

Do interest rate differentials drive the volatility of exchange rates? Evidence from an extended stochastic volatility model

December 2021

M. Ulm^{‡†} and J. Hambuckers^{†*}

Abstract

We study the link between the volatility of exchange rates and interest rate differentials (IRD), motivated by the importance of currency carry trade activities in exchange rate dynamics. We examine this link by means of an extended stochastic volatility model, for which we detail an efficient estimation strategy based on Gaussian mixture sampling and a linearization of the volatility process. We apply this approach to six currency pairs over the period from January 1999 to December 2017. Our results suggest that changes in IRD affect volatility differently for low and high-interest-rate currencies. The volatility reacts strongly and positively to increases in the low interest rate, an effect consistent with the unwinding of carry trade positions. In contrast, the response to a raise in the high interest rate is negative and substantially smaller. In general, we find that the informational content of the interest rate differentials regarding the volatility of exchange rate is greater during and after the global financial crisis, compared to the pre-crisis period.

JEL classification: C22, C11, C58

Keywords: stochastic volatility, exchange rate volatility, uncovered interest rate parity, currency carry trading.

[‡]Georg-August-University Goettingen, Department of Economics, Humboldtallee 3, 37073 Göttingen, Germany.

[†]University of Liège, HEC Liège Management School, Department of Finance, Rue Louvrex 14, 4000 Liège, Belgium.

*Corresponding author: jhambuckers@uliege.be

1 Introduction

The empirical failure of the uncovered interest rate parity condition (UIP) has become an active area of research during the past decades (see Rossi, 2013, for a review). UIP hypothesizes that any profit obtained from a difference in interest rates in two economies should be offset by an adverse movement in the exchange rate of their currencies. In turn, it implies that interest rate differentials (IRD) should serve as useful predictors of the future variations in spot foreign exchange rates. However, this result is only rarely supported by the data (Meese and Rogoff, 1983). A classic example of this failure can be found in the positive excess profit obtained from currency carry trading strategies¹ whereby this situation is ruled out by UIP (Burnside et al., 2011). Hence, the empirical evidence suggests that exchange rates and interest rates are connected differently than supposed by UIP.

To reconcile economic theory with empirical evidence, the literature follows two strands. On the one hand, the lack of support for UIP is explained by instabilities in the dependence between interest rates and exchange rates. For instance, Fama (1984) and, more recently, Li et al. (2011) and Christiansen et al. (2011) highlight the role of time-varying risk premia. Ismailov and Rossi (2018) examine the impact of uncertainty that might blur the relationship between exchange rates and IRD. Accominotti et al. (2019) emphasize the effect of the currency regime on the profitability of a carry trade strategy. On the other hand, several authors investigate the possibility that economic fundamentals and exchange rates exhibit high-order dependencies instead of a restricted dependence on the mean. For example, Chung and Hong (2007) study the link between IRD and high-order conditional moments of variations in foreign exchange rates (FX). They then demonstrate the usefulness of IRD as a predictor of the sign of changes in exchange rate. Brunnermeier et al. (2009) highlight the link between IRD and the skewness of exchange rate returns, suggesting that carry trade strategies are profitable because of the existence of a crash risk premium. Chernov et al. (2018) disentangle true currency crashes from normal volatility episodes, relying on a stochastic volatility (SV) model with a jump probability depending on interest rates.

Our study contributes to this second strand of the literature by examining the link

¹Carry trades describe investment strategies that comprise borrowing in low-interest-rate currencies and investing in high-interest-rate currencies to realize a profit from the IRD

between IRD and FX volatility². Our motivations are rooted in the role played by the financial channel, in particular carry trading activities, in the determination of exchange rates in imperfect financial markets. As highlighted by Froot and Ramadorai (2005), capital flows, financial conditions, and financiers' positions play an important role in the variations of currency returns. Gabaix and Maggiori (2015) formalize this mechanism in a theoretical model in which capital flows, influenced by relative levels of interest rates, affect both the level and the volatility of exchange rates via alterations in the balance sheets of financiers. A similar approach is used by de Rato (2007), Brunnermeier et al. (2009), and Farhi and Gabaix (2016) to describe exchange rate dynamics: if the IRD changes, speculators start unwinding their carry trade positions, leading to movements in capital flows. These capital flows, in turn, result in sudden exchange rate movements that materialize in an increase of the conditional volatility of these variations. Thus, considering that the estimated daily volume of FX transactions amounts to US \$5.1 trillion³, we should expect the relevance of the currency trading channel for short-term (i.e. daily) exchange rate movements and, by extension, of a link between IRD and FX volatility.

Our first contribution is therefore to investigate the following questions: are the daily FX volatility dynamics driven by IRD, as postulated by Gabaix and Maggiori (2015)? And if so, is the link observed empirically consistent with the theoretical effects implied by the carry trade hypothesis? To the best of our knowledge, the empirical study of this connection has been neglected in the literature, despite its importance in deciding upon effective monetary policy (Ames et al., 2017), or more generally in providing insights exploitable by market participants to adapt their hedging and investment strategies. Some exceptions are Omrane and Hafner (2015) and Omrane and Savaser (2017); however they focus on the link between FX volatility and news announcements, not IRD. The paper the closest in spirit with ours is that of Ichiue and Koyama (2011), which studies the intertwining between IRD, FX volatility, and the failure of UIP. They discover a significant link between IRD and FX volatility at a monthly horizon using a regime-switching linear regression approach. Our study goes one step further by studying the IRD-FX volatility connection at a higher frequency (i.e., daily) and in the context of an heteroscedastic time series model, whereas Ichiue and Koyama (2011) assume a constant (regime-specific) volatility. In doing so, we account for the confounding effects stemming from the persistence of high volatility episodes.

²By *FX volatility*, we refer to the volatility of variations in exchange rates.

³Bank of International Settlement, 2016

Our second contribution is methodological, and comprises detailing an extended stochastic volatility model (referred to as SVX model) to study the IRD-FX volatility relationship. Although the assumption of SV models for exchange rate returns is common (Bates, 1996; Chernov et al., 2018; Mahieu and Schotman, 1998) and was found superior to GARCH models for macro-financial variables (Chan and Grant, 2016; Clark and Ravazzolo, 2015), the specification of the volatility process is mainly restricted to random walks or autoregressive processes. As a consequence, we cannot use SV models to investigate the effect of IRD changes on the volatility process, in the idea of a GARCH-X model (Han and Kristensen, 2014). This is particularly damageable since, contrary to GARCH-X models, SV models account for the potential confounding effect of contemporaneous volatility shocks. In addition, SV models can be exploited in the broader framework of dynamic stochastic general equilibrium (DSGE) models to assess monetary policy interventions (Justiniano and Primiceri, 2008), or easily embedded in traditional continuous-time asset pricing models to solve asset allocation problems (Kim et al., 1998). Thus, to overcome this shortcoming, we introduce a further specification, in which the volatility term is extended to account for covariate effects, and which allows us to test for high-order dependencies between economic quantities. For estimation purposes, we detail a Gaussian mixture sampling approach relying on the linearization of the SV process and its reformulation into a linear state-space form with Gaussian error terms. The advantage of the latter is the straightforward way of incorporating the fixed regression coefficients into the state vector while simultaneously relying on a classical Bayesian estimation strategy. Consequently, we can also easily quantify estimation uncertainty and assess the significance of our findings. Since the properties of this approach are unknown in the context of a SV model, we conduct a realistic simulation study to stress the adequacy of the estimation procedure.

We use the proposed model to study the daily returns of six currencies vis-a-vis the US Dollar and find that changes in IRD have a significant impact on FX volatility for the majority of the considered currencies. To account for regime-specific effects of IRD, we discriminate between movements in IRD that imply a widening or a tightening of the differential. Our results reveal that the signs and the magnitudes of the impacts differ according to the IRD being negative or positive (with respect to one of the two currencies). If the interest rate in the economy of the funding currency (i.e. the currency with the lower interest rate) increases, FX volatility is found to increase significantly. For currency traders, the decrease in absolute IRD implies that the low-interest-rate currency becomes less attractive as a means for funding trades. Carry trade positions start to be

unwound, which impacts on FX volatility via alterations in capital flows. In contrast, if the interest rate in the low-interest-rate economy further decreases, FX volatility reacts negatively. Overall, the impact of the IRD becomes substantially stronger during and after the global financial crisis than in the pre-crisis period. This is particularly the case for the currencies whose interest rates differ most strongly from the US rate. The empirical results hold when we additionally control for the influence of past returns and the asymmetry of positive and negative shocks. For these two variables, we find significant evidence of an effect with the expected signs. Lagged returns impact negatively on volatility while negative shocks have a significant and positive influence.

With these results, we confirm empirically the link between IRD and FX volatility postulated by the theoretical framework of Gabaix and Maggiori (2015). They also deliver new insights for active investors (e.g. hedge funds) who trade at a high frequency on FX markets: when witnessing a decrease in absolute IRD, market participants should expect an increase (*ceteris paribus*) of short-term volatility, which can lead to the revision of hedging positions or volatility timing strategies. Finally, our results are useful to monetary authorities wishing to assess policy decisions targeting short-term interest rates: by narrowing the IRD between economies, short-term volatility of the corresponding currency pairs is more likely to increase, potentially leading to destabilizing market movements that must be accounted for.

In Section 2, we introduce our SVX model, outline its formulation into a state-space form, and detail the sampling steps for estimation. In Section 3, we conduct a simulation study to investigate the reliability of the proposed estimation approach. Section 4 contains the description of the data set. We report and discuss the empirical results in Section 5. We conclude in Section 6.

2 Empirical Model

In the following, we introduce the SVX model, its reformulation into a linear state-space model, and the estimation procedure. Subsequently, we discuss the role of IRD as determinants of FX volatility. Let E_t denote the daily nominal exchange rate. The mean adjusted exchange rate returns in percentage terms $r_t = 100 \times \log(E_t/E_{t-1})$ evolve

according to

$$r_t = \sigma_t \xi_t, \quad \xi_t \stackrel{iid}{\sim} N(0, 1), \quad (1)$$

$$\sigma_t^2 = \exp\{h_t + \mathbf{x}_{t-1}\boldsymbol{\beta}\}, \quad (2)$$

$$h_t = \mu_h + \phi(h_{t-1} - \mu_h) + \nu_t, \quad \nu_t \stackrel{iid}{\sim} N(0, \omega^2), \quad (3)$$

with $N(0, 1)$ denoting the standardized Gaussian distribution. Without loss of generality, we assume that the daily exchange rate data are characterized by an absence of dynamics in the mean. The error terms ξ_t and ν_t are assumed to be independent. The volatility of the returns is denoted σ_t , and its logarithmized version is assumed to change over time according to a stochastic process being the sum of two components. The first one, h_t , describes the logarithmized volatility, which follows a stationary autoregressive process with $|\phi| < 1$ and unconditional mean μ_h . The second term, $\mathbf{x}_{t-1}\boldsymbol{\beta}$, captures the impact of covariates on FX volatility. The specification in (2) takes into account the conclusion reached by, e.g., Ahmed and Valente (2015) that exchange rate volatility can be decomposed into distinct parts capturing either the persistence in the volatility or the effects of changes in macroeconomic conditions. Our focus centers on the impact of the (lagged) Libor rate for the US and the Libor rate for the country under consideration⁴. Moreover, as common in the empirical literature on asset returns, we additionally include past returns into the set of covariates. A detailed discussion on the considered explanatory variables follows in Subsection 2.2.

2.1 Formulation as a linear state-space model and estimation procedure

Our objective is now to obtain a reformulation of model (1) such that the Kalman filter and sampler can be used validly for estimation purposes. Indeed, our ability to use the Kalman filter is particularly important because this technique deals easily with the presence of covariates in addition to the time-series dynamics. A fundamental requirement of the Kalman filter is that the estimated model can be formulated as a linear state-space model with normally distributed error terms. However, in its current form, equations (1) and (2) are related in a non-linear fashion. Therefore, we need to

⁴Notice that, from January 2022 on, Libor will be discontinued by the Financial Conduct Authority and replaced by risk-free rate benchmarks such as the Secured Overnight Financing Rate (SOFR) or the European short term rate (ESFR) in cross-currency derivative contracts. These rates should be favored in future work.

linearize the stochastic volatility process to be able to estimate h_t and $\boldsymbol{\beta}$ with standard Gaussian mixture algorithms based on the Kalman filter.

To do so, we apply the procedure of Omori et al. (2007) and Kim et al. (1998) to our model. We start by squaring and taking the logarithm of (1):

$$\log \left[\left(\sqrt{\exp\{h_t + \mathbf{x}_{t-1}\boldsymbol{\beta}\}} \xi_t \right)^2 \right] = h_t + \mathbf{x}_{t-1}\boldsymbol{\beta} + \log \xi_t^2. \quad (4)$$

Defining $\kappa_t = \log \xi_t^2$, we approximate the non-linear stochastic volatility component by the following offset mixture model:

$$g_t = h_t + \mathbf{x}_{t-1}\boldsymbol{\beta} + \kappa_t, \quad (5)$$

with $g_t = \log(r_t^2 + c)$ and $c = 0.001$. The inclusion of a small positive constant in the definition of g_t enforces strict positivity and avoids numerical instabilities when r_t^2 is too close to 0, as suggested in Omori et al. (2007) and Kim et al. (1998).

From the normality assumption in (1), it follows that κ_t has a $\log\text{-}\chi_1^2$ distribution. As argued by Omori et al. (2007), it is of interest to approximate this distribution by a mixture of $K = 10$ Gaussian distributions for sampling efficiency at the estimation stage. The resulting density is therefore defined by:

$$f(\kappa_t) = \sum_{k=1}^K q_k f_N(\kappa_t | m_k, v_k^2) \quad \text{for } t = 1, \dots, T, \quad (6)$$

with f_N denoting a Gaussian density. The weights q_k , for $k = 1, \dots, K$, indicate the mixture probabilities. The means and variances of the K normal densities are given by the constants m_k and v_k^2 , respectively. The corresponding values for q_k , m_k and v_k^2 are taken from Omori et al. (2007), and were shown to ensure a good approximation of the $\log\text{-}\chi_1^2$ distribution⁵. The respective mixture component k for $k = 1, \dots, K$ is determined with the help of a stochastic indicator s_t for each observation. We sample the indicator s_t independently from a uniform distribution for each κ_t , such that $\kappa_t | (s_t = k) \sim N(m_k, v_k^2)$, where k indicates the respective elements of the ten-component mixture distribution. Thus, equation (5) can be rewritten as

$$g_t - m_{t s_t} = h_t + \mathbf{x}_{t-1}\boldsymbol{\beta} + \tilde{\kappa}_t, \quad (7)$$

⁵Notice that, contrary to Omori et al. (2007), Kim et al. (1998) explicitly define m_k as $m_k^* - 1.2704$. The term -1.2704 is equal to the expectation of a $\log\text{-}\chi_1^2$ distribution and ensures that (6) has the same expectation as a $\log\text{-}\chi_1^2$ distribution.

where $\tilde{\kappa}_t \sim N(0, v_{tk}^2)$. Using this reformulation of (1), we now have the capacity to estimate the corresponding regression coefficients with classical sampling procedures simply by including additional explanatory variables into the state vector of the Kalman filter. These results directly follow from the discussion in Durbin and Koopman (2012) on generalized state-space models. To the best of our knowledge, this approach has not yet been used to investigate the link between volatility and covariates.

Equation (7) represents the linear measurement equation with normally distributed errors for the Kalman filter. Throughout our analysis, we assume that $\beta = \beta_t = \beta_{t-1}$ at each time, i.e., that our additional explanatory variables (lagged IRD and past returns) have a time-invariant relationship with FX volatility. However, for technical reasons at the estimation stage, we cannot formally enforce a strict equality, but only that $\beta_t \approx \beta_{t-1}$ for each t . Therefore, to maintain consistency with the general notation in state-space models, we add in the following representation a time index t to the regression coefficient β but set the variance of β_t to a very small value⁶, such that the above approximation holds. Let the vector of unobserved states be denoted by $\alpha_t = (h_t \beta_t)'$. Combining equation (7) and (3), the respective state-space form for filtering α_t is obtained as:

$$y_t = Z_t \alpha_t + \tilde{\kappa}_t \quad \tilde{\kappa}_t \sim N(0, H_t) \quad (8)$$

$$\alpha_t = d + T_t \alpha_{t-1} + R_t \nu_t \quad \nu_t \sim N(0, Q_t), \quad (9)$$

where the matrices are defined according to:

$$y_t = g_t - m_{tk}, \quad Z_t = \begin{bmatrix} 1 & \mathbf{x}_{t-1} \end{bmatrix}, \quad \alpha_t = \begin{pmatrix} h_t \\ \beta_t \end{pmatrix}, \quad H_t = v_{tk}^2, \quad (10)$$

$$d = \begin{pmatrix} \mu_h \\ \mathbf{0} \end{pmatrix}, \quad T_t = \begin{bmatrix} \phi & \mathbf{0}' \\ \mathbf{0} & I_N \end{bmatrix}, \quad R_t = \begin{bmatrix} 1 & \mathbf{0}' \\ \mathbf{0} & I_N \end{bmatrix}, \quad Q_t = \begin{bmatrix} \omega^2 & \mathbf{0}' \\ \mathbf{0} & Q^* \end{bmatrix}. \quad (11)$$

The specification of T_t and Q_t ensures that $\beta_t \approx \beta_{t-1} = \beta$. In particular, the entries in the $N \times N$ diagonal matrix Q^* , denoting the variances of the regression coefficients, are set to $Q^* = 10^{-10} I_N$, i.e. very close to 0, such that the effects of the covariates are virtually constant over time. Finally, let α , \mathbf{y} , and X denote the vector and matrix of time series observations α_t , y_t , and \mathbf{x}_{t-1} , respectively. The joint posterior distribution is obtained by sequentially sampling from the conditional posterior distributions of the parameters with the following steps:

1. Sample the unobserved states α from $p(\alpha | \mathbf{y}, X, \mu_h, \phi, \omega^2)$ by means of Kalman

⁶We cannot set the variance of β_t exactly equal to zero.

filtering and smoothing.

2. Sample ω^2 from $p(\omega^2|\mathbf{h}, \phi, \mu_h)$.
3. Sample ϕ from $p(\phi|\mathbf{h}, \omega^2, \mu_h)$ by means of an independent chain Metropolis-Hastings (MH) step.
4. Sample μ_h from $p(\mu_h|\mathbf{h}, \phi, \omega^2)$.

For the first iteration, step 1 starts with a sequence for the SV process $\exp\{h_t^0 + \mathbf{x}_{t-1}\boldsymbol{\beta}_t^0\}\xi_t$ and $t = 1, \dots, T$. Furthermore, the regression coefficients $\boldsymbol{\beta}$ are all set to 0 to initialize the estimation. Subsequently, the sampler repeats steps 1 to 4 for 35,000 iterations where the first 10,000 are discarded as a burn-in sample. Final estimates and uncertainty quantification are obtained from these posterior draws.

We give a detailed version of the sampling steps, including the choices of the prior distributions, in Appendix B. Since this approach of including regression effects is non-standard in the context of SV models, we also conduct a simulation study in Section 3 to study the finite-sample properties of our estimation method.

2.2 Link between interest rates, carry trades, and FX volatility

Our econometric set-up relies on the hypothesis that speculative trading strategies build a transmission channel for the impact of interest rate changes on FX volatility. In this subsection, we discuss this idea more in detail. In addition, we define the other variables included in our vector \mathbf{x}_{t-1} of covariates.

A starting point is given by Ichiue and Koyama (2011), who explore the relationship between exchange rate returns, volatility, and IRD by means of regime-switching models. In particular, they extend the traditional UIP regression, such that the slope coefficient depends on the prevailing state of FX volatility. In contrast, we focus on the following direct mechanism: any variation in the interest rate in one economy leads to a change in the IRD between home and foreign economies. As a consequence, positions in carry trading strategies might be revised, leading to movements in the demand and supply of assets in different currencies. Ames et al. (2017), for instance, reveal that carry trading causes speculative in- and outflows of capital. Based on the assumption of imperfect financial markets, Gabaix and Maggiori (2015) demonstrate how such large-scale capital flows set exchange rates via alterations on financiers' balance sheets. The latter lead them to price currency risk differently, thus impacting on both the level and the volatility of exchange rates.

Recent contributions on carry trade activities categorize currencies into low-interest-rate, so-called “funding”, currencies and high-interest-rate, otherwise known as “investment” currencies (Ames et al., 2017). We compute the IRD as the difference between the short-term interest rate i^H in the home country and the respective US interest rate i^{US} . We argue that the consequences of a change in the interest rate in the home country for FX volatility differ according to i^H being the interest rate for the funding or the investment currency. An increase in the high-interest-rate currency does not change its role as an investment currency. Hence, the change in the IRD does not lead to a large unwinding of carry trade positions. In contrast, if the monetary authorities in the economy with the low interest rate start increasing interest rates, the IRD narrows and the currency becomes less attractive as a funding currency for the carry trades. Consequently, traders might start reversing their trading positions, which impacts on the volatility of the currency. Therefore, we assume that the sign and the magnitude of an effect of IRD depends on the prevalent “regime”, i.e. if the home currency is the funding or the investment currency. Consequently, we explicitly take into account the sign of the IRD by separating the effect into positive and negative IRD by means of the two variables

$$IRD_{t-1}^+ = \begin{cases} IRD_{t-1} & \text{if } i_{t-1}^H - i_{t-1}^{US} > 0 \\ 0 & \text{if } i_{t-1}^H - i_{t-1}^{US} \leq 0. \end{cases} \quad (12)$$

and

$$IRD_{t-1}^- = \begin{cases} IRD_{t-1} & \text{if } i_{t-1}^H - i_{t-1}^{US} < 0 \\ 0 & \text{if } i_{t-1}^H - i_{t-1}^{US} \geq 0. \end{cases} \quad (13)$$

Hence, for positive IRD_{t-1}^+ , the home currency is the investment currency. On the contrary, IRD_{t-1}^- captures the case when the home currency is the funding currency. We include both variables in our set of covariates \mathbf{x}_{t-1} .

From this model, we are able to deduce the signs to be expected for the effects of IRD_{t-1}^+ and IRD_{t-1}^- on FX volatility, according to the financial channel described above. For IRD_{t-1}^+ , we would expect negative variations to be linked with the unwinding of carry trade positions, leading to an increase in volatility. On the contrary, for IRD_{t-1}^- , positive variations are expected to be linked with the unwinding of the positions and an increase in volatility. Thus, we would expect opposite signs of the effects, namely a negative effect for IRD_{t-1}^+ and a positive effect for IRD_{t-1}^- reflecting the statement that smaller absolute IRD are associated with a higher FX volatility. Table 1 summarizes the signs of the different expected effects.

Funding ?	Channel	$i_{t-1}^H \nearrow$	$i_{t-1}^H \searrow$	$i_{t-1}^{US} \nearrow$	$i_{t-1}^{US} \searrow$
IRD_{t-1}^+	Unwind CT ? σ_t^2	No -	Yes +	Yes +	No -
IRD_{t-1}^-	Unwind CT ? σ_t^2	Yes +	No -	No -	Yes +

Table 1: Summary of the expected effects of a change in the home and the US rates on the unwinding of carry trades (CT) and FX volatility, given that the home currency is either the investment currency (IRD_{t-1}^+) or the funding currency (IRD_{t-1}^-).

In Section 5, we discuss the estimated changes in σ_t^2 associated with variations in IRD_{t-1}^+ and IRD_{t-1}^- , and interpret them in light of the carry trade hypothesis. However, channels other than carry trades could be responsible for the IRD - FX volatility connection. For example, Valchen (2020) considers the role played by fluctuations in bond convenience yield differentials as a compensation for a decrease in liquidity. Burnside et al. (2009) propose an explanation based on market microstructure with adverse selection of the market makers. Bacchetta and Van Wincoop (2010) highlight the role played by order flows and investors' heterogeneity. Precisely identifying the proper channel is a complex question beyond the scope of this paper.

Nevertheless, to motivate further our econometric set-up, we investigate whether our assumptions on the link between IRD and carry trade activity (CTA) are also justified from an empirical standpoint. Few empirical results exist beyond the cross-sectional analysis performed by Brunnermeier et al. (2009). Two exceptions are Anzuini and Fornari (2012) and Hutchison and Sushko (2013). Anzuini and Fornari (2012) find significant changes in CTA following an interest rate shock, confirming the results of Brunnermeier et al. (2009), while Hutchison and Sushko (2013) link CTA to risk reversal, noting the strong contemporaneous correlation of this measure with variations in interest rates. Although limited, empirical evidence is therefore in favor of a link between IRD and CTA.

To bring additional evidence to this question, we conduct a currency-by-currency regression analysis of the link between IRD and CTA for the six currencies studied in Section 5. As a proxy for CTA, we use the same quantity as Brunnermeier et al. (2009), Hutchison and Sushko (2013) and Anzuini and Fornari (2012), namely the net (long minus short) future position of non-commercial traders in the home currency (with respect

to the U.S. Dollar), expressed as a fraction of total open interests of non-commercial traders. Weekly⁷ data between 1999 and 2017 are obtained from the Commodity and Future Trading Commission (CFTC) website. Positive values are indicative of the home currency being the investment currency. As noted by Brunnermeier et al. (2009), this measure is not perfect, since over-the-counter trades are neglected; but the CFTC data are the best (if not the only) publicly available. To each CTA value, our associated predictor variable is the IRD (i.e., the difference in three-month Libor rates of the home currency and the U.S. Dollar, see Section 4 for more details) observed one week earlier. For each currency, we have between 964 and 991 data points. Our regression equation for each currency is the following:

$$CTA_t = \alpha_0^{(CTA)} + \beta^{(CTA)} IRD_{t-1} + \epsilon_t^{(CTA)}, \quad (14)$$

with CTA_t being the carry trade activity at time t , $\alpha_0^{(CTA)}$ and $\beta^{(CTA)}$ the regression parameters, and $\epsilon_t^{(CTA)}$ the error term. On Figure 1, we display the estimated regression functions obtained with OLS. For all currencies, we observe a positive relationship between IRD and CTA. It suggests that, when the home currency is the investment currency, CTA decreases when IRD moves closer to 0 (i.e. traders unwind their long positions in the home currency). Similarly, for the home currency being the funding currency, traders unwind their short positions in the home currency when IRD moves closer to 0. An additional kernel local linear regression estimate of the relationship (dashed red on the same figure) does not reveal important nonlinearities (two exceptions are for Euro, and for large IRD values of the Canadian Dollar). Using Newey-West robust standard errors of the estimated regression coefficients, we find the slope coefficients to be significantly different from 0 for all currencies except the Euro (detailed results, as well as a cointegration analysis, can be found in Table 13 in Appendix A). Repeating this analysis when splitting our samples into positive and negative IRD values or controlling for endogeneity with the instrument-free method of Park and Gupta (2012), our conclusions are qualitatively alike although less clear for the Euro and the Canadian Dollar.

To summarize these results, we reveal additional evidence of a link between IRD and CTA. In addition, we find signs of the estimated effects that are consistent with the carry trade channel outlined in Table 1: when absolute IRD narrows down, an unwinding of carry trades (i.e., a decrease in the absolute value of our CTA measure) is more

⁷Daily data are not publicly reported. We consider nevertheless that differences between weekly and daily processes are sufficiently small such that our results here can be extrapolated to daily dynamics.

likely to occur. This suggests that, if the IRD-FX volatility link outlined in Table 1 is empirically confirmed, the carry trade mechanism described above is a coherent channel. One alternative hypothesis, though, is that lagged IRD, CTA, and FX volatility are all exposed to similar macroeconomic shocks. In this respect, a potential alternative channel is suggested from the empirical findings of Hutchison and Sushko (2013) and the theoretical model of Bacchetta and Van Wincoop (2010): macro-news related to trade balance raises the perceived risk of a sharp appreciation of the currency under consideration, pushing some market participants to hedge their trades. Such trades are confused in the short run with signals about future value of the fundamentals, leading to a rise in volatility. Since our focus is on the link between IRD and FX volatility, we leave the clear identification of the channel to further research.

Finally, notice that to control for the potential influence of other factors, we broaden the set of covariates \mathbf{x}_{t-1} for the volatility term. Similar to common ARCH specifications for asset returns, we include the lagged return r_{t-1} as a third covariate. Moreover, for GARCH models, the empirical literature finds that shocks of equal magnitude but opposite signs often have different impacts on volatility. Therefore, to account for a potential asymmetry, we take an approach similar to that for threshold GARCH models: as a fourth explanatory variable, we incorporate the term $S_{t-1} r_{t-1}$ into \mathbf{x}_{t-1} where S_{t-1} denotes an indicator function, which takes the value 1 when $r_{t-1} < 0$ and 0 otherwise.

3 Simulation study

In this section, we conduct a simulation study to assess the reliability of the proposed estimation procedure, and by extension, of the results of our empirical analysis. The data-generating processes (DGP) are given by the empirical model detailed in Section 2. We consider two distinct DGP. In DGP1, we generate a return series r_t whose volatility only depends on a persistent part h_t without any influence of covariates. For DGP2, we consider the impact of three explanatory variables that resemble the macroeconomic covariates from our empirical study. For both DGPs we estimate a standard AR(1) specification and the SVX specification for the volatility term. Of particular interest in the simulation study are the two following aspects: i) the proper identification of zero and non-zero effects for β in the SVX model and ii) the comparison of parameter estimates and model fit for the SVX and the AR(1) specifications. We report posterior estimates for the parameters, as well as a summary on root mean square errors (RMSE) and coverage ratios for the regression parameters. Moreover, we provide details on the

computation of the test statistics used for model selection in Section 5.3.

3.1 Simulation set-up

The true parameter values for the DGPs are given in Table 2. For DGP1, we generate the return series excluding covariate effects. The parameters μ_h , ϕ_h , and ω^2 are set to -0.1 , 0.90 , and 0.04 , respectively. These values are in line with the estimates obtained in our empirical analysis of six currencies. For DGP2, we leave these values unchanged but we additionally assume the impact of three covariates with the following regression effects: $\beta_1 = 0.7$, $\beta_2 = -0.3$, and $\beta_3 = 0$. Similarly to the other parameters, these values are chosen close to those obtained in the empirical study. In addition, we simulate explanatory variables reflecting the characteristics of the IRD and lagged-return time series. We draw \mathbf{x}_1 from a Gaussian distribution but set it to 0 for a number of successive time periods, as has been observed for several IRD time series after the crisis. The variable \mathbf{x}_2 contains the one-period lagged return time series. Finally, the third covariate is specified as a Gaussian random variable whose impact is assumed to be 0. For estimation, we choose the same parameters of the prior distributions as in the empirical study (see Table 15 in Appendix B). Overall, the choices of these DGPs are intended to match the features of the financial time series used throughout this paper.

Table 2: True parameter values for DGP1 and DGP2

Parameter	μ_h	ϕ	ω^2	β_1	β_2	β_3
DGP 1	-0.1	0.90	0.04	-	-	-
DGP 2	-0.1	0.90	0.04	0.7	-0.3	0

3.2 Simulation results

In the following, we sum up the results of the simulation study. First, we focus on the proper estimation of the regression coefficients β . Second, we report the estimates and model diagnostics for the SVX model for three sample sizes $T \in \{1000, 2500, 5000\}$. Last, we compare the results from estimating both an AR(1) and a SVX model for DGP1 and DGP2 for 200 samples, respectively.

The histograms in Figure 2 display the estimates of the posterior mean for every regression coefficient β_1 , β_2 , and β_3 in 200 samples. The distributions of the estimates are

unimodal with most probability mass at the true parameter value, suggesting a good estimation. For $\beta_3 = 0$, posterior mean estimates vary slightly more.

As a next indicator, we consider the coverage ratios of the true parameter value and the highest posterior density interval (HPDI) of the respective estimates. The entries in Table 3 denote the proportion of samples for which the true parameter is covered by the 90% HPDI in 200 samples. Note that for DGP1, the true values for β are equal to zero while for DGP2, it holds that $\beta = (0.7 \ -0.3 \ 0)'$. For instance, the first entry of 0.945 implies that out of 200 estimated samples from DGP1, in 189 cases the true parameter value of $\beta_1 = 0$ is within the bounds of the estimated HPDI. Analogously, for DGP2, the entry of 0.960 indicates that the true parameter value of $\beta_1 = 0.7$ is included in the HPDI in 192 samples. For the regression coefficients, coverage rates are high, suggesting very good coverage of the true parameters in the HPDI. Overall, we conjecture that zero regression effects are correctly identified with high probability.

Table 3: Coverage rates for parameters

Parameter	β_1	β_2	β_3	μ	ϕ	ω^2
DGP1	0.945	0.925	0.925	1	1	0
DGP2	0.960	0.935	0.925	1	1	0

Proportions of samples for which the true parameter lies within the 90% HPDI across the 200 samples.

After examining the characteristics in repeated samples, we focus next on estimates for different sample sizes for DGP2. Table 4 contains estimates of the posterior distributions for the sample sizes $T \in \{1000, 2500, 5000\}$. For each parameter, we report the estimate of the posterior mean, the posterior standard deviation, and the respective 5% and 95% quantiles of the posterior distribution (HPDI). Moreover, we indicate the RMSE for increasing sample size. In general, the posterior mean estimates are close to the true value, with the HPDI narrowing with increasing sample size T . Similarly, we can observe decreasing posterior standard deviations and RMSEs for all parameters with increasing sample sizes. In particular, for the regression coefficients, RMSEs are small for a sample size of $T = 1000$ and shrink to zero for $T = 5000$.

In Figure 3, we display the trajectories for the estimated volatility $\exp\{h_t\}$ for $T = \{1000, 2500, 5000\}$. Additionally, we include the 90% HPDI around the posterior mean to grasp the variation in the posterior estimates. The estimated series captures the true process well and 90% of the values in the posterior distribution lie in a very narrow band

Table 4: Posterior estimates for SVX model with different sample sizes

Parameter	β_1	β_2	β_3	μ	ϕ	ω^2
$T = 1000$						
Mean	0.69	-0.23	0.02	-0.05	0.84	0.22
Stddev	0.062	0.052	0.055	0.466	0.108	0.366
HPDI	[0.58 0.79]	[-0.32 -0.15]	[-0.07 0.11]	[-0.56 0.40]	[0.62 0.98]	[0.08 0.6]
RMSE	0.0005	0.0022	0.0005	0.0014	0.0019	0.0056
$T = 2500$						
Mean	0.63	-0.29	-0.01	-0.04	0.78	0.18
Stddev	0.044	0.034	0.033	0.341	0.136	0.355
HPDI	[0.56 0.70]	[-0.34 -0.23]	[-0.07 0.04]	[-0.30 0.20]	[0.49 0.96]	[0.07 0.46]
RMSE	0.0014	0.0003	0.0002	0.0012	0.0025	0.0028
$T = 5000$						
Mean	0.70	-0.30	-0.02	-0.02	0.82	0.11
Stddev	0.029	0.020	0.023	0.224	0.075	0.082
HPDI	[0.65 0.75]	[-0.33 -0.27]	[-0.06 0.02]	[-0.17 0.13]	[0.69 0.94]	[0.07 0.22]
RMSE	0.0000	0.0000	0.0003	0.0012	0.0011	0.0010

The first row of each panel contains estimates of the posterior mean across 5000 Gibbs iterations. In the second row we report the parameter's posterior standard deviation (Stddev), followed by the 5% and 95 % quantiles of the posterior distribution in brackets in the third row. Finally, the fourth row contains the root mean squared error (RMSE) associated with the different parameters.

around the true process.

We conclude by examining the consequences of estimating the respective inappropriate model for DGP1 and DGP2. To this end, we simulate data according to DGP1 and estimate it first with the correct model (AR(1)). Subsequently, we estimate it with the “misspecified” SVX model that incorporates irrelevant covariates. We proceed similarly for DGP2, whereby this time the SVX model with covariates is the correct one, and the AR(1) model is misspecified. In Table 5, we report RMSE for the estimates of \mathbf{h} ($RMSE_h$). Moreover, we provide the Ljung-Box statistics $Q(\ell)$ and $Q^2(\ell)$ for a lag length of $\ell \approx \log(T) = 7$ for the standardized (squared) residuals as model diagnostics. Neglecting regression effects yields a higher RMSE for the log-volatility \mathbf{h} . On the contrary, estimating the SVX model including irrelevant covariates does not affect the quality of the estimation results, since effects are properly diagnosed as being 0.

Table 5: RMSE for h and Ljung-Box statistics

True DGP	Model	$RMSE_h$	Rejection $Q(7)$	Rejection $Q^2(7)$
DGP 1	AR(1)	0.0120	0.005	0.110
	SVX	0.0120	0.055	0.265
DGP 2	AR(1)	0.0147	0.065	0.705
	SVX	0.0120	0.060	0.280

Average RMSE for the time series of h_t , in 200 samples. The 4th and 5th columns contain the average rate of rejecting the null hypothesis of no autocorrelation in the standardized (resp. squared) residuals with $\alpha = 0.05$.

For the Ljung-Box statistic $Q(7)$, the average rate of rejecting the null hypothesis is close to the theoretical significance level of 0.05. For the squared residuals however, the empirical average rejection rate of the null hypothesis is substantially higher. The test might be oversized for the SVX models, i.e., the probability of rejecting the null falsely is higher than the theoretical level. To obtain more appropriate critical values for testing in the empirical study, we compute the empirical quantiles of the Ljung-Box Statistics in the Monte Carlo simulations. In particular, we obtain as 0.95-quantiles of the empirical distributions of simulated $Q^2(7)$ statistics the values 20.8807 and 19.3094 for the AR(1) and SVX models, respectively.

In Figure 4, we report the RMSEs in 200 samples for each parameter for the two DGP estimated in both model specifications. The first row captures the RMSEs for an estimation of the AR(1) specification. Neglecting the impact of the regressors leads to

higher RMSEs for μ , ϕ_h and ω^2 as shown in boxplot (b). The RMSEs in boxplots (c) and (d) are very similar. Hence, misleadingly estimating a SVX model for DGP1 does not lead to greater RMSEs.

The boxplots in Figure 5 display the estimates of the posterior means for all parameters in the 200 samples and a sample size of $T = 1000$. Overall, the posterior mean estimates do not vary substantially across the samples and are close to the true parameter values. In boxplot (b), we observe more variation in the posterior estimates for the AR(1) model. In contrast, a misspecified SVX model yields estimated regression coefficients close to zero, i.e., the true value (boxplot (d)). Hence, irrelevant covariates are detected to have zero impact on the return volatility. For DGP2, estimates are close to the true parameter values (boxplot (c)).

4 Data

We analyze daily FX data against the US Dollar (USD) for six currencies, namely the Australian (AUD) and Canadian Dollar (CAD), the Swiss Franc (CHF), the Euro (EUR), the Japanese Yen (JPY), and the British Pound (GBP), respectively. The exchange rates are expressed as the value of one unit of the national currency in USD. In this study, we confine the analysis to industrialized economies with floating exchange rates. Moreover, the independence of monetary authorities in these economies has the advantage of not facing major endogeneity problems in the course of estimation. All data are available on the website of the Federal Reserve of St. Louis. The sample period starts on January 4th, 1999 and ends on December 29th, 2017. After deleting time periods for which the spot exchange rates are missing, the sample includes 4,773 observations. To compute interest rate differentials, we use daily three-month Libor rates for CHF, EUR, GBP and JPY. For the analysis on CAD and AUD, respectively, we employ the three-month Euro Libor rates for the two countries, provided by the Financial Times and collected through Datastream.

Figure 6 contains the trajectories of the three-month Libor for the complete sample period for the CHF, EUR, GBP, USD, and JPY. Except for the JPY, interest rates vary substantially until the beginning of 2009. In the middle of 2009, they drop to zero for all currencies and move substantially less afterwards. These observations allow us to presume the existence of structural changes consecutive to the financial crisis. To account for this effect, we conduct analyses on two subperiods. In line with the NBER recession statistics, we define the pre-crisis period to run from 01/04/1999 to

06/30/2009 and the post-crisis period from 07/01/2009 to 12/29/2017. Table 6 contains descriptive statistics for all exchange rate returns, displaying substantial skewness and excess kurtosis for all returns.

Table 6: Descriptive statistics for exchange rate returns

	AUD	CAD	GBP	EUR	CHF	JPY
Mean	0.0047	-0.0040	0.0043	-0.0003	-0.0071	0.0001
Standard deviation	0.7984	0.5692	0.5961	0.6259	0.7014	0.6591
Skewness	-0.7896	-0.0904	0.7633	-0.1098	-1.0997	-0.2815
Kurtosis	15.1902	8.3064	14.7346	5.1875	35.0091	6.7895

Empirical mean, standard deviation, skewness, and kurtosis obtained from the daily exchange rate returns in percentage terms (r_t) for the complete sample period ranging from January 4, 1999 to December 29, 2017.

5 Empirical results

In this section, we interpret the estimation results for model (1) for the six currency pairs and the two subsamples.

We start with a summary of important diagnostics for the Markov Chain Monte Carlo (MCMC) sampling. Subsequently, we interpret the economic results and display the estimated trajectories for the unobserved volatility processes. Finally, we compare our SVX specification with a simpler alternative, and conduct several robustness checks.

5.1 MCMC diagnostics

After thinning the MCMC iterations, we compute posterior estimates based on 5,000 MCMC draws. We check for convergence of the Markov chains of the parameters by means of the Geweke convergence diagnostics (Geweke, 1992). In Table 7, we report the absolute value of the convergence diagnostic (CD) for each parameter in the two subsamples. The distribution of the test statistics under the null is asymptotically standard normal. Convergence of the chain is achieved if CD is less than 1.96 in absolute value for all parameters. We conclude that the posterior estimates rely on parameter draws that stem from the stationary Markov distribution. For the time-varying parameter h_t , we report the overall maximum convergence diagnostic, that is, the maximum out of the CD for each point in time. For a limited number of points in time, the critical value is exceeded, which we consider to be only a minor problem. Furthermore, the efficiency

Table 7: Geweke convergence diagnostics CD

Parameter	IRD_{t-1}^+	IRD_{t-1}^-	r_{t-1}	$S_{t-1}r_{t-1}$	h_t	μ_h	ϕ	ω^2
Subsample	01/04/1999 – 06/30/2009							
Currency								
AUD	0.28	0.27	0.28	0.43	2.98	1.10	0.80	0.93
CAD	1.40	1.42	0.10	0.03	2.05	0.10	0.37	0.28
GBP	0.41	0.93	0.41	0.21	2.56	1.08	0.95	0.35
EUR	0.43	0.27	0.47	0.33	1.13	0.72	0.35	0.71
CHF	–	0.94	1.10	0.72	2.37	0.57	1.20	1.47
JPY	–	1.32	0.05	0.14	2.11	0.24	1.28	1.70
Subsample	07/01/2009 – 12/29/2017							
Currency								
AUD	0.53	–	0.99	0.67	2.86	0.53	1.13	1.26
CAD	0.74	0.47	0.47	0.71	2.58	0.70	0.48	0.63
GBP	0.73	0.42	0.59	0.72	2.06	1.00	0.08	0.05
EUR	0.24	0.66	0.05	0.08	1.52	0.10	0.38	0.36
CHF	–	0.12	0.08	0.56	2.55	1.37	1.01	0.55
JPY	–	0.82	0.22	0.68	1.98	0.01	0.44	0.49

Table entries denote the absolute values of the Geweke (1992) convergence diagnostic CD for each parameter of the model. IRD_{t-1}^+ , IRD_{t-1}^- , r_{t-1} and $S_{t-1}r_{t-1}$ refer to the regression parameters associated with positive IRD, negative IRD, lagged return, and lagged negative shocks, respectively. For the time-varying state h_t , CD denotes the maximum absolute value of CD for each point in time. The hyphens indicate that the parameter was not estimated for the corresponding currency due to a lack of observations.

Table 8: Acceptance rates of $\hat{\phi}$ in MH-step

Acceptance rate	AUD	CAD	GBP	EUR	CHF	JPY
Subsample 1	0.9590	0.9592	0.9591	0.9588	0.9668	0.9580
Subsample 2	0.9520	0.9634	0.9522	0.9608	0.9604	0.9572

Table entries denote the acceptance rates of ϕ , computed as the number of accepted $\hat{\phi}$ from the proposal density divided by the number of Gibbs iterations after the burn-in. Subsample 1 covers the period 01/04/1999 – 06/30/2009, whereas Subsample 2 refers to the period 07/01/2009 – 12/29/2017.

of the sampler is of interest. The acceptance rate of the MH-step provides information on how often a new draw $\hat{\phi}$ from the proposal distribution is accepted. Low rates of acceptance imply that many candidate draws are rejected, and the chain remains stuck at a particular point for repeated iterations. Table 8 contains the acceptance rates. The high numbers indicate that the proposal draw $\hat{\phi}$ is not accepted only in a small number of iterations. An additional measure of algorithm reliability is given by the numerical standard errors of the posterior estimates. The numerical standard errors for the regression coefficients are low and range between 0.002 - 0.02, which suggests a reasonably good approximation of the posterior mean.

5.2 Impact of interest rate differentials on FX volatility

In the following section, we discuss the results of our empirical analysis. In Table 9, we provide a summary of these results by presenting the estimated regression coefficients associated with the variables supposed to affect FX volatility⁸.

First, we focus on the results for IRD_{t-1}^+ , which captures the effect of variations in IRD if $i_{t-1}^H - i_{t-1}^{US} > 0$. We find that an increase in IRD_{t-1}^+ implies a widening in the interest rate differentials originating from either a rise in i_{t-1}^H or a decrease in i_{t-1}^{US} . However, the crucial aspect is that the domestic currency remains the investment currency and the US Dollar remains the funding currency. Although the increase in the differential might attract additional traders, there is no incentive to unwind trading positions. Indeed, we find that if the interest rate differential widens, volatility tends to decrease. However, this effect is not significant. The 90% HPDI, indicated in the brackets below

⁸In the course of the analysis, we also examined specifications incorporating only past returns and the indicator of negative shocks. The coefficient estimates are similar across all the specifications considered. Moreover, the results for the model parameters μ_h, ϕ , and ω^2 are not affected. These estimated volatility parameters are available upon demand.

the posterior means, contain the value zero for all currencies. For the post-crisis period, in contrast, a widening in IRD_{t-1}^+ impacts more strongly on the volatility of currencies. For the CAD and GBP, the effect becomes significant and its magnitude increases substantially in comparison to the pre-crisis period.

The variable IRD_{t-1}^- contains the IRD for $i_{t-1}^H - i_{t-1}^{US} < 0$. In this regime, the home currency is the funding currency, and the US Dollar is the investment currency. Since IRD_{t-1}^- is negative, an increase in IRD_{t-1}^- implies a narrowing of the rate differential. This might be caused by either an increase in i^H or a decrease in i^{US} . For carry traders, an increase in i^H implies that the funding currency becomes less attractive for borrowing because of the higher interest rates. Vice versa, a decrease in i^{US} reduces the profit that can be made from lending in the US Dollar. As a consequence, speculative carry traders may start alternating trading positions, which leads to sudden shifts in the demand and supply for the foreign currency. Hence, we expect that an increase in IRD_{t-1}^- leads to a rise in exchange rate volatility. In the pre-crisis period, we find a significant positive effect for the CAD, the EUR, and the CHF. The magnitude of the effect increases in the post-crisis period for all currencies. Again, it is the volatilities of the CAD and the GBP that are particularly strongly influenced by changes in the currencies' Libor rates. The magnitude of the effect, however, also becomes larger for the EUR and the CHF. This finding is in line with the results of Omrane and Savaser (2017), who demonstrate that the reaction of currency volatility to news related to interest rates is larger during crisis times, and brings empirical confirmation of the theoretical framework proposed by Gabaix and Maggiori (2015). In addition, it shows that the IRD-FX volatility connection detected by Ichiue and Koyama (2011) at a monthly frequency holds also at the daily frequency. This result has thus several implications for active market participants (like hedge funds) trading frequently in sophisticated products, and market makers concerned by volatility hedging: information can be extracted from short-term IRD to better anticipate changes in volatility. These anticipations can be exploited to revise hedging position or volatility-timing investment strategies (Fleming et al., 2001).

Furthermore, rows 3 and 4 in Table 9 contain the results for the two additional covariates. The regression coefficients for the lagged return r_{t-1} and for the asymmetric effect of negative shocks $S_{t-1}r_{t-1}$ are consistent and significant across currencies and time periods. A lagged appreciation of the home currency tends to decrease the volatility, which is line with results from GARCH specifications for exchange rate returns. In contrast, negative shocks (i.e. a depreciation of the home currency) raise FX volatility substantially.

Table 9: Posterior distribution estimates and HPDI of the regression coefficients for the two subsamples

Return volatility	AUD	CAD	GBP	EUR	CHF	JPY
Subsample	01/04/1999 – 06/30/2009					
IRD_{t-1}^+	Mean HPDI	-0.04 [-0.20 0.11]	-0.43 [-0.85 0.09]	-0.15 [-0.35 0.11]	-0.17 [-0.44 0.15]	-
IRD_{t-1}^-	Mean HPDI	0.17 [-1.40 1.25]	0.68 [0.01 1.16]	0.45 [-1.02 1.52]	0.34 [0.09 0.54]	0.20 [0.07 0.28]
r_{t-1}	Mean HPDI	-0.43 [-0.59 -0.28]	-0.27 [-0.57 -0.01]	-0.32 [-0.54 -0.12]	-0.54 [-0.74 -0.35]	-0.48 [-0.65 -0.31]
$S_{t-1} r_{t-1}$	Mean HPDI	0.67 [0.41 0.96]	0.79 [0.37 1.32]	0.79 [0.41 1.20]	1.18 [0.86 1.54]	0.93 [0.65 1.23]
Subsample	07/01/2009 – 12/29/2017					
IRD_{t-1}^+	Mean HPDI	-0.03 [-0.17 0.12]	-1.20 [-1.59 -0.53]	-1.99 [-3.01 -0.53]	0.19 [-0.48 0.98]	-
IRD_{t-1}^-	Mean HPDI	-	2.44 [0.88 3.64]	1.04 [0.51 1.32]	0.56 [0.08 0.98]	0.57 [0.28 0.92]
r_{t-1}	Mean HPDI	-0.57 [-0.82 -0.35]	-0.36 [-0.57 -0.01]	-0.23 [-0.43 -0.06]	-0.55 [-0.80 -0.31]	-0.41 [-0.61 -0.22]
$S_{t-1} r_{t-1}$	Mean HPDI	0.88 [0.51 1.32]	1.01 [0.37 1.32]	0.79 [0.41 1.22]	1.16 [0.73 1.64]	0.49 [0.23 0.77]

Posterior mean estimates for positive and negative IRDs, lagged returns, and the negative lagged return, over the two different time periods. In brackets, we report the 90% HPDI defined by the 5th and 95th percentiles of the posterior distribution. Bold face entries denote estimates for which the corresponding 90% HPDI does not include 0. The hyphens for CHF, JPY, and AUD indicate that the corresponding variable is excluded because values are equal to zero throughout the respective sample period.

Finally, we compare the smoothed volatility estimates for $\exp\{h_t\}$ from the AR(1) process and the extended volatility specification. In Figure 7, we display the posterior mean and the 5% and 95% percentiles of $\exp\{h_t\}$ from the AR(1) specification and the extended model. The posterior mean of the latter is covered within the HPDI of the AR(1) specification for the majority of observations. However, the posterior mean estimates deviate during times of economic turmoil and high volatility. For the CHF, estimates differ around the time of the drastic drop in Swiss Libor rates in 2015. For the EUR and the JPY, the differences are more clear cut at the end of the second sample, when the IRD with i^{US} began to rise again. For the GBP, deviations are more pronounced during and in the aftermath of the global financial crisis. In summary, Figure 7 illustrates that the autoregressive component h_t tends to be underestimated in the simple specification, especially during periods of high FX volatility.

Overall, these results are in line with the expected effects implied by the carry trade hypothesis, and summarized in Table 1: a decrease in IRD_{t-1}^+ or an increase in IRD_{t-1}^- are associated with an increase in FX volatility, following an unwinding of the corresponding carry trade positions.

5.3 Model comparison and goodness-of-fit

The lack of a closed form for the likelihood function in SV models complicates the computation of likelihood-based model criteria to compare alternative specifications. The high number of time-varying parameters in the SV model also leads to the likelihood function being a high dimensional integral, which has to be approximated by means of simulation. For instance, alternative estimation methods for SV models such as efficient importance sampling (Liesenfeld and Richard, 2006) simulate an approximated likelihood function in the course of estimation. Our proposed extension of the SV model, in contrast, is analyzed by means of Gaussian mixture sampling which does not incorporate the evaluation of a likelihood function. Thus, to assess the fit of the model, we adopt a different point of view, as suggested in Bauwens et al. (2012). We compare the different volatility specifications with respect to the features they are supposed to be capturing. In particular, we concentrate on the ability to encompass the heteroscedasticity in the exchange rate data.

As a first means of evaluating the model, we compare the Ljung-Box-statistics $Q^2(7)$ to test the null of no autocorrelation in the squared standardized residuals from each model. Following Tsay (2005), we set the lag length to compute the test statistics to $\ell \approx \log(T)$. The preferred model is the one that enables the greater “whitening” of the residuals.

In Tables 10 and 11, we report residual diagnostics for subsamples 1 and 2 separately. We limit the comparison of results to estimates from a simple AR(1) specification for the volatility term and the specification including the covariate effects. Results differ across currencies and for the pre- and post-crisis subsamples. Although we reject the null hypothesis of no autocorrelation, including the covariates decreases the degree of autocorrelation in the squared residuals in the pre-crisis period. In particular for the AUD, the GBP, and the JPY, the Ljung-Box statistics are in favor of taking the IRD into account. Additionally, we report statistics for skewness and excess kurtosis. In the pre-crisis sample, the inclusion of IRD leads to a decrease in skewness for the CAD and the JPY. Moreover, after the crisis, the skewness of the residuals decreases nearly for all exchange rate returns. The close relationship of IRD and the skewness of exchange rates via currency carry trades has been emphasized, among others, by Brunnermeier et al. (2009).

Table 10: Residual diagnostics for subsample 1

FX volatility	AUD	CAD	GBP	EUR	CHF	JPY
Model	AR(1) without regression effects (SV)					
$Q^2(7)$	55.801	23.261	31.778	46.519	31.343	18.376
p -value	.000	.002	.000	.000	.000	.010
skewness	-0.1414	-0.0400	0.0625	-0.0317	-0.096	-0.138
kurtosis	2.5835	2.546	2.608	2.7038	2.7974	2.870
Model	Including regression effects (SVX)					
$Q^2(7)$	29.301	23.984	30.709	18.527	24.294	49.073
p -value	.000	.000	.000	.010	.000	.000
skewness	-0.1299	-0.0436	0.0560	-0.0273	-0.1236	-0.046
kurtosis	2.562	2.670	2.561	2.635	2.697	3.231

Residual diagnostics for the SV model with an AR(1) specification (upper panel) and the SVX specification (lower panel), obtained over the period 01/04/1999 – 06/30/2009. Table entries denote the test statistics and the corresponding p -values for the Ljung-Box test on the squared standardized residuals ($Q^2(7)$).

As an additional means to compare the in-sample fit, we compute the RMSE between the estimated variance term $\hat{\sigma}_t^2$ and the true variance σ_t^2 of the exchange rate returns. Thereby, we approximate the true variance by means of the most simple measure of realized volatility, the demeaned squared returns r_t^2 . The results in Table 12 demonstrate that the model with regression effects yields lower RMSEs than the AR(1) specification across all currencies and subperiods. Moreover, we conduct a Diebold-Mariano test for

Table 11: Residual diagnostics for subsample 2

FX volatility	AUD	CAD	GBP	EUR	CHF	JPY
Model	AR(1) without regression effects (SV)					
$Q^2(7)$	27.913	19.027	20.432	17.218	13.651	11.532
p -value	.000	.008	.005	.016	.058	.117
skewness	-0.0475	-0.0006	0.0928	0.0292	-0.0670	0.026
kurtosis	2.4939	2.7805	2.8461	2.845	3.027	3.018
Model	Including regression effects (SVX)					
$Q^2(7)$	8.807	20.013	13.529	29.739	17.457	7.069
p -value	.267	.006	.06	.000	.015	.422
skewness	-0.032	-0.053	0.070	-0.010	-0.045	0.019
kurtosis	2.553	2.744	2.607	2.677	2.843	3.039

Residual diagnostics for the SV model with an AR(1) specification (upper panel) and the SVX specification (lower panel), obtained over the period 07/01/2009 – 12/29/2017. Table entries denote the test statistics and the corresponding p -values for the Ljung-Box test on the squared standardized residuals ($Q^2(7)$).

equality of forecast accuracy, comparing the errors from the AR(1) and the extended specification for the volatility term. The errors are computed based on the filtered estimates of volatility from the Kalman filter. Hence, they may be interpreted in a similar fashion to pseudo-out-of-sample forecasts for volatility with an expanding estimation window. For the majority of currencies, we reject the null hypothesis of equality in forecast accuracy. Finally, we also conduct the conditional predictive ability test of Giacomini and White (2006) to compare the two forecasting methods. We similarly reject the null hypothesis of equal predictive ability in favor of our extended volatility specification for most currencies (Table 12).

5.4 Robustness checks

To assess the robustness of our results with respect to minor changes in the specification of the model, we re-estimate its parameters, adding various control variables likely to influence currency volatility. First, we include the GVIX-7 currency volatility index of JP Morgan in our set of predictors, to control for market-wide volatility effects. Then, we use (lagged) local stock market returns, computed from national stock indices. For each currency, we use a national reference index, namely the ASX, STOXX600, SSMI, FTSE100, GSPTSE, and NIKKEI225 indices. Finally, we control for a potential

Table 12: RMSE, Diebold-Mariano and Giacomini-White test statistics

		AUD	CAD	GBP	EUR	CHF	JPY
Subsample		01/04/1999 – 06/30/2009					
RMSE	AR (1)	2.3087	1.0602	1.0139	0.8432	0.9140	1.0280
	SVX	2.2960	0.8612	0.9239	0.7537	0.8586	0.9736
DM		0.08	6.29***	4.32***	9.16***	5.38***	7.33***
GW		0.99	52.24***	30.78***	121.86***	32.98***	40.20***
Subsample		07/01/2009 – 12/29/2017					
RMSE	AR(1)	0.8952	0.6866	1.3995	0.6923	3.6753	0.8531
	SVX	0.8492	0.5019	1.3306	0.5883	3.6129	0.8015
DM		4.11***	17.97***	1.59*	22.15***	0.99	8.70***
GW		27.85***	277.71***	3.18	223.73***	2.57	108.72***

Table entries denote the RMSE, computed as $RMSE_j = \frac{1}{T} \sum_{t=1}^T (\sigma_t^2 - \hat{\sigma}_{t,j}^2)^2$, where $j \in \{AR(1), SVX\}$ (AR(1) refers here to a SV model with a simple AR(1) specification). DM denotes the test statistic of equal forecast accuracy of Diebold and Mariano (1995), while GW denotes the test statistics for the predictive ability test of Giacomini and White (2006). *** (resp. *) indicates tests significant at the 1% (resp. 10%) test level. Upper panel: pre- and in-crisis period. Lower panel: post-crisis period.

day-of-the-week effect, since we would expect volatility to be higher on days of low trading volumes characterized by a low liquidity (typically Mondays and Fridays, see, e.g., Chordia et al., 2001). We thus include a dummy variable taking the value of 1 if the trading day is Tuesday, Wednesday, or Thursday, and 0 otherwise.

The results of these extended regressions are displayed in Tables 16 to 21 in Appendix C. We first include the additional control variables one at a time. For these regressions, the signs, magnitude, and significant characters of IRD effects are almost unchanged for all currencies. A similar comment applies to the parameters associated with r_{t-1} and $S_{t-1}r_{t-1}$. One notable exception is the specification including the GVIX for the EUR, where IRD_{t-1}^+ sees a significant increase in its regression coefficient for the post-crisis period. However, GVIX and IRD_{t-1}^+ are strongly and positively serially correlated (with a serial correlation around 60%) and the coefficient associated with the GVIX is significantly negative. It suggests that, in periods of positive IRD, EUR volatility decreases if global currency market volatility is large, a result to connect with the weaker evidence of a significant IRD-CTA connection for the EUR in Section 2.2.

In most cases though, the effects of the GVIX and of stock market returns are not found to be different from 0 (with the aforementioned exception). On the contrary, a day-of-the-week effect is found to be significant and negative in all currencies except the AUD.

It implies a higher volatility on days characterized by Chordia et al. (2001) as exhibiting low liquidity (i.e., Mondays and Fridays). Nevertheless, our estimates of the IRD effects are not affected in these specifications.

Then, we include the three control variables simultaneously. The sizes of the credible intervals tend to increase, but estimated signs and magnitudes of the effects stay globally unchanged, confirming the robustness of our findings.

6 Conclusions

We study the link between the volatility of exchange rate returns and past interest rate differentials. To do so, we first detail an econometric approach based on an extended stochastic volatility model, in which explanatory variables can easily enter the volatility equation. We demonstrate how the corresponding regression coefficients can be straightforwardly incorporated into the state vector of a classical Kalman filter, and estimated with a Gaussian mixture sampling approach. In a simulation study, we show that regression coefficients and model parameters are properly estimated with this strategy. Empirically, we detect that changes in the IRD between economies is associated with variations in FX volatility of most currencies of developed countries. Consistent with previous findings, we also find FX volatility to be highly persistent over time.

Currency carry trading activities are proposed as the respective transmission channel. Indeed, these investment strategies exploit the difference in interest rates to borrow in the low-interest-rate currency and invest in the high-interest-rate currency. Thus, a change in the IRD might cause the sudden unwinding of trading positions and, thereby, impact on FX volatility. We find that the influence of IRD becomes stronger in the aftermath of the global financial crisis. Controlling for the sign of the differential, we reveal that the effect of an interest rate hike differs for low and high-interest-rate currencies. For monetary authorities in the low-interest-rate economy, our results imply that *ceteris paribus* (c.p) raising the domestic interest rate is accompanied by an increase in FX volatility. The same effect would be documented by a c.p. decrease in the high interest rate. These results empirically confirm the theoretical framework of Gabaix and Maggiori (2015). These findings also have a number of implications for market participants: if, say, a hedge fund anticipates a decrease in absolute IRD, it must also anticipate an increase in (short-term) FX volatility that can have important implications for directional bets. Notice that we assume IRD to be related to FX volatility via movements of carry trading positions. Considering a more direct measure of the currency trading volumes might

therefore yield a more precise analysis of the relationship between currency trading and FX volatility. However, these order-flow data are not easily publicly available at a daily frequency. Finally, this work provides tools to investigate the determinants of currency volatility dynamics in emerging market economies, in which particularly high interest rates have attracted large amounts of capital inflow over the last decade. Sudden unwinding of carry trade positions and, subsequently, rising FX volatility might have severe consequences for the macro-financial stability of these economies. Additional evidence of this link will be necessary to develop and establish adequate regulatory policies.

Acknowledgments

J. Hambuckers (JH) acknowledges the financial support of the National Bank of Belgium, project REFEX. Parts of this research were conducted when JH was affiliated to the University of Göttingen, Chair of Statistics. The authors thank the Editor (R. Valkanov), an associate editor and an anonymous referee for helpful comments that improved the paper.

A Appendix: Carry trade channel

In Table 1, we provide the result of the IRD - CTA regressions discussed in Subsection 2.2. Standard OLS estimates exhibit the expected signs, except for the EUR in the full sample and when $IRD_{t-1} \leq 0$ (only the latter being significant). Results obtained with the instrument-free method of Park and Gupta (2012) controlling for endogeneity are more complicated to analyze. For the full sample, the CHF, GBP, and JPY have estimates that are significant and of the expected signs, whereas the AUD and EUR do not. However, decomposing between periods in which IRD_{t-1} is either null or negative, and positive, we see a clearer picture: estimates are all of the expected sign, except for the EUR when $IRD_{t-1} \leq 0$. The CAD, CHF, GBP, and JPY exhibit positive and significant coefficients in at least one of the considered specifications, whereas the AUD has always positive but not significant estimates. The EUR displays contradictory results, in line with the full sample results.

Finally, to assess whether our results may be spurious due to the non-stationarity of IRD and CTA (which seem to behave as integrated or near-integrated processes), we use the cointegration test of Johansen (1991). As specification for the vector error correction

(VEC) model of the test, we use a lag of order 1 and a deterministic trend. For all currencies, we find evidence of IRD and CTA to be cointegrated of order 1 at least, validating the results of our OLS analysis (Table 14).

Full sample										
Currency	$\hat{\alpha}_0^{(CTA)}$	Std.	t-stat.	$\hat{\beta}^{(CTA)}$	Std.	t-stat.	$\hat{\beta}_{IF}^{(CTA)}$	LB	UB	T
AUD	0.067	0.074	0.909	0.111	0.023	4.797***	-0.144**	-0.229	-0.012	964
CAD	0.094	0.041	2.308**	0.162	0.050	3.224***	-0.132	-0.344	0.084	961
CHF	0.079	0.051	1.537	0.146	0.026	5.635***	0.165**	0.106	0.215	989
EUR	0.088	0.038	2.327**	0.030	0.028	1.066	-0.285**	-0.379	-0.112	991
GBP	-0.169	0.042	-4.058***	0.128	0.032	4.038***	0.218**	0.144	0.286	991
JPY	-0.002	0.055	-0.039	0.063	0.019	3.381***	0.109**	0.067	0.140	991
$IRD_{t-1} \leq 0$										
Currency	$\hat{\alpha}_0^{(CTA)}$	Std.	t-stat.	$\hat{\beta}^{(CTA)}$	Std.	t-stat.	$\hat{\beta}_{IF}^{(CTA)}$	LB	UB	T
AUD	0.501	0.168	2.988***	1.250	0.343	3.642***	1.117	-0.573	2.927	94
CAD	0.127	0.042	3.020***	0.115	0.061	1.876*	0.115	-0.194	0.434	392
CHF	0.082	0.023	3.558***	0.147	0.012	12.152***	0.151**	0.087	0.207	962
EUR	-0.177	0.031	-5.723***	-0.156	0.023	-6.776***	-0.423**	-0.539	-0.260	546
GBP	-0.102	0.058	-1.752*	0.170	0.086	1.973**	0.569	-0.232	0.984	228
JPY	-0.019	0.023	-0.837	0.058	0.008	7.271***	0.153**	0.111	0.188	963
$IRD_{t-1} > 0$										
Currency	$\hat{\alpha}_0^{(CTA)}$	Std.	t-stat.	$\hat{\beta}^{(CTA)}$	Std.	t-stat.	$\hat{\beta}_{IF}^{(CTA)}$	LB	UB	T
AUD	0.123	0.031	3.987***	0.093	0.009	9.895***	0.002	-0.108	0.118	870
CAD	-0.150	0.038	-3.993***	0.437	0.037	11.700***	1.312***	0.954	1.567	599
CHF	-	-	-	-	-	-	-	-	-	27
EUR	-0.027	0.039	-0.687	0.216	0.039	5.615***	0.171	-0.063	0.550	445
GBP	-0.208	0.022	-9.465***	0.149	0.016	9.416***	0.194**	0.127	0.263	763
JPY	-	-	-	-	-	-	-	-	-	28

Table 13: Estimated regression coefficients, standard error, and t-statistics for equation (14). For the full-sample regression, we report Newey-West HAC estimators of the standard errors. For the subsamples, we use White's robust standard errors. $\hat{\beta}_{IF}^{(CTA)}$ denotes the estimates obtained with the instrument-free method of Park and Gupta (2012) controlling for endogeneity. LB and UB denote bounds of the 95% percentile-bootstrap confidence intervals for $\hat{\beta}_{IF}^{(CTA)}$. *, ** and *** denote t-statistics and regression coefficients found significantly different from 0 at the 10%, 5% and 1% confidence level, respectively. Middle and lower panels report the results obtained with subsamples based on the sign of IRD_{t-1} . For the EUR and JPY, since $T < 30$ for $IRD_{t-1} > 0$, we did not report the results.

Cointegration test								
Currency	$r(0)$	stat.	crit. val.	p-val.	$r(1)$	stat.	crit. val.	p-val.
AUD		30.90***	15.49	0.001		3.13*	3.84	0.077
CAD		34.47***	15.49	0.001		5.30**	3.84	0.022
CHF		68.78***	15.49	0.001		2.89*	3.84	0.089
EUR		22.78***	15.49	0.004		3.09*	3.84	0.078
GBP		34.27***	15.49	0.001		1.53	3.84	0.243
JPY		26.69***	15.49	0.001		1.75	3.84	0.186

Table 14: Cointegration test of Johansen (1991) assuming a lag of order 1 and a deterministic linear trend in the underlying VEC model. $r(\cdot)$ denotes the cointegration rank under the tested null hypothesis (0 or 1). We report the test statistics, the critical value at the 5% test level, and the p-value. *, ** and *** denote a rejection of the null hypothesis at levels 10%, 5% and 1%, respectively.

B Appendix: Detailed sampling steps and prior assumptions

In the following, we describe the sampling steps in a more detailed outline including the parameters of the posterior densities. We summarize the information on the respective prior distribution in Table 15.

- We start by sampling the latent states in $\boldsymbol{\alpha}$ by means of the Kalman filter and smoother according to the state-space model in (8). The corresponding matrices in equation (10) are subsequently updated to the most recent parameter estimates for every sweep of the sampler.
- The variance parameter ω^2 is drawn from the inverse gamma distribution $IG(\nu_0 + T/2, S_0 + \lambda)$ with $\lambda = ((h_1)^2(1 - \phi^2) + \sum_{t=2}^T (h_t - \phi h_{t-1})^2) / 2$.
- Next, the AR coefficient ϕ is sampled by means of an independent chain Metropolis-Hastings step with proposal density $N(\hat{\phi}, D_\phi) \mathbb{1}(|\phi| < 1)$ with $\mathbb{1}$ being the indicator function. The parameters for the proposal density are given by

$$D_\phi = \left(V_\phi^{-1} + \mathbf{X}'_\phi \mathbf{X}_\phi / \omega^2 \right)^{-1} \quad \text{and} \quad \hat{\phi} = D_\phi \left(V_\phi^{-1} \phi_0 + \mathbf{X}'_\phi \mathbf{z}_\phi / \omega^2 \right), \quad (15)$$

where $\mathbf{X}_\phi = (h_1, \dots, h_{T-1})'$ and $\mathbf{z}_\phi = (h_2, \dots, h_T)'$. Moreover, ϕ_0 and V_ϕ^{-1} denote the prior parameters for the proposal density.

- Lastly, μ_h is sampled from $N(\hat{\mu}_h, D_{\mu_h})$, where

$$D_{\mu_h} = \left(V_{\mu_h}^{-1} + \mathbf{X}'_{\mu_h} \boldsymbol{\Sigma}_h^{-1} \mathbf{X}_{\mu_h} \right)^{-1} \quad \text{and} \quad \hat{\mu}_h = D_{\mu_h} \left(V_{\mu_h}^{-1} \mu_{h0} + \mathbf{X}'_{\mu_h} \boldsymbol{\Sigma}_h^{-1} \mathbf{z}_{\mu_h} \right), \quad (16)$$

where $\mathbf{X}_{\mu_h} = (1, 1 - \phi, \dots, 1 - \phi)'$, $\mathbf{z}_{\mu_h} = (h_1, h_2 - \phi h_1, \dots, h_T - \phi h_{T-1})'$, and $\boldsymbol{\Sigma}_h = \text{diag}(\omega^2 / (1 - \phi^2), \omega^2, \dots, \omega^2)$.

The hyperparameters for the prior distribution are given in Table 15. Since a priori we are ignorant of the distribution of the regression coefficients, we initiate the Kalman filter with a diffuse prior for $\boldsymbol{\beta}$. In contrast, since h_t follows a stationary AR(1) process, we draw h_1 from a normal distribution with the unconditional mean 0 and variance $\omega^2 / (1 - \phi^2)$.

Table 15: Parameters for prior distributions

Parameter	Distribution	Hyperparameters	
μ_h	N	$\mu_0 = 1$	$V_\mu = 10$
ϕ	N	$\phi_0 = 0.90$	$V_\phi = 0.01$
ω^2	IG	$\nu_0 = 5$	$S_h = 0.04 * 4$

C Appendix: Robustness checks

In this section, we display the results of the robustness checks discussed in Section 5.4. Tables 16 to 21 give the estimated parameters and HPDI at level 90% for each of the specifications, on a currency-by-currency basis.

Table 16: Posterior distribution of regression coefficients with additional control variables for AUD.

		AUD			
Subsample		01/04/1999 – 06/30/2009			
		I	II	III	IV
IRD_{t-1}^+	Mean	-0.10	-0.04	-0.03	-0.10
	HPDI	[-0.26 0.04]	[-0.21 0.11]	[-0.19 0.11]	[-0.24 0.05]
IRD_{t-1}^-	Mean	0.67	0.17	0.13	0.75
	HPDI	[-1.07 1.81]	[-1.39 1.25]	[-1.39 1.21]	[-0.83 1.82]
r_{t-1}	Mean	-0.43	-0.43	-0.44	-0.43
	HPDI	[-0.59 - 0.27]	[-0.60 - 0.29]	[-0.61 - 0.28]	[-0.65 - 0.26]
$S_{t-1} r_{t-1}$	Mean	0.66	0.68	0.68	0.68
	HPDI	[0.39 0.95]	[0.42 0.98]	[0.40 0.98]	[0.37 1.06]
$GVIX_{t-1}$	Mean	0.04	-	-	0.04
	HPDI	[-0.01 0.11]			[-0.01 0.11]
r_{t-1}^{ASX}	Mean	-	-0.20	-	-0.24
	HPDI		[-3.41 3.16]		[-3.78 3.28]
WD_t	Mean	-	-	-0.06	-0.09
	HPDI			[-0.17 0.06]	[-0.24 0.05]
Subsample		07/01/2009 – 12/29/2017			
		I	II	III	IV
IRD_{t-1}^+	Mean	-0.07	-0.03	-0.054	-0.07
	HPDI	[-0.31 0.13]	[-0.17 0.12]	[-0.21 0.11]	[-0.30 0.12]
IRD_{t-1}^-	Mean	-	-	-	-
	HPDI				
r_{t-1}	Mean	-0.59	-0.58	-0.57	-0.57
	HPDI	[-0.85 - 0.36]	[-0.83 - 0.36]	[-0.83 - 0.34]	[-0.84 - 0.33]
$S_{t-1} r_{t-1}$	Mean	0.91	0.89	0.87	0.88
	HPDI	[0.52 1.38]	[0.49 1.34]	[0.49 1.34]	[0.48 1.37]
$GVIX_{t-1}$	Mean	0.01	-	-	0.01
	HPDI	[-0.06 0.11]			[-0.06 0.11]
r_{t-1}^{ASX}	Mean	-	-2.34	-	-2.42
	HPDI		[-8.39 4.19]		[-8.87 4.01]
WD_t	Mean	-	-	0.18	0.19
	HPDI			[0.03 0.34]	[0.04 0.33]

Posterior mean estimates for our baseline SVX specification, augmented with three control variables for AUD: lagged GVIX-7 levels ($GVIX_{t-1}$), lagged national stock index return (r_{t-1}^{ASX}), and an indicator variable for the week-of-the-day effect (WD_t) taking value 0 if the return was recorded on a Monday or Friday, 0 otherwise. In brackets, we report the 90% HPDI obtained from the 5th and 95th percentiles of the posterior distribution. Bold face entries denote estimates for which the 90% HPDI does not include 0. Hyphens indicate that the corresponding variable is excluded from the regression.

Table 17: Posterior distribution of regression coefficients with additional control variables for CAD.

		CAD			
Subsample		01/04/1999 – 06/30/2009			
		I	II	III	IV
IRD_{t-1}^+	Mean	-0.31	-0.43	-0.41	-0.30
	HPDI	[-0.71 0.19]	[-0.83 0.08]	[-0.80 0.06]	[-0.71 0.16]
IRD_{t-1}^-	Mean	0.50	0.68	0.65	0.50
	HPDI	[-0.11 1.00]	[0.047 1.15]	[0.02 1.11]	[-0.08 0.99]
r_{t-1}	Mean	-0.51	-0.53	-0.54	-0.50
	HPDI	[-0.76 - 0.29]	[-0.78 - 0.31]	[-0.81 - 0.30]	[-0.79 - 0.28]
$S_{t-1} r_{t-1}$	Mean	0.75	0.79	0.80	0.74
	HPDI	[0.34 1.25]	[0.38 1.26]	[0.37 1.31]	[0.31 1.30]
GVIX $_{t-1}$	Mean	-0.03	-	-	-0.03
	HPDI	[-0.09 0.04]			[-0.08 0.03]
r_{t-1}^{TSX}	Mean	-	-1.43	-	-1.55
	HPDI		[-6.10 3.03]		[-5.71 2.39]
WD $_t$	Mean	-	-	-0.19	-0.17
	HPDI			[-0.32 - 0.06]	[-0.29 - 0.04]
Subsample		07/01/2009 – 12/29/2017			
		I	II	III	IV
IRD_{t-1}^+	Mean	-1.14	-1.21	-1.08	-1.13
	HPDI	[-1.57 - 0.49]	[-1.59 - 0.61]	[-1.47 - 0.43]	[-1.53 - 0.49]
IRD_{t-1}^-	Mean	2.30	2.45	2.25	2.31
	HPDI	[0.79 3.54]	[0.92 3.64]	[0.75 3.43]	[0.80 3.49]
r_{t-1}	Mean	-0.63	-0.65	-0.65	-0.65
	HPDI	[-0.90 - 0.40]	[-0.90 - 0.42]	[-0.90 - 0.43]	[-0.92 - 0.40]
$S_{t-1} r_{t-1}$	Mean	0.97	1.01	0.99	0.98
	HPDI	[0.56 1.46]	[0.60 1.47]	[0.58 1.42]	[0.56 1.49]
GVIX $_{t-1}$	Mean	-0.00	-	-	0.01
	HPDI	[-0.05 0.07]			[-0.03 0.08]
r_{t-1}^{TSX}	Mean	-	2.57	-	3.30
	HPDI		[-4.77 9.82]		[-3.71 10.16]
WD $_t$	Mean	-	-	-0.31	-0.32
	HPDI			[-0.43 - 0.17]	[-0.46 - 0.19]

Posterior mean estimates for our baseline SVX specification, augmented with three control variables for CAD: lagged GVIX-7 levels (GVIX $_{t-1}$), lagged national stock index return (r_{t-1}^{TSX}), and an indicator variable for the week-of-the-day effect (WD $_t$) taking value 0 if the return was recorded on a Monday or Friday, 0 otherwise. In brackets, we report the 90% HPDI obtained from the 5th and 95th percentiles of the posterior distribution. Bold face entries denote estimates for which the 90% HPDI does not include 0. Hyphens indicate that the corresponding variable is excluded from the regression.

Table 18: Posterior distribution of regression coefficients with additional control variables for GBP.

		GBP			
Subsample		01/04/1999 – 06/30/2009			
		I	II	III	IV
IRD_{t-1}^+	Mean	-0.072	-0.16	-0.14	-0.08
	HPDI	[-0.28 0.16]	[-0.35 0.07]	[-0.34 0.09]	[-0.28 0.14]
IRD_{t-1}^-	Mean	0.11	0.44	0.42	0.14
	HPDI	[-1.28 1.24]	[-1.05 1.54]	[-1.07 1.52]	[-1.13 1.20]
r_{t-1}	Mean	-0.43	-0.45	-0.46	-0.43
	HPDI	[-0.68 - 0.21]	[-0.68 - 0.23]	[-0.71 - 0.22]	[-0.67 - 0.21]
$S_{t-1} r_{t-1}$	Mean	0.73	0.76	0.78	0.72
	HPDI	[0.36 1.17]	[0.37 1.16]	[0.37 1.20]	[0.35 1.12]
GVIX $_{t-1}$	Mean	-0.04	-	-	-0.03
	HPDI	[-0.09 0.02]			[-0.09 0.03]
r_{t-1}^{FTSE}	Mean	-	-3.87	-	-3.72
	HPDI		[-7.93 0.33]		[-7.68 0.32]
WD $_t$	Mean	-	-	-0.02	0.00
	HPDI			[-0.17 0.12]	[-0.13 0.12]
Subsample		07/01/2009 – 12/29/2017			
		I	II	III	IV
IRD_{t-1}^+	Mean	-2.23	-2.03	-1.74	-2.24
	HPDI	[-3.62 - 0.74]	[-3.08 - 0.59]	[-2.75 - 0.36]	[-3.57 - 0.83]
IRD_{t-1}^-	Mean	1.10	1.06	0.93	1.08
	HPDI	[0.24 2.02]	[0.22 1.87]	[0.14 1.79]	[0.31 1.91]
r_{t-1}	Mean	-0.56	-0.55	-0.55	-0.57
	HPDI	[-0.86 - 0.30]	[-0.87 - 0.28]	[-0.84 - 0.29]	[-0.87 - 0.30]
$S_{t-1} r_{t-1}$	Mean	0.80	0.79	0.74	0.79
	HPDI	[0.41 1.25]	[0.39 1.30]	[0.37 1.19]	[0.39 1.27]
GVIX $_{t-1}$	Mean	0.02	-	-	0.03
	HPDI	[-0.04 0.09]			[-0.03 0.09]
r_{t-1}^{FTSE}	Mean	-	-4.16	-	-4.10
	HPDI		[-10.03 2.11]		[-10.20 1.84]
WD $_t$	Mean	-	-	-0.33	-0.35
	HPDI			[-0.46 - 0.19]	[-0.49 - 0.22]

Posterior mean estimates for our baseline SVX specification, augmented with three control variables for GBP: lagged GVIX-7 levels (GVIX $_{t-1}$), lagged national stock index return (r_{t-1}^{FTSE}), and an indicator variable for the week-of-the-day effect (WD $_t$) taking value 0 if the return has been recorded on a Monday or Friday, 0 otherwise. In brackets, we report the 90% HPDI obtained from the 5th and 95th percentiles of the posterior distribution. Bold face entries denote estimates for which the 90% HPDI does not include 0. Hyphens indicate that the corresponding variable is excluded from the regression.

Table 19: Posterior distribution of regression coefficients with additional control variables for EUR.

		EUR			
Subsample		01/04/1999 – 06/30/2009			
		I	II	III	IV
IRD_{t-1}^+	Mean	-0.27	-0.17	-0.13	-0.27
	HPDI	[-0.56 0.09]	[-0.44 0.17]	[-0.41 0.18]	[-0.55 0.08]
IRD_{t-1}^-	Mean	0.40	0.34	0.32	0.40
	HPDI	[0.11 0.62]	[0.09 0.53]	[0.08 0.51]	[0.13 0.60]
r_{t-1}	Mean	-0.67	-0.64	-0.65	-0.68
	HPDI	[-0.97 - 0.48]	[-0.84 - 0.46]	[0.87 - 0.45]	[-0.96 - 0.48]
$S_{t-1} r_{t-1}$	Mean	1.24	1.19	1.19	1.25
	HPDI	[0.89 1.76]	[0.86 1.57]	[0.86 1.57]	[0.89 1.76]
GVIX $_{t-1}$	Mean	0.01	-	-	0.02
	HPDI	[-0.05 0.06]			[-0.05 0.06]
r_{t-1}^{STOXX}	Mean	-	-3.17	-	-3.35
	HPDI		[-7.01 0.60]		[-7.30 0.47]
WD $_t$	Mean	-	-	-0.18	-0.21
	HPDI			[-0.31 - 0.07]	[-0.35 - 0.08]
Subsample		07/01/2009 – 12/29/2017			
		I	II	III	IV
IRD_{t-1}^+	Mean	0.82	0.17	0.38	0.83
	HPDI	[0.09 1.59]	[-0.50 0.98]	[-0.25 1.08]	[0.13 1.61]
IRD_{t-1}^-	Mean	0.33	0.55	0.43	0.28
	HPDI	[-0.11 0.70]	[0.08 0.97]	[0.01 0.83]	[-0.13 0.66]
r_{t-1}	Mean	-0.56	-0.61	-0.55	-0.51
	HPDI	[-0.94 - 0.30]	[-0.89 - 0.36]	[-0.83 - 0.30]	[-0.85 - 0.25]
$S_{t-1} r_{t-1}$	Mean	1.06	1.16	1.07	1.00
	HPDI	[0.61 1.76]	[0.73 1.64]	[0.64 1.55]	[0.55 1.58]
GVIX $_{t-1}$	Mean	-0.08	-	-	-0.06
	HPDI	[-0.14 - 0.02]			[-0.13 - 0.01]
r_{t-1}^{STOXX}	Mean	-	-3.86	-	-3.41
	HPDI		[-9.52 2.21]		[-8.62 2.40]
WD $_t$	Mean	-	-	-0.54	-0.49
	HPDI			[-0.72 - 0.37]	[-0.62 - 0.36]

Posterior mean estimates for the baseline SVX specification, with three control variables, for EUR: lagged GVIX-7 levels (GVIX $_{t-1}$), lagged national stock index return (r_{t-1}^{STOXX}), and an indicator variable for the week-of-the-day effect (WD $_t$) taking value 0 if the return has been recorded on a Monday or Friday, 0 otherwise. In brackets, we report the 90% HPDI obtained from the 5th and 95th percentiles of the posterior distribution. Bold face entries denote estimates for which the 90% HPDI does not include 0. Hyphens indicate that the corresponding variable is excluded from the regression.

Table 20: Posterior distribution of regression coefficients with additional control variables for CHF.

		CHF			
Subsample		01/04/1999 – 06/30/2009			
		I	II	III	IV
IRD_{t-1}^+	Mean	-	-	-	-
	HPDI				
IRD_{t-1}^-	Mean	0.22	0.20	0.19	0.22
	HPDI	[0.11 0.30]	[0.08 0.28]	[0.07 0.26]	[0.12 0.29]
r_{t-1}	Mean	-0.47	-0.45	-0.45	-0.47
	HPDI	[-0.66 - 0.30]	[-0.63 - 0.30]	[-0.62 - 0.29]	[-0.67 - 0.30]
$S_{t-1} r_{t-1}$	Mean	0.96	0.93	0.90	0.95
	HPDI	[0.66 1.30]	[0.65 1.25]	[0.63 1.22]	[0.64 1.31]
$GVIX_{t-1}$	Mean	0.01	-	-	0.02
	HPDI	[-0.02 0.04]			[-0.01 0.04]
r_{t-1}^{SSMI}	Mean	-	1.50	-	1.41
	HPDI		[-2.76 5.78]		[-2.87 5.75]
WD_t	Mean	-	-	-0.17	-0.20
	HPDI			[-0.28 - 0.06]	[-0.33 - 0.09]
Subsample		07/01/2009 – 12/29/2017			
		I	II	III	IV
IRD_{t-1}^+	Mean	-	-	-	-
	HPDI				
IRD_{t-1}^-	Mean	0.47	0.55	0.49	0.43
	HPDI	[0.19 0.79]	[0.25 0.86]	[0.20 0.81]	[0.15 0.70]
r_{t-1}	Mean	-0.07	-0.06	-0.08	-0.05
	HPDI	[-0.19 0.05]	[-0.19 0.07]	[-0.21 0.05]	[-0.17 0.07]
$S_{t-1} r_{t-1}$	Mean	0.45	0.47	0.48	0.43
	HPDI	[0.19 0.71]	[0.20 0.75]	[0.21 0.77]	[0.17 0.70]
$GVIX_{t-1}$	Mean	-0.04	-	-	-0.02
	HPDI	[-0.08 0.01]			[-0.07 0.02]
r_{t-1}^{SSMI}	Mean	-	-4.27	-	-3.75
	HPDI		[-11.02 2.30]		[-10.35 2.80]
WD_t	Mean	-	-	-0.45	-0.42
	HPDI			[-0.58 - 0.31]	[-0.55 - 0.29]

Posterior mean estimates for the baseline SVX specification, with three control variables, for CHF: lagged GVIX-7 levels ($GVIX_{t-1}$), lagged national stock index return (r_{t-1}^{SSMI}), and an indicator variable for the week-of-the-day effect (WD_t) taking value 0 if the return has been recorded on a Monday or Friday, 0 otherwise. In brackets, we report the 90% HPDI obtained from the 5th and 95th percentiles of the posterior distribution. Bold face entries denote estimates for which the 90% HPDI does not include 0. Hyphens indicate that the corresponding variable is excluded from the regression.

Table 21: Posterior distribution of regression coefficients with additional control variables for JPY.

		JPY			
Subsample		01/04/1999 – 06/30/2009			
		I	II	III	IV
IRD_{t-1}^+	Mean HPDI	-	-	-	-
IRD_{t-1}^-	Mean HPDI	0.14 [0.05 0.21]	0.12 [0.00 0.19]	0.11 [-0.01 0.18]	0.13 [0.05 0.21]
r_{t-1}	Mean HPDI	-0.25 [-0.42 - 0.08]	-0.24 [-0.42 - 0.07]	-0.24 [-0.42 - 0.08]	-0.25 [-0.43 - 0.08]
$S_{t-1} r_{t-1}$	Mean HPDI	0.69 [0.40 1.01]	0.68 [0.37 1.00]	0.67 [0.38 1.01]	0.68 [0.37 1.01]
$GVIX_{t-1}$	Mean HPDI	0.02 [-0.02 0.07]	-	-	0.02 [-0.02 0.07]
r_{t-1}^{NIKKEI}	Mean HPDI	-	-0.32 [-3.66 3.12]	-	-0.53 [-3.85 2.96]
WD_t	Mean HPDI	-	-	-0.04 [-0.16 0.07]	-0.05 [-0.17 0.08]
Subsample		07/01/2009 – 12/29/2017			
		I	II	III	IV
IRD_{t-1}^+	Mean HPDI	-	-	-	-
IRD_{t-1}^-	Mean HPDI	0.33 [-0.13 0.85]	0.51 [0.00 1.01]	0.38 [-0.08 0.88]	0.28 [-0.15 0.76]
r_{t-1}	Mean HPDI	-0.25 [-0.48 - 0.03]	-0.33 [-0.57 - 0.09]	-0.32 [-0.59 - 0.09]	-0.26 [-0.49 - 0.05]
$S_{t-1} r_{t-1}$	Mean HPDI	0.49 [0.11 0.89]	0.64 [0.24 1.07]	0.64 [0.23 1.09]	0.51 [0.15 0.89]
$GVIX_{t-1}$	Mean HPDI	-0.05 [-0.10 - 0.00]	-	-	-0.04 [-0.08 0.01]
r_{t-1}^{NIKKEI}	Mean HPDI	-	1.24 [-3.51 5.99]	-	0.79 [-3.61 5.07]
WD_t	Mean HPDI	-	-	-0.41 [-0.57 - 0.24]	-0.34 [-0.47 - 0.21]

Posterior mean estimates for the baseline SVX specification, with three control variables, for JPY: lagged $GVIX$ -7 levels ($GVIX_{t-1}$), lagged national stock index return (r_{t-1}^{NIKKEI}), and an indicator variable for the week-of-the-day effect (WD_t) taking value 0 if the return has been recorded on a Monday or Friday, 0 otherwise. In brackets, we report the 90% HPDI obtained from the 5th and 95th percentiles of the posterior distribution. Bold face entries denote estimates for which the 90% HPDI does not include 0. Hyphens indicate that the corresponding variable is excluded from the regression.

Figures

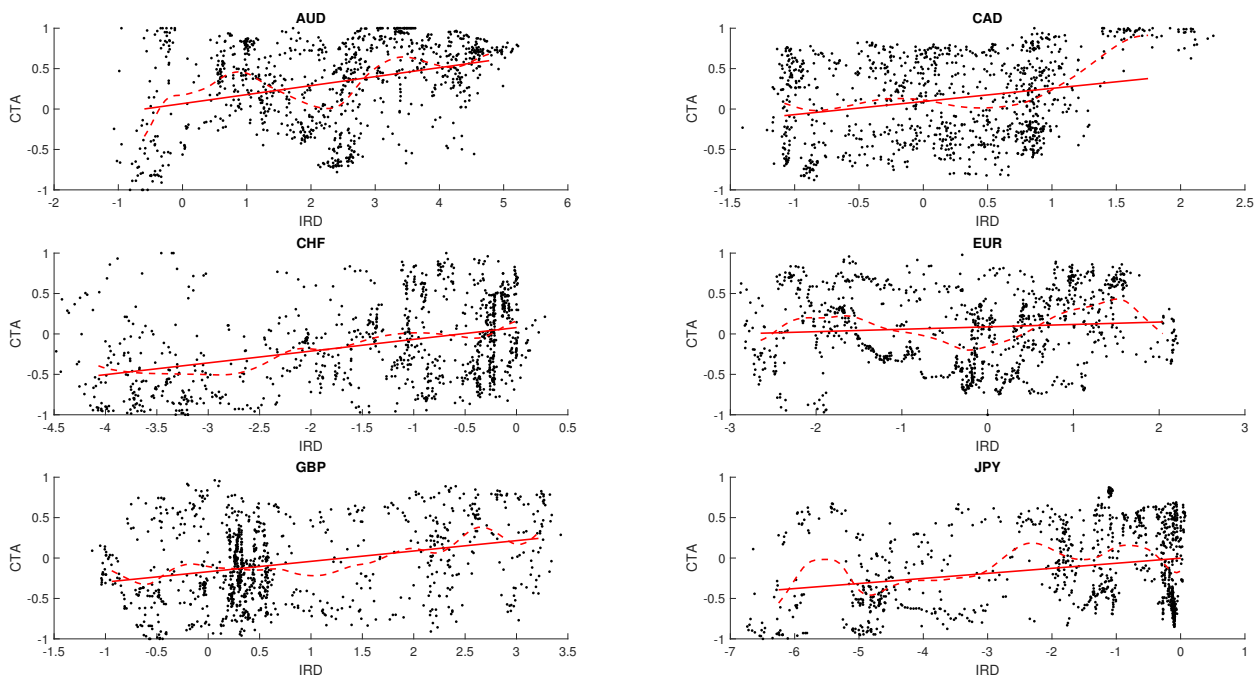


Figure 1: Link between weekly IRD and CTA, for the Australian and Canadian Dollar, the Swiss Franc, the Euro, the British Pound and the Japanese Yen. Solid red: linear regression function obtained with OLS. Dashed red: smoothed local linear kernel (LLK) estimate. To avoid boundary effects, we only display the LLK estimate for values of the predictor between the empirical quantiles at levels 2.5% and 97.5%.




Figure 2: Posterior mean estimates of the regression effects (β) across 200 samples. The true parameter value is indicated by the vertical black line.




Figure 3: Estimated volatility process with the posterior mean displayed in blue and the 90% HPDI as gray shadowed area for $t = 1, \dots, T$. The true volatility process $\exp\{h_t\}$ is indicated by the cyan line.

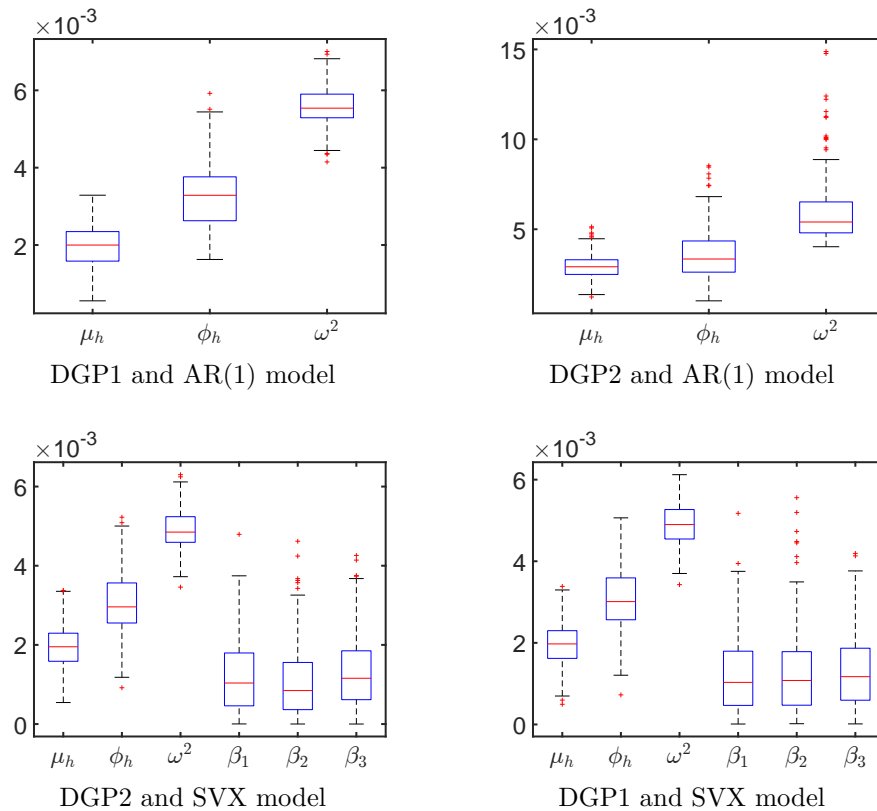


Figure 4: RMSEs for all parameters over 200 simulated samples. The first row contains the RMSEs for the SV model with an AR(1) specification. The graphs in the second row display the RMSEs for the SVX model.

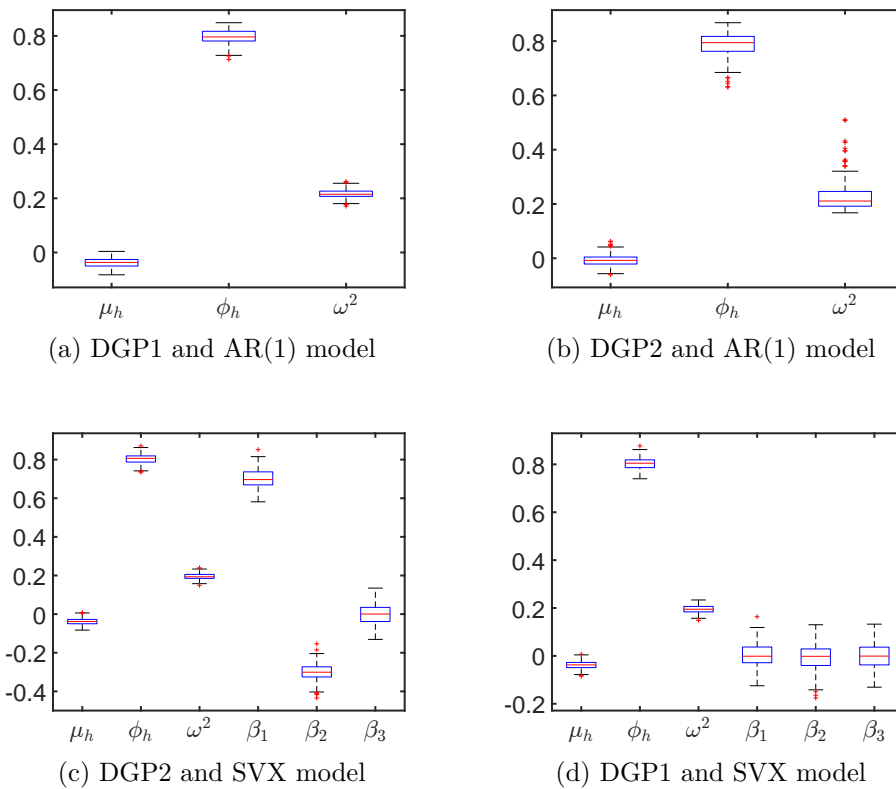


Figure 5: Posterior mean estimates for parameters in 200 samples. The first row contains the results from estimating an AR(1) model for the volatility while the second row displays estimates from the SVX model.

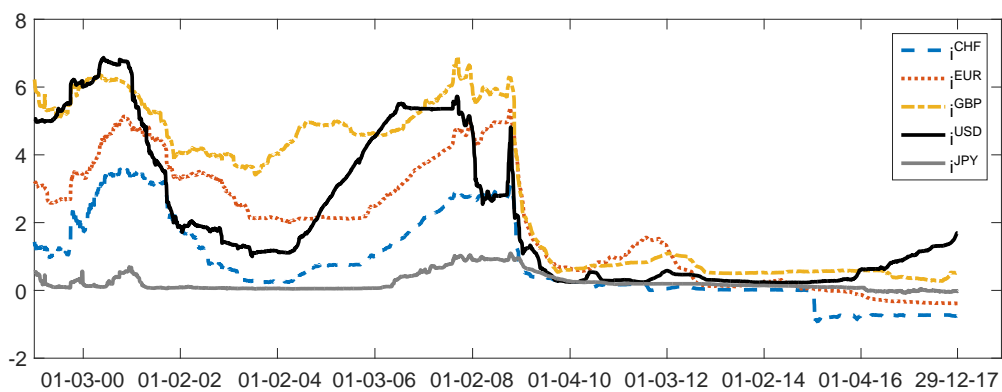


Figure 6: Daily 3-month Libor rates from January, 4th 1999 to December, 29th 2017.

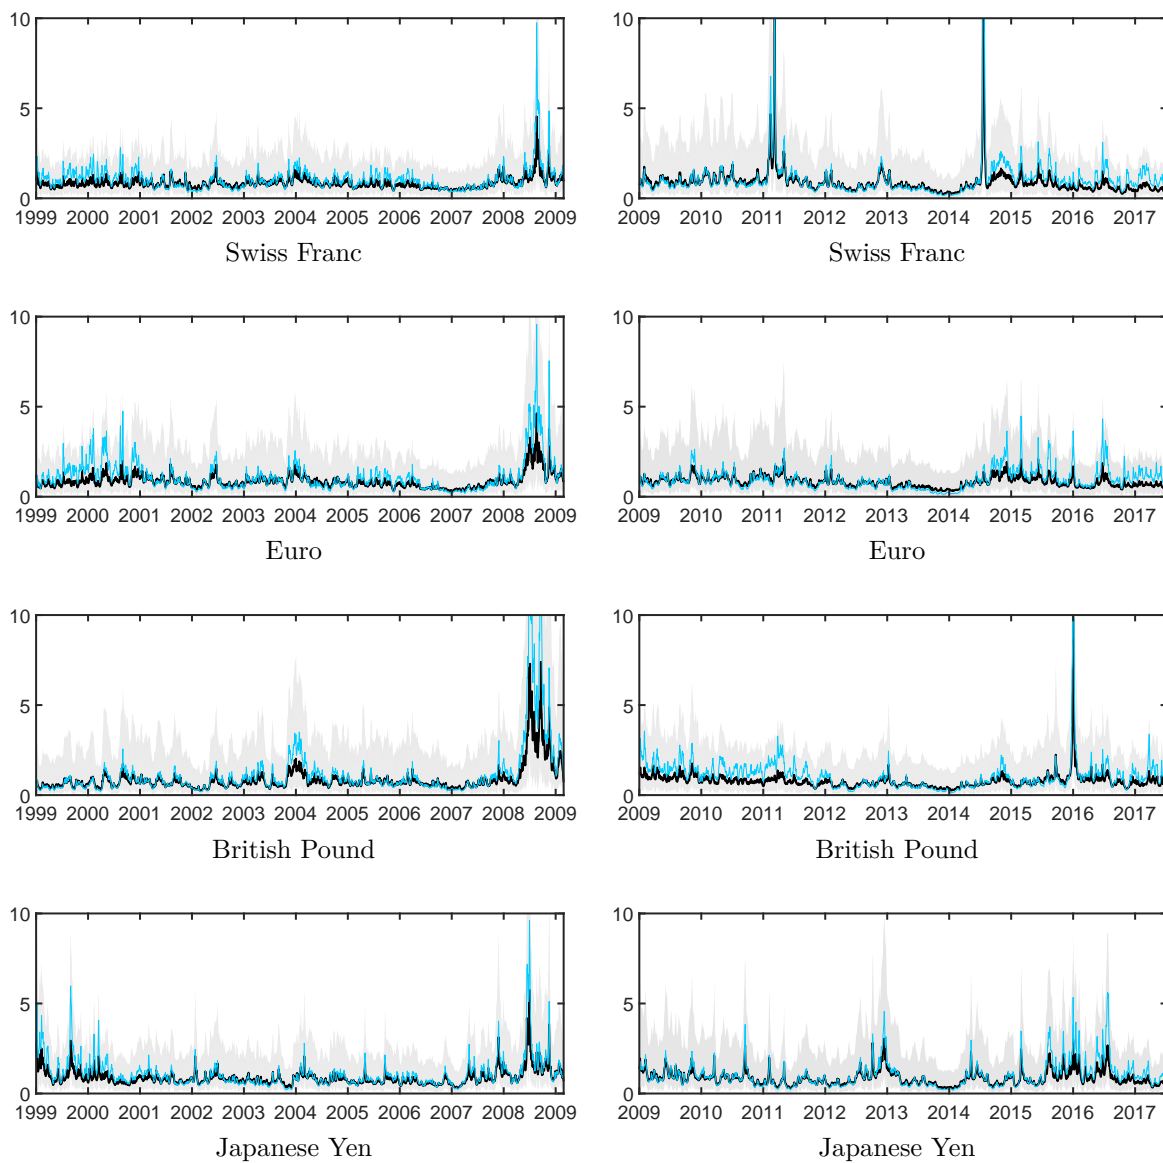


Figure 7: Smoothed estimates of volatility for the two subsamples. The black line and the gray shadowed area indicate the posterior mean and the 90 % credible interval for the volatility $\exp\{h_t\}$ in the AR(1) model. The cyan line indicates the posterior mean of volatility $\exp\{h_t\}$ in the extended model. The left column of the figure displays estimates for the subperiod 01/04/1999 – 06/30/2009, the right column the results for 07/01/2009 – 12/29/2017.

References

- Accominotti, O., J. Cen, D. Chambers, and I. W. Marsh (2019). Currency regimes and the carry trade. *Journal of Financial and Quantitative Analysis* 54(5), 2233–2260.
- Ahmed, S. and G. Valente (2015). Understanding the price of volatility risk in carry trades. *Journal of Banking and Finance* 57, 118–129.
- Ames, M., G. Bagnarosa, and G. W. Peters (2017). Violations of uncovered interest rate parity and international exchange rate dependencies. *Journal of International Money and Finance* 73(A), 162–187.
- Anzuini, A. and F. Fornari (2012). Macroeconomic determinants of carry trade activity. *Review of International Economics* 20(3), 468–488.
- Bacchetta, P. and E. Van Wincoop (2010). Can information heterogeneity explain the exchange rate determination puzzle? *American Economic Review* 96(3), 552–576.
- Bates, D. S. (1996). Jumps and stochastic volatility: Exchange rate processes implicit in deutsche mark options. *The Review of Financial Studies* 9(1), 69–107.
- Bauwens, L., C. Hafner, and S. Laurent (2012). *Handbook of volatility modeling and their applications*. John Wiley and Sons.
- Brunnermeier, M., S. Nagel, and L. Pedersen (2009). Carry trades and currency crashes. In D. Acemoglu, K. Rogoff, and M. Woodford (Eds.), *NBER Macroeconomic Annual 2008*, pp. 313–347.
- Burnside, C., M. Eichenbaum, and S. Rebelo (2009). Understanding the forward premium puzzle: A microstructure approach. *American Economic Journal: Macroeconomics* 1(2), 127–154.
- Burnside, C., M. Eichenbaum, and S. Rebelo (2011). Carry trade and momentum in currency markets. *Annual Review of Financial Economics* 3(1), 511–535.
- Chan, J. C. and A. L. Grant (2016). Modeling energy price dynamics: GARCH versus stochastic volatility. *Energy Economics* 54, 182 – 189.

- Chernov, M., J. Graveline, and I. Zviadadze (2018). Crash risk in currency returns. *Journal of Financial and Quantitative Analysis* 53(1), 137–170.
- Chordia, T., R. Roll, and A. Subrahmanyam (2001). Market liquidity and trading activity. *The Journal of Finance* 56(2), 501–530.
- Christiansen, C., A. Rinaldo, and P. Söderling (2011). The Time-Varying Systematic Risk of Carry Trade Strategies. *The Journal of Financial and Quantitative Analysis* 46(4), 1107–1125.
- Chung, J. and Y. Hong (2007). Model-free evaluation of directional predictability in foreign exchange markets. *Journal of Applied Econometrics* 22(5), 855–889.
- Clark, T. E. and F. Ravazzolo (2015). Macroeconomic forecasting performance under alternative specifications of time-varying volatility. *Journal of Applied Econometrics* 30(4), 551–575.
- de Rato, R. (2007). Speech by Rodrigo de Rato, Managing Director at the IMF at the Harvard Business School Alumni Dinner.
- Diebold, F. and R. Mariano (1995). Comparing predictive accuracy. *Journal of Business and Economic Statistics* 13(3), 253–263.
- Durbin, J. and S. Koopman (2012). *Time Series Analysis by State Space Methods*. Oxford Statistical Science Series.
- Fama, E. (1984). Forward and spot exchange rates. *Journal of Monetary Economics* 14(3), 319–338.
- Farhi, E. and X. Gabaix (2016, 10). Rare disasters and exchange rates. *The Quarterly Journal of Economics* 131(1), 1–52.
- Fleming, J., C. Kirby, and B. Ostdiek (2001). The economic value of volatility timing. *The Journal of Finance* 56(1), 329–352.
- Froot, K. and T. Ramadorai (2005). Currency returns, intrinsic value, and institutional-investor flows. *The Journal of Finance* 60(3), 1535–1566.
- Gabaix, X. and M. Maggiori (2015). International liquidity and exchange rate dynamics. *The Quarterly Journal of Economics* 130(3), 1369 – 1420.

- Geweke, J. (1992). Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments. In J. M. Bernardo, J. Berger, A. Dawid, and A. Smith (Eds.), *Bayesian Statistics 4*, pp. 169–193. Oxford University Press.
- Giacomini, R. and H. White (2006). Tests of conditional predictive ability. *Econometrica* 74(6), 1545–1578.
- Han, H. and K. Kristensen (2014). Asymptotic Theory for the QMLE in GARCH-X Models With Stationary and Nonstationary Covariates. *Journal of Business & Economic Statistics* 32(3), 416–429.
- Hutchison, M. and S. Sushko (2013). Impact of macro-economic surprises on carry trade activity. *Journal of Banking and Finance* 37, 1133–1147.
- Ichiue, H. and K. Koyama (2011). Regime switches in exchange rate volatility and uncovered interest rate parity. *Journal of International Money and Finance* 30(7), 1436–1450.
- Ismailov, A. and B. Rossi (2018). Uncertainty and deviations from uncovered interest rate parity. *Journal of International Money and Finance* 88, 242–259.
- Johansen, S. (1991). Estimation and hypothesis testing of cointegration vectors in gaussian vector autoregressive models. *Econometrica* 59(6), 1551–1580.
- Justiniano, A. and G. Primiceri (2008, June). The time-varying volatility of macroeconomic fluctuations. *American Economic Review* 98(3), 604–41.
- Kim, S., N. Shephard, and S. Chib (1998). Stochastic volatility: Likelihood inference and comparison with ARCH Models. *The Review of Economic Studies* 65(3), 361–393.
- Li, D., A. Ghoshray, and B. Morley (2011). Uncovered interest parity and the risk premium. *University of Bath Working Paper*.
- Liesenfeld, R. and J. Richard (2006). Classical and Bayesian analysis of univariate and multivariate stochastic volatility models. *Econometric Reviews* 25(2-3), 335–360.
- Mahieu, R. J. and P. C. Schotman (1998). An empirical application of stochastic volatility models. *Journal of Applied Econometrics* 13(4), 333–360.
- Meese, R. and K. Rogoff (1983). Empirical exchange rate models of the seventies: Do they fit out of sample? *Journal of International Economics* 14, 3–24.

- Omori, Y., S. Chib, N. Shephard, and J. Nakajima (2007). Stochastic volatility with leverage: Fast and efficient likelihood inference. *Journal of Econometrics* 140(2), 425–449.
- Omrane, W. and C. Hafner (2015). Macroeconomic news surprises and volatility spillover in foreign exchange markets. *Empirical Economics* 48(2), 577– 607.
- Omrane, W. and T. Savaser (2017). Exchange rate volatility response to macroeconomic news during the global financial crisis. *International Review of Financial Analysis* 52, 130–143.
- Park, S. and S. Gupta (2012). Handling endogenous regressors by joint estimation using copulas. *Journal of Marketing Science* 31(4), 567–586.
- Rossi, B. (2013). Exchange rate predictability. *Journal of Economic Literature* 51(4), 1063–1119.
- Tsay, R. S. (2005). *Analysis of Financial Time Series*. John Wiley & Sons.
- Valchen, R. (2020). Bond convenience yields and exchange rate dynamics. *American Economic Journal: Macroeconomics* 12(2), 124–166.

