

Crash Risk in Currency Returns

Mikhail Chernov, Jeremy Graveline, and Irina Zviadadze*

Abstract

We develop an empirical model of bilateral exchange rates. It includes normal shocks with stochastic variance and jumps in an exchange rate and in its variance. The probability of a jump in an exchange rate corresponding to depreciation (appreciation) of the U.S. dollar is increasing in the domestic (foreign) interest rate. The probability of a jump in variance is increasing in the variance only. Jumps in exchange rates are associated with announcements; jumps in variance are not. On average, jumps account for 25% of currency risk. The dollar carry index retains these features. Options suggest that jump risk is priced.

I. 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.

*Chernov, mikhail.chernov@anderson.ucla.edu, University of California–Los Angeles (UCLA) Anderson School of Management, Center for Economic and Policy Research (CEPR), and National Bureau of Economic Research (NBER); Graveline, jeremy.graveline@blackrock.com, BlackRock; and Zviadadze (corresponding author), irina.zviadadze@hhs.se, Stockholm School of Economics and CEPR. We are grateful to an anonymous referee and Hendrik Bessembinder (the editor), as well as many people who gave us advice, including Yacine Ait-Sahalia, David Backus, Magnus Dahlquist, Itamar Drechsler, Bjorn Eraker, Peter Feldhutter, Diego Garcia, Francisco Gomes, Michael Johannes, Igor Makarov, Christopher Neely, 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, and Stan Zin, as well as participants in seminars at and conferences sponsored by the American Economic Association (AEA), BI Oslo, University of Chicago, Ecole des Hautes Etudes Commerciales du Nord (EDHEC), the Federal Reserve, Goldman Sachs, Imperial College, London Business School (LBS), London School of Economics (LSE), Montreal Structured Finance and Derivatives Institute (IFSID) and Bank of Canada, Moscow Higher School of Economics (HSE), the New York Federal Reserve Bank, New York University, the Society for Economic Dynamics (SED), the Society for Financial Econometrics (SoFiE), Stockholm University, the Swedish House of Finance, Tinbergen Institute, the University of British Columbia, UCLA, the University of North Carolina, and the University of Pennsylvania. Zviadadze gratefully acknowledges support from the AXA Research Fund and Jan Wallander and Tom Hedelius Foundation.

Therefore, the distribution of risks is an important ingredient in understanding currency premiums.

Some researchers have suggested that investors in currency markets require high returns on average as compensation for crash risk. The merit of this explanation hinges on the magnitude and probability of large moves in currency markets. The goal of our article is to offer an empirical study that uses an affine model to systematically quantify crashes and document when they occur historically and what their determinants are.

The main advantage of the affine framework is that it offers tractable asset valuation. Therefore, our preferred formulation can be used for option pricing, as well as for developing macro-based equilibrium models of currencies. Our model incorporates both crash risk and regular, or normal, risk. 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. Using data on implied volatilities in addition to exchange rates themselves is 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 modeling 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), and 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 normal (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), Brunermeier, Nagel, and Pedersen (2008), Carr and Wu (2007), and Neely (1999)).

Third, changes in the at-the-money implied volatility of a typical exchange rate exhibit an unconditional skewness of 1 and a kurtosis of 10 or more. To accommodate this property, our model allows for jumps in the variance of normal shocks to exchange rates. The importance of such jumps for modeling equity returns has been emphasized by Broadie, Chernov, and Johannes (2007), Duffie, Pan, and Singleton (2000), Eraker, Johannes, and Polson (2003), among others. To our knowledge, our article is the first to investigate the role of jumps in the volatility of exchange rates.

Although both are designed to generate large movements, a jump in an exchange rate is qualitatively different from a jump in its variance. A jump in an exchange rate is a model of a direct extreme shock. Because extreme events are rare, one does not necessarily expect there to be many subsequent large innovations in

¹In this article, 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 modeling the economy (see, e.g., Backus, Chernov, and Martin (2011) for an example of how asset prices could be used to infer the behavior of consumption).

the near future. By contrast, when there is a jump in the variance, one expects there to be many large subsequent changes generated by normal innovations 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: the Australian dollar, the Swiss franc, the British pound, and the Japanese yen. Crashes are rare, and volatility is persistent; thus, it is important to use a long time span of data. Short samples are likely to either over- or underrepresent jumps and periods of high or low volatility, leading to biased estimates of the required probabilities.

We rely on Bayesian Markov chain Monte Carlo (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 jumps in both exchange rates and 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 foreign exchange (FX). These 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 U.S. dollar, is positively related to the domestic (U.S.) interest rate. The probability of a jump down, which corresponds to an appreciation of the U.S. 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 they are quite distinct economically. We find that most of the jumps in FX coincide with important macroeconomic or political announcements. In contrast, jumps in variance cluster at 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 et al. (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 (2014). Although diversification reduces jump frequency, it also reduces the volatility of the normal risks. As a result, the relative importance of jumps stays unchanged. We also estimate our model on the excess return to the dollar carry trade, which switches the position of the dollar index from long to short, depending on the difference between the U.S. 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. Although such analysis is outside of the scope of this article, 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 is discussed in the Internet Appendix (available at www.jfqa.org).

II. Preliminaries

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

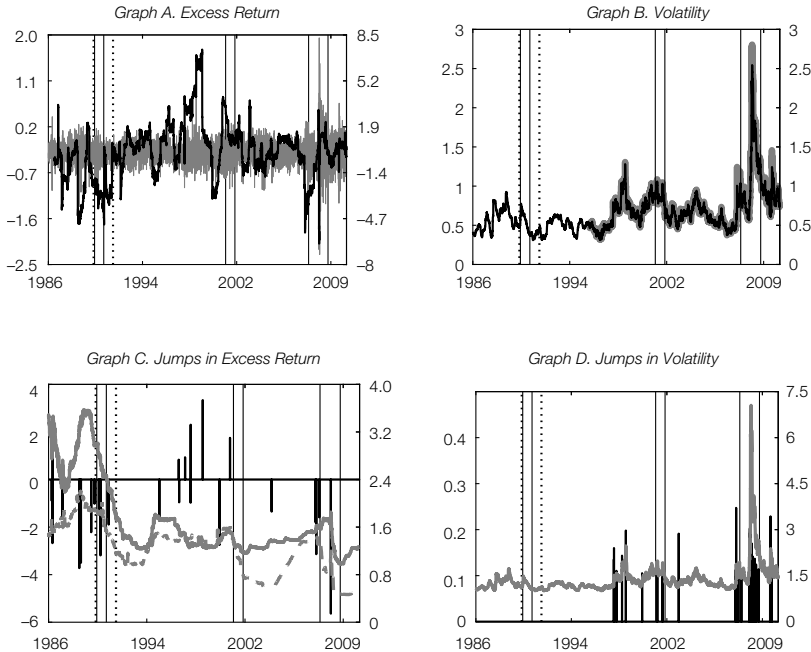
A. Excess Returns

In this article we always treat U.S. dollars (USD) 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 1 USD buys $1/S_t$ units of foreign currency at time t , which grows at the foreign risk-free rate to $1/S_t \times \exp(\tilde{r}_t)$ units at time $t + 1$, and can be exchanged for $S_{t+1}/S_t \times \exp(\tilde{r}_t)$ USD. 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} (Graph A), and implied volatilities (Graph B) for the currencies we consider in this article. We have selected four currencies, the Australian dollar (AUD), the Swiss franc (CHF), the British pound (GBP), and the Japanese yen (JPY), based on the availability of daily data and cross-sectional and time-series variation in the interest rate differential. We use 1-month London Interbank Offered Rate (LIBOR) to proxy for interest rates. Using 1-month rather than overnight rates implicitly assumes a flat term structure at the very short end of the LIBOR curve and allows

FIGURE 1
AUD Data and Estimated States

In Graph A of Figure 1, the gray line displays observed log currency returns (right scale); the black line displays the conditional 6-month skewness of the log currency returns (left scale). In Graph B, the thick gray line shows observed 1-month at-the-money implied volatility; the thin black line shows the estimated spot volatility $\sqrt{V_t}$. The bottom graphs display estimated jump sizes in returns (Graph C) and volatility (Graph D) with jump intensities. In Graph C, the gray solid line is the intensity of the jump down, the gray dashed line is the intensity of the jump up, and the black bars are the jumps themselves. In Graph D, the gray line is the intensity of the jump in volatility; the black 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%. Vertical bars with a dashed border indicate recessions in Australia; vertical bars with a solid black border indicate recessions in the United States.



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.

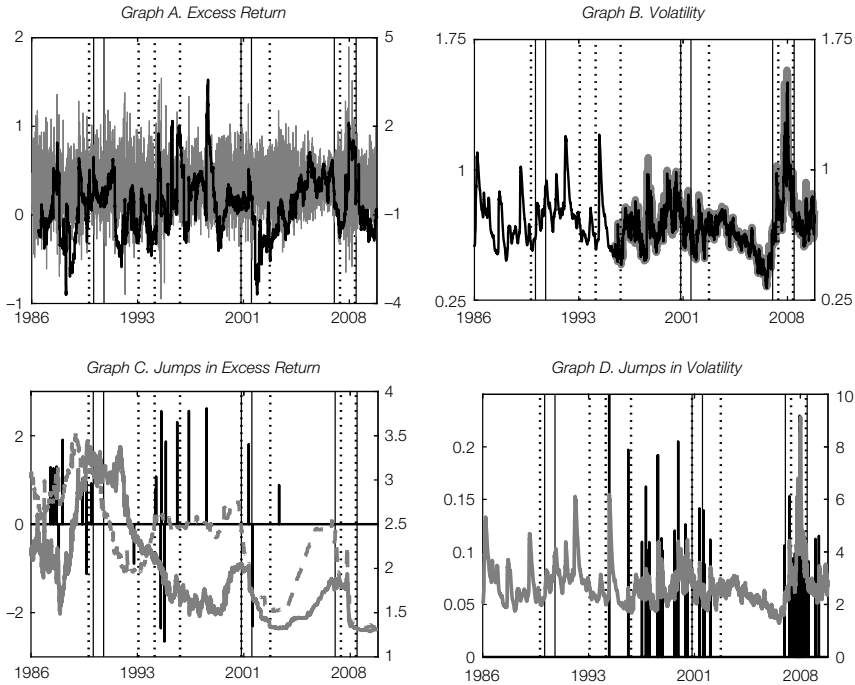
B. Properties of Excess Returns

In Table 1 we provide summary statistics of daily log excess returns and changes in the 1-month at-the-money implied volatility. Means are close to 0 at a daily frequency. Therefore, these summary statistics inform us primarily about the properties of shocks.

All currencies have a 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 suggestive estimate of conditional skewness by computing a 6-month rolling window. The time series of these estimates are displayed in Graph A of Figures 1–4. Depending on the currency, conditional skewness ranges from –2

FIGURE 2
CHF Data and Estimated States

In Graph A of Figure 2, the gray line displays observed log currency returns (right scale); the black line displays the conditional 6-month skewness of the log currency returns (left scale). In Graph B, the thick gray line shows observed 1-month at-the-money implied volatility; the thin black line shows the estimated spot volatility $\sqrt{V_t}$. The bottom graphs display estimated jump sizes in returns (Graph C) and volatility (Graph D) with jump intensities. In Graph C, the gray solid line is the intensity of the jump down, the gray dashed line is the intensity of the jump up, and the black bars are the jumps themselves. In Graph D, the gray line is the intensity of the jump in volatility; the black 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%. Vertical bars with a dashed border indicate recessions in Switzerland; vertical bars with a solid border indicate recessions in the United States.



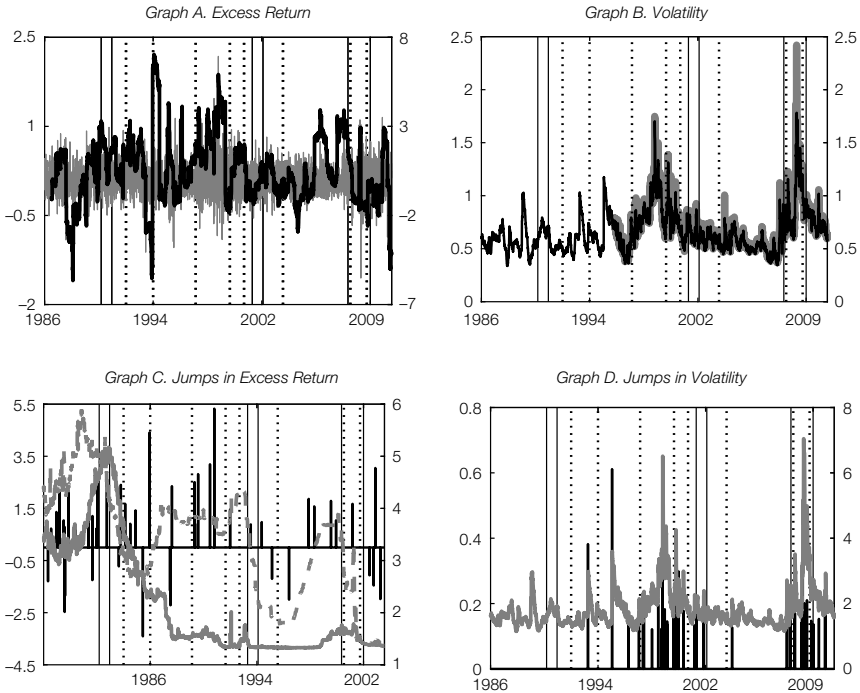
to 2. Thus, excess returns are not only fat tailed but also asymmetric, with the degree of asymmetry changing over time.

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 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 incorporate.

As a reference, we report the same summary statistics for the Standard & Poor's (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 compared with currencies. In particular, negative unconditional skewness is evident (in fact, a measure of conditional skewness rarely becomes positive). In contrast, changes in the volatility index (VIX), a cousin of implied variance, display weaker non-normalities than those of currencies. These statistics suggest

FIGURE 3
GBP Data and Estimated States

In Graph A of Figure 3, the gray line displays observed log currency returns (right scale); the black line displays the conditional 6-month skewness of the log currency returns (left scale). In Graph B, the thick gray line shows observed 1-month at-the-money implied volatility; the thin black line shows the estimated spot volatility $\sqrt{v_t}$. The bottom graphs display estimated jump sizes in returns (Graph C) and volatility (Graph D) with jump intensities. In Graph C, the gray solid line is the intensity of the jump down, the gray dashed line is the intensity of the jump up, and the black bars are the jumps themselves. In Graph D, the gray line is the intensity of the jump in volatility; the black 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%. Vertical bars with a dashed border indicate recessions in the United Kingdom; vertical bars with a solid border indicate recessions in the United States.



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.

C. Risks and Expected Excess Returns

We can generically represent excess returns as follows:

$$(1) \quad y_{t+1} = E_t(y_{t+1}) + \text{SHOCKS.}$$

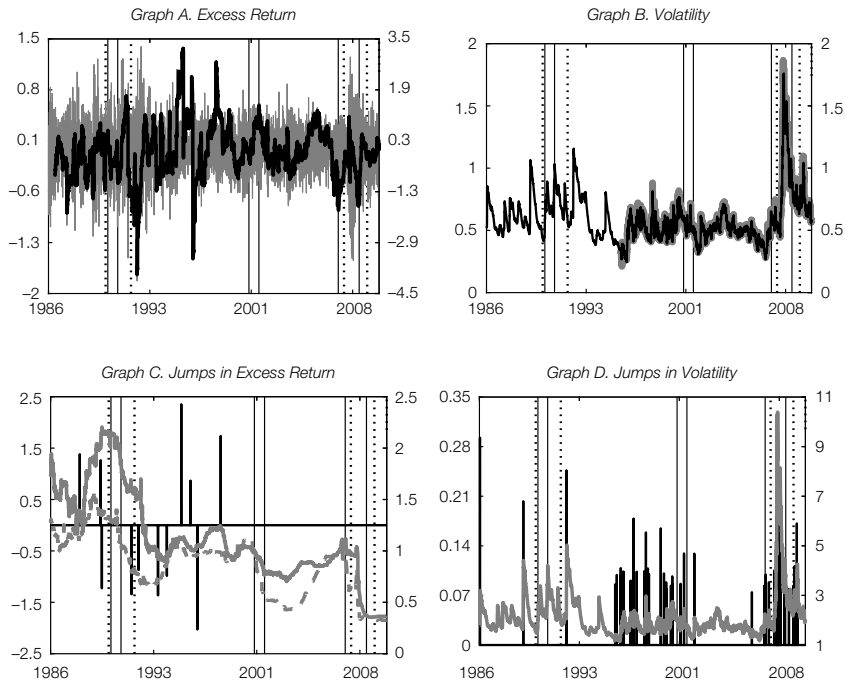
Most research has focused on conditional expected excess returns $E_t(y_{t+1})$. For example, if currencies do not carry a risk premium, then uncovered interest rate parity (UIP) holds, and $E_t(y_{t+1})=0$. However, Bilson (1981), Fama (1984), and Tryon (1979) establish that the regression

$$(2) \quad s_{t+1} - s_t = a_1 + a_2(f_t - s_t) + \text{SHOCKS,}$$

where f_t is the log of the 1-month forward exchange rate, typically yields estimates of a_2 of approximately -2 . If covered interest rate parity (i.e., no arbitrage)

FIGURE 4
JPY Data and Estimated States

In Graph A of Figure 4, the gray line displays observed log currency returns (right scale); the black line displays the conditional 6-month skewness of the log currency returns (left scale). In Graph B, the thick gray line shows observed 1-month at-the-money implied volatility; the thin black line shows the estimated spot volatility $\sqrt{v_t}$. The bottom graphs display estimated jump sizes in returns (Graph C) and volatility (Graph D) with jump intensities. In Graph C, the gray solid line is the intensity of the jump down, the gray dashed line is the intensity of the jump up, and the black bars are the jumps themselves. In Graph D, the gray line is the intensity of the jump in volatility; the black 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%. Vertical bars with a dashed border indicate recessions in Japan; vertical bars with a solid border indicate recessions in the United States.



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 the following:

$$(3) \quad y_{t+1} = a_1 + (a_2 - 1)(r_t - \tilde{r}_t) + \text{SHOCKS},$$

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

Asset pricing theory relates expected excess returns to compensation for bearing risks; that is, it relates “ $E_t(y_{t+1})$ ” to the “shocks” in equation (1). In the language of pricing kernels, expected excess returns are determined by the covariation of currency risks with a pricing kernel. We do not empirically test any specific asset pricing theories (pricing kernels) in this article, 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 in the Internet Appendix (available at www.jfqa.org) an example of two theories that can lead to identical expected

TABLE 1
Properties of Excess Log Currency and S&P 500 Returns and Changes in Implied Volatility

Table 1 shows descriptive statistics for daily log currency and log Standard & Poor's (S&P) 500 excess returns and changes in implied volatility, in percentages, per day: AUD return from Sept. 25, 1986, to Dec. 31, 2010, and AUD IV from Dec. 6, 1995, to Dec. 31, 2010; CHF return from Jan. 3, 1986, to Dec. 31, 2010, and CHF IV from May 8, 1996, to Dec. 31, 2010; GBP return from Jan. 3, 1986, to Dec. 31, 2010, and GBP IV from May 8, 1996, to Dec. 31, 2010; JPY return from July 2, 1986, to Dec. 31, 2010, and JPY IV from Dec. 5, 1995, to Dec. 31, 2010. The currency index return is from Jan. 2, 1986, to Dec. 31, 2010; S&P 500 return from Jan. 3, 1986, to Dec. 31, 2010; and VIX from Jan. 2, 1996, to Dec. 31, 2010. The interest rates used to compute returns are 1-month London Interbank Offered Rate (LIBOR) rates. Source: Bloomberg.

Asset	Metric	Mean	Std. Dev.	Skewness	Kurtosis	No. of Obs.
AUD	Excess return	0.0186	0.7435	-0.3870	13.7202	6,332
	$\Delta\sqrt{IV}$	0.0109	3.7661	0.9077	9.7290	3,933
CHF	Excess return	0.0057	0.7232	0.1194	4.7841	6,521
	$\Delta\sqrt{IV}$	0.0073	3.8057	0.9966	9.8095	3,823
GBP	Excess return	0.0096	0.6197	-0.2337	5.6832	6,521
	$\Delta\sqrt{IV}$	0.0142	4.0001	1.3884	44.2683	3,823
JPY	Excess return	0.0003	0.6950	0.3626	8.0878	6,393
	$\Delta\sqrt{IV}$	-0.0045	4.8277	1.0395	10.7764	3,934
Currency Index	Excess return	0.0086	0.4664	0.0342	4.6448	6,393
S&P 500	Excess return	0.0090	1.1803	-1.3584	32.9968	6,522
	$\Delta\sqrt{VIX}$	0.0089	5.8997	0.5096	6.7502	3,914

excess returns despite their different shock structures. 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 uncovered interest rate parity (UIP) regressions result by allowing for linear dependence of means on the domestic and foreign interest rates and includes the variance of FX returns as an extra variable.² Because we are working with daily returns, the magnitude of the conditional mean 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. As such, to avoid overfitting, we did not include any other variables in the drift of the exchange rate.³

Although our focus is on careful modeling of currency risks themselves, our conclusions should have implications for modeling of economic channels leading to the observed risk premiums. As highlighted by our examples in the Internet Appendix, successful equilibrium models should be able to replicate not only the measured risk premiums but also 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.

²This 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), Brandt and Santa-Clara (2002), Graveline (2006), Lustig et al. (2011), and Menkhoff et al. (2012).

³Recent 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 (2015), Lustig et al. (2011), Lustig and Verdelhan (2007), and Menkhoff et al. (2012)).

III. Empirical Model

We start by presenting our empirical model in Section III.A. In Section III.B, we discuss how we arrived at the assumed functional forms.

A. Currency Dynamics

We model bilateral log excess FX returns as

$$(4) \quad y_{t+1} \equiv (s_{t+1} - s_t) - (r_t - \tilde{r}_t) = \mu_t + v_t^{1/2} w_{t+1}^s + z_{t+1}^u - z_{t+1}^d,$$

where w_{t+1}^s is a standard normal shock (i.e., 0 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.⁴ Consistent with the empirical literature estimating conditional expected currency returns, we assume that $\mu_t = \mu_0 + \mu_r r_t + \tilde{\mu}_r \tilde{r}_t + \mu_v v_t$. Details regarding the properties of μ_t are provided in the Internet Appendix.

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

$$(5) \quad v_{t+1} = (1 - \nu)v + \nu v_t + \sigma_v v_t^{1/2} w_{t+1}^v + z_{t+1}^v,$$

which itself can jump with intensity h_t^v .⁵ The normal shocks to excess returns w^s and to conditional spot variance w^v have a correlation coefficient $\text{corr}(w^s, w^v) = \rho$. Finally, to ensure positivity of the variance when jumps are present, we only allow for upward jumps so that z_{t+1}^v 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

$$(6) \quad \text{Prob}(j_{t+1}^k = j) = e^{-h_t^k} (h_t^k)^j / j!, \quad k = u, d, v.$$

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

$$(7) \quad h_t^k = h_0^k + h_r^k r_t + \tilde{h}_r^k \tilde{r}_t + h_v^k v_t, \quad k = u, d, v.$$

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

$$(8) \quad z_t^k | j \sim \text{Gamma}(j, \theta_k), \quad k = u, d, v.$$

⁴This 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 normal variables, are jumps in discrete time. We model small jumps via normal 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.

⁵In continuous time, the Feller condition $\sigma_v^2 < 2\nu(1 - \nu)$ ensures that the variance stays positive if there are no jumps. A formal modeling of this process in discrete time is achieved via a Poisson mixture of gamma distributions (e.g., Gouriéroux and Jasiak (2006), Le, Singleton, and Dai (2010)). We use a direct discretization of the continuous-time counterpart because this is what most researchers do at the estimation stage anyway. We ensure that the variance stays positive at the estimation stage by a careful design of the simulation strategy.

Intuitively, because we consider daily data, a Bernoulli distribution is a very good approximation to our model because it is reasonable to assume no more than one jump per day. Then, the probability of a jump is $1 - e^{-h_t^k} \approx h_t^k$, and the distribution of the jump size is exponential with mean parameter θ_k .⁶

A popular alternative to our modeling approach is to specify the pricing kernel expressed in the units of domestic and foreign currencies. A ratio of the two would imply the exchange rate. One cannot identify how the multiple sources of risk contemplated in this article are priced using exchange rate data alone. Therefore, there is no need to use pricing kernel language in the context of the questions that we are addressing. The Internet Appendix outlines conditions when our approach (direct specification of the dynamics of y) is equivalent to the pricing kernel formulation. The Internet Appendix also illustrates why one cannot identify prices of risk using FX data only.

We model each exchange rate in isolation from others. This approach speaks to the large fraction of currency analysis, such as UIP regressions, option valuation, or equilibrium modeling, that is conducted on a currency-by-currency basis. We are able to identify the normal and non-normal shocks and how they should be modeled. 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 Internet Appendix discusses how one may approach modeling of the joint behavior of multiple currencies using our bilateral models as a foundation.

This discussion in the Internet Appendix has a number of implications. First, our bilateral models are a useful input into building joint behavior. Second, joint modeling leads to the proliferation of jump parameters because possible interactions between jumps in different currencies can be quite involved. Third, there are no economic restrictions gained by considering exchange rates jointly. For all these reasons, we leave a rather complicated analysis of joint behavior for future research.

To identify dynamics of variance better, we complement our data on exchange 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 Brandt and Santa-Clara (2002), Chernov and Ghysels (2000), Jones (2003), Pan (2002), and Pastorello, Renault, and Touzi (2000)). Our empirical specification is motivated by the fact that the risk-adjusted expectation of the average future variance realized over an option's lifetime is a linear function of state variables v_t , r_t , and \tilde{r}_t in our model. See the Internet Appendix for details.

Because implied variance is not exactly equal to the risk-adjusted expectation of the future variance, we treat the former as a noisy and biased observation of the conditional spot variance v_t :

$$(9) \quad IV_t = \alpha_{iv} + \beta_{iv} v_t + \sigma_{iv} v_t \sqrt{\lambda_t} \epsilon_t,$$

⁶Our choice of the variance jump size distribution is frequently used when modeling variance to ensure its positivity, as discussed previously. 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.

where IV_t is computed using prices of a 1-month at-the-money option and is expressed in daily terms; ϵ_t has a standard normal distribution, $N(0, 1)$; and λ_t has an inverse gamma distribution, $\mathcal{IG}(v/2, v/2)$. So the product $\sqrt{\lambda_t}\epsilon_t$ is t_v distributed (Cheung (2008), Jacquier, Polson, and Rossi (2004)).⁷ We have considered a version of equation (9) with non-zero loadings on r_t and \tilde{r}_t , but this specification did not find empirical support.⁸

Many authors use implied variance or model-free implied variance (MFIV) in empirical work by interpreting it as a very accurate approximation of the risk-adjusted expectation of the average future variance realized. 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 0 to retain the simple relationship. The importance of careful accounting for jumps is manifested more clearly in the literature on MFIV, such as VIX for the S&P 500, where analytic expressions are feasible. Martin (2013) shows that in the presence of jumps, VIX is equal to the risk-adjusted expected variance plus additional terms reflecting the higher-order risk-adjusted cumulants of returns. Our empirical specification allows us to avoid these difficulties in the explicit connection between implied variance or MFIV and risk-adjusted expected future variance. The cost of such an approach is our inability to estimate the risk-adjusted parameters of the model.

We conclude with a discussion of our approach to modeling 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 1-day r_t and \tilde{r}_t can be reasonably proxied with short-term yields. We view this feature as a strength of our approach because explicitly modeling the behavior of spot interest rates entails significant effort.⁹ Further, interest rates and currencies have a low conditional correlation, and variability in interest rates is much smaller than that in currencies. Thus, elaborate modeling and estimation of interest rate dynamics do not appear to be a first-order concern for the questions that we are addressing.

Nonetheless, we will need to use a model to compute the expectation of future variance that depends on the distribution of interest rates. In order to obtain reasonable quantities, we assume the simplest possible model for the interest rates. The foreign interest rate is $\tilde{r}_{t+1} = (1 - \tilde{b}_r)\tilde{a}_r + \tilde{b}_r\tilde{r}_t + \tilde{\sigma}_r\tilde{r}_t^{1/2}\tilde{w}_{t+1}^r$, and the domestic is the same without tildes. A square-root process for interest rates is subject to the 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

⁷Jones (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_w , a normal error with volatility $\sigma_w v_t$, and the Student t -error described previously. We find that heavy-tailed t_3 works very well.

⁸Interest rates have low variability relative to exchange rates and their variance. Therefore, it could be the case that the contribution of interest rates to the variation in implied variance cannot be empirically distinguished from the measurement error if the former is reasonably small.

⁹There is a separate area of the 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.

variation in parameters confirm our intuition that they have minimal impact on the role of normal and non-normal currency risks.

B. 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 previously described 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 et al. (2003), Eraker (2004), and Jones (2003)) and some of the few models of currencies (Bates (1996), Jorion (1988), and Maheu and McCurdy (2008)). The salient features of equity data are the presence of substantial moves up and down and a pronounced negative skewness in the return distribution. Therefore, jumps in equity returns are often modeled via a single compound Poisson process with a normally distributed non-zero-mean size. However, in contrast to equity returns, currency returns have very mild skewness over long samples, which suggests a zero-mean symmetric distribution for jump sizes.

Further, Bates (1996), Campa, Chang, and Reider (1998), and Carr and Wu (2007) emphasize the time-varying and sign-switching nature of the risk-adjusted skewness of exchange rates. The key to modeling 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 the potential to generate the requisite variation via either time-varying jump intensities or time-varying jump size distributions, or both. We do not explore time-varying jump means because 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. 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 is hard to distinguish the down and up components empirically. 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 is again in the equity literature, where all jumps in returns and variance are guided by the same (or at least correlated) Poisson processes. We find that the model with correlated Poisson processes fits the data poorly.

IV. 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 jump sizes) are a natural by-product of the estimation procedure. The Internet 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. The complicated dependencies of jump intensities on state variables and the sheer number of separate Poisson processes were 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 (equation (5)). At this stage, we select the best model by checking the significance of parameters on the basis of both credible intervals and Bayes odds ratios. Specifically, the parameters of concern are the ones controlling the jump intensity in equation (7) 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)).¹⁰ 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 of equations (4) and (5) and combine it with equation (9) 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 equation (7) 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. Such a result does not necessarily preclude the dependence of the FX distribution's asymmetry on the interest rate differential because the behavior of the FX skewness is dramatically affected by the 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 the expected jump would be equal to $\theta h_r (r_t - \tilde{r}_t)$.

¹⁰The recent literature on equity returns also finds support for self-exciting jumps (see, e.g., Ait-Sahalia, Cacho-Diaz, and Laeven (2015), Nowotny (2011), and Santa-Clara and Yan (2010)).

Thus, the excess return asymmetries will be directly driven by the interest rate differential, as noted by Brunnermeier et al. (2008).¹¹ The final version of this model that incorporates all the unrejected null hypotheses is referred to as the preferred model.

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. The Internet Appendix describes the details.

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 (the so-called automatic Occam's razor). The computations are described in the Internet Appendix.

V. Results

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

A. Statistical Properties of Currency Risks

Tables 2–5 report the parameter estimates. The corresponding model diagnostics are reported in the Internet Appendix. Table 6 displays the results of specification tests on the basis of Bayes odds ratios. Table 7 summarizes the parameters of the calibrated interest rates.

The results exhibit a lot of similarities across the different currencies.¹² As we move from models with stochastic variance to models with 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 ν decline. Taking AUD as an example, ν declines from 0.9943 to 0.9855. This seemingly small change translates into a drop in the half-life of the conditionally normal component, $\ln(2/(1-\nu))$, 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-persistence 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. IV *var* (time-series average of the

¹¹An alternative way to generate such a relation between skewness and the interest rate differential is to hardwire the linear dependence of jump intensity on the latter. Because the interest rate differential may switch signs, the probability of a jump may become negative.

¹²This is not the focus of our article, but for completeness, we provide a discussion of the theoretical and empirical implications of our model for UIP regressions in the Internet Appendix.

TABLE 2
AUD Parameter Estimates

In Table 2, the estimates correspond to daily excess currency returns, in percentages. The 95% credible intervals are reported in parentheses. The preferred model is as follows:

$$\begin{aligned}
 Y_{t+1} &= \mu_0 + \mu_r(r_t - \bar{r}_t) + \mu_v V_t + V_t^{1/2} W_{t+1}^s + Z_{t+1}^u - Z_{t+1}^d, \\
 V_{t+1} &= (1 - \nu)V + \nu V_t + \sigma_v V_t^{1/2} W_{t+1}^v + Z_{t+1}^v, \\
 h_t^u &= h_0 + h_r r_t, \quad h_t^d = h_0 + h_r \bar{r}_t, \quad h_t^v = h_0^v + h_v V_t, \\
 z_t^u | j &\sim \text{Gamma}(j, \theta), \quad z_t^d | j \sim \text{Gamma}(j, \theta), \quad z_t^v | j \sim \text{Gamma}(j, \theta).
 \end{aligned}$$

"Risk-adjusted" is the preferred model calibrated to 1-month options (with parameters that are denoted by symbols fixed at their time-series values).

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

variance of the error term in equation (9) 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_r . Bayes odds ratios strongly favor 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 the Internet Appendix 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

TABLE 3
CHF Parameter Estimates

In Table 3, the estimates correspond to daily excess currency returns, in percentages. The 95% credible intervals are reported in parentheses. The preferred model is as follows:

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

“Risk-adjusted” is the preferred model calibrated to 1-month options (with parameters that are denoted by symbols fixed at their time-series values).

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

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 different currencies). Jumps in variance are self-exciting, so a jump increases the likelihood of another jump. When the variance jumps by θ_v , the 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 of Brunnermeier et al. (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

TABLE 4
GBP Parameter Estimates

In Table 4, the estimates correspond to daily excess currency returns, in percentages. The 95% credible intervals are reported in parentheses. The preferred model is as follows:

$$\begin{aligned}
 Y_{t+1} &= \mu_0 + \mu_r(r_t - \tilde{r}_t) + \mu_v V_t + V_t^{1/2} W_{t+1}^s + Z_{t+1}^u - Z_{t+1}^d, \\
 V_{t+1} &= (1 - \nu)V + \nu V_t + \sigma_v V_t^{1/2} W_{t+1}^v + Z_{t+1}^v, \\
 h_t^u &= h_0 + h_r r_t, \quad h_t^d = h_0 + h_r \tilde{r}_t, \quad h_t^v = h_0^v + h_v V_t, \\
 z_t^u | j &\sim \mathcal{G}\text{amma}(j, \theta), \quad z_t^d | j \sim \mathcal{G}\text{amma}(j, \theta), \quad z_t^v | j \sim \mathcal{G}\text{amma}(j, \theta).
 \end{aligned}$$

"Risk-adjusted" is the preferred model calibrated to 1-month options (with parameters that are denoted by symbols fixed at their time-series values).

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

trade investor who borrows money in USD and invests in AUD. This investor can lose 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 a 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.

The probability of USD depreciation, h_r^u , depends positively on the U.S. interest rate. This property implies that when the interest rate differential is positive,

TABLE 5
JPY Parameter Estimates

In Table 5, the estimates correspond to daily excess currency returns, in percentages. The 95% credible intervals are reported in parentheses. The preferred model is as follows:

$$\begin{aligned}
 Y_{t+1} &= \mu_0 + \mu_r(r_t - \bar{r}_t) + \mu_v V_t + V_t^{1/2} W_{t+1}^s + Z_{t+1}^u - Z_{t+1}^d, \\
 V_{t+1} &= (1 - \nu)V + \nu V_t + \sigma_v V_t^{1/2} W_{t+1}^v + Z_{t+1}^v, \\
 h_t^u &= h_0 + h_r r_t, \quad h_t^d = h_0 + h_r \bar{r}_t, \quad h_t^v = h_0^v + h_v V_t, \\
 z_t^u | j &\sim \mathcal{G}amma(j, \theta), \quad z_t^d | j \sim \mathcal{G}amma(j, \theta), \quad z_t^v | j \sim \mathcal{G}amma(j, \theta).
 \end{aligned}$$

“Risk-adjusted” is the preferred model calibrated to 1-month options (with parameters that are denoted by symbols fixed at their time-series values).

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

TABLE 6
Log Odds Ratios

Table 6 presents a statistical model assessment. We compare the SV ($\theta=0, \theta_v=0$), SVJ ($\theta=0$), and preferred models pair-wise. 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).

Model	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

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

TABLE 7
Calibration of the Interest Rates

In Table 7, we calibrate processes for domestic (U.S.) interest rates as

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

and foreign interest rates as

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

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

Parameter	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

rate. Indeed, the third conditional cumulant can be obtained by taking the third derivative of the cumulant-generating function of log currency returns. For example, in the preferred model, $\kappa_{3r}(s_{t+1} - s_t) = 6\theta^3 h_r (r_t - \tilde{r}_t)$.

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

Parameters reflecting the average jump size have a different interpretation compared with the jump in variance. The former is a jump in exchange rate expressed as a percentage of that exchange rate. The latter is an absolute magnitude of a 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 “2-sigma” event. The 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.

B. 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, Table 8, and Tables 1–4 of the Internet Appendix.

For each exchange rate, the figures display the time series of data (excess returns and implied volatilities) complemented by 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 because 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 1 if their actual value is 0.5 or higher, and 0 otherwise. Then we plot the estimated jump sizes on the days with the assigned value of 1. Our strategy yields 219 jumps overall across all currencies. Approximately 25% of these jumps take place simultaneously in at least two currencies. We refer

TABLE 8
Decomposition of the Total Risk

In Table 8, 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).

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

to such jumps as common. There are 8 episodes when FX and variance jump at the same time. We overlay the plots of the estimated jump sizes with the ex ante jump probabilities h_t^v , h_t^u , and h_t^d .

To enhance interpretation of the plots, we 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 that the mean is approximately 0. Thus, a “regular” excess return can be within the range of $\pm 2\%$ ($\mu \pm 3\sigma$). Graph A 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 also time-varying and uncertain. The estimation procedure takes this uncertainty into account by producing the ex post probability of a jump taking place and estimated jump sizes. The magnitude of estimated jumps 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 Oct. 28, 2008, the plot of excess returns shows a big drop of 5.5%. The model shows no jumps on that day. Such a big move is generated via a normal component because of the jump in variance. Volatility jumps by $(v_{t-1} + z_t^v)^{1/2} - v_{t-1}^{1/2} = 0.18\%$ (2.9% annualized), on average, on each of the

3 days between Oct. 22 and 24. Each day, the jump in variance is increasing the probability of a jump the following day. Over these 3 days, volatility moves from 1.35% (21.5% annualized) to 1.8% (28% annualized). To put this number into perspective, the long-run volatility mean is $v_j^{1/2} = 0.66\%$ (10.4% annualized). Thus, the increase in volatility over these 3 days is roughly equal to the average level of volatility. Moreover, there is no significant news associated with either Oct. 22–24 or Oct. 28. Thus, we attribute these events to market uncertainty.

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 takes place after 2000. In fact, even the famous “Black Wednesday,” the GBP devaluation on Sept. 16, 1992, is attributed primarily to a jump in variance on Sept. 8. The volatility moves from 0.93% (14.7% annualized) to 1.13% (17.9% annualized), and then it drifts up to 1.15% (18.3% annualized) on the day of the crash. These values of volatility are high; 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_j 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 the exchange rate is only 26% (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 experiences a jump, we check if there is 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 complementary to the one employed in studies of the impact of news surprises on financial assets because 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 the overall market anxiety that is sometimes evident in the news. The Internet 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. In stark contrast, only 23% of jumps in variance across all currency pairs are associated with such events.¹³ 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 a critical difference between jumps in FX and variance. Tables 1–4 of the Internet Appendix provide a summary of the types of announcements associated with jumps.

¹³We avoid double counting by accounting for common jump events in different currency pairs only once, and we do not count simultaneous jumps in FX and variance (only 8 incidents across all currencies).

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 1980s and 1990s. Later in the sample, AUD, as a high-yield commodity currency, is vulnerable to events around the Asian crisis and credit crunch, 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 U.S. economy and third countries' economies. The exceptions are an interest rate reduction by the Swiss National Bank on Nov. 21, 2008, and news indicating monetary policy uncertainty on Mar. 12, 2009. The British pound has the largest number of jumps in variance; hence, most of the jump events (72%) are driven by economic uncertainty, such as uncertainty about the monetary system in Europe (exchange rate mechanism, euro launch) and spillovers from the U.S. 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 realized jump returns on the yen tend to be positive in the recent financial crisis.

C. 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:

$$(10) \quad L_t(S_{t+n}/S_t) = \ln E_t(e^{s_{t+n} - s_t}) - E_t(s_{t+n} - s_t).$$

Entropy is a loaded term. Our use of entropy is closest to the way it is used by Martin (2013), who analyzes the entropy of equity index returns. Entropy's connection to the cumulants of log FX returns makes it attractive for our purposes. Specifically, the definition in equation (10) implies that

$$(11) \quad L_t(S_{t+n}/S_t) = k_t(1; s_{t+n} - s_t) - \kappa_{1t}(s_{t+n} - s_t)$$

$$(12) \quad = \kappa_{2t}(s_{t+n} - s_t)/2! + \kappa_{3t}(s_{t+n} - s_t)/3! \\ + \kappa_{4t}(s_{t+n} - s_t)/4! + \dots,$$

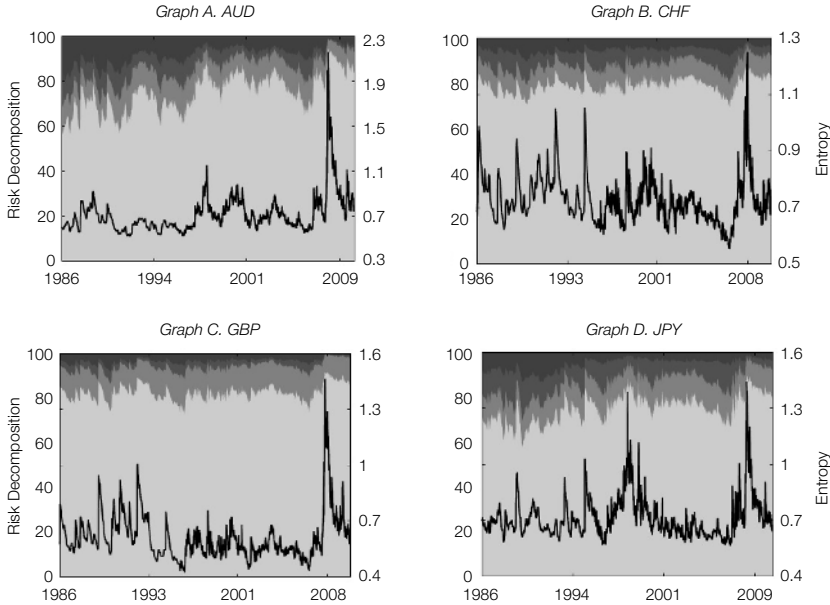
where $k_t(1; s_{t+n} - s_t)$ is the conditional cumulant-generating function of $s_{t+n} - s_t$ evaluated at the argument equal to 1, $\kappa_{jt}(s_{t+n} - s_t)$ is the j th conditional cumulant of $s_{t+n} - s_t$, and we use the fact that $k_t(1; s_{t+n} - s_t)$ is equal to the infinite sum of scaled cumulants. The significance of the property in equation (12) is that if currency changes are normally distributed, then entropy is equal to half of their variance (the first term in the sum). All the higher-order terms arise from non-normalities. Thus, entropy captures the tail behavior of returns in a compact form and serves as a natural generalization of variance as a risk measure.

FIGURE 5
 Conditional Decomposition of the Total Risk for Quarterly Returns

Figure 5 displays the cumulative contribution of the different risks to the total risk of excess returns (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 j th cumulant of log FX returns. The investment horizon is $n=63$ days. The contributions of the down jumps in FX, of the up jumps in FX, of the jumps in variance, and of the normal shocks are displayed in gray, from darkest to lightest gray. The black line shows $\sqrt{2L_t/n}$ in percentages (right axis). This quantity has an interpretation of 1-period volatility in the case of normally distributed returns.



We decompose the total risk of currency returns into the contribution of the jump and normal components. The Internet Appendix explains how we compute the full entropy on the basis of equation (11). We can compute the individual components by zeroing out the rest.¹⁴ Figure 5 displays the contributions of these components for the investment horizon of 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 percentages. Thus, we plot $\sqrt{2L_t/n}$. Finally, Table 8 reports summary statistics of the relative contributions of the different components for horizons of 1 month and 3 months.

¹⁴If 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.

The regular, or normal, risk is the most prominent regardless of the horizon. The average total contribution of jumps at a 1-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, 6.82%–9.17%] is higher than that of a jump in variance [the range, 3.76%–4.83%] and has higher time variation at a 1-month horizon (GBP is the exception because jumps in variance contribute 4.37% compared with 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 as we extend the investment horizon to 3 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, the individual contribution of jumps in variance [the range, 8.94%–11.48%] is higher than that of jumps in FX [the range, 2.78%–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 toward 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. 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.

VI. 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 offer initial thoughts on this topic, leaving an in-depth analysis for future research.

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. There are two cases in which it could be useful to study a portfolio in order to establish whether a particular risk is priced. First, a portfolio could be systematic or diversified in an equilibrium sense, and therefore, economic theory would suggest that all of the risk that it is exposed to should be priced. A value-weighted portfolio of all traded assets is an example of such a portfolio in the capital asset pricing model (CAPM) world. Second, it can be fruitful to construct a portfolio that is exposed to a single risk and nothing else. If that portfolio has a significant excess return, then it can be taken as empirical evidence that the risk must be priced.

It is difficult to examine the first type of portfolio because currencies are not claims on tangible assets. Therefore, we are not aware of an existing theory that would deliver a systematic currency portfolio. Nevertheless, we study the empirical properties of the dollar index introduced by Lustig et al. (2014). With some stretch of the imagination, one could relate this portfolio to a systematic

portfolio in a CAPM world where the only traded assets are sovereign bonds of different countries. If we treat the U.S. bond as risk-free (and not just default-free), then a portfolio of bonds in other currencies, weighted by the supply of these bonds, could be considered as systematic. In practice, the equal-weighted dollar index should be close to such an idealized value-weighted portfolio. Although we readily admit that this argument is quick and dirty, this is the best motivation we could have come up with for a proxy for a systematic portfolio. All other currency portfolios considered in the literature seem further away from being systematic.

Given the complexity of the risks identified in the bilateral exchange rates, one cannot construct portfolios of the second type using currencies alone. In this situation, a more direct way to assess the valuation of crash risk is to use a model to estimate 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 article. Instead, we provide a calibration of jump risk premiums on the basis of average implied volatility smiles.

A. Crash Risk of the Dollar Index

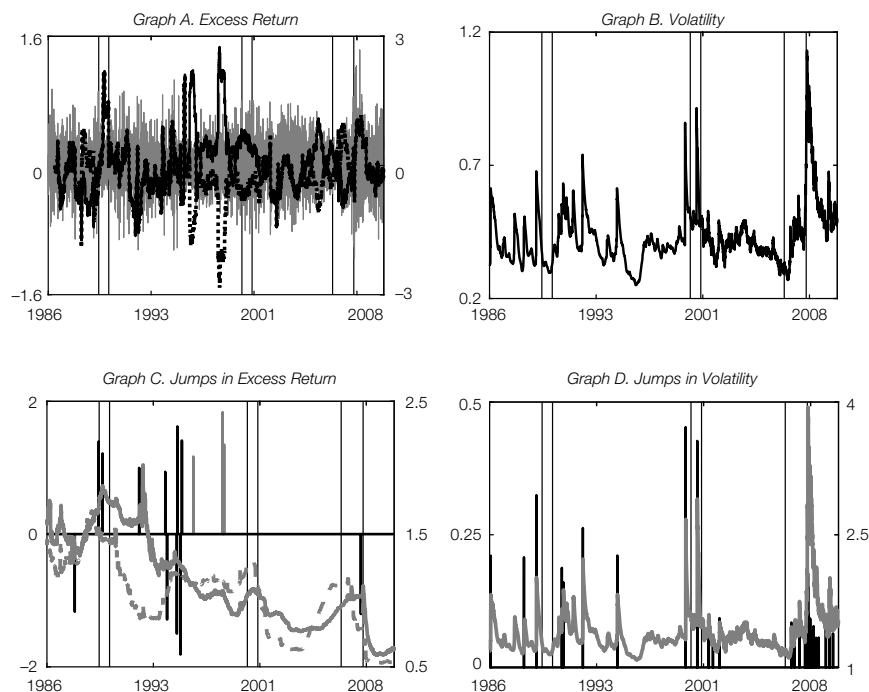
We closely follow Lustig et al. (2014) in constructing the dollar index. We use daily exchange rates on the same 21 currency pairs against USD from Jan. 1, 1986, to Dec. 31, 2010. We replace 11 European currencies with the euro starting from Jan. 1, 1999. We construct an equally weighted index of currencies against USD: An investor buys a \$1 equivalent of each foreign currency in the basket. The 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. Graph A of Figure 6 reveals that skewness is time-varying within a range similar to that of individual FX rates.

Options on the index are not traded, so we do not have an option-implied variance to use in equation (9) 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 1-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 Feb. 1, 2001, to Dec. 31, 2010. The choice of currencies and data span is limited by the 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} from equation (9) may not be in the neighborhood 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 article, to measure the first-order effect of jumps on currency indices.

We use the 1-month LIBOR or the 1-month Eurocurrency rate as a proxy for the 1-period interest rate. If these are unavailable, we use their overnight or 3-month counterparts. The foreign interest rate is an equal-weighted index of the

FIGURE 6
Index and Estimated States

In Graph A of Figure 6, the thin gray line displays observed log currency returns (right scale); the solid black line displays the conditional 6-month skewness of the log currency returns (left scale); the dashed black line displays the conditional 6-month skewness of the log currency returns of the dollar carry trade. In Graph B, the thin black line shows the estimated spot volatility \sqrt{V}_t . The bottom graphs display estimated jump sizes in returns (Graph C) and volatility (Graph D) with jump intensities. In Graph C, the gray solid line is the intensity of the jump down, the gray dashed line is the intensity of the jump up, the black bars are the jumps themselves, and the gray bars are jumps that change the direction if there is rebalancing based on the interest rate differential. In Graph D, the gray line is the intensity of the jump in volatility; the black 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%. Vertical bars with a thin solid border indicate recessions in the United States.



respective individual rates. We calibrate the same model as for the individual interest rates. See Table 7.

Table 9 reports the parameter estimates (the corresponding model diagnostics are in the Internet Appendix). Table 6 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.

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 year). 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 slightly smaller than the average jump size for AUD (1.69) and is larger than those for

TABLE 9
Index Parameter Estimates

In Table 9, the estimates correspond to daily excess returns on the currency index, in percentages. The 95% credible intervals are reported in parentheses. The preferred model is as follows:

$$\begin{aligned}
 Y_{t+1} &= \mu_0 + \mu_r(r_t - \bar{r}_t) + \mu_v V_t + V_t^{1/2} W_{t+1}^s + Z_{t+1}^u - Z_{t+1}^d, \\
 V_{t+1} &= (1 - \nu)V + \nu V_t + \sigma_v V_t^{1/2} W_{t+1}^v + Z_{t+1}^v, \\
 h_t^u &= h_0 + h_r r_t, \quad h_t^d = h_0 + h_r \bar{r}_t, \quad h_t^v = h_0^v + h_v V_t, \\
 z_t^u | j &\sim \text{Gamma}(j, \theta), \quad z_t^d | j \sim \text{Gamma}(j, \theta), \quad z_t^v | j \sim \text{Gamma}(j, \theta_v).
 \end{aligned}$$

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

other currencies. The 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 6 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 Aug. 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: Aug. 8, 2008.

Table 5 of the Internet Appendix 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 common jumps do not get diversified away in an index.

The findings are consistent with the notion of diversification in that the normal 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 8 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 as 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 compared with individual FX rates. The role of jumps in variance increases with the investment horizon, similarly to individual currency pairs.

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 et al. (2014) argue that investing in the basket of foreign currencies when the corresponding average foreign interest rate exceeds the U.S. rate and in the USD otherwise is i) more profitable than a conventional carry portfolio (high minus low) and ii) 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-à-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 the dollar carry 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 favor 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 gray in Graph C of Figure 6). So the realized Sharpe ratio increases at the cost of making returns negatively skewed; the skewness of the rebalanced index is -0.10 , whereas that of the long-only index is 0.03 . The plot of the time-varying skewness of the dollar carry portfolio in Graph A shows a higher incidence of negative values compared with 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 the speculator, at least in our sample. Conditioning on interest rates has no effect on jumps in variance.

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

B. 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 scope of this article. 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 Oct. 16, 2003, to Dec. 31, 2010, and for CHF from Jan. 5, 2006, to Dec. 31, 2010; we start as early as data on five strikes become available (0.25 and 0.10 risk reversals and butterflies, at-the-money (ATM) options). We compute the average smiles for 1-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 7 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 with an asterisk *. The “volatility of volatility” σ_v and the “leverage” parameter ρ are the same under both probabilities by convention, a by-product 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 prices of domestic and foreign 1-period risk-free bonds, respectively. Dividing by q_t^1 and taking logs, we obtain the following:

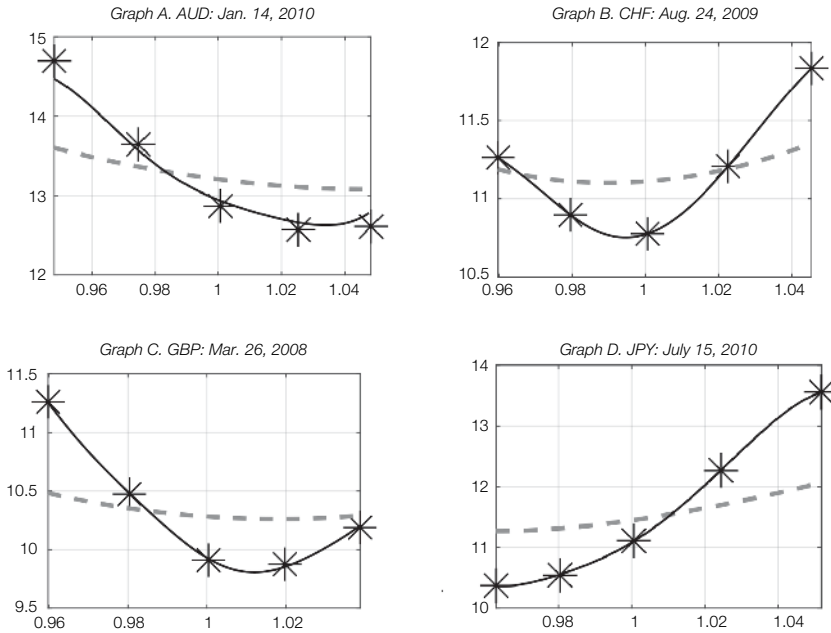
$$\begin{aligned} r_t - \tilde{r}_t &= \ln E_t^*(e^{s_{t+1} - s_t}) \\ &= \mu_t^* + v_t/2 + h_t^{u*}(E_t^* \exp(z_{t+1}^u) - 1) \\ &\quad + h_t^{d*}(E_t^* \exp(z_{t+1}^d) - 1) \\ &= \mu_t^* + v_t/2 + h_t^{u*}((1 - \theta_u^*)^{-1} - 1) + h_t^{d*}((1 + \theta_d^*)^{-1} - 1). \end{aligned}$$

Thus,

$$(13) \quad \mu_t^* = r_t - \tilde{r}_t - v_t/2 - h_t^{u*}\theta_u^*/(1 - \theta_u^*) + h_t^{d*}\theta_d^*/(1 + \theta_d^*).$$

FIGURE 7
Implied Volatility

In Figure 7, 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 gray lines display theoretical smiles when only compensation for normal risks is allowed (we optimize over v^* and v^* only). The solid black lines depict theoretical implied volatilities when jump risk premiums are allowed as well (we optimize over v^* , v^* , $h_0^{v^*}$, $h_0^{u^*}$, and $h_0^{d^*}$).



In our first exercise, we do not allow for any jump risk premiums. That is, we try to fit the smiles in Figure 7 by minimizing the sum of squared implied volatility errors and by allowing for deviation between v and v^* , and v and v^* only. Our best fit is depicted with dashed gray 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 et al. (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 1 day per currency. Therefore, we cannot separate differences between risk-adjusted parameters that control the constant and the time-varying parts of jump intensities. Therefore, we only allow for $h_0^{v^*} \neq h_0^v$ and $h_0^{d^*} \neq h_0^d$. Although we find the constant components of $h_0^{u^*}$ and h_0^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 7 depicts the results of the second exercise with solid lines. The smiles can be fitted perfectly. Although 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 1 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 (13).¹⁶ To interpret these parameters, consider GBP as an example. 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 (see the Internet Appendix). 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 a 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 in the next period. The price of this security is

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

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 (a.k.a. the risk premium) is $rpz_t^k \equiv \ln[E_t(I_{J_{t+1}=0}^k)/P_t^k] - r_t = h_t^{k*} - h_t^k$. The risk premiums range (in percentages 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 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 this endeavor for future research.

VII. 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

¹⁶Readers may wonder why the loading of -0.5 on v , in equation (13) 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 percentages. Thus, in order to get a sensible result for log returns expressed in percentages, the easiest solution is to use -0.005 and keep other quantities as they are.

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 two important extensions of our analysis. First, currency options can be used to estimate the pricing kernel. This would allow one to characterize how the risks documented in this article are valued in the marketplace. Second, our evidence suggests informally that common jump risks are shared across different currencies. It would be useful to establish a model that explicitly incorporates common and country-specific jump components.

References

- Ait-Sahalia, Y.; J. Cacho-Diaz; and R. J. A. Laeven. "Modeling Financial Contagion Using Mutually Exciting Jump Processes." *Journal of Financial Economics*, 117 (2015), 585–600.
- Alvarez, F., and U. Jermann. "Using Asset Prices to Measure the Persistence of the Marginal Utility of Wealth." *Econometrica*, 73 (2005), 1977–2016.
- Andersen, T. G.; L. Benzoni; and J. Lund. "An Empirical Investigation of Continuous-Time Equity Returns." *Journal of Finance*, 57 (2002), 1239–1284.
- Bacchetta, P., and E. Van Wincoop. "Can Information Heterogeneity Explain the Exchange Rate Determination Puzzle?" *American Economic Review*, 96 (2006), 552–576.
- Backus, D.; M. Chernov; and I. Martin. "Disasters Implied by Equity Index Options." *Journal of Finance*, 66 (2011), 1969–2012.
- Baillie, R. T., and T. Bollerslev. "The Message in Daily Exchange Rates: A Conditional-Variance Tale." *Journal of Business and Economic Statistics*, 7 (1989), 297–305.
- Bakshi, G.; P. Carr; and L. Wu. "Stochastic Risk Premiums, Stochastic Skewness in Currency Options, and Stochastic Discount Factors in International Economics." *Journal of Financial Economics*, 87 (2008), 132–156.
- Barroso, P., and P. Santa-Clara. "Beyond the Carry Trade: Optimal Currency Portfolios." *Journal of Financial and Quantitative Analysis*, 50 (2015), 1037–1056.
- Bates, D. K. "Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options." *Review of Financial Studies*, 9 (1996), 69–107.
- Bekaert, G., and R. J. Hodrick. "Characterizing Predictable Components in Excess Returns on Equity and Foreign Exchange Markets." *Journal of Finance*, 47 (1992), 467–509.
- Bekaert, G., and R. J. Hodrick. "On Biases in the Measurement of Foreign Exchange Risk Premiums." *Journal of International Money and Finance*, 12 (1993), 115–138.
- Bilson, J. F. O. "The 'Speculative Efficiency' Hypothesis." *Journal of Business*, 54 (1981), 435–451.
- Brandt, M. W., and P. Santa-Clara. "Simulated Likelihood Estimations of Diffusions with an Application to Exchange Rate Dynamics in Incomplete Markets." *Journal of Financial Economics*, 63 (2002), 161–210.
- Brennan, M. J., and Y. Xia. "International Capital Markets and Foreign Exchange Risk." *Review of Financial Studies*, 19 (2006), 753–795.
- Broadie, M.; M. Chernov; and M. Johannes. "Model Specification and Risk Premia: Evidence from Futures Options." *Journal of Finance*, 62 (2007), 1453–1490.
- Brunnermeier, M. K.; S. Nagel; and L. H. Pedersen. "Carry Trades and Currency Crashes." In *NBER Macroeconomics Annual*, Vol. 23, D. Acemoglu, K. Rogoff, and M. Woodford, eds. Chicago, IL: University of Chicago Press (2008).
- Campa, J. M.; P. H. K. Chang; and R. L. Reider. "Implied Exchange Rate Distributions: Evidence from OTC Option Markets." *Journal of International Money and Finance*, 17 (1998), 117–160.
- Carr, P., and L. Wu. "Stochastic Skew in Currency Options." *Journal of Financial Economics*, 86 (2007), 213–247.
- Chernov, M. "On the Role of Risk Premia in Volatility Forecasting." *Journal of Business and Economic Statistics*, 25 (2007), 411–426.
- Chernov, M.; A. R. Gallant; E. Ghysels; and G. Tauchen. "Alternative Models for Stock Price Dynamics." *Journal of Econometrics*, 116 (2003), 225–257.
- Chernov, M., and E. Ghysels. "A Study Towards a Unified Approach to the Joint Estimation of Objective and Risk Neutral Measures for the Purposes of Options Valuations." *Journal of Financial Economics*, 56 (2000), 407–458.
- Cheung, S. "An Empirical Analysis of Joint Time-Series of Returns and the Term-Structure of Option Prices." Working Paper, Columbia University (2008).

- Duffie, D.; J. Pan; and K. Singleton. "Transform Analysis and Asset Pricing for Affine Jump-Diffusions." *Econometrica*, 68 (2000), 1343–1376.
- Engel, C., and J. D. Hamilton. "Long Swings in the Dollar: Are They in the Data and Do Markets Know It?" *American Economic Review*, 80 (1990), 689–713.
- Engle, R. F.; T. Ito; and W.-L. Lin. "Meteor Showers or Heat Waves? Heteroskedastic Intra-Daily Volatility in the Foreign Exchange Market." *Econometrica*, 58 (1990), 525–542.
- Eraker, B.; M. Johannes; and N. Polson. "The Impact of Jumps in Volatility and Returns." *Journal of Finance*, 58 (2003), 1269–1300.
- Eraker, B. "Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices." *Journal of Finance*, 58 (2004), 1269–1300.
- Fama, E. F. "Forward and Spot Exchange Rates." *Journal of Monetary Economics*, 14 (1984), 319–338.
- Gourieroux, C., and J. Jasiak. "Autoregressive Gamma Processes." *Journal of Forecasting*, 25 (2006), 129–152.
- Graveline, J. J. "Exchange Rate Volatility and the Forward Premium Anomaly." Working Paper, University of Minnesota (2006).
- Harvey, C. R., and R. D. Huang. "Volatility in the Foreign Currency Futures Market." *Review of Financial Studies*, 4 (1991), 543–569.
- Hawkes, A. G. "Spectra of Some Self-Exciting and Mutually Exciting Point Processes." *Biometrika*, 58 (1971), 83–90.
- Jacquier, E.; N. G. Polson; and P. E. Rossi. "Bayesian Analysis of Stochastic Volatility Models with Fat-Tails and Correlated Errors." *Journal of Econometrics*, 122 (2004), 185–212.
- Jarrow, R., and D. Madan. "Option Pricing Using the Term Structure of Interest Rates to Hedge Systematic Discontinuities in Asset Returns." *Mathematical Finance*, 5 (1995), 311–336.
- Johannes, M., and N. G. Polson. "MCMC Methods for Continuous-Time Financial Econometrics." In *Handbook of Financial Econometrics*, Y. Ait-Sahalia and L. Hansen, eds. Amsterdam, Netherlands: North-Holland (2009).
- Jones, C. S. "The Dynamics of Stochastic Volatility: Evidence from Underlying and Options Markets." *Journal of Econometrics*, 116 (2003), 181–224.
- Jorion, P. "On Jump Processes in the Foreign Exchange and Stock Markets." *Review of Financial Studies*, 1 (1988), 427–445.
- Kass, R. E., and A. E. Raftery. "Bayes Factors." *Journal of the American Statistical Association*, 90 (1995), 773–795.
- Le, A.; K. J. Singleton; and Q. Dai. "Discrete-Time Affine^Q Term Structure Models with Generalized Market Prices of Risk." *Review of Financial Studies*, 23 (2010), 2184–2227.
- Lustig, H.; N. Roussanov; and A. Verdelhan. "Common Risk Factors in Currency Markets." *Review of Financial Studies*, 24 (2011), 3731–3777.
- Lustig, H.; N. Roussanov; and A. Verdelhan. "Countercyclical Currency Risk Premia." *Journal of Financial Economics*, 111 (2014), 527–553.
- Lustig, H., and A. Verdelhan. "The Cross-Section of Foreign Currency Risk Premia and US Consumption Growth Risk." *American Economic Review*, 97 (2007), 89–117.
- Maheu, J. M., and T. H. McCurdy. "Modeling Foreign Exchange Rates with Jumps." In *Forecasting in the Presence of Structural Breaks and Model Uncertainty*, Vol. 3, D. E. Rapach and M. E. Wohar, eds. Bingley, UK: Emerald (2008).
- Martin, I. "Simple Variance Swaps." NBER Working Paper (2013).
- Menkhoff, L.; L. Sarno; M. Schmelzing; and A. Schrimpf. "Carry Trades and Global Foreign Exchange Volatility." *Journal of Finance*, 67 (2012), 681–718.
- Neely, C. J. "Target Zones and Conditional Volatility: The Role of Realignments." *Journal of Empirical Finance*, 6 (1999), 177–192.
- Nowotny, M. "Disaster Begets Crisis: The Role of Contagion in Financial Markets." Working Paper, Boston University (2011).
- Pan, J. "The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study." *Journal of Financial Economics*, 63 (2002), 3–50.
- Pastorello, S.; E. Renault; and N. Touzi. "Statistical Inference for Random-Variance Option Pricing." *Journal of Business and Economic Studies*, 18 (2000), 358–367.
- Platen, E., and R. Rebolledo. "Weak Convergence of Semimartingales and Discretization Methods." *Stochastic Processes and Their Applications*, 20 (1985), 41–58.
- Santa-Clara, P., and S. Yan. "Crashes, Volatility, and the Equity Premium: Lessons from S&P 500 Options." *Review of Economics and Statistics*, 92 (2010), 435–451.
- Tryon, R. "Testing for Rational Expectations in Foreign Exchange Markets." Working Paper, Federal Reserve (1979).