag{3.7}$$

As discussed in the previous section, we assume that

$$\mu_t = \mu_0 + \mu_r r_t + \tilde{\mu}_r \tilde{r}_t + \mu_v v_t. \tag{3.8}$$

The resulting expected excess return is

$$E_t\left[y_{t+1}\right] = \mu_0^* + \mu_r^* r_t + \tilde{\mu}_r^* \tilde{r}_t + \mu_v^* v_t \tag{3.9}$$

where

$$\mu_0^* = \mu_0 + h_0^u \theta_u - h_0^d \theta_d \,, \tag{3.10a}$$
$$\mu_r^* = \mu_r + h_r^u \theta_u - h_r^d \theta_d \,, \tag{3.10b}$$
$$\tilde{\mu}_r^* = \tilde{\mu}_r + \tilde{h}_r^u \theta_u - \tilde{h}_r^d \theta_d \,, \tag{3.10c}$$
$$\mu_v^* = \mu_v + h_v^u \theta_u - h_v^d \theta_d \,. \tag{3.10d}$$

Thus, our risk premium encompasses the UIP regressions which set

$$\tilde{\mu}_r^* = -\mu_r^* \,, \tag{3.11}$$
$$\mu_v^* = 0 \,. \tag{3.12}$$

We conclude with a discussion of our approach to modelling interest rates. We do not need an explicit model of interest rates to estimate our model of FX excess returns if we are willing to assume that one-day r_t and $\tilde{r}_t$ can be reasonably proxied with short-term yields. We

[9] Jones (2003) makes a strong case for heteroscedastic measurement errors in implied variance. His specification sets $\lambda_t = 1$. Cheung (2008) generalizes the specification to the Student $t-$error. We tried using a normal error with volatility σ_{iv}, a normal error with volatility $\sigma_{iv} v_t$, and the Student $t-$error described above. We find that heavy-tailed t_3 works very well.

[10] The error specification in (3.6) is very flexible. Therefore, it could be the case that the contribution of interest rates to the variation in implied variance cannot be empirically distinguished from the error, if the former is reasonably small.

view this feature as a strength of our approach because explicitly modelling the behaviour of spot interest rates entails significant effort. There is a separate literature dedicated to this task and state-of-the-art models rely on five factors for capturing interest rate dynamics. These studies are typically conducted with monthly or quarterly data, so they do not take into account higher-frequency movements in interest rates which are susceptible to jumps themselves (e.g., Johannes, 2004; Piazzesi, 2005). Moreover, interest rates and currencies have low conditional correlation and variability in interest rates is much smaller than that in currencies. In summary, elaborate modelling and estimation of interest rate dynamics does not appear to be a first order concern for the questions that we are addressing.

Nonetheless, we use the estimated model to compute some useful objects (expectations of future variance, or expected excess returns over multiple horizon) that depend on the distribution of interest rates. In order to obtain reasonable quantities, we assume the simplest possible model for the interest rates:

$$r_{t+1} = (1-b_r)a_r + b_r r_t + \sigma_r r_t^{1/2} w_{t+1}^r, \tag{3.13a}$$

$$\tilde{r}_{t+1} = (1-\tilde{b}_r)\tilde{a}_r + \tilde{b}_r r_t + \tilde{\sigma}_r \tilde{r}_t^{1/2} \tilde{w}_{t+1}^r. \tag{3.13b}$$

A square root process for interest rates is subject to caveats in discrete time that we previously discussed for the variance process. We calibrate the models to match the mean, variance and serial correlation of the respective observed short-term interest rates. Our computations with reasonable variation in parameters confirm our intuition that they have minimal impact on the role of normal and non-normal currency risks.

3.2 Qualitative Features of the Model

In this section we explain how we arrived at the specified functional form of the model. We evaluated too many models to provide a detailed account of our analysis, so we briefly summarize the results that led us to the above specification. Our initial specifications were motivated by the well-developed literature on equity returns (Andersen, Benzoni, and Lund, 2002; Chernov, Gallant, Ghysels, and Tauchen, 2003; Eraker, Johannes, and Polson, 2003; Eraker, 2004; Jones, 2003) and some of the few models of currencies (Bates, 1996; Johnson, 2002; Jorion, 1988; Maheu and McCurdy, 2008). The salient features of equity data are presence of substantial moves up and down and a pronounced negative skewness in the return distribution. Therefore, jumps in equity returns are often modelled via a single compound Poisson process with a normally distributed size of non-zero mean. However, in contrast to equity returns, currency returns have very mild skewness over long samples, which suggests a zero-mean normal distribution for jump sizes.

Further, Bates (1996), Campa, Chang, and Reider (1998), Carr and Wu (2007), and Johnson (2002) emphasize the time-varying and sign-switching nature of the risk-adjusted skewness of exchange rates. The key to modelling this feature successfully is to allow the conditional expected jump to vary over time. A single jump process with a zero mean jump size implies a zero conditional expected jump. Two jump processes have a potential to generate the requisite variation either via time-varying jump intensities, or time-varying jump size distributions, or both. We do not explore time-varying jump means as such specifications do not allow for tractable option valuation in the affine framework, and we eventually want our model to be used to analyze option prices. As can be seen from the expression for the currency risk premium (3.7), the conditional jump expectation is $h_t^u\theta_u - h_t^d\theta_d$, which is capable of producing the needed variation. We have also considered normally distributed jump sizes in excess returns with means of the jump size distribution having opposing signs. However, because normal distributions have infinite support, it was hard to distinguish empirically the down and up components. The exponential distribution does not have this issue because the support is on the positive line.

Another interesting feature of our specification is that we allow not only for two different Poisson processes in currency returns, but also for a third one in the variance. Our starting point was again in the equity literature where all jumps in returns and variance are guided by the same (or at least correlated) Poisson processes. We found that the model with correlated Poisson processes fit the data poorly.

4 Empirical Approach

We use Bayesian MCMC to estimate the model. This method has been successfully implemented in many applications (see Johannes and Polson, 2009 for a review). For our purposes, a key advantage of this approach over other methodologies is that unobserved variables – e.g., variance, jump times, and sizes – are a natural by-product of the estimation procedure. The online Appendix describes all the details of the implementation.

It is worth pointing out how we distinguish jumps and normal shocks in the model. Formally, all shocks are discontinuous in our discrete-time formulation. We think of jumps as relatively infrequent events with relatively large variance. We use priors on jump arrival and jump size parameters to express this view.

It proved to be extremely fruitful to use option implied variances in our estimation. Ignoring information in option prices made it very hard to settle on a particular model. Parameters were estimated imprecisely and the algorithm had poor convergence properties – both are manifestations of the data being not sufficiently informative about the model. We had a similar experience when estimating the most general model, even when using the options data. Complicated dependencies of jump intensities on state variables, and the sheer number of separate Poisson processes was too much for the available data.

As such, we pursue the following model selection strategy. First, we treat implied variances as observed spot variances and estimate the model of variance (3.2). At this stage we select the best model by checking the significance of parameters on the basis of both confidence intervals and Bayes odds ratios. Specifically, the parameters of concern are the ones controlling the jump intensity in equation (3.4) for $k = v$. It turns out that, regardless of the currency, only the loading on variance is significant. In other words, the probability of jumps in the variance is affected by the variance itself. Thus, jumps in the variance are self-exciting (Hawkes, 1971).[11] Pinning down the model of variance is an extremely useful step in our estimation procedure.

Second, we use the lessons from the estimation exercise on the basis of implied variance alone to guide us in a formal search in the context of our full model. That is, we take the model (3.1), (3.2) and combine it with equation (3.6) that recognizes implied variances as noisy observations of the spot variance. As a benchmark, we estimate the stochastic variance model with no jumps. Next, we estimate a model with jumps in variance but no jumps in exchange rates ($h_t^u = h_t^d = 0$). We refer to this model as stochastic variance with jumps.

Finally, we allow for the full model with jumps in both exchange rates and variance. Here, we focus on the significance of the parameters controlling the jump intensities in (3.4) for $k = u$ and d. We are not reporting all the details here, but we find that $\tilde{h}_r^u$, h_v^u, h_r^d and h_v^d are insignificant. Thus, the probability of jumps up in the exchange rate is driven by the domestic rates only, and the probability of jumps down in the exchange rate is driven by the foreign rate only. We also test if some interesting parameters, or combinations of parameters, are equal to zero. First, we can test the UIP regression restrictions on the risk premiums in Eq. (3.11) (whether interest rates affect the risk premium as a differential) and Eq. (3.12) (whether the variance affects the risk premium). As we noted earlier, the behaviour of the FX skewness is dramatically affected by expected effect of jumps, which is equal to $h_t^u\theta_u - h_t^d\theta_d$. Here we are interested in testing whether $\theta_u = \theta_d = \theta$, $h_0^u = h_0^d = h_0$, and $h_r^u = \tilde{h}_r^d = h_r$. These hypotheses are interesting because if they cannot be jointly rejected then expected jump would be equal to $\theta h_r(r_t - \tilde{r}_t)$. Thus, the excess return asymmetries will be directly driven by the interest rate differential as noted in Brunnermeier, Nagel, and Pedersen (2008). The final version of this model that incorporates all the unrejected null hypotheses is referred to as the preferred.

We implement a series of informal diagnostics and specification tests to establish the preferred model. The diagnostics test the null hypothesis that the shocks to the observable excess return, w^s, and implied variance, ε, should be normal under the null of a given model. We can construct the posterior distribution of these shocks and evaluate how they change from model to model and whether they are normal. The procedure is described in the online Appendix.

One has to exhibit caution when interpreting the evidence on normality of ε. The variance of the error term in the implied variance equation (3.6), $\sigma_{iv}^2 v_t^2 \lambda_t$ is very flexible. If a model is

[11]The recent literature on equity returns also finds support for self-exciting jumps. See, for example, Aït-Sahalia, Cacho-Diaz, and Laeven (2011); Carr and Wu (2011); Nowotny (2011); Santa-Clara and Yan (2010).

misspecified, λ_t will adjust so that the ε is close to a normal variable. Therefore, diagnostics of ε are not enough. We should be tracking the size of the variance of the error term. A better specified model should have smaller variance. We keep track of the time-series average of this variance – which we refer to as $IVvar$ – and report its posterior distribution.

Bayes odds ratios offer a formal specification test of the models. The test produces a number that measures the relative odds of two models given the data (the posterior distribution of the null model is in the denominator of the ratio). Following Kass and Raftery (1995), we interpret a log odds ratio that is greater than 3 as strong evidence against the null. Odds ratios do not necessarily select more complex models because the ratios contain a penalty for using more parameters (so-called automatic Occam's razor). The computations are described in the online Appendix.

5 Results

We start by highlighting statistical properties of the estimated models. Next, we study economic implications.

5.1 Statistical Properties of Currency Risks

Tables 2 - 5 report the parameter estimates and Tables 6 - 9 report the corresponding model diagnostics. Table 10 displays the results of specification tests on the basis of Bayes odds ratios. Table 11 summarizes parameters of the calibrated interest rate processes.

The results exhibit a lot of similarities across the different currencies. As we move from models with stochastic variance to stochastic variance with jumps, we observe a change in two key parameters: both the persistence ν of variance and the long-run mean of its conditionally normal component v decline. Taking AUD as an example, ν declines from 0.9943 to 0.9855. This seemingly small change translates into drop in the half-life of the conditionally normal component, $\log 2/(1-\nu)$, from 122 to 48 days. The high persistence of variance in the model without jumps is a sign of misspecification. Variance has to take high values occasionally to generate the observed exchange rates in the data. In the absence of jumps, variance builds up to the high values gradually via the high persistent channel. Additionally, in the case of GBP only, the volatility of variance σ_v declines significantly from 0.0321 to 0.0272. High σ_v helps the misspecified model with stochastic variance in generating high values of variance. The diagnostics support this interpretation. $IVvar$ drops by 50% across all currencies; this change is statistically significant. As expected, diagnostics for ε show that it is close to a normal variable for both models because of the flexibility in λ_t. Bayes odds ratios strongly favour stochastic variance with jumps.

Continuing with AUD, its volatility ($\sqrt{v}$) declines from 0.70% to 0.53% per day (11.19% to 8.43% per year). This happens because the total variance has contributions from the regular and jump components in the model with stochastic variance with jumps. When there are no jumps in FX, the long-run variance is equal to $v_J = [(1-\nu)v + h_0^v\theta_v]/[1-\nu-h_v^v\theta_v]$. See Appendix A.2 for more details. This expression produces the average volatility of 0.65%, much closer to the figure in the model with stochastic variance.

To aid in interpreting the parameters controlling jumps in variance, consider the impact of a jump in variance. Suppose the current variance is at its long-run mean and the variance jumps by the average amount θ_v. Then in the case of AUD, the resulting volatility will move from 0.65% to $(v_J + \theta_v)^{1/2} = 0.90\%$, a nearly 40% increase in volatility (this increase ranges from 20% to 40% for the different currencies). The average jump intensity is equal to $h_0^v + h_v^v v_J = 0.0053$ jumps per day, or 1.34 per year (this number ranges from 1.34 to 2.61 for the the different currencies). Jumps in variance are self-exciting, so that a jump increases the likelihood of another jump. When the variance jumps by θ_v, intensity changes to 1.71 per year for AUD (the range is from 1.71 to 3.41 for all the currencies).

Also note that ρ, the "leverage effect," has the same sign as the average interest rate differential. It is positive for JPY and CHF, and negative for AUD and GBP. This result is consistent with the analysis in Brunnermeier, Nagel, and Pedersen (2008) and the common wisdom among market participants that investors who are long carry are essentially short volatility. For example, consider the position of a carry trade investor who borrows money in USD and invests in AUD. This investor can loose money when the AUD depreciates against the USD. We estimate that ρ is negative for this currency pair, so the volatility of this exchange rate tends to increase during times when the AUD depreciates.

The preferred version of the full model is the one with all of the aforementioned restrictions imposed ($\theta_u = \theta_d = \theta$, $h_0^u = h_0^d = h_0$, and $h_r^u = \tilde{h}_r^d = h_r$). That is, the size of the jumps in FX up and down are symmetric and their intensities have numerically identical functional form (but they depend on different interest rates). As a result, the overall structure of jump arrivals differs from the one used in popular models of S&P 500 returns, where jumps in variance and the index are simultaneous.

Parameters reflecting the average jump size have a different interpretation as compared to jump in variance. The latter is a jump in the level of the variable, while the former is the jump in return: it reflects by how much the return changes at the moment of the jump. Thus, on average, AUD returns increase (decline) by 1.69% when there is a jump up (down). To put this number into perspective, the daily volatility of AUD returns is 0.74% a day (Table 1). Thus, an average jump exceeds a "two sigma" event. Average intensities of down and up jumps are similar to each other and to those of variances for a given currency: they range from 0.87 to 3.35 jumps per year (we use sample averages of interest rates to compute average h_t^u and h_t^d).

The diagnostics of residuals w^s indicate that the major improvement in moving from stochastic variance with jumps to the preferred model comes from a statistically significant drop in kurtosis from roughly 4 to 3.5 across all currencies. The absolute value of skewness of w experiences a significant drop for all currencies except for GBP, where it was insignificantly different from zero in the model with stochastic variance with jumps already. Serial correlation is slightly negative for all currencies except for GBP (where it is zero in the model with stochastic variance with jumps already), and the change from one model to another is insignificant. $IVvar$ does not change appreciably because we did not change our model for variance. Bayes odds ratios strongly favor the preferred model. In summary, the preferred model is clearly superior, but there are some residual non-normalities left in the fitted shocks to exchange rates. We leave improvements to future research.

The expected excess return in (3.9) can be simplified for the preferred model to

$$E_t(y_{t+1}) = \mu_0 + (\mu_r + h_r\theta)r_t + (\tilde{\mu}_r - h_r\theta)\tilde{r}_t + \mu_v v_t. \tag{5.1}$$

Thus, by testing if $\mu_r = -\tilde{\mu}_r$ and $\mu_v = 0$, we test the UIP regression specification (3.11) - (3.12) of currency excess returns across all three models. For all currency pairs, we cannot reject that $\mu_r = -\tilde{\mu}_r$ at the conventional significance levels. Moreover, $\mu_r \approx -3$ for all currencies, which is consistent with our earlier discussion of UIP regression results. In addition, the loading on the variance μ_v is significantly negative in all currencies except for JPY which has a significantly positive estimate. The tiny serial correlation of the residuals w^s suggests that this model is adequate in capturing the conditional mean of excess returns and, therefore, potentially omitted variables cannot materially affect our conclusions about the structure of currency risks.

The probability of USD depreciation, h_t^u, depends positively on the US interest rate. This result seems to contradict many people's intuition about the relationship between interest rates and changes in FX. It is important to note that this connection is applicable to jumps only. Parameter estimates and expression (5.1) imply that, all else equal, the USD appreciates, on average, when the difference between the U.S. and foreign interest rate is higher. However, when the interest rate differential is positive, our model says that the probability of a large depreciation in the USD (a jump up) is higher than the probability of a large appreciation (a jump down). Another way to illustrate this effect is to focus on the asymmetry of the conditional distribution of an exchange rate. Indeed, the third conditional cumulant can be obtained by taking the third derivative of the cumulant-generating function of log currency returns (provided in Appendix A.3). For example,

$$\kappa_{3t}(s_{t+1} - s_t) = 6\theta^3 h_r(r_t - \tilde{r}_t)$$

in the preferred model. We see that loadings of jump intensities on the different interest rates translates into time-varying skewness controlled by the interest rate differential.

5.2 Jumps and Events

In this section we study the economic properties of the documented jumps. We ask basic questions regarding the structure of the jump components, examine whether jumps can be related to important economic events, and gauge their impact on the overall risk of currency trades. Our discussion is based on Figures 1 - 4 and Table 12.

For each exchange rate, the figures display the time series of data (excess returns and implied volatilities) complemented with the estimated unobservable model components: spot volatility $v_t^{1/2}$ and jumps. Regarding the latter, the model produces an estimate of a jump size and an ex-post probability that a jump actually took place on a specific day. However, all of this information is not easy to digest as there are a lot of small jumps and jumps with a small ex-post probability of taking place. Thus, to simplify the reporting, we stratify the jump probabilities by assigning them a value of one if their actual value is 0.5 or higher, and zero otherwise. Then we plot the estimated jumps sizes on the days with the assigned value of one. Our strategy yields 219 jumps overall across all currencies. Approximately 25% of these jumps take place simultaneously in at least two currencies. We refer to such jumps as common. There are only 8 episodes when FX and variance jumped at the same time. We overlay the plots of the estimated jumps sizes with the state-dependent ex-ante jump probabilities h_t^v, h_t^u, and h_t^d.

To interpret the plots better, we have to reference the jump magnitudes against the summary statistics available in Table 1. Let us use JPY in Figure 4 as an example. Table 1 tells us that the volatility of JPY is 0.7% per day and the mean is approximately zero. Thus, a "regular" excess return can be within the range of $\pm 2\%$ ($\mu \pm 3\sigma$). The upper left panel of Figure 4 shows that there are quite a few days when excess returns are outside of this range. In practice, the volatility is time-varying and unobserved. Therefore, the "regular" range is time-varying also and uncertain. The estimation procedure takes this uncertainty into account by producing the ex-post probability of a jump taking place. We arbitrarily select the level of uncertainty about jumps that we are comfortable with by discarding the jumps with such probabilities less than 50%. The bottom left panel confirms this by showing the estimated jumps in excess returns. Their magnitude ranges from 2% to 6% in absolute value. Interestingly, the larger jumps coincide with spikes in the moving-window estimates of skewness across all currencies.

However, not all of the big spikes in excess returns are attributable to jumps in FX. For example, on October 28, 2008, the plot of excess returns shows a big drop of 5.5%. The model tells us that there were no jumps on that day. The model is capable of generating such a big move via a normal component because of the jump in variance. Volatility jumped by $(v_{t-1} + z_t^v)^{1/2} - v_{t-1}^{1/2} = 0.18\%$ (2.9% annualized), on average, on each of three days between October 22 and 24. Each day the jump in variance was increasing the probability of a jump the following day. Over these three days volatility moved from 1.35% (21.5% annualized) to 1.8% (28% annualized). To put this number into a perspective, the long-run volatility mean is $v_J^{1/2} = 0.66\%$ (10.4% annualized). Thus, the increase in volatility over these three days was roughly equal to the average level of volatility. Moreover, there is no significant news associated with either October 22-24 or October 28. Thus, we attribute these events to pure uncertainty in the markets.

GBP generates large movements via jumps in variance in the most transparent way. The exchange rate itself exhibits only 11 jumps throughout the sample, none of which take place after 2000. In fact, even the famous "Black Wednesday" – the GBP devaluation on September 16, 1992 – is attributed primarily to a jump in variance on September 8. The volatility moved from 0.93% (14.7% annualised) to 1.13% (17.9% annualised), then it drifted up to 1.15% (18.3% annualised) on the day of the crash. These values of volatility are high, as the average volatility of GBP is $v_J^{1/2} = 0.55\%$. Nonetheless, this level of volatility is still insufficient to generate the whole drop of -4.10%. Of course, these rough computations assume that v_t is known with certainty. It is not, and a small deviation in the estimate may be able to attribute the whole return to a normal shock in FX. This is why the estimated probability of a jump in exchange rate is only 26% on this day (and the estimated jump size is -0.42%).

As a next step, we check if the jumps we detect are related to economic, political or financial events. For each day a currency has experienced a jump, we check if there were significant news. To this end, we search Factiva for articles that explicitly relate movements in currency markets to events rather than collect a comprehensive set of macro-economic releases on the days of jumps. This strategy is complimentary to the one employed in studies of the news impact on financial assets (see, e.g., Andersen, Bollerslev, Diebold, and Vega, 2003 for FX). Usually one measures news surprise by computing a standardized difference between an announcement's expectation and realization and then checks, usually at intra-day frequency, if the surprise had an impact on values of financial assets. Our approach does not require us taking a stand on measuring a surprise. In addition, we are careful in distinguishing announcements, a clear public release of a fact, from uncertainty: anticipation, comments on the current economic situation and overall market anxiety that is sometimes evident in the news. The online Appendix provides a jump-by-jump description of all events.

Across all currency pairs, we find that 65% of jumps in FX are associated with announcements of economic, political, or financial events. By stark contrast, only 23% of jumps in variance across all currency pairs are associated with such events.[12] Perhaps, the fact that many jumps in FX are associated with announcements is not so surprising. After all, announcements take place on many days of a year. What is surprising is that there are so few jumps in variance that coincide with announcements. On such days we see a lot of discussion and interpretation of past events and, perhaps, anticipation of the future ones. Thus, jumps in variance appear to reflect economic uncertainty rather than the revelation of a specific fact. Thus, there is is a critical difference between jumps in FX and variance. Table 12 provides a summary of the types of announcements associated with jumps.

We find that economically different types of announcements and uncertainty matter for different currency pairs. Australia, a country with a large current account deficit, experiences a number of interest rate cuts and interventions. These, together with the trade news, lead to a number of jumps in AUD in the eighties and nineties. Later in the sample,

[12]We avoid double counting by accounting for common jump events only once and do not count simultaneous jumps in FX and variance (only 8 incidences across all currencies).

AUD, as a high yield commodity currency, is vulnerable to Asian crisis and credit crunch events, which results in numerous realizations of large negative returns. In contrast to all other currencies, jumps in CHF are driven solely by foreign announcements and uncertainty about the US and third countries' economies. The exceptions are an interest rate cut by the Swiss National Bank on November 21, 2008, and a monetary policy uncertainty on March 12, 2009. The British pound has the largest number of jumps in variance and hence most of the jump events (72%) are driven by economic uncertainty – uncertainty about monetary system in Europe (Exchange Rate Mechanism, Euro launch) and spillovers from the US financial market and banking sector in Europe. Similar to AUD, JPY experiences many jumps in prices caused by either trade news or central bank interventions in the beginning of the sample, whereas more recent jump episodes are related to turbulence in world financial markets associated with the credit crisis. In contrast to AUD, however, Japan runs a huge trade surplus and the realised jump returns on yen tend to be positive in the recent financial crisis.

While we do not formally model the joint behaviour of exchange rates, we find many simultaneous jumps across the currencies, particularly jumps in variance. Therefore, it appears promising to extend the existing research on common and currency-specific factors affecting risk premiums (e.g., Lustig, Roussanov, and Verdelhan, 2013) by allowing for common and currency-specific jump risk.

The plots of jump intensities provide some insight into why jumps in variance are so prominent. Probabilities of jumps in FX are moving together with interest rates, which experienced a secular decline in our sample in all countries. As we highlighted earlier, the probability of a jump in variance is primarily driven by the variance itself. This probability went through a couple of cycles of high and low values, as volatility is less persistent than interest rates.

5.3 Decomposing the total risk

Are these risks quantitatively important? Jumps in FX and variance should affect the tail events the most. The properties of tails are captured by high-order moments or cumulants. One could report each of these statistics separately and measure how they are affected by the various shocks present in our model. We choose to summarize all this information by one number and measure the total risk corresponding to investment horizon n using the entropy of changes in FX:

$$L_t(S_{t+n}/S_t) = \log E_t(e^{s_{t+n}-s_t}) - E_t(s_{t+n} - s_t). \tag{5.2}$$

Entropy is a loaded term. Our use of entropy is similar to that of Backus, Chernov, and Martin (2011) and Backus, Chernov, and Zin (2013), who characterize entropy of the pricing kernel, and is closest to the way it is used in Martin (2013) who uses entropy of equity index returns. Entropy's connection to the cumulants of log FX returns makes it attractive for our purposes.

Specifically, the definition (5.2) implies that

$$\begin{aligned} L_t(S_{t+n}/S_t) &= k_t(1; s_{t+n} - s_t) - \kappa_{1t}(s_{t+n} - s_t) & (5.3) \\ &= \kappa_{2t}(s_{t+n} - s_t)/2! + \kappa_{3t}(s_{t+n} - s_t)/3! + \kappa_{4t}(s_{t+n} - s_t)/4! + \ldots, & (5.4) \end{aligned}$$

where $k_t(1; s_{t+n} - s_t)$ is the conditional cumulant-generating function of $s_{t+n} - s_t$ evaluated at the argument equal to 1, $\kappa_{jt}(s_{t+n} - s_t)$ is the jth conditional cumulant of $s_{t+n} - s_t$, and we used the fact that $k_t(1; s_{t+n} - s_t)$ is equal to the infinite sum of $\kappa_{jt}(s_{t+n} - s_t)/j!$ over j starting with $j = 1$. The significance of the property (5.4) is that if currency changes are normally distributed, then entropy is equal to a half of their variance (the first term in the sum). All the higher-order terms arise from non-normalities. Thus, entropy captures tail behaviour of returns in a compact form. As a result, we view entropy as a natural generalization of variance as a risk measure. For this reason, Alvarez and Jermann (2005) explicitly refer to L_t as generalized variance.

We decompose the total risk of currency returns into the contribution of the jump and normal components. Appendix A.3 explains how we compute the full entropy on the basis of equation (5.3). We can compute the individual components by zeroing out the rest.[13] Figures 5 and 6 display the contributions of these components for the investment horizons of 1 months ($n = 21$) and 3 months ($n = 63$). We overlay these contributions with a time-series plot of entropy computed treating USD as the domestic currency. We scale entropy to ensure that it is equal to variance in the normal case and to adjust for the horizon. Finally, to make the number easily interpretable, we take the square-root and express it in percent. Thus, we plot $\sqrt{2L_t/n}$. Finally, Table 13 supports the plots by reporting summary statistics of the relative contributions of the different components.

We start by characterizing the contribution of the different components at a given point in time. We see that the regular, or normal, risk is the most prominent regardless of the horizon. The average total contribution of jumps at a one-month horizon ranges from 11.03 % for GBP to 20.19% for AUD. The contribution of a jump in FX (up or down) [the range is between 6.82% and 9.17%] is higher than that of a jump in variance [the range is between 3.76% and 4.83%] and has higher time-variation at a one-month horizon (GBP is the exception as jumps in variance contribute 4.37% as compared to 2.98% for jumps up and 3.68% for jumps down). Therefore, in the short-term the risk of a jump in variance has the smallest contribution to the overall currency risk. However, this conclusion changes

[13] If two variables x_s and y_s are conditionally independent, then $L_t(x_s y_s) = L_t(x_s) + L_t(y_s)$. Therefore, our decomposition approach correctly separates the contributions of the two jumps in currencies. Because the probability of a jump in variance depends on the variance itself, the normal shock to variance and the jump are conditionally independent only over one period, $n = 1$. When $n > 1$ our procedure attributes all the covariance terms, which are positive because of the estimated functional form of jump probabilities, to the jump in variance. We think that this approach is sensible because the presence of these covariance terms is due to jumps.

as we extend the investment horizon to three months. The average total contribution of jumps at this horizon increases – the range is from 17.71% for GBP to 25.19% for JPY. In this case individual contribution of jumps in variance [the range is 8.94% to 11.48%] is higher than those of jumps in FX [the range is 2.78% to 8.57%] (in the case of GBP, the contribution of a jump in variance is greater than the sum of jumps up and down).

The contribution of jumps in FX declines towards the end of our sample, thereby making the contribution of jumps in variance more important. We connect this result to the secular decline in the probability of FX jumps that we highlighted earlier. This effect diminishes the expected contribution of such jumps to the overall risk. In contrast, the probability of jumps in variance is less persistent and, therefore, exhibits mean-reversion in our sample.

The time-series variation in total risk resembles the time-series variation in the spot variance v_t. This is not surprising because L_t is a linear function of v_t (Appendix A.3). Thus, whenever the spot variance moves, and especially when it jumps, we observe a clear move in entropy. We conclude that the risk of jumps in variance are at least as important as the risk of jumps in FX, even though the nature of the two types of jumps is different.

6 Importance of crash risk for valuation

The large amount of jump risk in bilateral exchange rates prompts a natural question of whether this risk is priced. Answering this question has important implications for the literature that establishes high returns to currency speculation. In particular, one should be able to attribute a specific fraction of documented risk premiums to compensation for bearing crash risk. The full answer to this question requires an explicit model of the pricing kernel and the use of assets that are sensitive to jump risk for estimation of risk premiums. We leave an in-depth analysis of this issue for future research.

Herein, we offer two perspectives on the question. First, one might argue that portfolios of exchange rates could diversify crash risks and, therefore, their importance for pricing is less relevant than it may seem on the basis of bilateral rates. Moreover, the carry trade literature emphasizes the importance of the timing of the long vs short position in a given currency. Arguably, perfect timing may make crashes always go in a speculator's favor. In particular, Lustig, Roussanov, and Verdelhan (2013) construct one such portfolio and the corresponding trade (the dollar index and the dollar carry, respectively). The authors show that it is more profitable than the more traditional carry trades. This index fits nicely with our focus on foreign currencies vis-a-vis the U.S. dollar, so we repeat our analysis for the dollar index and investigate its risk structure.

A more direct way to assess the valuation of crash risk is to measure the magnitude of risk premiums using prices of assets that are sensitive to jump risk. In this regard, out-of-the-money options are particularly informative about the price of jump risks (i.e., the covariance of the pricing kernel with jumps). A thorough analysis along these lines is outside of the scope of this paper. Instead, we provide calibration of jump risk premiums on the basis of average implied volatility smiles.

6.1 Crash risk of the dollar index

We follow closely Lustig, Roussanov, and Verdelhan (2013) in constructing the dollar index. We use daily exchange rates on twenty one currency pairs against USD: Australian dollar, Austrian schilling, Belgian franc, Canadian dollar, Danish krone, French franc, Finnish markka, Deutsche mark, Greek drachma, Italian lira, Irish pound, Japanese yen, Dutch guilder, New Zealand dollar, Norwegian krone, Portuguese escudo, Spanish peseta, Swedish krona, Swiss franc, British pound, and Euro from January 1, 1986 to December 31, 2010. We replace eleven European currencies with the Euro starting from January 1, 2002. We construct an equally-weighted index of currencies against USD: an investor buys a one-dollar equivalent of each foreign currency in the basket. Summary statistics in Table 1 indicate that, consistent with the notion of risk diversification, the excess returns of the index are less volatile, skewed and leptokurtic than those of bilateral exchange rates. As is the case with individual FX rates, Figure 7(a) reveals that skewness is time-varying. The range of possible values is comparable with those of specific country pairs.

Options on the index are not traded, so we do not have an option-implied variance to use in equation (3.6) as a noisy observation of the variance of the log return on the currency index. Instead, we use the first principal component of variances implied from at-the-money one-month options on the bilateral exchange rates. Specifically, we collect implied variances for AUD, CAD, CHF, DKK, EUR, GBP, JPY, NOK, NZD, and SEK from February 1, 2001 to December 31, 2010.[14] The choice of currencies and data span are limited by availability of Bloomberg data. To be clear, there is no obvious theoretical link between the principal component and the index volatility. The principal component is a rescaled version of implied variances, so its units do not correspond to the variance level. The values of α_{iv} and β_{iv} may not be in the neighbourhood of 0 and 1, respectively. Moreover, the noise term could be reasonably large. Given the absence of option-implied volatility on the currency index, our empirical approach is simply our best effort, within the scope of this paper, to measure the first order effect of jumps on currency indices.

We use the one-month LIBOR or the one-month Eurocurrency rate as a proxy for the one-period interest rate. If these are unavailable, we use their overnight or 3 months counterparts. The foreign interest rate is an equal-weighted index of the respective individual rates. We calibrate the same model for the index as for the individual interest rates, see Table 11.

Table 14 reports the parameter estimates and Table 15 reports the corresponding model diagnostics. Table 10 displays the results of specification tests on the basis of Bayes odds ratios. All the diagnostics show that the (restricted) largest model is strongly preferred.

[14] Options on the Euro started trading before the currency was circulated physically.

The currency index has a smaller long-run variance and a smaller average jump in variance than bilateral exchange rates. In relative terms, jumps in variance matter as much for the index as they do for individual currency pairs. When current variance is at its long-run mean and jumps by the average jump size θ_v, volatility increases by 35%, from 0.44% to 0.60% per day (6.93 to 9.52 per annum). The average jump intensity is equal to $h_0^v + h_v^v v_J = 0.0050$ jumps per day or 1.25 jumps per year, which is less frequent than that of individual exchange rates.

The average jump size in FX for the index is $\theta = 1.64$. It is only slighter smaller than the average jump size for AUD (1.69) and is larger than those for other currencies. Average intensities of up and down jumps are 1.03 and 1.14 per year, respectively. This is on the low end of jump intensities of individual currency pairs.

Figure 7 displays estimated jumps, ex-ante jump probabilities and volatility states. We estimate 5 jumps down, 10 jumps up and 38 jumps in volatility. All jumps up and down take place before 1998 with the exception of the one on August 8, 2008. In contrast, jumps in variance dominate after 2000 and especially cluster in 2008. There is only one episode when a jump in return and a jump in variance coincide: August 8, 2008.

Table 12 reports how many jumps in our original four bilateral exchange rates are common with the index. Among the jumps in FX, 12 jumps are common (they coincide mainly with jump events in CHF and JPY). Among the jumps in variance, 31 jumps are common. This suggests that many of the jumps that are common between two or three currencies do not get diversified away in an index.

The findings are consistent with the notion of diversification in that the regular risks are smaller, and the jump risks are less frequent. In order to establish the relative importance of the different risks we decompose the entropy of exchange rate changes. Table 13 shows that normal risk is the main component of risks across different investment horizons. Still, jumps contribute jointly on average 30% and 25% of variation at monthly and quarterly horizons, correspondingly. The role of jumps in FX is in fact higher for the currency index than for individual currencies. This is because the mean jump size is on the high end and the jump intensity is about the same for the index for currency pairs. In addition, the overall variation of the index returns is smaller than for individual currency returns. These two effects create a more prominent role for the jump risk as compared to individual FX rates. Similar to individual currency pairs, the role of jumps in variance increases with the investment horizon. Overall, the role of jumps in variance for the currency index is about the same as for CHF.

Thus, combining different currencies into a portfolio does not necessarily reduce the prominence of crash risk. This is not the end of the analysis, however. Lustig, Roussanov, and Verdelhan (2013) argue that investing in the basket of foreign currencies when the corresponding average foreign interest rate exceeds the US rate and in the USD otherwise is (1) more profitable than a conventional carry portfolio (high-minus-low) and (2) uncorrelated with it. Whereas a high-minus-low strategy is well known to suffer at times of turbulence in financial markets, and hence many authors conjecture that it is subject to crash risk, it is less clear what happens with the dollar carry portfolio.

Consider the relative profitability of the original currency index (long-only portfolio) vis-a-vis the dollar carry portfolio. If one uses the realized Sharpe ratio as a measure of performance, the dollar carry portfolio yields a resounding success: its annualized Sharpe ratio is 0.81, compared to 0.33 for the long-only currency index.[15] The Sharpe ratio does not reflect the departures from normality, however.

To evaluate the crash risk of such a strategy we ask whether negative jump risk disappears as a result of flipping position in the USD. We focus our interest on the downward jumps because upward jumps work in favour of the investor and jumps in variance do not disappear because of rebalancing. The dollar carry portfolio experiences 9 jumps down and 6 jumps up, versus the original currency index that has 5 jumps down and 10 jumps up (the jumps in FX that switch signs as a result of rebalancing are marked with green in Figure 7(c)). So the realized Sharpe ratio increases at the cost of making returns negatively skewed – the skewness of the rebalanced index is -0.10 while that of the long-only index is 0.03. A plot of time-varying skewness of the dollar carry portfolio in Figure 7(a) shows a higher incidence of negative values as compared to the long-only index.

To summarize, indexation of currencies does reduce the occurrences of crashes. Indexation diversifies regular risks as well. As a result, the relative contribution of crash risk to the overall riskiness of the dollar index is similar to that of individual exchange rates. Conditioning on the interest rate differential in the dollar carry portfolio appears to align normal shocks in favor of a speculator. However jumps in currencies end up going against her, at least in our sample. Conditioning on interest rates has no effect on jumps in variance.

6.2 Crash risk premiums implicit in options

The large amount of jump risk prompts a natural question of whether this risk is priced. A qualitative answer to this question is borderline obvious. Given that diversification does not remove jumps, one needs an infinite number of financial instruments to hedge the jump risk perfectly (e.g., Jarrow and Madan, 1995). Thus, one would anticipate jump risk premiums in the currency markets. Answering this question quantitatively has important implications for the carry trade literature. In particular, one should be able to attribute a specific fraction of the carry risk premium to compensation for bearing crash risk.

The full answer to this question requires an explicit model of the pricing kernel and the use of assets that are sensitive to jump risk for estimation of risk premiums. In this regard, out-of-the-money options are particularly informative about the price of jump risks (i.e., the covariance of the pricing kernel with jumps). However, such analysis is outside of the

[15] Our numbers are somewhat higher than reported previously in the literature. We can potentially improve performance of the dollar carry portfolio because we are considering daily data and ignoring transaction costs.

scope of this paper. Instead, we provide a back-of-the-envelope computation, which we view as preliminary evidence of priced jump risk.

We check what kind of assumptions about risk premiums are needed to fit a "representative" implied volatility smile. We use option data for our original three bilateral exchange rates (AUD, GBP, and JPY) from October 16, 2003 to December 31, 2010, and for CHF from January 5, 2006 to December 31, 2010 – we start as early as data on five strikes become available (0.25 and 0.10 risk reversals and butterflies, ATM options). We compute the average smiles for one-month options on each exchange rate. Then for each currency, we find a specific day such that implied volatility on this day closely resembles the average smile for the respective currency. We do so because we can use the state variables from that day (the observed domestic and foreign interest rates and the estimated variance). Figure 8 depicts the resulting data points with asterisks.

Next, we construct theoretical options prices and the respective implied volatilities using different assumptions about risk premiums. In our reduced-form framework the simplest way to specify risk premiums is to assume that parameters under the risk-adjusted probability are different from the parameters estimated from the time-series of returns. We distinguish risk-adjusted parameters by an asterisk $*$. The "volatility of volatility" σ_v and the "leverage" parameter ρ are the same under both probabilities by convention, a byproduct of the model's continuous-time origins. The risk-adjusted drift μ_t^* of returns (rather than excess returns) is constrained by the law of one price. Specifically, we have

$$\tilde{q}_t^1 = q_t^1 E_t^*(S_{t+1}/S_t),$$

where q_t^1 and $\tilde{q}_t^1$ are domestic and foreign one-period risk-free bonds, respectively. Dividing by q_t^1 and taking logs we obtain

$$\begin{aligned} r_t - \tilde{r}_t &= \log E_t^*(e^{s_{t+1}-s_t}) = \mu_t^* + v_t/2 + h_t^{u*}(E_t^* \exp(z_{t+1}^u) - 1) + h_t^{d*}(E_t^* \exp(z_{t+1}^d) - 1) \\ &= \mu_t^* + v_t/2 + h_t^{u*}((1-\theta_u^*)^{-1} - 1) + h_t^{d*}((1+\theta_d^*)^{-1} - 1). \end{aligned}$$

Thus,

$$\mu_t^* = r_t - \tilde{r}_t - v_t/2 - h_t^{u*}\theta_u^*/(1-\theta_u^*) + h_t^{d*}\theta_d^*/(1+\theta_d^*). \tag{6.1}$$

The terms following the interest rate differential reflect the "convexity adjustment" associated with log-returns.

In our first exercise, we do not allow for any jump risk premiums. That is, we try to fit the smiles in Figure 8 by minimizing the sum of squared implied volatility errors and by allowing for deviation between v and v^*, and ν and ν^* only. Our best fit is depicted with green lines. It is evident that without jump risk premiums we cannot get anywhere close to the observed smiles. The theoretical risk-adjusted probabilities of extreme outcomes are too low relative to those in the data.

In our second exercise, we allow jump risk to be priced as well. However, we have too much flexibility in selecting the risk-adjusted jump parameters. First, the jump mean and jump intensity are not identified separately in practice because they approximately enter option pricing formulas as a product (e.g., Broadie, Chernov, and Johannes, 2007). Therefore, we constrain the jump means to be the same under both probabilities. Second, jump intensities could be time-varying, but we are looking at only one day per currency. Therefore, we cannot separate differences between risk-adjusted parameters that control the constant and time-varying part of jump intensities. Therefore, we only allow for $h_0^{v*} \neq h_0^v$, and $h_0^* \neq h_0$. While we find the constant components of h_t^u and h_t^d to be the same under the actual probability, this does not have to be the case under the risk-adjusted probability. Thus, instead of calibrating just h_0^* we distinguish h_0^{u*} and h_0^{d*}.

Figure 8 depicts the results of the second exercise with solid lines. The smiles can be fitted perfectly. While our calibration is limited in its scope, it clearly suggests that jump risks are priced. In what follows, we provide some discussion to help our readers with interpretation of the jump risk premium parameters. Before we do so, we wish to emphasize that the key conclusion is that without jump risk premiums, one cannot fit options even in the simplest case of one representative day with only one maturity. The flip side of the simplicity of this setting is that the calibrated jump risk premiums may be unreliable.

Tables 2 – 5 report the calibrated values of the risk-adjusted parameters. Parameters μ_0^*, μ_r^* and μ_v^* are not calibrated: they are determined by the law of one price via equation (6.1).[16] To interpret these parameters, consider GBP as an example. First, the risk-adjusted parameters are larger than their time-series counterparts. Option investors assign higher probabilities of extreme outcomes when valuing assets that are sensitive to these outcomes. Second, the smile is skewed to the left which indicates that, relatively speaking, out-of-the-money calls are less expensive than out-of-the-money puts. Therefore, extreme negative outcomes (GBP depreciation) are more likely than extreme positive outcomes. The risk-adjusted jump intensities of $h_0^{u*} = 0.0329$ and $h_0^{d*} = 0.0723$ reflect this asymmetry. The effect of a jump in variance is more subtle. As jump intensity increases, so does the overall level of variance (Appendix A.2). As a result, the relative contribution of jumps in currencies declines, resulting in fewer extreme negative outcomes and thereby skewing the smile to the right. Therefore, the risk-adjusted intensity of jump in variance cannot be too high to match the observed smile.

It is customary to measure risk premiums as the difference between risk-adjusted and actual parameters. To see why, consider a hypothetical security that pays \$1 if there are no jumps next period. The price of this security is

$$P_t^k = e^{-r_t} E_t^*(I_{j_{t+1}^k = 0}) = e^{-r_t} P_t^*(j_{t+1}^k = 0) = e^{-r_t} e^{-h_t^{k*}},$$

where $k = u, d, v$, I_x is the indicator function of event x, and P_t^* is the risk-adjusted conditional cumulative distribution function. The excess log expected return on this security,

[16] Readers may wonder why the loading of -0.5 on v_t in (6.1) is replaced by -0.005 in the Tables. The adjustment of -0.5 corresponds to all quantities being expressed in decimals. We express $v_t^{1/2}$ in percent. Thus, in order to get a sensible result for log-returns expressed in percent, the easiest solution is to use -0.005 and keep other quantities as they are.

a.k.a. the risk premium, is:

$$rpz_t^k \equiv \log \frac{E_t(I_{j_{t+1}^k=0})}{P_t^k} - r_t = h_t^{k*} - h_t^k.$$

The risk premiums range (in percent per day) from 0.02 (JPY) to 0.07 (GBP) for $k = d$, from 0.01 (AUD) to 0.07 (JPY) for $k = u$, and from 0.01 (AUD) to 0.03 (JPY) for $k = v$. All risk premiums are positive reflecting the aforementioned higher risk-adjusted probability of jumps. The variance jump premiums are smaller than those of currency jumps, but they are of the same order of magnitude. Whether the magnitudes of these risk premiums are reliable and sensible is hard to establish without more elaborate empirical work and a reference equilibrium model, respectively. We leave these endeavours for future research.

7 Conclusion

We explore sources of risk affecting currency returns. We find a large time-varying component that is attributable to jump risk. The most interesting part of this finding is that jumps in currency variance play an important role, especially at long (quarterly) investment horizons, yet there is no obvious link between macroeconomic fundamentals and these jumps. We interpret this evidence as a manifestation of economic uncertainty.

We see at least two important directions in which our analysis can be extended. First, the prices of financial assets associated with currencies (e.g., bonds, options) can be used to estimate the pricing kernel. This would allow one to characterize how the risks documented in this paper are valued in the marketplace. Second, existing research presents evidence of common and currency-specific factors affecting risk premiums. Our evidence suggests informally that common jump risks are shared across different currencies. It would be useful to establish a model of joint currency behaviour that explicitly incorporates common and country-specific jump components and how they contribute to risk premiums.

References

Aït-Sahalia, Yacine, Julio Cacho-Diaz, and Roger Laeven, 2011, Modeling financial contagion using mutually exciting jump processes, working paper, Princeton University.

Aït-Sahalia, Yacine, and Robert Kimmel, 2007, Maximum likelihood estimation of stochastic volatility models, *Journal of Financial Economics* 83, 413–452.

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

Andersen, Torben G., Luca Benzoni, and Jesper Lund, 2002, An empirical investigation of continuous-time equity return models, *Journal of Finance* 57, 1239–1284.

Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Clara Vega, 2003, Micro effects of macro announcements: Real-time price discovery in foreign exchange, *American Economic Review* 93, 38–62.

Bacchetta, Philippe, and Eric van Wincoop, 2006, Can information heterogeneity explain the exchange rate determination puzzle, *American Economic Review* 96, 552–576.

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

Backus, David, Mikhail Chernov, and Stanley Zin, 2013, Sources of entropy in representative agent models, Journal of Finance, forthcoming.

Backus, David K., Federico Gavazzoni, Christopher Telmer, and Stanley E. Zin, 2010, Monetary policy and the uncovered interest rate parity puzzle, Working paper, NBER.

Baillie, Richard T., and Tim Bollerslev, 1989, The message in daily exchange rates: A conditional-variance tale, *Journal of Business and Economic Statistics* 7, 297–305.

Bakshi, Gurdip, Peter Carr, and Liuren Wu, 2008, Stochastic risk premiums, stochastic skewness in currency options, and stochastic discount factors in international economies, *Journal of Financial Economics* 87, 132–156.

Bansal, Ravi, and Ivan Shaliastovich, 2012, A long-run risks explanation of predictability puzzles in bond and currency markets, *Review of Financial Studies* 26, 1–33.

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

Barroso, Pedro, and Pedro Santa-Clara, 2011, Beyond the carry trade optimal currency portfolios, Working paper, NOVA School of Business and Economics.

Bates, David S., 1996, Jumps and stochastic volatility: Exchange rate processes implicit in deutsche mark options, *Review of Financial Studies* 9, 69–107.

Beber, Alessandro, Francis Breedon, and Andrea Buraschi, 2010, Differences in beliefs and currency risk premiums, *Journal of Financial Economics* 98, 415 – 438.

Bekaert, Geert, 1995, The time variation of expected returns and volatility in foreign-exchange markets, *Journal of Business and Economic Statistics* 13, 397–408.

———, 1996, The time variation of risk and return in foreign exchange markets: a general equilibrium perspective, *Review of Financial Studies* 9, 427–470.

———, and Robert J. Hodrick, 1992, Characterizing predictable components in excess returns on equity and foreign exchange markets, *Journal of Finance* 47, 467–509.

———, 1993, On biases in the measurement of foreign exchange risk premiums, *Journal of International Money and Finance* 12, 115–138.

Benzoni, Luca, Pierre Collin-Dufresne, and Robert S. Goldstein, 2011, Explaining asset pricing puzzles associated with the 1987 market crash, *Journal of Financial Economics* 101, 552–573.

Bilson, John, 1981, The "speculative efficiency" hypothesis, *Journal of Business* 54, 435–451.

Brandt, Michael W., and Pedro Santa-Clara, 2002, Simulated likelihood estimation of diffusions with an application to exchange rate dynamics in incomplete markets, *Journal of Financial Economics* 63, 161–210.

Brennan, Michael J., and Yihong Xia, 2006, International capital markets and foreign exchange risk, *Review of Financial Studies* 19, 753–795.

Broadie, Mark, Mikhail Chernov, and Michael Johannes, 2007, Model specification and risk premia: Evidence from futures options, *Journal of Finance* 62, 1453–1490.

Brunnermeier, Markus K., Stephan Nagel, and Lasse Pedersen, 2008, Carry trades and currency crashes, in Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, ed.: *NBER Macroeconomics Annual, 23(1), 313-347* (University of Chicago Press).

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

Campa, Jose M., P.H. Kevin Chang, and Robert L. Reider, 1998, Implied exchange rate distributions: Evidence from OTC option markets, *Journal of International Money and Finance* 17, 117–160.

Carr, Peter, and Liuren Wu, 2007, Stochastic skew in currency options, *Journal of Financial Economics* 86, 213–247.

———, 2011, Leverage effect, volatility feedback, and self-exciting market disruptions, Working paper, Baruch College.

Chernov, Mikhail, 2007, On the role of risk premia in volatility forecasting, *Journal of Business and Economic Statistics* 25, 411–426.

———, A. Ronald Gallant, Eric Ghysels, and George Tauchen, 2003, Alternative models for stock price dynamics, *Journal of Econometrics* 116, 225–257.

Chernov, Mikhail, and Eric Ghysels, 2000, A study towards a unified approach to the joint estimation of objective and risk neutral measures for the purposes of options valuation, *Journal of Financial Economics* 56, 407–458.

Cheung, Sam, 2008, An empirical analysis of joint time-series of returns and the term-structure of option prices, Working paper, Columbia University.

Colacito, Riccardo, 2009, Six anomalies looking for a model. A consumption based explanation of international finance puzzles, Working Paper, University of North Carolina at Chapel Hill.

———, and Mariano M. Croce, 2013, International asset pricing with recursive preferences, Journal of Finance, forthcoming.

Domowitz, Ian, and Craig S. Hakkio, 1985, Conditional variance and the risk premium in the foreign exchange market, *Journal of International Economics* 19, 47–66.

Drechsler, Itamar, and Amir Yaron, 2011, What's vol got to do with it, *Review of Financial Studies* 24, 1–45.

Duffie, Darrell, Jun Pan, and Kenneth Singleton, 2000, Transform analysis and asset pricing for affine jump-diffusions, *Econometrica* 68, 1343–1376.

Engel, Charles, and James D. Hamilton, 1990, Long swings in the dollar: Are they in the data and do markets know it?, *American Economic Review* 80, 689–713.

Engle, Robert F., Takatoshi Ito, and Wen-Ling Lin, 1990, Meteor showers or heat waves? heteroskedastic intra-daily volatility in the foreign exchange market, *Econometrica* 58, 525–542.

Eraker, Bjorn, 2004, Do stock prices and volatility jump? Reconciling evidence from spot and option prices, *Journal of Finance* 59, 1367–1404.

———, Michael Johannes, and Nicholas Polson, 2003, The impact of jumps in volatility and returns, *Journal of Finance* 58, 1269–1300.

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

Farhi, Emmanuel, Samuel P. Fraiberger, Xavier Gabaix, Romain Ranciere, and Adrien Verdelhan, 2009, Crash risk in currency markets, Working Paper, NBER.

Farhi, Emmanuel, and Xavier Gabaix, 2008, Rare disasters and exchange rates, Working paper, NBER.

Gourieroux, C., and J. Jasiak, 2006, Autoregressive gamma processes, *Journal of Forecasting* 25, 129–152.

Gourio, Francois, Michael Siemer, and Adrien Verdelhan, 2013, International risk cycles, *Journal of International Economics* 89, 471–484.

Graveline, Jeremy J., 2006, Exchange rate volatility and the forward premium anomaly, Working Paper, University of Minnesota.

Guo, Kai, 2007, Exchange rates and asset prices in an open economy with rare disasters, Working Paper, Harvard University.

Hansen, Lars Peter, John C. Heaton, and Nan Li, 2008, Consumption strikes back? Measuring long-run risk, *Journal of Political Economy* 116, 260–302.

Harvey, Campbell R., and Roger D. Huang, 1991, Volatility in the foreign currency futures market, *Review of Financial Studies* 4, 543–569.

Hawkes, Alan G., 1971, Spectra of some self-exciting and mutually exciting point processes, *Biometrika* 58, 83–90.

Jacquier, Eric, Nicholas G. Polson, and Peter E. Rossi, 2004, Bayesian analysis of stochastic volatility models with fat-tails and correlated errors, *Journal of Econometrics* 122, 185–212.

Jarrow, Robert, and Dilip Madan, 1995, Option pricing using the term structure of interest rates to hedge systematic discontinuities in asset returns, *Mathematical Finance* 5, 311–336.

Johannes, Michael, 2004, The statistical and economic role of jumps in continuous-time interest rate models, *Journal of Finance* 59, 227–260.

———, and Nicholas G. Polson, 2009, MCMC methods for continuous-time financial econometrics, in Yacine Aït-Sahalia, and Lars Hansen, ed.: *Handbook of Financial Econometrics, 1-72* (Elsevier: Oxford).

Johnson, Timothy C., 2002, Volatility, momentum, and time-varying skewness in foreign exchange returns, *Journal of Business and Economic Statistics* 20, 390–411.

Jones, Christopher S., 2003, The dynamics of stochastic volatility: Evidence from underlying and options markets, *Journal of Econometrics* 116, 181–224.

Jordà, Òscar, and Alan M. Taylor, 2012, The carry trade and fundamentals: Nothing to fear but feer itself, *Journal of International Economics* 88, 74–90.

Jorion, Philippe, 1988, On jump processes in the foreign exchange and stock markets, *Review of Financial Studies* 1, 427–445.

Jurek, Jakub W., 2009, Crash-neutral currency carry trades, Working Paper, Princeton University.

Kass, Robert E., and Adrien E. Raftery, 1995, Bayes factors, *Journal of the American Statistical Association* 90, 773–795.

Le, Anh, Kenneth J. Singleton, and Qiang Dai, 2010, Discrete-time affine term structure models with generalized market prices of risk, *Review of Financial Studies* 23, 2184–2227.

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

——— , 2013, Countercyclical currency risk premia, Journal of Financial Economics, forthcoming.

Lustig, Hanno, and Adrien Verdelhan, 2007, The cross section of foreign currency risk premia and consumption growth risk, *The American Economic Review* 97, 89–117.

Maheu, John M., and Thomas H. McCurdy, 2008, Modeling foreign exchange rates with jumps, in David E. Rapach, and Mark E. Wohar, ed.: *Forecasting in the Presence of Structural Breaks and Model Uncertainty*, vol. 3 chap. 12, pp. 449–475.

Martin, Ian, 2013, Simple variance swaps, Working paper, NBER.

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

Nowotny, Michael, 2011, Disaster begets crisis: The role of contagion in financial markets, Working paper, Boston University.

Pan, Jun, 2002, The jump-risk premia implicit in options: Evidence from an integrated time-series study, *Journal of Financial Economics* 63, 3–50.

Pastorello, Sergio, Eric Renault, and Nizar Touzi, 2000, Statistical inference for random-variance option pricing, *Journal of Business and Economic Statistics* 18, 358–367.

Piazzesi, Monika, 2005, Bond yields and the Federal Reserve, *Journal of Political Economy* 113, 311–344.

Plantin, Guillaume, and Hyun Song Shin, 2011, Carry trades, monetary policy and speculative dynamics, Working Paper, CEPR.

Platen, E., and R. Rebolledo, 1985, Weak convergence of semimartingales and discretisation methods, *Stochastic Processes and their Applications* 20, 41–58.

Ready, Robert, Nikolai Roussanov, and Colin Ward, 2013, Commodity trade and the carry trade: a tale of two countries, Working paper, NBER.

Santa-Clara, Pedro, and Shu Yan, 2010, Crashes, volatility, and the equity premium: Lessons from S&P 500 options, *Review of Economics and Statistics* 92, 435–451.

Sarno, Lucio, Paul G. Schneider, and Christian Wagner, 2012, Properties of foreign exchange risk premiums, *Journal of Financial Economics* 105, 279–310.

Tryon, Ralph, 1979, Testing for rational expectations in foreign exchange markets, Working paper, Federal Reserve.

Verdelhan, Adrien, 2011, The share of systematic variation in bilateral exchange rates, Working Paper, MIT.

Wachter, Jessica, 2013, Can time-varying risk of rare disasters explain aggregate stock market volatility?, *Journal of Finance* 68, 987–1035.

Table 1
Properties of excess log currency and S&P 500 returns and changes in implied volatility

		Mean	Std Dev	Skewness	Kurtosis	Nobs
AUD	Excess return	0.0186	0.7435	-0.3870	13.7202	6332
	$\Delta\sqrt{IV}$	0.0109	3.7661	0.9077	9.7290	3933
CHF	Excess return	0.0057	0.7232	0.1194	4.7841	6521
	$\Delta\sqrt{IV}$	0.0073	3.8057	0.9966	9.8095	3823
GBP	Excess return	0.0096	0.6197	-0.2337	5.6832	6521
	$\Delta\sqrt{IV}$	0.0142	4.0001	1.3884	44.2683	3823
JPY	Excess return	0.0003	0.6950	0.3626	8.0878	6393
	$\Delta\sqrt{IV}$	-0.0045	4.8277	1.0395	10.7764	3934
Currency Index	Excess return	0.0086	0.4664	0.0342	4.6448	6393
S&P 500	Excess return	0.0090	1.1803	-1.3584	32.9968	6522
	$\Delta\sqrt{VIX}$	0.0089	5.8997	0.5096	6.7502	3914

Notes. Descriptive statistics for daily log currency and log S&P 500 excess returns and changes in implied volatility, in percent, per day: AUD return from September 25, 1986, to December 31, 2010, and AUD IV from December 6, 1995, to December 31, 2010; CHF return from January 3, 1986, to December 31, 2010, and CHF IV from May 8, 1996, to December 31, 2010; GBP return from January 3, 1986, to December 31, 2010, and GBP IV from May 8, 1996, to December 31, 2010; JPY return from July 2, 1986, to December 31, 2010, and JPY IV from December 5, 1995, to December 31, 2010; Currency Index return is from January 2, 1986, to December 31, 2010; S&P 500 return from January 3, 1986, to December 31, 2010 and VIX from January 2, 1996, to December 31, 2010. The interest rates used to compute returns are one-month LIBOR rates. Source: Bloomberg.

Table 2
AUD Parameter Estimates

	SV ($\theta = 0$, $\theta_v = 0$)	SVJ ($\theta = 0$)	Preferred	Risk-adjusted
μ_0	-0.0004	0.0003	0.0014	0.0640
	(-0.0181,0.0173)	(-0.0175, 0.0182)	(-0.0166, 0.0194)	
μ_r	-2.4893	-2.5200	-2.7643	$-0.2988^u/0.2888^d$
	(-3.5895, -1.3910)	(-3.6170, -1.4190)	(-3.8613, -1.6801)	
μ_v	-0.0150	-0.0152	-0.0147	-0.0050
	(-0.0247, -0.0053)	(-0.0249, -0.0056)	(-0.0244, -0.0051)	
v	0.4968	0.2819	0.2819	0.2165
	(0.2903, 0.8984)	(0.2101, 0.3728)	(0.2110, 0.3734)	
ν	0.9943	0.9855	0.9857	0.8198
	(0.9925, 0.9961)	(0.9837, 0.9873)	(0.9838, 0.9875)	
σ_v	0.0391	0.0343	0.0342	σ_v
	(0.0379, 0.0404)	(0.0330, 0.0357)	(0.0329, 0.0356)	
ρ	-0.2924	-0.2770	-0.2839	ρ
	(-0.3279, -0.2563)	(-0.3156, -0.2378)	(-0.3237, -0.2435)	
θ_v		0.3864	0.3837	θ_v
		(0.3392, 0.4406)	(0.3367, 0.4362)	
θ			1.6910	θ
			(1.5208, 1.8779)	
h_0^v		0.0037	0.0036	0.0098
		(0.0029, 0.0040)	(0.0028, 0.0040)	
h_v		0.0038	0.0038	h_v
		(0.0031, 0.0040)	(0.0031, 0.0040)	
h_0			0.0017	$0.0136^u/0.0525^d$
			(0.0000, 0.0038)	
h_r			0.1737	h_r
			(0.1177, 0.1992)	
α_{iv}	0.0033	0.0027	0.0033	
	(0.0009, 0.0064)	(-0.0002, 0.0056)	(0.0005, 0.0063)	
β_{iv}	1.0006	1.0021	1.0022	
	(0.9958, 1.0054)	(0.9983, 1.0059)	(0.9984, 1.0061)	

Notes. The estimates correspond to daily excess currency returns, in percent. The 95% confidence intervals are reported in parentheses. The preferred model is:

$$\begin{aligned} y_{t+1} &= \mu_0 + \mu_r(r_t - \tilde{r}_t) + \mu_v v_t + v_t^{1/2} w_{t+1}^s + z_{t+1}^u - z_{t+1}^d \\ v_{t+1} &= (1-\nu)v + \nu v_t + \sigma_v v_t^{1/2} w_{t+1}^v + z_{t+1}^v \\ h_t^u &= h_0 + h_r r_t, \ h_t^d = h_0 + h_r \tilde{r}_t, \ h_t^v = h_0^v + h_v v_t \\ z_t^u | j &\sim \mathcal{G}\text{amma}(j, \theta), \ z_t^d | j \sim \mathcal{G}\text{amma}(j, \theta), \ z_t^v | j \sim \mathcal{G}\text{amma}(j, \theta_v) \end{aligned}$$

"Risk-adjusted" is the preferred model calibrated to one-month options (with parameters that are denoted by symbols fixed at their time-series values).

Table 3
CHF Parameter Estimates

	SV ($\theta = 0$, $\theta_v = 0$)	SVJ ($\theta = 0$)	Preferred	Risk-adjusted
μ_0	0.0340	0.0353	0.0324	-0.0138
	(0.0150, 0.0531)	(0.0163, 0.0543)	(0.0132, 0.0516)	
μ_r	-2.9851	-3.0674	-3.2501	$-0.2995^u/0.2914^d$
	(-4.3345, -1.6354)	(-4.4174, -1.7169)	(-4.5952, -1.8996)	
μ_v	-0.0198	-0.0199	-0.0199	-0.0050
	(-0.0333, -0.0064)	(-0.0335, -0.0065)	(-0.0334, -0.0064)	
v	0.5088	0.3502	0.3427	0.2377
	(0.3136, 0.8364)	(0.2564, 0.4741)	(0.2507, 0.4639)	
ν	0.9891	0.9789	0.9785	0.7324
	(0.9863, 0.9919)	(0.9758, 0.9818)	(0.9755, 0.9815)	
σ_v	0.0388	0.0337	0.0337	σ_v
	(0.0373, 0.0404)	(0.0321, 0.0352)	(0.0321, 0.0352)	
ρ	0.0875	0.0856	0.0856	ρ
	(0.0480, 0.1271)	(0.0439, 0.1273)	(0.0416, 0.1298)	
θ_v		0.2205	0.2145	θ_v
		(0.1845, 0.2622)	(0.1804, 0.2550)	
θ			1.3582	θ
			(1.1771, 1.5744)	
h_0^v		0.0037	0.0037	0.0212
		(0.0029, 0.0040)	(0.0028, 0.0040)	
h_v		0.0145	0.0144	h_v
		(0.0131, 0.0150)	(0.0130, 0.0150)	
h_0			0.0051	$0.0376^u/0.0283^d$
			(0.0011, 0.0078)	
h_r			0.2175	h_r
			(0.0615, 0.2973)	
α_{iv}	0.0061	0.0041	0.0056	
	(0.0009, 0.0113)	(-0.0013, 0.0093)	(0.0003, 0.0104)	
β_{iv}	0.9919	0.9934	0.9946	
	(0.9753, 1.0091)	(0.9795, 1.0071)	(0.9802, 1.0096)	

Notes. The estimates correspond to daily excess currency returns, in percent. The 95% confidence intervals are reported in parentheses. The preferred model is:

$$\begin{aligned}
y_{t+1} &= \mu_0 + \mu_r(r_t - \tilde{r}_t) + \mu_v v_t + v_t^{1/2} w_{t+1}^s + z_{t+1}^u - z_{t+1}^d \\
v_{t+1} &= (1-\nu)v + \nu v_t + \sigma_v v_t^{1/2} w_{t+1}^v + z_{t+1}^v \\
h_t^u &= h_0 + h_r r_t,\ h_t^d = h_0 + h_r \tilde{r}_t,\ h_t^v = h_0^v + h_v v_t \\
z_t^u | j &\sim \mathcal{G}\text{amma}(j,\theta),\ z_t^d | j \sim \mathcal{G}\text{amma}(j,\theta),\ z_t^v | j \sim \mathcal{G}\text{amma}(j,\theta_v)
\end{aligned}$$

"Risk-adjusted" is the preferred model calibrated to one-month options (with parameters that are denoted by symbols fixed at their time-series values).

Table 4
GBP Parameter Estimates

	SV ($\theta = 0$, $\theta_v = 0$)	SVJ ($\theta = 0$)	Preferred	Risk-adjusted
μ_0	0.0348	0.0360	0.0337	0.0445
	(0.0128, 0.0565)	(0.0138, 0.0584)	(0.0116, 0.0556)	
μ_r	-3.0632	-3.1692	-3.1897	-0.1445^u/0.1412^d
	(-4.4127, -1.7209)	(-4.5046, -1.8351)	(-4.5138, -1.8673)	
μ_v	-0.1326	-0.1377	-0.1341	-0.0050
	(-0.1928, -0.0720)	(-0.1986, -0.0773)	(-0.1952, -0.0731)	
v	0.3773	0.2227	0.2180	0.1015
	(0.1855, 0.8002)	(0.1631, 0.2989)	(0.1619, 0.2909)	
ν	0.9941	0.9810	0.9809	0.7544
	(0.9919, 0.9963)	(0.9786, 0.9834)	(0.9786, 0.9833)	
σ_v	0.0321	0.0272	0.0273	σ_v
	(0.0311, 0.0332)	(0.0262, 0.0283)	(0.0263, 0.0284)	
ρ	-0.1341	-0.1295	-0.1303	ρ
	(-0.1713, -0.0965)	(-0.1692, -0.0896)	(-0.1709, -0.0895)	
θ_v		0.1953	0.1959	θ_v
		(0.1728, 0.2206)	(0.1731, 0.2211)	
θ			1.1680	θ
			(0.9593, 1.4127)	
h_0^v		0.0038	0.0038	0.0133
		(0.0033, 0.0040)	(0.0033, 0.0040)	
h_v		0.0121	0.0121	h_v
		(0.0110, 0.0125)	(0.0110, 0.0125)	
h_0			0.0012	0.0329^u/0.0723^d
			(0.0001, 0.0020)	
h_r			0.1223	h_r
			(0.0634, 0.1491)	
α_{iv}	0.0109	0.0063	0.0089	
	(0.0063, 0.0155)	(0.0009, 0.0109)	(0.0039, 0.0137)	
β_{iv}	0.9905	0.9951	0.9940	
	(0.9855, 0.9955)	(0.9905, 0.9996)	(0.9895, 0.9985)	

Notes. The estimates correspond to daily excess currency returns, in percent. The 95% confidence intervals are reported in parentheses. The preferred model is:

$$
\begin{aligned}
y_{t+1} &= \mu_0 + \mu_r(r_t - \tilde{r}_t) + \mu_v v_t + v_t^{1/2} w_{t+1}^s + z_{t+1}^u - z_{t+1}^d \\
v_{t+1} &= (1-\nu)v + \nu v_t + \sigma_v v_t^{1/2} w_{t+1}^v + z_{t+1}^v \\
h_t^u &= h_0 + h_r r_t,\ h_t^d = h_0 + h_r \tilde{r}_t,\ h_t^v = h_0^v + h_v v_t \\
z_t^u | j &\sim \mathcal{G}\text{amma}(j, \theta),\ z_t^d | j \sim \mathcal{G}\text{amma}(j, \theta),\ z_t^v | j \sim \mathcal{G}\text{amma}(j, \theta_v)
\end{aligned}
$$

"Risk-adjusted" is the preferred model calibrated to one-month options (with parameters that are denoted by symbols fixed at their time-series values).

Table 5
JPY Parameter Estimates

	SV ($\theta = 0$, $\theta_v = 0$)	SVJ ($\theta = 0$)	Preferred	Risk-adjusted
μ_0	0.0253	0.0253	0.0203	-0.0717
	(0.0064, 0.0441)	(0.0062, 0.0443)	(0.0018, 0.0389)	
μ_r	-3.1861	-3.2046	-3.4590	$-0.5561^u/0.5426^d$
	(-4.4200, -1.9526)	(-4.4540, -1.9531)	(-4.6992, -2.2266)	
μ_v	0.0151	0.0152	0.0152	-0.0050
	(0.0054, 0.0248)	(0.0055, 0.0249)	(0.0054, 0.0248)	
v	0.4816	0.3143	0.3012	0.2382
	(0.2926, 0.8111)	(0.2328, 0.4223)	(0.2207, 0.4079)	
ν	0.9896	0.9762	0.9771	0.7022
	(0.9868, 0.9924)	(0.9730, 0.9794)	(0.9739, 0.9802)	
σ_v	0.0496	0.0438	0.0436	σ_v
	(0.0476, 0.0516)	(0.0419, 0.0458)	(0.0417, 0.0455)	
ρ	0.3681	0.3505	0.3631	ρ
	(0.3316, 0.4040)	(0.3098, 0.3902)	(0.3205, 0.4047)	
θ_v		0.3917	0.3771	θ_v
		(0.3313, 0.4622)	(0.3198, 0.4447)	
θ			1.2351	θ
			(1.0847, 1.4071)	
h_0^v		0.0037	0.0037	0.0331
		(0.0031, 0.0040)	(0.0029, 0.0040)	
h_v		0.0077	0.0076	h_v
		(0.0068, 0.0080)	(0.0067, 0.0080)	
h_0			0.0052	$0.0774^u/0.0206^d$
			(0.0034, 0.0060)	
h_r			0.4447	h_r
			(0.3133, 0.4984)	
α_{iv}	0.0140	0.0116	0.0159	
	(0.0086, 0.0193)	(0.0062, 0.0169)	(0.0099, 0.0214)	
β_{iv}	1.0052	1.0083	1.0248	
	(0.9871, 1.0248)	(0.9916, 1.0256)	(1.0069, 1.0431)	

Notes. The estimates correspond to daily excess currency returns, in percent. The 95% confidence intervals are reported in parentheses. The preferred model is:

$$\begin{aligned} y_{t+1} &= \mu_0 + \mu_r(r_t - \tilde{r}_t) + \mu_v v_t + v_t^{1/2} w_{t+1}^s + z_{t+1}^u - z_{t+1}^d \\ v_{t+1} &= (1-\nu)v + \nu v_t + \sigma_v v_t^{1/2} w_{t+1}^v + z_{t+1}^v \\ h_t^u &= h_0 + h_r r_t,\ h_t^d = h_0 + h_r \tilde{r}_t,\ h_t^v = h_0^v + h_v v_t \\ z_t^u | j &\sim \mathcal{G}\text{amma}(j, \theta),\ z_t^d | j \sim \mathcal{G}\text{amma}(j, \theta),\ z_t^v | j \sim \mathcal{G}\text{amma}(j, \theta_v) \end{aligned}$$

"Risk-adjusted" is the preferred model calibrated to one-month options (with parameters that are denoted by symbols fixed at their time-series values).

Table 6
Model diagnostics for AUD

	SV ($\theta = 0$, $\theta_v = 0$)	SVJ ($\theta = 0$)	Preferred
$skewness^C$	-0.3080	-0.3074	-0.2004
	(-0.3308, -0.2860)	(-0.3304, -0.2855)	(-0.2408, -0.1599)
$kurtosis^C$	4.1472	4.0822	3.4892
	(4.0677, 4.2366)	(4.0006, 4.1810)	(3.3802, 3.6055)
$autocorrelation^C$	-0.0281	-0.0271	-0.0324
	(-0.0311, -0.0252)	(-0.0303, -0.0241)	(-0.0406, -0.0242)
$skewness^{IV}$	0.0402	0.0303	0.0310
	(-0.0373, 0.1181)	(-0.0466, 0.1070)	(-0.0459, 0.1080)
$kurtosis^{IV}$	3.0618	3.0385	3.0375
	(2.9103, 3.2314)	(2.8902, 3.2034)	(2.8896, 3.2033)
$autocorrelation^{IV}$	0.1043	0.0634	0.0637
	(0.0749, 0.1336)	(0.0331, 0.0937)	(0.0334, 0.0940)
$IV var$	0.0064	0.0034	0.0034
	(0.0041, 0.0122)	(0.0021, 0.0070)	(0.0021, 0.0070)

Notes. Posterior means and 95% confidence intervals (reported in parentheses) for the residuals from the currency return and from the IV equations. Superscript C stands for the residuals from the currency return equation, superscript IV stands for the residuals from the IV equation.

Table 7
Model diagnostics for CHF

	SV ($\theta = 0$, $\theta_v = 0$)	SVJ ($\theta = 0$)	Preferred
$skewness^C$	0.1178	0.1282	0.0586
	(0.0994, 0.1365)	(0.1078, 0.1486)	(0.0182, 0.0983)
$kurtosis^C$	3.9497	3.9438	3.4333
	(3.8825, 4.0198)	(3.8919, 4.0011)	(3.3373, 3.5405)
$autocorrelation^C$	-0.0203	-0.0198	-0.0272
	(-0.0227, -0.0179)	(-0.0226, -0.0170)	(-0.0352, -0.0192)
$skewness^{IV}$	0.0224	0.0201	0.0210
	(-0.0574, 0.1022)	(-0.0585, 0.0985)	(-0.0573, 0.0995)
$kurtosis^{IV}$	3.0648	3.0399	3.0406
	(2.9091, 3.2378)	(2.8887, 3.2097)	(2.8890, 3.2094)
$autocorrelation^{IV}$	0.0777	0.0565	0.0564
	(0.0459, 0.1094)	(0.0247, 0.0883)	(0.0246, 0.0881)
$IV var$	0.0010	0.0006	0.0006
	(0.0007, 0.0017)	(0.0004, 0.0011)	(0.0004, 0.0011)

Notes. Posterior means and 95% confidence intervals (reported in parentheses) for the residuals from the currency return and from the IV equations. Superscript C stands for the residuals from the currency return equation, superscript IV stands for the residuals from the IV equation.

Table 8
Model diagnostics for GBP

	SV ($\theta = 0$, $\theta_v = 0$)	SVJ ($\theta = 0$)	Preferred
$skewness^C$	-0.0407	-0.0211	-0.0232
	(-0.0606, -0.0202)	(-0.0436, 0.0012)	(-0.0609, 0.0143)
$kurtosis^C$	3.9181	3.8540	3.4947
	(3.8427, 4.0061)	(3.7784, 3.9423)	(3.4006, 3.5969)
$autocorrelation^C$	0.0009	0.0006	-0.0027
	(-0.0024, 0.0040)	(-0.0038, 0.0047)	(-0.0094, 0.0037)
$skewness^{IV}$	0.0352	0.0212	0.0215
	(-0.0443, 0.1146)	(-0.0565, 0.0995)	(-0.0568, 0.0998)
$kurtosis^{IV}$	3.0710	3.0293	3.0296
	(2.9160, 3.2461)	(2.8798, 3.1972)	(2.8786, 3.1977)
$autocorrelation^{IV}$	0.0791	0.0510	0.0510
	(0.0483, 0.1096)	(0.0204, 0.0814)	(0.0204, 0.0815)
$IVvar$	0.0011	0.0004	0.0004
	(0.0007, 0.0019)	(0.0003, 0.0008)	(0.0003, 0.0008)

Notes. Posterior means and 95% confidence intervals (reported in parentheses) for the residuals from the currency return and from the IV equations. Superscript C stands for the residuals from the currency return equation, superscript IV stands for the residuals from the IV equation.

Table 9
Model diagnostics for JPY

	SV ($\theta = 0$, $\theta_v = 0$)	SVJ ($\theta = 0$)	Preferred
$skewness^C$	0.3348	0.3360	0.1298
	(0.3060, 0.3650)	(0.3038, 0.3668)	(0.0799, 0.1800)
$kurtosis^C$	4.8254	4.7148	3.6054
	(4.7109, 4.9645)	(4.5982, 4.8361)	(3.4829, 3.7445)
$autocorrelation^C$	-0.0146	-0.0140	-0.0221
	(-0.0176 -0.0116)	(-0.0174, -0.0108)	(-0.0312, -0.0131)
$skewness^{IV}$	0.0568	0.0278	0.0311
	(-0.0210, 0.1349)	(-0.0495, 0.1054)	(-0.0465, 0.1087)
$kurtosis^{IV}$	3.0707	3.0430	3.0423
	(2.9175, 3.2420)	(2.8940, 3.2100)	(2.8923, 3.2098)
$autocorrelation^{IV}$	0.1042	0.0758	0.0768
	(0.0733, 0.1349)	(0.0443, 0.1070)	(0.0453, 0.1083)
$IVvar$	0.0061	0.0029	0.0037
	(0.0036, 0.0125)	(0.0017, 0.0059)	(0.0021, 0.0078)

Notes. Posterior means and 95% confidence intervals (reported in parentheses) for the residuals from the currency return and from the IV equations. Superscript C stands for the residuals from the currency return equation, superscript IV stands for the residuals from the IV equation.

Table 10
Log-Bayes-Odds Ratios

	AUD	CHF	GBP	JPY	Index
SVJ/SV	22.03	52.05	44.50	34.89	24.12
Preferred/SVJ	26.36	18.77	13.43	61.22	10.13

Notes. Statistical model assesement. We compare the SV ($\theta = 0$, $\theta_v = 0$), SVJ ($\theta = 0$) and the preferred models pairwise. In the first row, we consider the SV and SVJ models and quantify evidence against the SV model; in the second row, we consider the SVJ and the preferred models and quantify evidence against the SVJ model. Focusing on the AUD, the first number tells us the odds are 22.03 to 1 that the SVJ model (the alternative) is preferable to the SV model (the null).

Table 11
Calibration of the interest rates

	AUD	CHF	GBP	JPY	Index
a_r	0.0181	0.0184	0.0184	0.0182	0.0185
$\tilde{a}_r$	0.0291	0.0121	0.0269	0.0077	0.0222
b_r	0.9999	0.9997	0.9997	0.9998	0.9998
$\tilde{b}_r$	0.9992	0.9995	0.9998	0.9997	0.9998
σ_r	0.0012	0.0016	0.0016	0.0014	0.0015
$\tilde{\sigma}_r$	0.0035	0.0030	0.0018	0.0027	0.0015

Notes. We calibrate processes for domestic (US)

$$r_{t+1} = (1 - b_r)a_r + b_r r_t + \sigma_r r_t^{1/2} w_{t+1}^r$$

and foreign interest rates

$$\tilde{r}_{t+1} = (1 - \tilde{b}_r)\tilde{a}_r + \tilde{b}_r \tilde{r}_t + \tilde{\sigma}_r \tilde{r}_t^{1/2} \tilde{w}_{t+1}^r$$

Parameters correspond to daily interest rates in percent. There are five versions of the parameters corresponding to the US interest rate. This is because the foreign data have different starting dates, and we calibrated the US rate in the matching samples.

Table 12
Summary of announcements associated with jumps

	Up	Jumps Down	Vol	Up	Jumps Down	Vol
		AUD			**CHF**	
Total	6	21	34	17	6	45
Common	2	5	17	9	3	23
Common with index	1	2	14	6	3	19
Type of announcement						
Trade	1	1	1	4	0	1
Political	0	1	0	3	0	1
Intervention	0	4	1	2	2	2
Macroeconomic	3	2	4	1	3	4
Monetary policy	1	6	5	1	0	1
		GBP			**JPY**	
Total	5	6	56	34	13	52
Common	3	1	25	12	4	19
Common with index	3	1	18	6	2	13
Type of announcement						
Trade	1	0	1	9	2	3
Political	0	1	4	2	1	1
Intervention	2	1	0	7	5	1
Macroeconomic	0	2	4	3	2	5
Monetary policy	0	0	3	1	0	5

Notes: We assess how many jumps in the preferred model are associated with economic, political or financial announcements. In the first three rows we report the total number of jumps, how many of these are common (at least two currencies jump on the same day), and how many of these are common with the dollar index. The subsequent five rows report counts of jumps that took place on the days of announcements.

Table 13
Decomposition of the total risk

	Jump Up		Jump Down		Jump in Vol		Normal	
	$n = 21$	$n = 63$	$n = 21$	$n = 63$	$n = 21$	$n = 63$	$n = 21$	$n = 63$
				AUD				
Mean	6.90	6.37	9.17	8.40	4.12	10.18	79.81	75.05
Std	3.70	2.97	5.22	4.13	1.10	1.99	9.60	8.38
Min	0.31	0.38	0.60	0.74	1.58	4.45	55.17	56.33
Max	15.84	12.41	26.24	21.34	7.14	14.73	97.45	94.32
				CHF				
Mean	6.82	6.66	5.48	5.37	3.76	8.94	83.94	79.04
Std	2.56	2.03	1.82	1.50	0.27	0.39	4.32	3.34
Min	1.02	1.31	0.98	1.26	3.13	8.01	66.78	69.49
Max	16.34	11.88	12.04	10.55	4.84	9.91	94.87	89.38
				GBP				
Mean	2.98	2.78	3.68	3.45	4.37	11.48	88.96	82.29
Std	1.51	1.19	1.68	1.38	0.80	1.35	3.73	3.42
Min	0.18	0.23	0.25	0.32	2.64	7.79	74.21	72.60
Max	8.84	5.83	9.87	8.45	7.70	15.48	96.87	91.58
				JPY				
Mean	9.10	8.57	5.80	5.43	4.83	11.19	80.27	74.82
Std	4.18	3.32	3.53	2.96	0.71	1.08	7.56	6.10
Min	0.93	1.23	0.69	0.91	3.05	7.96	54.70	57.34
Max	22.04	18.05	17.68	13.89	6.82	13.61	65.20	89.67
				Index				
Mean	10.39	10.05	10.96	10.63	3.70	8.85	74.95	70.48
Std	3.96	3.28	4.01	3.38	0.57	0.98	8.05	6.78
Min	1.37	1.76	1.91	2.48	1.93	5.24	55.85	55.56
Max	19.44	17.65	21.65	19.11	5.18	10.87	94.58	90.12

Notes. We report summary statistics of the percentage contribution of the different risks to the total risk of currency returns (horizon $n = 21$ and $n = 63$ days).

Table 14
Index Parameter Estimates

	SV ($\theta = 0$, $\theta_v = 0$)	SVJ ($\theta = 0$)	Preferred
μ_0	0.0062	0.0059	0.0050
	(-0.0042, 0.0166)	(-0.0045, 0.0163)	(-0.0055, 0.0154)
μ_r	-2.3987	-2.5243	-2.6140
	(-3.4799, -1.3162)	(-3.6003, -1.4557)	(-3.6949, -1.5336)
μ_v	-0.0208	-0.0209	-0.0209
	(-0.0345, -0.0072)	(-0.0345, -0.0073)	(-0.0346, -0.0073)
v	0.1935	0.1357	0.1335
	(0.0854, 0.4543)	(0.0911, 0.1994)	(0.0991, 0.1783)
ν	0.9941	0.9832	0.9801
	(0.9916, 0.9965)	(0.9801, 0.9859)	(0.9775, 0.9826)
σ_v	0.0199	0.0168	0.0168
	(0.0192, 0.0206)	(0.0162, 0.0174)	(0.0161, 0.0174)
ρ	-0.0552	-0.0530	-0.0508
	(-0.0944, -0.0159)	(-0.0944, -0.0115)	(-0.0932, -0.0083)
θ_v		0.1687	0.1714
		(0.1521, 0.1870)	(0.1546, 0.1902)
θ			1.6385
			(1.4635, 1.8334)
h_0^v		0.0037	0.0037
		(0.0030, 0.0040)	(0.0031, 0.0040)
h_v		0.0092	0.0093
		(0.0073, 0.0100)	(0.0075, 0.0100)
h_0			0.0020
			(0.0001, 0.0052)
h_r			0.1141
			(0.0113, 0.2542)

Notes. The estimates correspond to daily excess returns on currency index, in percent. The 95% confidence intervals are reported in parentheses. The preferred model is:

$$
\begin{aligned}
y_{t+1} &= \mu_0 + \mu_r(r_t - \tilde{r}_t) + \mu_v v_t + v_t^{1/2} w_{t+1}^s + z_{t+1}^u - z_{t+1}^d \\
v_{t+1} &= (1-\nu)v + \nu v_t + \sigma_v v_t^{1/2} w_{t+1}^v + z_{t+1}^v \\
h_t^u &= h_0 + h_r r_t,\ h_t^d = h_0 + h_r \tilde{r}_t,\ h_t^v = h_0^v + h_v v_t \\
z_t^u | j &\sim \mathcal{G}\text{amma}(j, \theta),\ z_t^d | j \sim \mathcal{G}\text{amma}(j, \theta),\ z_t^v | j \sim \mathcal{G}\text{amma}(j, \theta_v)
\end{aligned}
$$

Table 15
Model diagnostics for Currency Index

	SV ($\theta = 0$, $\theta_v = 0$)	SVJ ($\theta = 0$)	Preferred
$skewness^C$	0.0144	0.0222	0.0049
	(-0.0053, 0.0342)	(0.0032, 0.0456)	(-0.0292, 0.0389)
$kurtosis^C$	3.6921	3.6386	3.4022
	(3.6292, 3.7611)	(3.5790, 3.7057)	(3.3053, 3.5100)
$autocorrelation^C$	0.0020	0.0005	-0.0081
	(-0.0010, 0.0049)	(-0.0027, 0.0041)	(-0.0155, -0.0012)
$skewness^{IV}$	0.0220	0.0212	0.0153
	(-0.0729, 0.1173)	(-0.0743, 0.1169)	(-0.0794, 0.1097)
$kurtosis^{IV}$	3.0422	3.0419	3.0206
	(2.8601, 3.2515)	(2.8598, 3.2506)	(2.8423, 3.2248)
$autocorrelation^{IV}$	0.0604	0.0508	0.0499
	(0.0240, 0.0966)	(0.0141, 0.0873)	(0.0138, 0.0859)
$IV var$	0.0518	0.0388	0.0239
	(0.0304, 0.1077)	(0.0229, 0.0810)	(0.0155, 0.0416)

Notes. Posterior means and 95% confidence intervals (reported in parentheses) for the residuals from the currency return and from the IV equations. Superscript C stands for the residuals from the currency return equation, superscript IV stands for the residuals from the IV equation.

Figure 1
AUD data and estimated states

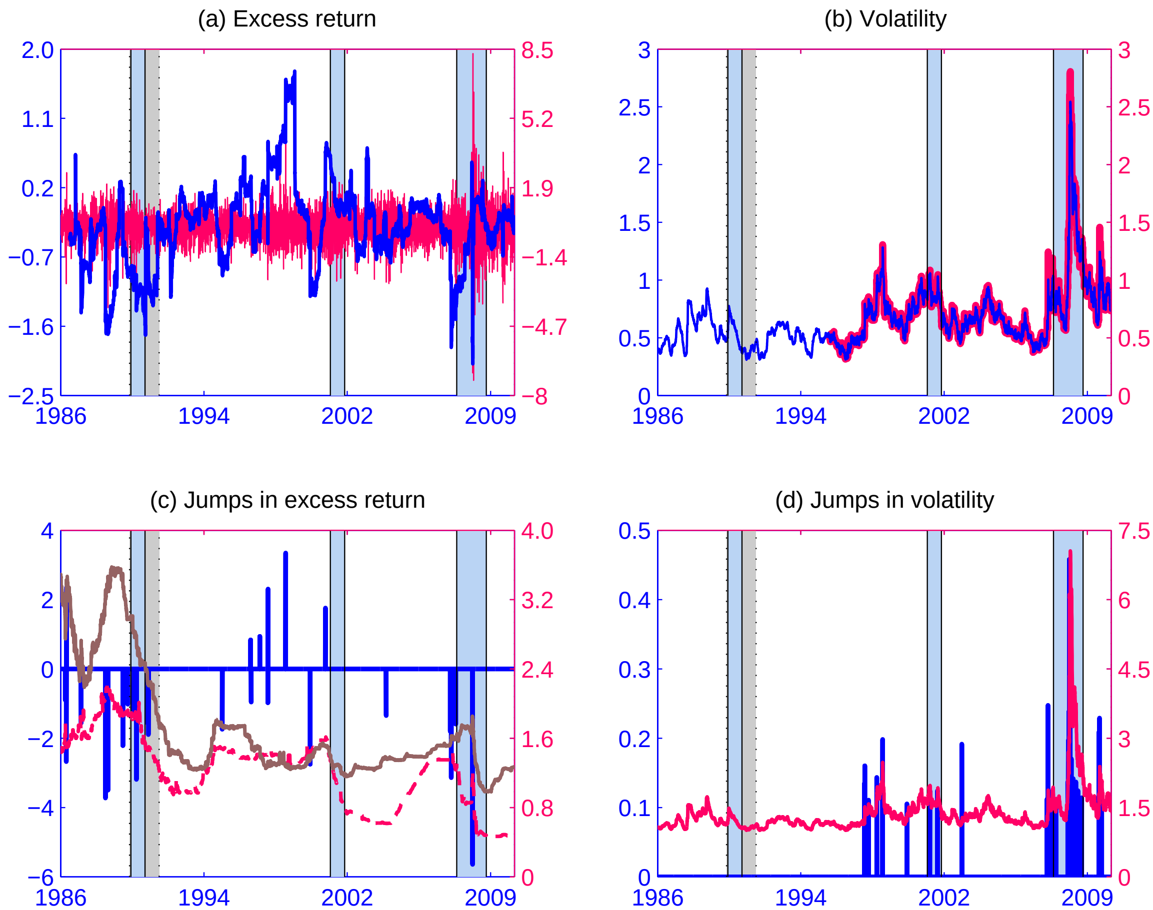

Notes. Panel (a): thin red line displays observed log currency returns; thick blue line displays the conditional 6-months skewness of the log currency returns. Panel (b): thick red line shows observed one-month at-the-money implied volatility; thin blue line shows the estimated spot volatility $\sqrt{v_t}$. The bottom panels display estimated jump sizes in returns (c) and volatility (d) with jump intensities. Panel (c): brown solid line is the intensity of the jump down; the red dashed line is the intensity of the jump up; blue bars are jumps themselves. Panel (d): the red line is the intensity of the jump in volatility; blue bars are jumps. We say that there was a jump at time t if the estimated probability of a jump on that day was above 50%. Gray vertical bars with the dashed border indicate recessions in Australia; light blue vertical lines with the the thin solid border indicate recessions in the US.

Figure 2
CHF data and estimated states

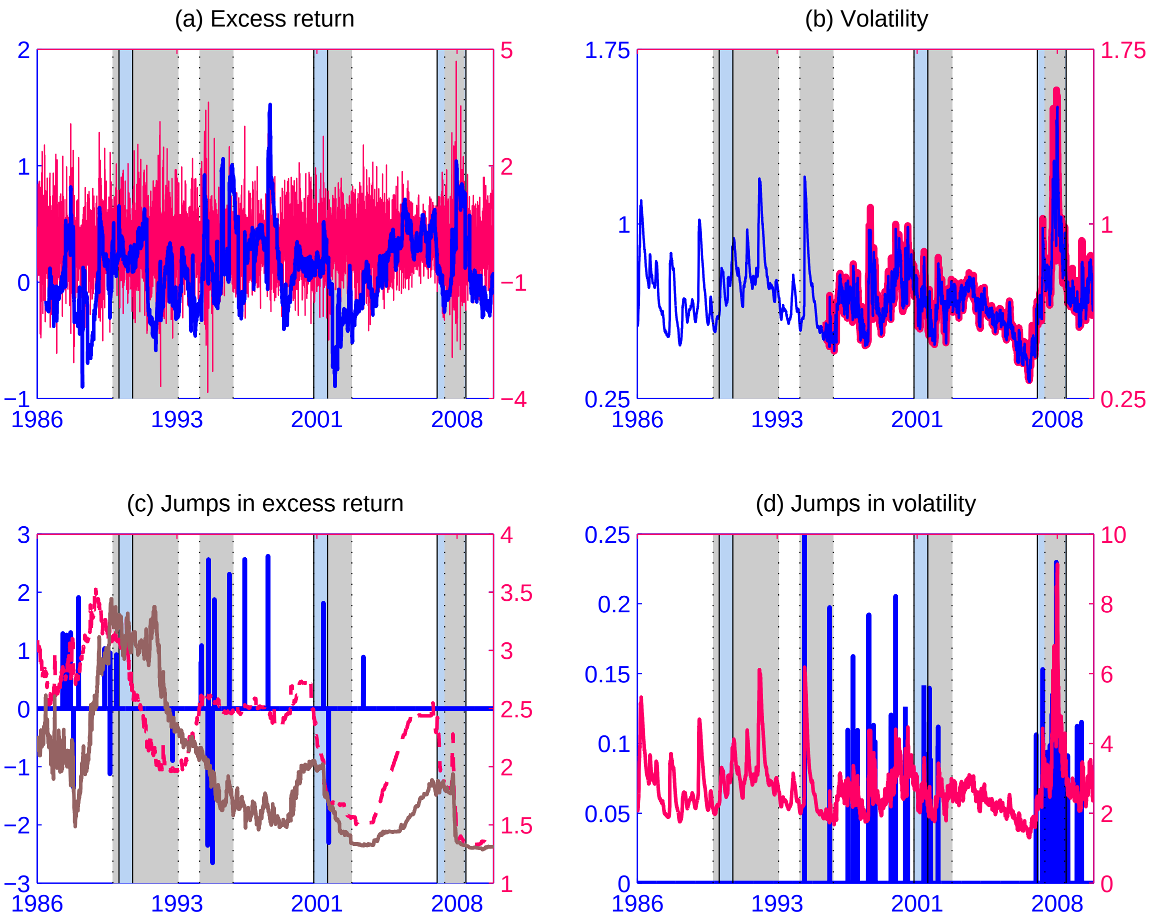

Notes. Panel (a): thin red line displays observed log currency returns; thick blue line displays the conditional 6-months skewness of the log currency returns. Panel (b): thick red line shows observed one-month at-the-money implied volatility; thin blue line shows the estimated spot volatility $\sqrt{v_t}$. The bottom panels display estimated jump sizes in returns (c) and volatility (d) with jump intensities. Panel (c): brown solid line is the intensity of the jump down; the red dashed line is the intensity of the jump up; blue bars are jumps themselves. Panel (d): the red line is the intensity of the jump in volatility; blue bars are jumps. We say that there was a jump at time t if the estimated probability of a jump on that day was above 50%. Gray vertical bars with the dashed border indicate recessions in Switzerland; light blue vertical lines with the thin solid border indicate recessions in the US.

Figure 3
GBP data and estimated states

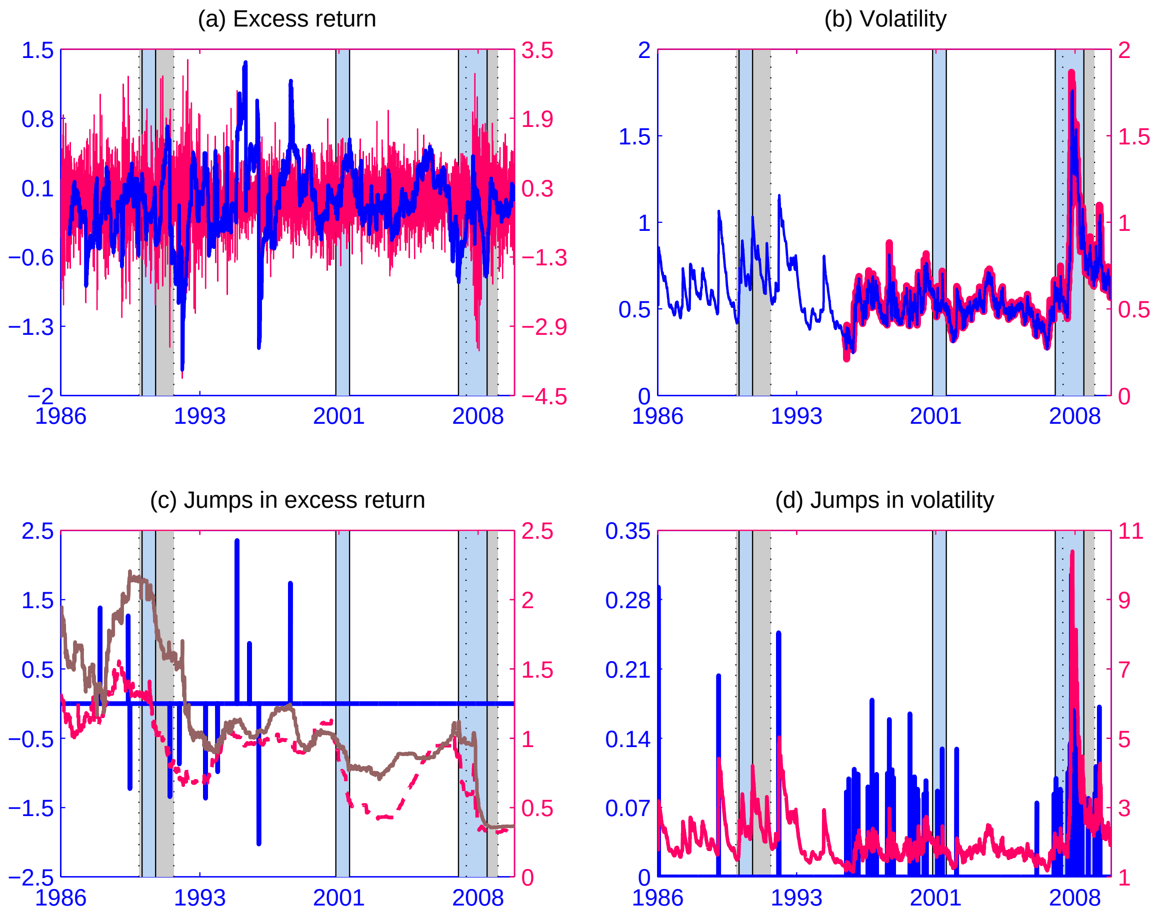

Notes. Panel (a): thin red line displays observed log currency returns; thick blue line displays the conditional 6-months skewness of the log currency returns. Panel (b): thick red line shows observed one-month at-the-money implied volatility; thin blue line shows the estimated spot volatility $\sqrt{v_t}$. The bottom panels display estimated jump sizes in returns (c) and volatility (d) with jump intensities. Panel (c): brown solid line is the intensity of the jump down; the red dashed line is the intensity of the jump up; blue bars are jumps themselves. Panel (d): the red line is the intensity of the jump in volatility; blue bars are jumps. We say that there was a jump at time t if the estimated probability of a jump on that day was above 50%. Gray vertical bars with the dashed border indicate recessions in the UK; light blue vertical lines with the thin solid border indicate recessions in the US.

Figure 4
JPY data and estimated states

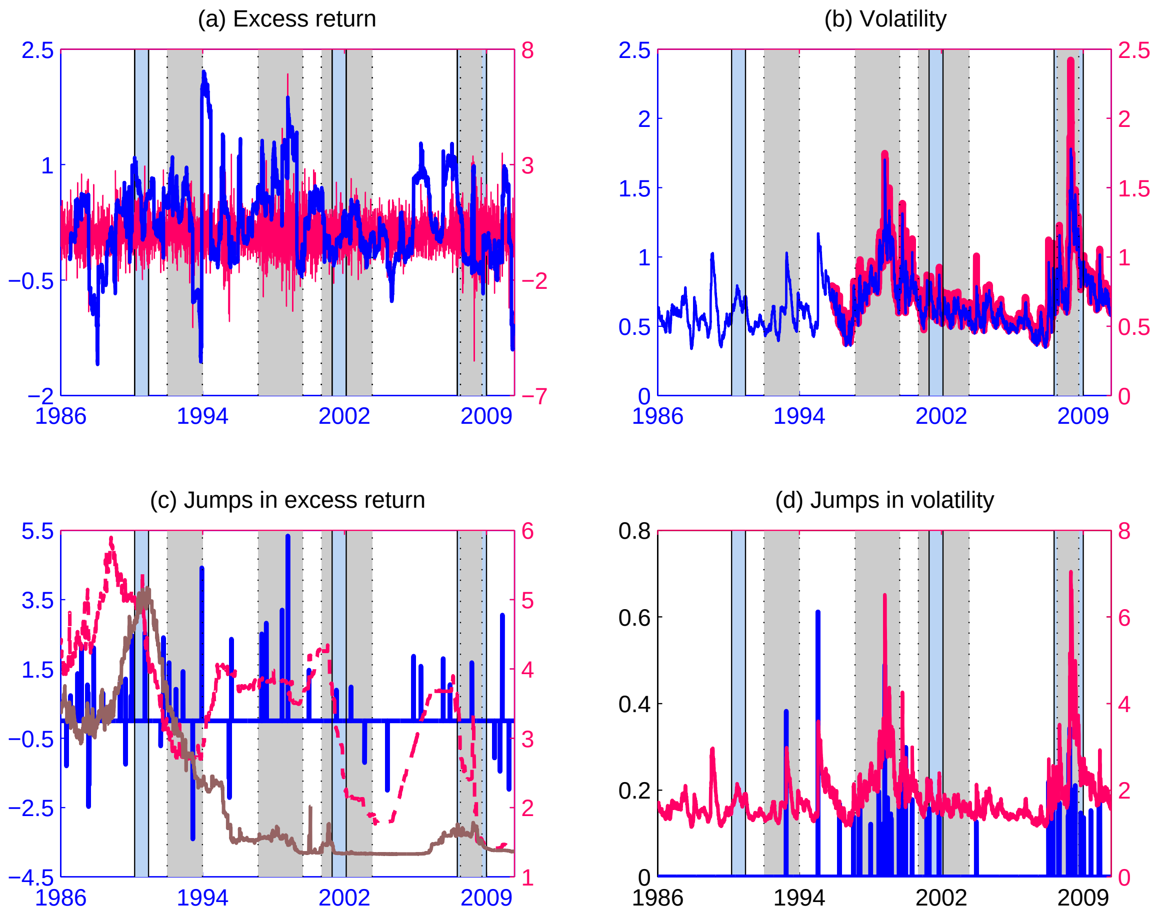

Notes. Panel (a): thin red line displays observed log currency returns; thick blue line displays the conditional 6-months skewness of the log currency returns. Panel (b): thick red line shows observed one-month at-the-money implied volatility; thin blue line shows the estimated spot volatility $\sqrt{v_t}$. The bottom panels display estimated jump sizes in returns (c) and volatility (d) with jump intensities. Panel (c): brown solid line is the intensity of the jump down; the red dashed line is the intensity of the jump up; blue bars are jumps themselves. Panel (d): the red line is the intensity of the jump in volatility; blue bars are jumps. We say that there was a jump at time t if the estimated probability of a jump on that day was above 50%. Gray vertical bars with the dashed border indicate recessions in Japan; light blue vertical lines with the thin solid border indicate recessions in the US.

Figure 5
Conditional decomposition of the total risk for monthly returns

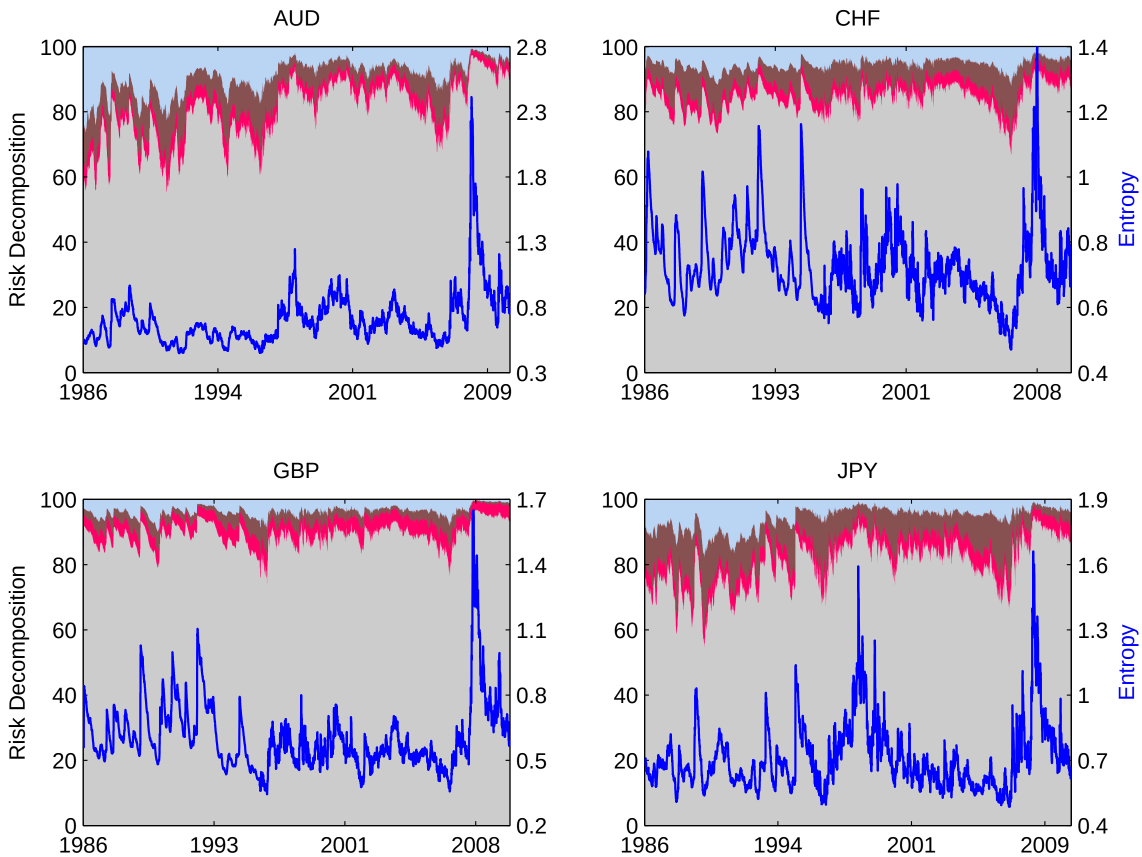

Notes. We display cumulative contribution of the different risks to the total risk of excess returns (the left axis). We measure the total amount of risk using entropy

$$L_t(S_{t+n}/S_t) = \kappa_{2t}(s_{t+n} - s_t)/2! + \kappa_{3t}(s_{t+n} - s_t)/3! + \kappa_{4t}(s_{t+n} - s_t)/4! + \dots,$$

where $\kappa_{jt}(s_{t+n} - s_t)$ is the jth cumulant of log FX returns. Investment horizon is $n = 21$ days. The contribution of the down jumps in FX is displayed in light blue, contribution of the up jumps in FX is in brown, and the contribution of the jumps in variance is in red. Gray area is the contribution of the normal shocks. The blue line shows $\sqrt{2L_t/n}$ in percent (the right axis). This quantity has an interpretation of one-period volatility in the case of normally distributed returns.

Figure 6
Conditional decomposition of the total risk for quarterly returns

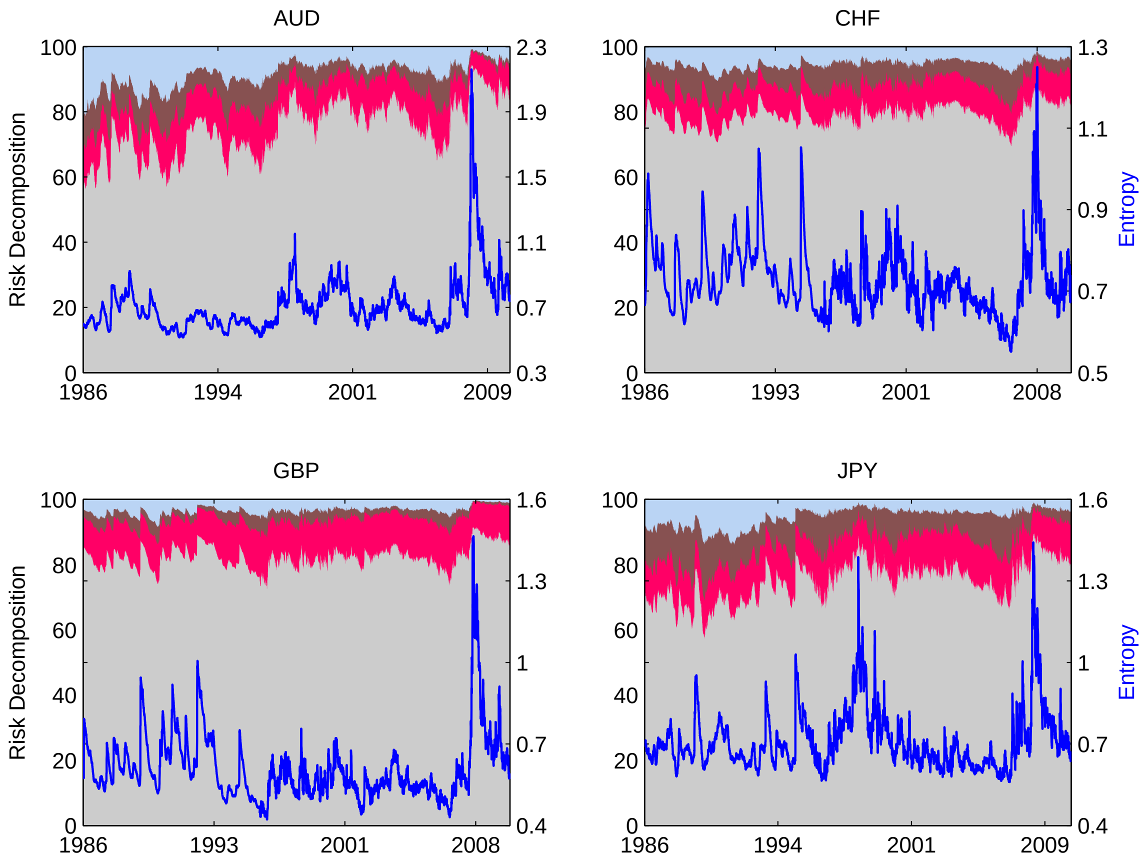

Notes. We display cumulative contribution of the different risks to the total risk of excess returns (the left axis). We measure the total amount of risk using entropy

$$L_t(S_{t+n}/S_t) = \kappa_{2t}(s_{t+n} - s_t)/2! + \kappa_{3t}(s_{t+n} - s_t)/3! + \kappa_{4t}(s_{t+n} - s_t)/4! + \dots,$$

where $\kappa_{jt}(s_{t+n} - s_t)$ is the jth cumulant of log FX returns. Investment horizon is $n = 63$ days. The contribution of the down jumps in FX is displayed in light blue, contribution of the up jumps in FX is in brown, and the contribution of the jumps in variance is in red. Gray area is the contribution of the normal shocks. The blue line shows $\sqrt{2L_t/n}$ in percent (the right axis). This quantity has an interpretation of one-period volatility in the case of normally distributed returns.

Figure 7
Index and estimated states

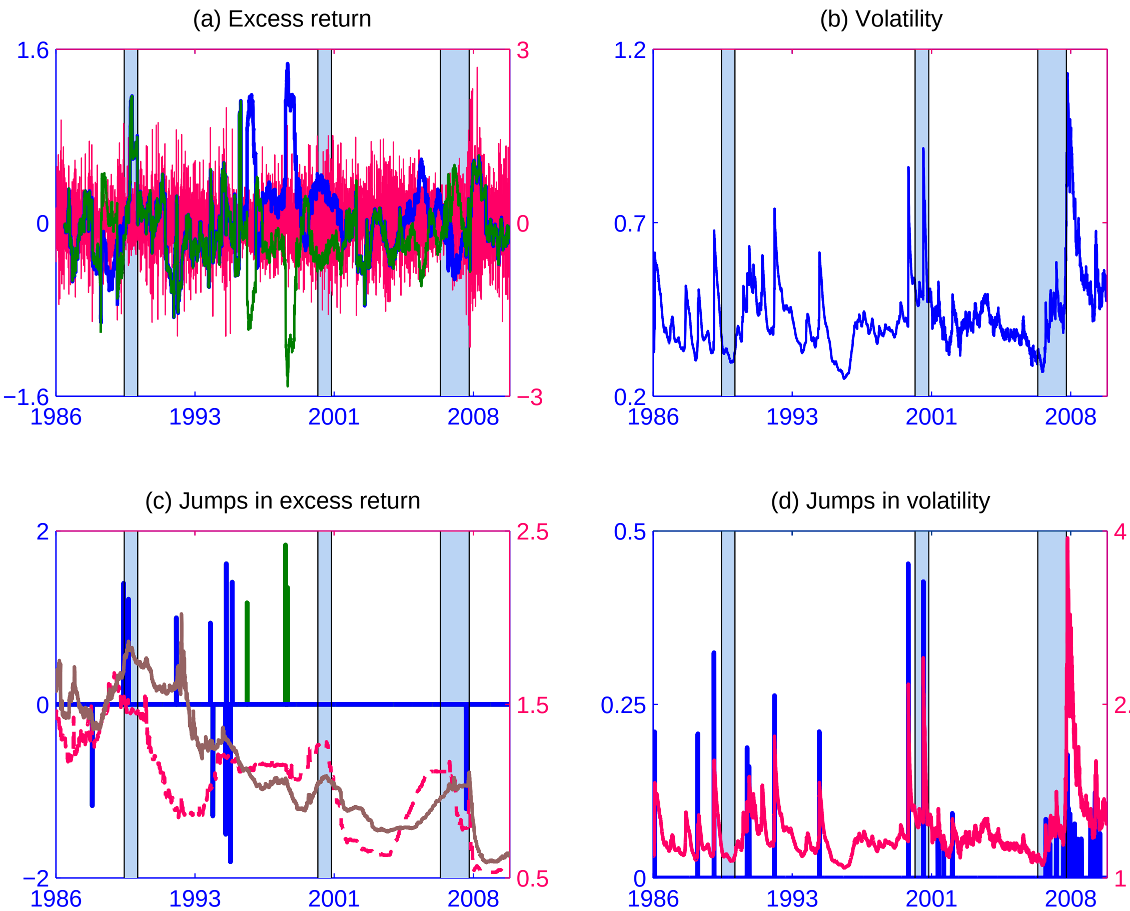

Notes. Panel (a): thin red line displays observed log currency returns; thick blue line displays the conditional 6-months skewness of the log currency returns; thick green line displays the conditional 6-months skewness of the log currency returns of the dollar carry trade. Panel (b): thin blue line shows the estimated spot volatility $\sqrt{v_t}$. The bottom panels display estimated jump sizes in returns (c) and volatility (d) with jump intensities. Panel (c): brown solid line is the intensity of the jump down; the red dashed line is the intensity of the jump up; blue bars are jumps themselves; green bars are jumps that change the direction if there is rebalancing based on the interest rate differential. Panel (d): the red line is the intensity of the jump in volatility; blue bars are jumps. We say that there was a jump at time t if the estimated probability of a jump on that day was above 50%. Light blue vertical lines with the thin solid border indicate recessions in the US.

Figure 8
Implied volatility

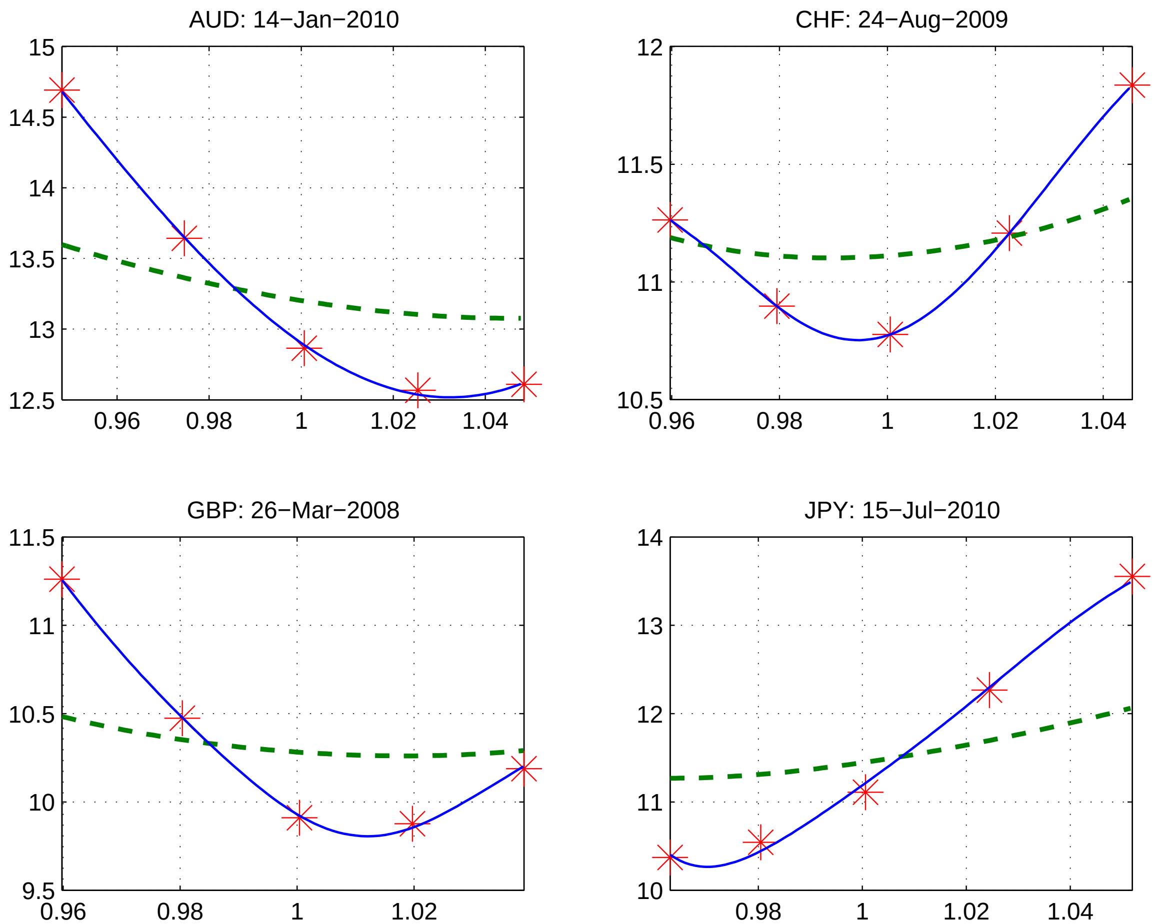

Notes. We check the ability of the model to generate the implied volatility (IV) smiles. The smiles on selected days are the closest to the sample average. The asterisks indicate observed implied volatilities. The dashed green lines display theoretical smiles when only compensation for normal risks is allowed (we optimize over v^* and ν^* only). The solid blue lines depict theoretical implied volatilities when jump risk premiums are allowed as well (we optimize over v^*, ν^*, h_0^{v*}, h_0^{u*}, and h_0^{d*}).

A Appendix

A.1 Long-Run Risk models with identical risk premium implications

In this section we provide two examples of models with critically different shocks to the respective endowment processes that, nonetheless, yield the same functional dependence of currency excess returns on observable variables. We rely on the Long-Run Risk framework of Bansal and Yaron (2004) and use various modelling elements inspired by Bansal and Shaliastovich (2012); Benzoni, Collin-Dufresne, and Goldstein (2011); Drechsler and Yaron (2011); Wachter (2013).

We neither make any claims about realism of these models nor attempt to distinguish them empirically. In fact, we try to construct the simplest models possible that deliver risk premiums dependent on the interest rate differential and the variance of changes in the exchange rate. Moreover, the models have implications for real exchange rates while we are studying the empirical behaviour of nominal exchange rates. Thus, these models serve for pure illustrative purposes.

We use recursive preferences and define utility from date t on

$$U_t = \left[(1-\beta)c_t^{\rho} + \beta\mu_t(U_{t+1})^{\rho}\right]^{1/\rho}, \tag{A.1}$$

and certainty equivalent function,

$$\mu_t(U_{t+1}) = \left[E_t(U_{t+1}^{\alpha})\right]^{1/\alpha}.$$

In standard terminology, $\rho < 1$ captures time preference (with intertemporal elasticity of substitution $1/(1-\rho)$) and $\alpha < 1$ captures risk aversion (with coefficient of relative risk aversion $1-\alpha$). The time aggregator and certainty equivalent functions are both homogeneous of degree one, which allows us to scale everything by current consumption. If we define scaled utility $u_t = U_t/c_t$, equation (A.1) becomes

$$u_t = \left[(1-\beta) + \beta\mu_t(g_{t+1}u_{t+1})^{\rho}\right]^{1/\rho}, \tag{A.2}$$

where $g_{t+1} = c_{t+1}/c_t$ is consumption growth. The pricing kernel is

$$\begin{aligned} m_{t+1} &= \beta(c_{t+1}/c_t)^{\rho-1}\left[U_{t+1}/\mu_t(U_{t+1})\right]^{\alpha-\rho} \\ &= \beta g_{t+1}^{\rho-1}\left[g_{t+1}u_{t+1}/\mu_t(g_{t+1}u_{t+1})\right]^{\alpha-\rho}. \end{aligned}$$

The relationship (A.2) serves, essentially, as a Bellman equation. Its loglinear approximation

$$\begin{aligned} \log u_t &= \rho^{-1}\log\left[(1-\beta) + \beta\mu_t(g_{t+1}u_{t+1})^{\rho}\right] \\ &= \rho^{-1}\log\left[(1-\beta) + \beta e^{\rho\log\mu_t(g_{t+1}u_{t+1})}\right] \\ &\approx b_0 + b_1\log\mu_t(g_{t+1}u_{t+1}) \end{aligned} \tag{A.3}$$

gives us transparent closed-form expressions for pricing kernels (Hansen, Heaton, and Li, 2008). The last line is a first-order approximation of $\log u_t$ in $\log \mu_t$ around the point $\log \mu_t = \log \mu$, with

$$\begin{aligned} b_1 &= \beta e^{\rho \log \mu} / [(1-\beta) + \beta e^{\rho \log \mu}] \\ b_0 &= \rho^{-1} \log[(1-\beta) + \beta e^{\rho \log \mu}] - b_1 \log \mu. \end{aligned}$$

The equation is exact when $\rho = 0$, in which case $b_0 = 0$ and $b_1 = \beta$. We focus on this case for the simplicity sake and to avoid the debate on the accuracy of the log-linear approximation as this subject is not the focus of our paper.

We extend this setting to two countries that we refer to as home (US) and foreign. The representative agents in each country have different risk aversion: $1-\alpha$ and $1-\tilde{\alpha}$, respectively. Similarly, all other foreign-country-specific objects, such as consumption growth, or pricing kernel are denoted by tilde $\tilde{}$.

A.1.1 Model 1: Stochastic Variance

The domestic consumption growth is

$$\begin{aligned} \log g_{t+1} &= \log g + x_t + k\sigma_{g,t}\eta_{t+1}, \\ x_{t+1} &= \gamma x_t + \sigma_{x,t} e_{t+1}, \\ \sigma^2_{g,t+1} &= (1-\nu_g)v_g + \nu_g \sigma^2_{g,t} + \sigma_{gw}\sigma_{g,t} w_{g,t+1}, \\ \sigma^2_{x,t+1} &= (1-\nu_x)v_x + \nu_x \sigma^2_{x,t} + \sigma_{xw}\sigma_{x,t} w_{x,t+1}. \end{aligned}$$

The foreign consumption growth is similar, except for the loading of the variance of consumption growth on $\sigma^2_{g,t}$:

$$\begin{aligned} \log \tilde{g}_{t+1} &= \log g + x_t + \tilde{k}\sigma_{g,t}\tilde{\eta}_{t+1}, \\ x_{t+1} &= \gamma x_t + \sigma_{x,t} e_{t+1}, \\ \sigma^2_{g,t+1} &= (1-\nu_g)v_g + \nu_g \sigma^2_{g,t} + \sigma_{gw}\sigma_{g,t} w_{g,t+1}, \\ \sigma^2_{x,t+1} &= (1-\nu_x)v_x + \nu_x \sigma^2_{x,t} + \sigma_{xw}\sigma_{x,t} w_{x,t+1}. \end{aligned}$$

For simplicity, we assume that all the shocks, η_t, $\tilde{\eta}_t$, e_t, $w_{g,t}$, and $w_{x,t}$, are independent.

The model is solved by guessing the value function for each country

$$\log u_t = \log u + p_x x_t + p_{\sigma g}\sigma^2_{g,t} + p_{\sigma x}\sigma^2_{x,t}.$$

and finding the coefficients p_x, $p_{\sigma g}$, and $p_{\sigma x}$ using the equation (A.3). The full solution of the model is presented in the online Appendix. Here we summarize the key implications. The log pricing kernel at home is

$$\begin{aligned} \log m_{t+1} &= \log \beta - \log g_{t+1} + \alpha[\log g_{t+1}u_{t+1} - \log \mu_t(g_{t+1}u_{t+1})] \\ &= \log \beta - \log g - x_t - \alpha^2(k^2 + p^2_{\sigma g}\sigma^2_{gw})\sigma^2_{g,t}/2 - \alpha^2(p^2_x + p^2_{\sigma x}\sigma^2_{xw})\sigma^2_{x,t}/2 \\ &+ (\alpha-1)k\sigma_{g,t}\eta_{t+1} + \alpha p_x \sigma_{x,t} e_{t+1} + \alpha p_{\sigma g}\sigma_{gw}\sigma_{g,t} w_{g,t+1} + \alpha p_{\sigma x}\sigma_{xw}\sigma_{x,t} w_{x,t+1}. \end{aligned}$$

The log pricing kernel abroad $\tilde{m}_t$ has a similar expression.

We can compute domestic interest rate $r_t = -\log E_t(m_{t+1})$. A similar expression applies to the foreign interest rate $\tilde{r}_t$. By no-arbitrage:

$$s_{t+1} - s_t = \log \tilde{m}_{t+1} - \log m_{t+1}. \tag{A.4}$$

Therefore, we can compute expected excess returns $E_t(s_{t+1} - s_t - (r_t - \tilde{r}_t))$ and conditional variance of the exchange rate changes $var_t(s_{t+1} - s_t)$. All of these objects depend on $\sigma^2_{g,t}$ and $\sigma^2_{x,t}$. Therefore, one can express the dependence of expected excess returns on $\sigma^2_{g,t}$ and $\sigma^2_{x,t}$ as a dependence on the interest rate differential and conditional variance of the exchange rate changes. The details are in the online Appendix.

As a result, we have:

$$E_t(s_{t+1} - s_t - (r_t - \tilde{r}_t)) = \delta_r(r_t - \tilde{r}_t) + \delta_v var_t(s_{t+1} - s_t),$$

where

$$\begin{aligned}
\delta_r &= (B_v s_g - A_v s_x)/(A_r B_v), \\
\delta_v &= s_x/B_v, \\
s_g &= [(\alpha-1)^2 k^2 - (\tilde{\alpha}-1)^2 \tilde{k}^2 + \alpha^2 p^2_{\sigma g}\sigma^2_{gw} - \tilde{\alpha}^2 \tilde{p}^2_{\sigma g}\sigma^2_{gw}]/2, \\
s_x &= [p^2_x(\alpha^2 - \tilde{\alpha}^2) + \alpha^2 p^2_{\sigma x}\sigma^2_{xw} - \tilde{\alpha}^2 \tilde{p}^2_{\sigma x}\sigma^2_{xw}]/2, \\
A_r &= [(2\alpha-1)k^2 - (2\tilde{\alpha}-1)\tilde{k}^2]/2 \\
A_v &= (\tilde{\alpha}-1)^2\tilde{k}^2 + (\alpha-1)^2 k^2 + (\tilde{\alpha}\tilde{p}_{\sigma g} - \alpha p_{\sigma g})^2\sigma^2_{gw}, \\
B_v &= p^2_x(\tilde{\alpha}-\alpha)^2 + (\tilde{\alpha}\tilde{p}_{\sigma x} - \alpha p_{\sigma x})^2\sigma^2_{xw}.
\end{aligned}$$

A.1.2 Model 2: Disasters

The domestic and foreign consumption growths are

$$\begin{aligned}
\log g_{t+1} &= \log g + x_t + \sigma_g \eta_{t+1} + z_{g,t+1}, \\
\log \tilde{g}_{t+1} &= \log g + x_t + \sigma_g \tilde{\eta}_{t+1} + z_{g,t+1}, \\
x_{t+1} &= \gamma x_t + \sigma_x e_{t+1} + z_{x,t+1},
\end{aligned}$$

where the jump sizes are drawn from the normal distributions

$$\begin{aligned}
z_{g,t+1}|j &\sim \mathcal{N}(j\mu_g, j\sigma^2_g), \\
z_{x,t+1}|j &\sim \mathcal{N}(j\mu_x, j\sigma^2_x),
\end{aligned}$$

and the jump arrival rate is controlled by a Poisson distribution

$$Prob(j_{t+1} = j) = \exp(-h_{k,t}) h^j_{k,t}/j!, \;\; k = g, x.$$

The jump intensities $h_{g,t}$ and $h_{x,t}$ are time-varying:

$$\begin{aligned} h_{g,t+1} &= (1-\nu_{hg})v_{hg} + \nu_{hg}h_{g,t} + \sigma_{hg}h_{g,t}^{1/2}\varepsilon_{hg,t+1}, \\ h_{x,t+1} &= (1-\nu_{hx})v_{hx} + \nu_{hx}h_{x,t} + \sigma_{hx}h_{x,t}^{1/2}\varepsilon_{hx,t+1}. \end{aligned}$$

The model is solved by guessing the value function for each country

$$\log u_t = \log u + p_x x_t + p_{hg}h_{g,t} + p_{hx}h_{x,t}.$$

and finding the coefficients p_x, p_{hg}, and p_{hx} using the equation (A.3) ($p_x = \beta/(1-\beta\gamma)$ when $\rho = 0$). The details are provided in the online Appendix.

The log pricing kernel at home is

$$\begin{aligned} \log m_{t+1} &= [\log\beta - \log g - \alpha^2\sigma_g^2/2 - \alpha^2 p_x^2\sigma_x^2/2] - x_t - \alpha^2 p_{hg}^2\sigma_{hg}^2 h_{g,t}/2 - \alpha^2 p_{hx}^2\sigma_{hx}^2 h_{x,t}/2 \\ &- (e^{\alpha\mu_g + (\alpha\sigma_g)^2/2} - 1)h_{g,t} - (e^{\alpha p_x\mu_x + (\alpha p_x\sigma_x)^2/2} - 1)h_{x,t} + (\alpha-1)\sigma_g\eta_{t+1} + \alpha p_x\sigma_x e_{t+1} \\ &+ \alpha p_{hg}\sigma_{hg}h_{g,t}^{1/2}\varepsilon_{hg,t+1} + \alpha p_{hx}\sigma_{hx}h_{xt}^{1/2}\varepsilon_{hx,t+1} + (\alpha-1)z_{g,t+1} + \alpha p_x z_{x,t+1}. \end{aligned}$$

The log pricing kernel abroad has a similar expression.

We can compute domestic interest rate $r_t = -\log E_t(m_{t+1})$. A similar expression applies to the foreign interest rate $\tilde{r}_t$. By no-arbitrage:

$$s_{t+1} - s_t = \log\tilde{m}_{t+1} - \log m_{t+1}. \tag{A.5}$$

Therefore, we can compute expected excess returns $E_t(s_{t+1} - s_t - (r_t - \tilde{r}_t))$ and conditional variance of the exchange rate changes $var_t(s_{t+1} - s_t)$. All of these objects depend on $h_{g,t}$ and $h_{x,t}$. Therefore, one can express the dependence of expected excess returns on $h_{g,t}$ and $h_{x,t}$ as a dependence on the interest rate differential and conditional variance of the exchange rate changes. The details are in the online Appendix.

As a result, we have:

$$E_t(s_{t+1} - s_t - (r_t - \tilde{r}_t)) = \delta_0 + \delta_r(r_t - \tilde{r}_t) + \delta_v var_t(s_{t+1} - s_t),$$

where

$$\begin{aligned} \delta_0 &= ((\alpha-1)^2 - (\tilde{\alpha}-1)^2)\sigma_g^2/2 + (\alpha^2 - \tilde{\alpha}^2)p_x^2\sigma_x^2/2 - s_g(\alpha-\tilde{\alpha})\sigma_g^2/A_r \\ &- s_x(A_r[((\tilde{\alpha}-1)^2 + (\alpha-1)^2)\sigma_g^2 + (\tilde{\alpha}-\alpha)^2 p_x^2\sigma_x^2] - A_v(\alpha-\tilde{\alpha})\sigma_g^2)/(A_r B_v), \\ \delta_r &= (B_v s_g - A_v s_x)/(A_r B_v), \\ \delta_v &= s_x/B_v, \\ s_g &= (\alpha^2 p_{hg}^2 - \tilde{\alpha}^2\tilde{p}_{hg}^2)\sigma_{hg}^2/2 + (e^{\mu_g(\alpha-1) + ((\alpha-1)\sigma_g)^2/2} - e^{(\tilde{\alpha}-1)\mu_g + ((\tilde{\alpha}-1)\sigma_g)^2/2}) + \mu_g(\tilde{\alpha}-\alpha), \\ s_x &= (\alpha^2 p_{hx}^2 - \tilde{\alpha}^2\tilde{p}_{hx}^2)\sigma_{hx}^2/2 + e^{\alpha p_x\mu_x + (\alpha p_x\sigma_x)^2/2} - e^{\tilde{\alpha}p_x\mu_x + (\tilde{\alpha}p_x\sigma_x)^2/2} + \mu_x p_x(\tilde{\alpha}-\alpha), \\ A_r &= e^{\alpha\mu_g + (\alpha\sigma_g)^2/2} - e^{\tilde{\alpha}\mu_g + (\tilde{\alpha}\sigma_g)^2/2} + e^{(\tilde{\alpha}-1)\mu_g + ((\tilde{\alpha}-1)\sigma_g)^2/2} - e^{(\alpha-1)\mu_g + ((\alpha-1)\sigma_g)^2/2} \\ A_v &= (\tilde{\alpha}\tilde{p}_{hg} - \alpha p_{hg})^2\sigma_{hg}^2 + (\tilde{\alpha}-\alpha)^2\mu_g^2 + \sigma_g^2, \\ B_v &= (\tilde{\alpha}\tilde{p}_{hx} - \alpha p_{hx})^2\sigma_{hx}^2 + p_x^2(\tilde{\alpha}-\alpha)^2\mu_x^2 + \sigma_x^2. \end{aligned}$$

A.2 Expected future variance

We do not consider the most general model to streamline the presentation. We focus on the empirically relevant case where intensity of jumps in variance depends on variance only, and jumps up (down) in FX depend on domestic (foreign) interest rate only. We start by computing expectation of the variance process in (3.2). Conditional expectation $E_t(v_{t+i}) \equiv v_{t,i}$ can be computed via a recursion. Note that $v_{t,0} = v_t$. Suppose we know $v_{t,i-1}$. Then

$$\begin{aligned} v_{t,i} &= (1-\nu)v + \nu v_{t,i-1} + \sigma_v E_t(E_{t+i-1}(v_{t+i-1}^{1/2} w_{t+i}^v)) + E_t(E_{t+i-1} z_{t+i}^v) \\ &= (1-\nu)v + \nu v_{t,i-1} + \theta_v h_0^v + \theta_v h_v^v v_{t,i-1} = (1-\nu)v + \theta_v h_0^v + (\nu + \theta_v h_v^v) v_{t,i-1}. \end{aligned}$$

We can solve this recursion analytically:

$$\begin{aligned} v_{t,i} &= [(1-\nu)v + \theta_v h_0^v](1 + (\nu + \theta_v h_v^v)) + (\nu + \theta_v h_v^v)^2 v_{t,i-2} \\ &= [(1-\nu)v + \theta_v h_0^v](1 - (\nu + \theta_v h_v^v)^i)/(1 - (\nu + \theta_v h_v^v)) + (\nu + \theta_v h_v^v)^i v_t. \end{aligned}$$

Next, we can compute expectation of average future v :

$$\begin{aligned} & E_t\left(\sum_{i=1}^n v_{t+i}\right)/n = 1/n \sum_{i=1}^n E_t v_{t+i} = 1/n \sum_{i=1}^n v_{t,i} \\ = \;& 1/n \sum_{i=1}^n [(1-\nu)v + \theta_v h_0^v](1 - (\nu + \theta_v h_v^v)^i)/(1 - (\nu + \theta_v h_v^v)) + 1/n \sum_{i=1}^n (\nu + \theta_v h_v^v)^i v_t \\ = \;& \frac{(1-\nu)v + \theta_v h_0^v}{1 - (\nu + \theta_v h_v^v)} \left[1 - \frac{\nu + \theta_v h_v^v}{n} \frac{1 - (\nu + \theta_v h_v^v)^n}{1 - (\nu + \theta_v h_v^v)}\right] + \frac{\nu + \theta_v h_v^v}{n} \frac{1 - (\nu + \theta_v h_v^v)^n}{1 - (\nu + \theta_v h_v^v)} v_t \\ \equiv \;& \underbrace{\frac{(1-\nu)v + \theta_v h_0^v}{1 - (\nu + \theta_v h_v^v)} [1 - \beta_n]}_{\alpha_n} + \beta_n v_t. \end{aligned}$$

Similarly, we can obtain conditional expectations of future interest rates:

$$r_{t,i} \equiv E_t(r_{t+i}) = a_r(1 - b_r^i) + b_r^i r_t,$$

and the expectations of average future interest rates:

$$\begin{aligned} & E_t\left(\sum_{i=1}^n r_{t+i}\right)/n = 1/n \sum_{i=1}^n E_t r_{t+i} = 1/n \sum_{i=1}^n r_{t,i} \\ = \;& a_r\left[1 - \frac{b_r}{n} \frac{1 - b_r^n}{1 - b_r}\right] + \frac{b_r}{n} \frac{1 - b_r^n}{1 - b_r} r_t \end{aligned}$$

and the similar expression holds for expectations associated with $\tilde{r}_t$.

Now, we can characterize the variance of excess returns:

$$v_t^y \equiv var_t(y_{t+1}) = v_t + 2h_t^u\theta_u^2 + 2h_t^d\theta_d^2.$$

Therefore, the conditional expectation of this variance can be computed on the basis of our results for the variance of the normal component v_t and the expectations of interest rates:

$$v_{t,i}^y \equiv E_t(v_{t+i}^y) = v_{t,i} + 2\theta_u^2 h_0^u + 2\theta_u^2 h_r^u E_t(r_{t+i}) + 2\theta_d^2 h_0^d + 2\theta_d^2 \tilde{h}_r^d E_t(\tilde{r}_{t+i}).$$

This expression implies that the unconditional expectation, or long-run mean, of the variance is:

$$\begin{aligned} v_J = \lim_{i\to\infty} v_{t,i}^y &= [(1-\nu)v + \theta_v h_0^v]/(1-(\nu+\theta_v h_v^v)) \\ &+ 2\theta_u^2 h_0^u + 2\theta_u^2 h_r^u a_r + 2\theta_d^2 h_0^d + 2\theta_d^2 \tilde{h}_r^d \tilde{a}_r. \end{aligned}$$

When there are no jumps, that is, $\theta_v = 0$, $\theta_u = 0$, and $\theta_d = 0$, then $v_J = v$.

Next, we compute $E_t(\sum_{i=1}^n v_{t+i}^y)/n$

$$\begin{aligned} & E_t\left(\sum_{i=1}^n v_{t+i}^y\right)/n = 1/n\sum_{i=1}^n E_t v_{t+i}^y = 1/n\sum_{i=1}^n v_{t,i}^y \\ = \quad & \alpha_n + 2\theta_u^2 h_0^u + 2\theta_u^2 h_r^u a_r\left[1 - \frac{b_r}{n}\frac{1-b_r^n}{1-b_r}\right] + 2\theta_d^2 h_0^d + 2\theta_d^2 \tilde{h}_r^d \tilde{a}_r\left[1 - \frac{\tilde{b}_r}{n}\frac{1-\tilde{b}_r^n}{1-\tilde{b}_r}\right] \\ + \quad & \beta_n v_t + 2\theta_u^2 h_r^u \frac{b_r}{n}\frac{1-b_r^n}{1-b_r} r_t + 2\theta_d^2 \tilde{h}_r^d \frac{\tilde{b}_r}{n}\frac{1-\tilde{b}_r^n}{1-\tilde{b}_r}\tilde{r}_t. \end{aligned}$$

A.3 Computing entropy

Entropy of currency changes over a horizon of n days is equal to:

$$L_t(S_{t+n}/S_t) = \log E_t(e^{x_{t,n}}) - E_t(x_{t,n}) = k_t(1; x_{t,n}) - \kappa_{1t}(x_{t,n}),$$

where $x_{t,n} = \log(S_{t+n}/S_t) = \sum_{i=t}^{t+n}(s_{i+1} - s_i)$, $k_t(s; x_{t,n})$ is a cumulant-generating function of $x_{t,n}$, and $\kappa_{1t}(x_{t,n})$ is the first cumulant of $x_{t,n}$. Thus, we need to compute the cumulant-generating function of $x_{t,n}$:

$$k_t(s; x_{t,n}) = \log E_t e^{s x_{t,n}}.$$

The first cumulant can be recovered as $\partial k_t(s; x_{t,n})/\partial s$ at $s = 0$. Denote the the drift of log currency changes by $\bar{\mu}_t = \mu_t + (r_t - \tilde{r}_t)$.

Guess

$$k_t(s; x_{t,n}) = A(n) + B_v(n)v_t + B_r(n)r_t + \tilde{B}_r(n)\tilde{r}_t.$$

Then

$$
\begin{aligned}
& A(n) + B_v(n)v_t + B_r(n)r_t + \tilde{B}_r(n)\tilde{r}_t \\
= \quad & k(s; x_{t,n}) = \log E_t[e^{sx_{t,1}} E_{t+1} e^{sx_{t+1,n-1}}] \\
= \quad & \log E_t[e^{sx_{t,1}} e^{A(n-1)+B_v(n-1)v_{t+1}+B_r(n-1)r_{t+1}+\tilde{B}_r(n-1)\tilde{r}_{t+1}}] \\
= \quad & A(n-1) + \log E_t e^{sx_{t,1}+B_v(n-1)v_{t+1}} + \log E_t e^{B_r(n-1)r_{t+1}+\tilde{B}_r(n-1)\tilde{r}_{t+1}} \\
= \quad & A(n-1) + s\bar{\mu}_t + B_v(n-1)((1-\nu)v + \nu v_t) \\
+ \quad & B_r(n-1)((1-b_r)a_r + b_r r_t) + \tilde{B}_r(n-1)((1-\tilde{b}_r)\tilde{a}_r + \tilde{b}_r\tilde{r}_t) \\
+ \quad & \log E_t e^{s(1-\rho^2)^{1/2}v_t^{1/2}w^s_{t+1}+s\rho v_t^{1/2}w^v_{t+1}+sz^u_{t+1}+sz^d_{t+1}+B_v(n-1)\sigma_v v_t^{1/2}w^v_{t+1}+B_v(n-1)z^v_{t+1}} \\
+ \quad & \log E_t e^{B_r(n-1)r_t^{1/2}\sigma_r w^r_{t+1}+\tilde{B}_r(n-1)\tilde{r}_t^{1/2}\tilde{\sigma}_r\tilde{w}^r_{t+1}} \\
= \quad & A(n-1) + s\bar{\mu}_t + B_v(n-1)((1-\nu)v + \nu v_t) \\
+ \quad & B_r(n-1)((1-b_r)a_r + b_r r_t) + \tilde{B}_r(n-1)((1-\tilde{b}_r)\tilde{a}_r + \tilde{b}_r\tilde{r}_t) \\
+ \quad & s^2 v_t/2 + v_t s\rho\sigma_v B_v(n-1) + B_v^2(n-1)\sigma_v^2 v_t/2 + h^u_t((1-s\theta_u)^{-1}-1) + h^d_t((1+s\theta_d)^{-1}-1) \\
+ \quad & h^v_t((1-B_v(n-1)\theta_v)^{-1}-1) + B_r^2(n-1)\sigma_r^2 r_t/2 + \tilde{B}_r^2(n-1)\tilde{\sigma}_r^2\tilde{r}_t/2 \\
= \quad & A(n-1) + s(\mu + (\mu_r+1)(r_t - \tilde{r}_t) + \mu_v v_t) + B_v(n-1)((1-\nu)v + \nu v_t) \\
+ \quad & B_r(n-1)((1-b_r)a_r + b_r r_t) + \tilde{B}_r(n-1)((1-\tilde{b}_r)\tilde{a}_r + \tilde{b}_r\tilde{r}_t) \\
+ \quad & s^2 v_t/2 + v_t s\rho\sigma_v B_v(n-1) + B_v^2(n-1)\sigma_v^2 v_t/2 + s\theta_u(h^u_0 + h^u_r r_t)/(1-s\theta_u) - s\theta_d(h^d_0 + \tilde{h}^d_r\tilde{r}_t)/(1+s\theta_d) \\
+ \quad & (h^v_0 + h^v_v v_t)B_v(n-1)\theta_v/(1-B_v(n-1)\theta_v) + B_r^2(n-1)\sigma_r^2 r_t/2 + \tilde{B}_r^2(n-1)\tilde{\sigma}_r^2\tilde{r}_t/2.
\end{aligned}
$$

Collect terms, match them with the corresponding terms in the first line, solve for the coefficients:

$$
\begin{aligned}
A(n) &= A(n-1) + s\mu + B_v(n-1)(1-\nu)v + s\theta_u h^u_0/(1-s\theta_u) - s\theta_d h^d_0/(1+s\theta_d) \\
&+ h^v_0 B_v(n-1)\theta_v/(1-\theta_v B_v(n-1)) + B_r(n-1)(1-b_r)a_r + \tilde{B}_r(n-1)(1-\tilde{b}_r)\tilde{a}_r \\
B_v(n) &= s\mu_v + B_v(n-1)\nu + s^2/2 + s\rho\sigma_v B_v(n-1) + B_v^2(n-1)\sigma_v^2/2 \\
&+ h^v_v B_v(n-1)\theta_v/(1-B_v(n-1)\theta_v), \\
B_r(n) &= s(\mu_r+1) + B_r(n-1)b_r + s\theta_u h^u_r/(1-s\theta_u) + B_r^2(n-1)\sigma_r^2/2, \\
\tilde{B}_r(n) &= -s(\mu_r+1) + \tilde{B}_r(n-1)\tilde{b}_r - s\theta_d\tilde{h}^d_r/(1+s\theta_d) + \tilde{B}_r^2(n-1)\tilde{\sigma}_r^2/2.
\end{aligned}
$$

To compute initial conditions for the above recursion, write down the cumulant generating function of a one-period return:

$$
k_t(s; x_{t,1}) = s\bar{\mu}_t + s^2 v_t/2 + (h^u_0 + h^u_r r_t)\frac{s\theta_u}{1-s\theta_u} - (h^d_0 + \tilde{h}^d_r\tilde{r}_t)\frac{s\theta_d}{1+s\theta_d}.
$$

Therefore,

$$\begin{aligned}
A(1) &= s\mu + h_0^u \frac{s\theta_u}{1 - s\theta_u} - h_0^d \frac{s\theta_d}{1 + s\theta_d}, \\
B_v(1) &= s\mu_v + s^2/2, \\
B_r(1) &= s(\mu_r + 1) + s\theta_u h_r^u/(1 - s\theta_u), \\
\tilde{B}_r(1) &= -s(\mu_r + 1) - s\theta_d \tilde{h}_r^d/(1 + s\theta_d).
\end{aligned}$$