EFA Eastern Finance Association

The Financial Review 54 (2019) 201–229

The Financial Review

Higher Moments and Exchange Rate Behavior

Siroos Khademalomoom

Department of Treasury and Finance, Victoria, Australia

Paresh Kumar Narayan

Deakin University

Susan Sunila Sharma*

Deakin University

Abstract

This paper uses 15-minute exchange rate returns data for the six most liquid currencies (i.e., the Australian dollar, British pound, Canadian dollar, Euro, Japanese yen, and Swiss franc) vis-à-vis the United States dollar to examine whether a GARCH model augmented with higher moments (HM-GARCH) performs better than a traditional GARCH (TG) model. Two findings are unraveled. First, the inclusion of odd/even moments in modeling the return/variance

**Corresponding author*: Centre for Financial Econometrics, Deakin Business School, Deakin University, 221 Burwood Highway, Burwood, VIC 3125, Australia; Phone: +61 3 9244 6871; Fax: +61 3 9244 6034; E-mail: s.sharma@deakin.edu.au.

The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Department of Treasury and Finance, Victoria, Australia.

This paper is original and has not been submitted elsewhere for publication. It is a chapter of the first author's PhD thesis undertaken at Deakin University, Melbourne, Australia. Earlier versions of this paper were presented at the Centre for Financial Econometrics—*Journal of Banking and Finance* 2015 conference on "Recent Developments in Financial Econometrics" at Deakin University, Geelong, Australia. We acknowledge helpful comments and suggestions provided by conference participants, seminar participants at Deakin University, and Professors Jonathan Batten and Niklas Wagner. Helpful comments and suggestions from the Editor (Dr. Richard Warr) and anonymous reviewers of this journal are duly acknowledged.

improves the statistical performance of the HM-GARCH model. Second, trading strategies that extract buy and sell trading signals based on exchange rate forecasts from HM-GARCH models are more profitable than those that depend on TG models.

Keywords: foreign exchange, high frequency, modeling, higher moments, trading strategy

JEL Classifications: C5, C58, F31, G15

1. Introduction

In this paper, we examine the role of high-order moments in influencing exchange rate behavior. We are not the first to explore the role of higher moments in understanding exchange rate behavior. There is a literature on this; see Aggarwal (1990), Harvey and Siddique (1999), and Mittnik and Paolella (2000). These studies show that higher moments improve the statistical performance of the models. However, these studies only consider up to the fourth moment. High-order moments in excess of the fourth moment have not been considered in terms of how they influence exchange rate behavior.[1]

There are several economic channels/mechanisms through which higher moments can impact exchange rate behavior. The first channel of effect is "liquidity spirals" that results from the theoretical model of Brunnermeier and Pedersen (2009). The basic idea of their model is that invested securities contain positive average returns and a negative skewness. They explain the source of this positive returns and negative skewness. Positive returns owe to the premium resulting from speculators' provision of liquidity while negative skewness is because investors make heavy losses and relatively mild gains from negative shocks (such as financial constraints) and positive shocks (such as liquidity), respectively. In other words, the impact of shocks is asymmetric, skewed heavily in favor of negative shocks. More specifically, their work implies that funding constraints determine market liquidity. When there are funding shocks, market liquidity declines leading to higher margins, which exerts

[1] The motivation for this is rooted in the fact that these high-order moments contain different types of risk-related information. The second-order moment (variance), for instance, represents volatility, the third-order moment (skewness) accounts for the probability of positive and negative values, and the fourth-order moment (kurtosis) is indicative of the relative importance of tails versus shoulders in causing dispersion. High kurtosis is linked to heavy tails, and low kurtosis corresponds to heavy shoulders. In other words, the fourth moment or volatility of variance indicates how uncertain the uncertainty is. The fifth moment (hyper-skewness) indicates the relative importance of tails versus the center in causing skewness and corresponds to a heavy tail. There is little movement in the mode when the hyper-skewness is high and a greater change in the shoulders when hyper-skewness is low. In other words, hyper-skewness measures the asymmetric sensitivity of the kurtosis. For instance, in the case of the stock market, the fifth moment can explain a symmetric or asymmetric distribution of fluctuations in assets with high/low volatilities. The sixth moment, hyper-kurtosis measures both the peakedness and tails relative to the normal distribution. See also Kostakis, Muhammad and Siganos (2012).

further pressure on speculators' funding constraints. They refer to this as "margin spiral." Moreover, a funding shock dents market liquidity, leading to speculators losses on their initial positions. Speculators are forced to sell more, thus causing a price decline. This they refer to as a "loss spiral." These liquidity driven spirals are reasons for extreme events and substantial losses, generating higher order moments in the currency market.

The second channel of effect on exchange rates from higher moments owes to carry trade as shown, for instance, in the work of Brunnermeier, Nagel and Pedersen (2008). They show that sudden exchange rate movements (not related to news) are likely due to unwinding of carry trades when speculators are faced with funding constraints. They show that crash risk is a feature of investment currencies—that is, positive interest rate differentials (carry) and conditional skewness of exchange rate movements are negatively related. Moreover, Brunnermeier Nagel and Pedersen (2008) argue that the speculators' positions (which are positively related to the carry) increase crash risk, which is consistent with Plantin and Shin's (2007) claim that carry trades can be destabilizing. Carry traders, because of crash risk, demand a risk premium (Brunnermeier, Nagel and Pedersen, 2008). Empirical evidence reported in Doskov and Swinkels (2015) shows that carry strategies result in substantial losses over the 1901–2012 period. In other words, carry trade lead not only to extreme events but can trigger substantial losses, thus gives rise to higher skewness and kurtosis in currency markets.

The third channel that motivates the presence of extreme events (higher order kurtosis) in currency markets can be traced to the price-contingent trading hypothesis, as argued in Osler and Savaser (2011). The main thesis of this hypothesis is that trading results when the price arrives at a specific level. The main source of price-contingent trading is algorithmic and technical trading, and dynamic option hedging—either one of these is responsible for extreme events in currency markets. Multiple factors contribute to such extreme events, including (1) fat tails in the distribution of order sizes (Osler and Savaser, 2011), (2) clustering of order executions (Osler, 2003), and (3) feedback between trading and returns (Genotte and Leland, 1990).

A common feature emerging from these theoretical motivations for the role of higher order moments in influencing exchange rate behavior is that the theory does not point out the precise higher order moment that matters. For example, liquidity spirals, carry trade, and price contingent trading all contribute to extreme events in the exchange rate market, leading to substantial losses, thus giving rise to skewness and kurtosis. Theory does not dwell on what precisely is the order of these moments. The reason because the precise order of moments would depend on the magnitude of effect on exchange rate from extreme events. Theory does not predict the magnitude of this effect. Therefore, the precise order and indeed whether even or odd moments impact mean and/or variance of exchange rate returns is purely an empirical issue and needs to be dealt with accordingly. To date, the implications of these higher moments on exchange rates have not been studied. In addition, what has also not been explored so far is the economic implication of these higher moments of exchange rate returns. The goal of this paper is to fill this research gap.

Our approaches to achieving the aim of our paper are threefold. In the first step, we compile a data set on exchange rates sourced from Thomson Reuters Tick History (TRTH). Specifically, we compile 15-minute exchange rate price data for six most traded currencies—Australian dollar (AUD), British pound (GBP), Canadian dollar (CAD), Euro (EUR), Japanese yen (JPY), and the Swiss franc (CHF)—vis-à-vis the United States dollar (USD). In the second step, using exchange rate returns, we compute the conditional higher moments of exchange rate returns, namely, skewness, kurtosis, hyper-skewness, and hyper-kurtosis.

Finally, we test the statistical performance of the exchange rate returns model that includes these high-order moments. The aim here is to explore the economic importance of the inclusion of these higher moments on the profitability of exchange rate trading. Specifically, to estimate the economic significance of high-order moments, we forecast exchange rate returns using a higher moment GARCH (HM-GARCH) model and utilize the sign and magnitude of the return forecasts to generate buy and sell signals. Profits are computed from these trading signals. The resulting profits are compared with profits obtained when exchange rate return forecasts are generated using a traditional GARCH (TG) model—that is, a model that excludes the higher moments.

Our approaches contribute to two fresh revelations with respect to the exchange rate market. Our first finding is statistical. We run GARCH-type regressions that include high-order moments in both the mean and variance equations. While broadly this relation was understood, what was not understood was how exactly the odd and even moments impacted exchange rate returns. Specifically, we find that including the odd moments in the mean equation and the even moments in the variance equation maximizes the statistical fit of the exchange rate returns model. An extensive in-sample (IS) and out-of-sample (OOS) period forecasting analysis also convincingly reveals that the proposed HM-GARCH model beats the TG model. Our second finding revolves around the economic importance of the statistical evidence favoring the role of higher moments in influencing exchange rate behavior. Using several IS and OOS periods, we forecast exchange rate returns of each of the six currencies using the HM-GARCH model. We find that buy and sell signals extracted from exchange rate forecasts based on the HM-GARCH model offer higher net (of transaction costs) profits compared to profits generated using exchange rate returns models that exclude higher moments (TG model). Both statistical and economic significance results are robust to different sample periods of data.

Our findings contribute to the literature in the following way. Our first finding that including higher moments improve the statistical performance of exchange rate return models is consistent with the exchange rate studies (see, among others, Aggarwal, 1990; Harvey and Siddique, 1999; Mittnik and Paolella, 2000) and also the equity market studies (see, among others, Rubinstein,1973; Kraus and Litzenberger, 1976, 1983; Brooks, Burke, Heravi and Persand, 2005; Harvey, Liechty, Liechty and

Müller, 2010; Brooks, Cerny and Miffre, 2012). However, our contribution is not that we show the statistical superiority of high-order moments in explaining exchange rate behavior; rather, we show, for the first time in this literature, that the superior statistical performance translates into profitable trading strategies.

In this regard, our study joins the literature on the profitability of exchange rate markets. In this literature, profitability of exchange rate trading has been analyzed from different perspectives, dominated mostly by the momentum trading strategies and technical trading rules. See, among others, Levich and Thomas (1993), Taylor (1994), Kho (1996), Lee and Mathur (1996), LeBaron (1999), Raj (2000), Okunev and White (2003), Mende and Menkhoff (2006), Harris and Yilmaz (2009), Menkhoff, Sarno, Schmeling and Schrimpf (2012), Narayan, Mishra, Narayan and Thuraisamy (2015), and Taylor and Allen (1992). The main message from this literature is that the exchange rate market is profitable. We contribute to this by showing that tracking high-order moments in forecasting exchange rate returns leads to profitable trading strategies. In other words, we do not use momentum trading strategies to show profitability of the market; rather, we forecast exchange rate returns using higher moments, and show that trading signals obtained on the basis of forecasted returns are also profitable. Our study, therefore, offers a fresh perspective on exchange rate profitability.

The rest of this paper is organized as follows. Section 2 contains a discussion on hypotheses development. Section 3 describes the estimation method. The data and empirical findings are presented in Section 4. A key feature of this section is the construction of a trading strategy that illustrates the economic significance of our statistical findings. A summary of the paper appears in the final section.

2. Hypotheses development

We propose two hypotheses that have not been investigated previously using high-frequency exchange rate data.

H1: The inclusion of high-order moments improves the forecasting performance of exchange rate models.

It is true that several studies show the importance of incorporating higher moments into models of the mean and volatility of returns. However, a key feature of this literature is that they mostly focus on the equity market, where the emphasis is on understanding the consequences of ignoring these higher moments. There are only a very few studies that are based on the currency market. Aggarwal (1990), for instance, examines the statistical distribution of major currencies. He suggests that in foreign currency forecasting and hedging practices, mean-variance portfolio analysis, foreign currency option pricing, and other research involving exchange rates, one should account for these significant deviations from normality (i.e., skewness and kurtosis). Harvey and Siddique (1999) find that many economic variables, such as exchange rate and stock index returns, have asymmetric distributions. They discover that the inclusion of skewness in models of financial data influences the persistence of the variance and causes the asymmetry in the variance to disappear. Premaratne and Bera (2000) claim that models (applied to a range of financial data) that include the mean, variance, skewness, and kurtosis simultaneously tend to outperform traditional models that include only the first two moments. Likewise, Mittnik and Paolella (2000) show that the third and the fourth moments are required to model the daily exchange rate returns of the East Asian currencies against the USD and demonstrate that the inclusion of these moments significantly improves the statistical performance of models.

The main message emanating from these studies is that high-order moments influence exchange rate behavior. A feature of these studies in arriving at this conclusion is that they are based on at most the fourth-order moment. The limitation, therefore, is that these are not the only moments that can potentially influence exchange rate behavior. This idea has roots in the work of Brooks, Burke, Heravi and Persand (2005), who argue that ignoring higher moments can lead to substantial approximation error not only in estimation but also in optimal decision and portfolio selection. More specifically, and therefore giving credence to our proposed hypothesis, they argue in favor of models that allow for dynamic higher moments since it is these models that describe the distributional properties of financial asset returns better than moment-free models, especially when modeling high-frequency data. We, therefore, hypothesize that the presence of higher moments (up to the sixth-order) will improve the statistical performance of the exchange rate model vis-à-vis a constant returns model.

H2: The inclusion of odd/even moments in modeling currency return/variance improves the statistical performance of the model.

Several studies have argued that high-order moments are important for risk assessment (see, e.g., Athayde and Flores, 2006; Malevergne and Sornette, 2006; Conrad, Dittmar and Ghysels, 2008). They show that the first four moments are important in investors' decisions. In particular, they show that higher odd moments and lower even moments tend to reduce risk. "Agents 'like' odd moments and 'dislike' even ones" (Athayde and Flores, 2006, p. 38). Malevergne and Sornette (2006) argue that while volatility is considered the main measure of risk, higher moments are often ignored, and including these higher moments could provide an opportunity to increase expected returns while minimizing substantial risk. Conrad, Dittmar and Ghysels (2008) find that every high moment matters and that investors' expectations of these moments can change over time.

To summarize, although these previously mentioned studies emphasize the importance of considering high-order moments in modeling the mean and volatility of returns, thus far, there is no empirical evidence on the impact of high-order moments (greater than the fourth moment—kurtosis) on asset returns and volatility. This motivates our second hypothesis—that is, we examine whether inclusion of odd (up to the fifth) and even (up to the sixth) moments in modeling currency returns/variance improves the statistical performance of the exchange rate model.

3. Estimation approach

In our estimation procedure, we first calculate the returns of the selected currencies against the USD using the mid-quote price at the beginning and end of each 15-minute interval. Our motivation for using 15-minute data has roots in the arguments presented in Andersen, Bollerslev and Diebold (2007). The point about data frequency here is that too high a frequency (such as 1 minute or 5 minutes) exposes the sample to market microstructure contamination and thus leads to biased variance estimates. Striking a balance in terms of the appropriate high-frequency data to use is therefore important: several studies (see Chevallier, 2011; Gnabo, Hvozdyk and Lahaye, 2014; Palandri, 2015), as a result, take 15-minute data for application purposes. We do the same.

We have: $EP_t = (P_t^{OB} + P_t^{OA} + P_t^{CB} + P_t^{CA})/4$, where EP_t is the estimated mid-quote price (foreign currency value of the USD) at time t and P_t^{OB}, P_t^{OA}, P_t^{CB}, and P_t^{CA} are the prices at the opening bid, opening ask, closing bid, and closing ask, respectively. Consequently, the return is: $R_t = \log(EP_t/EP_{t-1}) * 100$, where R_t denotes the 15-minute exchange rate return. The way in which we compute exchange rate returns represents a direct quote such that an increase in the exchange rate implies a depreciation of the local currency vis-à-vis the USD. Then, to ensure that our proposed hypotheses test is free of heteroskedasticity, we consider a GARCH-type model (with higher moments included)—assuming that the residuals follow a conditional Student's t distribution. The model has the following form:

$$R_t = \alpha_0 + \sum_{i=1}^{n} \alpha_i R_{t-i} + \sum_{j=0}^{m} \beta_j HM_{t-j} + \varepsilon_t, \tag{1}$$

$$h_t = \omega_0 + \sum_{i=1}^{p} \omega_i \varepsilon_{t-i}^2 + \sum_{j=1}^{q} \theta_j h_{t-j} + \sum_{k=0}^{d} \delta_k HM_{t-k}. \tag{2}$$

Equation (1) represents the mean equation, where α and β are the parameters to be estimated and ε_t represents the error term. The lagged values of returns ($\sum_{i=1}^{n} \alpha_i R_{t-i}$) are included to account for autocorrelation, and $\sum_{j=0}^{m} \beta_j HM_{t-j}$ represents the conditional higher moments (i.e., skewness, kurtosis, hyper-skewness, and hyper-kurtosis) and their lagged values. The optimal lag length is chosen using the Akaike information criterion (AIC).[2]

[2] The optimum lags based on both AIC and Bayesian information criterion (BIC) are two lags for skewness and hyper-skewness and one lag for kurtosis and hyper-kurtosis. The test statistics also reveal eight lags

Equation (2) represents the variance equation of exchange rate returns, which is specified as a function of news from previous periods (ε^2_{t-i}), the forecast variance from previous periods (h_{t-j}), and the conditional higher moments and their lagged values (HM_{t-k}).

Different methods have become popular for estimating conditional moments. Here, we use one of the most common definitions of the conditional skewness (S_t), kurtosis (K_t), hyper-skewness (HS_t), and hyper-kurtosis (HK_t):

$$S_t = \frac{E_{t-1}(r_t - \mu_t)^3}{h^{3/2}_{t-1}}, \quad K_t = \frac{E_{t-1}(r_t - \mu_t)^4}{h^2_{t-1}}, \quad HS_t = \frac{E_{t-1}(r_t - \mu_t)^5}{h^{5/2}_{t-1}},$$

$$HK_t = \frac{E_{t-1}(r_t - \mu_t)^6}{h^3_{t-1}}, \tag{3}$$

where r_t is the return at the time t, μ_t is the conditional mean, h_t is the variance at the time t, and E_{t-1} is the expected value at the time $t - 1$.[3]

4. Data and empirical findings

4.1. Data and preliminary results

We test our hypotheses based on the 15-minute exchange rate returns of the six most traded currencies, namely, the AUD, GBP, CAD, EUR, JPY, and CHF, vis-à-vis the USD. According to the Bank for International Settlements (2013) Triennial Survey, in the spot market, the USD was involved in 87% of transactions, followed by the EUR (33.4%), JPY (23%), GBP (11.8%), AUD (8.6%), CHF (5.2%), and CAD (4.6%). Overall, these six currencies were involved in approximately 87% of total transactions in the currency market.[4]

The data are sourced from TRTH and cover the period from January 1, 2004 (12:00 AM—GMT) to January 1, 2014 (11:45 PM—GMT), excluding weekends. This culminates into a total of 250,560 observations per currency.[5] A plot of the data is presented in Figure 1, and a summary of the data appears in Table 1. Several

for the lagged dependent variable and we therefore include this. Finally, p and q, according to AIC and BIC, are 1.

[3] We estimate the parameters using the GARCH model with Student's t error distribution and accordingly the conditional variance and residuals are extracted. Then, following the definitions as in Equation (3), we compute each of the conditional higher moments by dividing the estimated residual at time t by the variance at time $t - 1$ to the power of 3, 4, 5, and 6 for the skewness, kurtosis, hyper-skewness, and hyper-kurtosis, respectively. The Appendix presents a plot of these higher moments for each of the six exchange rate returns.

[4] Each transaction involves two currencies; therefore, percentages for all currencies should be divided by 2.

[5] In some cases, we have missing exchange rate data (less than 2% of the data per currency) at the 15-minute frequency. Our approach here is to use the linear interpolation technique to deal with missing values.

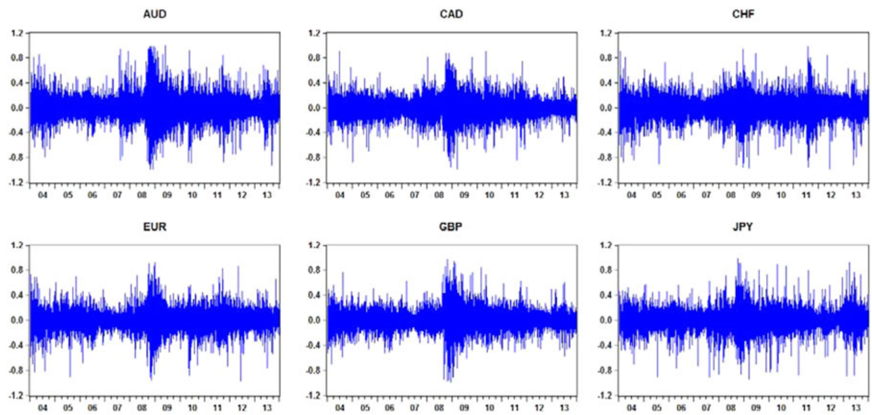

Figure 1

A plot of spot returns

This figure plots 15-minute spot returns of six currencies vis-à-vis the USD (namely, AUD, CAD, CHF, EUR, GBP, and JPY) over the period January 1, 2004 (12:00 AM—GMT) to January 1, 2014 (11:45 PM—GMT). Our data sample excludes weekends and we have a total of 250,560 observations per currency.

interesting features of currency returns are reflected in Figure 1. For example, some currencies (the EUR and GBP) are relatively stable whereas others are extremely volatile (the AUD). Moreover, the significant impact of the global financial crisis (GFC) can be observed in all currencies as they all experience strong patterns of volatility clustering during the GFC period.

Selected descriptive statistics of the currencies' returns are reported in Table 1. There are some points worth noting here. First, as mentioned earlier, while these six currencies are involved in approximately 87% of total transactions of the currency market, each of these currencies contain unique statistical characteristics that allow us to test our hypotheses on a data set containing unique currencies. For instance, during the period 2004–2014, we notice that all currencies (except the JPY) appreciated (negative mean) against the USD. The most volatile currency, based on the standard deviation of the exchange rate returns, turns out to be AUD, whereas the EUR and GBP are relatively less volatile than other currencies. In four of six currencies, the skewness is positive, suggesting that the chance that these currencies will further depreciate is high. By contrast, for the remaining two currencies (the CAD and JPY), the skewness is negative, implying that the chance of further appreciation is high. The large kurtosis values, for all currencies, is symptomatic of fat tails—a feature of high-frequency exchange rate returns (see Westerfield, 1977).

Second, the results of the Jarque-Bera test reveal that the null hypothesis of normality can be rejected at the 1% level of significance for all variables. This confirms that none of the returns are normally distributed. Table 1 also reports the

Table 1

Descriptive statistics of the returns

This table reports selected descriptive statistics for the exchange rate returns. These statistics include annualized mean returns, standard deviation, skewness, kurtosis, and the *p*-value of Jarque-Bera test, which tests the null hypothesis that exchange rate returns are normally distributed. Additionally, we also report the *t*-statistics from ADF unit root test model, which examines the null hypothesis of a "unit root" in the return series, followed by the first-order autoregression (AR(1)) coefficient and its *p*-value in parentheses. Finally, we also report results (*F*-statistics and its corresponding *p*-value in parentheses) obtained from the Lagrange multiplier test to examine the null hypothesis of "no ARCH" in return series.

	AUD	CAD	CHF	EUR	GBP	JPY
Panel A: 15-minute data						
Annualized mean	−3.994	−1.747	−0.749	−2.496	−2.995	5.990
Std dev	0.058	0.040	0.041	0.038	0.037	0.040
Skewness	0.087	−0.006	4.62E-04	0.029	0.019	−0.034
Kurtosis	6.699	4.769	4.809	4.993	5.213	4.743
JB *p*-value	0.000	0.000	0.000	0.000	0.000	0.000
ADF *t*-statistics	−173.35	−202.93	−201.69	−200.14	−199.63	−176.39
AR(1) coef.	0.40	0.35	0.36	0.39	0.38	0.38
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
ARCH *F*-statistics	30,412.97	19,346.09	19,614.70	21,800.74	23,187.19	19,179.28
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Panel B: 30-minute data						
Annualized mean	0.137	0.416	−0.416	−1.274	−0.824	2.647
Std dev	0.058	0.040	0.042	0.038	0.037	0.040
Skewness	0.103	3.68E-05	0.002	0.035	0.024	−0.034
Kurtosis	6.673	4.785	4.805	4.993	5.213	4.768
JB *p*-value	0.000	0.000	0.000	0.000	0.000	0.000
ADF *t*-statistics	−256.55	−211.94	−211.18	−209.91	−257.61	−212.00
AR(1) coef.	−0.023	−0.033	−0.03	−0.027	−0.028	−0.029
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
ARCH *F*-statistics	4,370.88	3,431.49	2,950.76	3,105.87	3,867.88	2,565.92
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Panel C: Hourly data						
Annualized mean	4.236	3.808	3.114	3.445	2.306	2.722
Std dev	0.058	0.040	0.042	0.039	0.038	0.041
Skewness	0.149	0.021	0.024	0.064	0.025	–0.025
Kurtosis	6.656	4.684	4.713	4.902	5.132	4.689
JB *p*-value	0.000	0.000	0.000	0.000	0.000	0.000
ADF *t*-statistics	−254.09	−256.14	−254.95	−253.86	−254.52	−253.66
AR(1) coef.	−0.015	−0.023	−0.019	−0.014	−0.017	−0.013
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
ARCH *F*-statistics	1,851.65	1,430.57	913.72	1,127.37	1,294.06	801.084
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)

results from the augmented Dickey-Fuller (ADF; Dickey and Fuller, 1979) unit root test, which examines the null hypothesis of a unit root against the alternative that the series is stationary. The test is implemented for a model with an intercept but no time trend, and the results reveal that all variables are stationary in their levels. Additionally, the first-order autoregressive coefficient, AR(1), suggests a persistent and serially correlated currency returns for all countries. Finally, the results from the Autoregressive Conditional Heteroskedasticity Lagrange Multiplier (ARCH-LM) test for heteroskedasticity indicate that the null hypothesis of "no ARCH" effect is easily rejected at the 1% level of significance. This is true for all currency returns.

To conclude, different characteristics of these most liquid currencies provide us with an opportunity to assess the proposed hypotheses on different currencies. Moreover, the preliminary results strongly suggest the existence of autocorrelation and heteroskedasticity in the regression model. Therefore, our approach of using a GARCH-type model, augmented with lagged values of returns, is ideal for testing our proposed hypotheses.

Before we test the proposed hypotheses though, we begin by understanding the correlation between the return (variance) and the higher moments of returns. As Table 2 reports, the third and fifth moments (i.e., skewness and hyper-skewness) show higher correlation with returns compared to the other moments. Likewise, when it comes to variance, the fourth and the sixth moments (i.e., kurtosis and hyper-kurtosis) show relatively higher correlation with variance of all currencies except for the AUD. Therefore, one may expect that the presence of skewness and hyper-skewness in the mean equation of the GARCH model and the presence of kurtosis and hyper-kurtosis in the variance equation of the GARCH model should have implications for the statistical performance of the models.

4.2. Hypotheses and findings

In this section, we test our two proposed hypotheses:

H1: The inclusion of high-order moments improves the forecasting performance of exchange rate models.

H2: The inclusion of odd/even moments in modeling currency return/variance improves the statistical performance of the model.

We first run our intraday regression model, Equation (1), for each of the selected currencies. In our sample, the results suggest that the GARCH model with Student's *t* error distribution is superior to the other GARCH-type models. Then, one at a time, we add each of the high-order moments (namely, conditional skewness, kurtosis, hyper-skewness, and hyper-kurtosis) with their lag values in both the mean and variance equations to gauge the optimal model using information criterion statistics.[6]

[6] We observe that both AIC and BIC show very similar statistics on the impact of the inclusion of higher moments in the mean and variance equations.

Table 2

Unconditional correlation between return/variance and higher moments

In this table, we report results on the unconditional correlation between exchange rate returns/volatility and higher moments. Panel A reports the unconditional correlation between return and higher moments (i.e., variance, skewness, kurtosis, hyper-skewness, and hyper-kurtosis) for each of the six currencies (i.e., AUD, CAD, CHF, EUR, GBP, and JPY). Panel B reports the unconditional correlation between variance and higher moments (i.e., third to sixth) of these currencies. *p*-Values are in parentheses.

Panel A: Unconditional correlation between return and higher moments

Return	Variance	Skewness	Kurtosis	Hyper-skew	Hyper-kurt
AUD	0.003	0.234	0.013	0.065	−0.010
	(0.192)	(0.000)	(0.000)	(0.000)	(0.000)
CAD	0.002	0.364	−0.002	0.114	0.003
	(0.328)	(0.000)	(0.272)	(0.000)	(0.112)
CHF	−0.003	0.343	−0.002	0.118	−0.004
	(0.155)	(0.000)	(0.405)	(0.000)	(0.044)
EUR	0.002	0.289	0.021	0.100	0.020
	(0.451)	(0.000)	(0.000)	(0.000)	(0.000)
GBP	−0.006	0.276	0.008	0.094	0.006
	(0.003)	(0.000)	(0.000)	(0.000)	(0.004)
JPY	0.003	0.359	−0.005	0.124	9.91E-05
	(0.183)	(0.000)	(0.007)	(0.000)	(0.960)

Panel B: Unconditional correlation between variance and higher moments

Variance	Skewness	Kurtosis	Hyper-skew	Hyper-kurt
AUD	−2.70E-04	1.89E-04	0.003	0.004
	(0.892)	(0.925)	(0.118)	(0.052)
CAD	0.001	−0.023	−0.001	−0.013
	(0.641)	(0.000)	(0.920)	(0.000)
CHF	−0.002	−0.021	−0.001	−0.013
	(0.277)	(0.000)	(0.496)	(0.000)
EUR	−0.004	−0.013	−0.002	−0.007
	(0.059)	(0.000)	(0.224)	(0.000)
GBP	−0.004	−0.011	−0.003	−0.007
	(0.032)	(0.000)	(0.200)	(0.000)
JPY	0.003	−0.016	0.002	−0.008
	(0.121)	(0.000)	(0.417)	(0.000)

We observe that intraday models with the third moment (skewness) included in the mean equation and the fourth and the sixth moments (kurtosis and hyper-kurtosis) included in the variance equation have significantly lower AIC and BIC statistics. Thus, our results suggest that the inclusion of higher moments in both the mean and variance equations significantly improves the statistical performance of the exchange rate models. More specifically, we find that the inclusion of odd moments (skewness) in the mean equation and even moments (kurtosis and hyper-kurtosis) in the variance equation improves the statistical performance of the models. According to the AIC and BIC statistics, for example, on average, the inclusion of higher moments improves the performance of the models by 13% compared to models that do not include these higher moments. Table 3 reports the performance of TG model and our proposed HM-GARCH model.

In short, our findings are consistent in that we find support for our first hypothesis—that is, the inclusion of higher moments improves the statistical performance of high-frequency currency return and volatility models. As reported in Table 3, our model performs best when we include the third moment (skewness—an odd moment) into the mean equation and when we include the fourth and sixth moments (i.e., kurtosis and hyper-kurtosis—even moments) into the variance equation. Therefore, we also conclude in support of our second hypothesis that odd/even moments explain the returns/variance better than the other moments. Our proposed GARCH-type regression model (with higher moments included) thus outperforms the standard GARCH model in estimating the currency returns and variance.

4.3. *Economic significance*

The goal of this section is to test whether our proposed model (HM-GARCH) has any economic significance for investors. More specifically, we test whether predicted returns obtained from our proposed model can be used in trading strategies to generate economically meaningful profits for investors in the currency market. We propose a higher moment-based trading strategy (HMTS) to test the profitability of exchange rate trading. First, based on forecasted returns from our HM-GARCH model, we compare the profitability of our HMTS with that of Anatolyev and Gerko's (2005) excess predictability (EP) trading strategy. Next, we compare the profitability of our HMTS based on the forecasted returns from the HM-GARCH model vis-à-vis the TG model.

To generate forecasts of exchange rate returns, we use nine sets of IS and OOS periods that represent the following combinations (in percentage) of the IS/OOS periods: 10/90, 20/80, 30/70, 40/60, 50/50, 60/40, 70/30, 80/20, and 90/10. For instance, in the case of 90% IS and 10% OOS, we use the first nine years of our data as an IS estimation window (2004–2013) to obtain a series of one-period ahead forecasts for the final year of our data sample (2013–2014). This approach of considering multiple IS and OOS periods is important for two reasons. First, it allows us to investigate whether the choice of the estimation and forecast sample periods has any impact on the profitability of the trading strategy. Second, this approach is needed because in the forecasting literature, studies have used a range of IS and OOS periods with the 50/50 split been the most popular (see, among others, Choi, Hauser and Kopecky, 1999; Olson, 2004; Campbell and Thompson, 2008; Narayan, Sharma and Thuraisamy, 2014). With this approach, therefore, we keep within the spirit of this literature on forecasting asset price returns.

Table 3

Inclusion of higher moments and the performance of the model

This table compares the statistical performance of the TG model with our proposed HM-GARCH model. In order to select the best model between TG and HM-GARCH, here we report values on Akaike information criterion (AIC), Bayesian information criterion (BIC), and the log-likelihood. In the last row of the table, we also report the average improvement of the models (after including the higher moments—HM-GARCH) using AIC and BIC.

	AUD	CAD	CHF	EUR	GBP	JPY
Panel A: Traditional GARCH model (TG)						
AIC	−3.444	−4.001	−3.932	−4.182	−4.279	−4.071
BIC	−3.444	−4.001	−3.931	−4.182	−4.279	−4.071
Likelihood	388,350	451,130	443,340	471,578	482,478	459,049
*Panel B: GARCH model with higher moments included (HM-GARCH)—*t *distribution*						
AIC	−3.842	−4.587	−4.555	−4.731	−4.739	−4.617
BIC	−3.842	−4.586	−4.555	−4.730	−4.738	−4.616
Likelihood	433,240	517,152	513,639	533,465	534,337	520,549
Avg. improvement	12%	15%	16%	13%	11%	13%
Panel C: GARCH model with higher moments included (HM-GARCH)—normal distribution						
AIC	−3.868	−4.478	−4.422	−4.647	−3.906	−4.774
BIC	−3.868	−4.477	−4.421	−4.646	−3.905	−4.774
Likelihood	436,181	504,918	498,623	523,982	440,381	538,313
Avg. improvement	12%	12%	12%	11%	−9%	17%
Panel D: GARCH model with higher moments included (HM-GARCH)—GE distribution						
AIC	−2.713	−4.541	−4.396	−4.596	−4.837	−4.678
BIC	−2.712	−4.540	−4.396	−4.595	−4.836	−4.678
Likelihood	305,917	511,965	495,707	518,234	545,410	527,505
Avg. improvement	−21%	13%	12%	10%	13%	15%
*Panel E: EGARCH model with higher moments included (HM-EGARCH)—*t *distribution*						
AIC	−3.307	−4.237	−4.160	−4.282	−4.450	−4.291
BIC	−3.306	−4.236	−4.159	−4.281	−4.449	−4.291
Likelihood	372,863	477,757	469,072	482,807	501,753	483,868
Avg. improvement	−4%	6%	6%	2%	4%	5%
*Panel F: Power-GARCH model with higher moments included (HM-Power-GARCH)—*t *distribution*						
AIC	−3.836	−4.571	−4.511	−4.727	−4.835	−4.661
BIC	−3.835	−4.571	−4.510	−4.727	−4.834	−4.660
Likelihood	432,494	515,445	508,590	533,033	545,130	525,536
Avg. improvement	11%	14%	15%	13%	13%	14%
*Panel G: GARCH model with higher moments (only include hyper-skewness and hyper-kurtosis), (HM-GARCH)—*t *distribution*						
AIC	−3.826	−4.289	−4.078	−4.377	−4.685	−4.360
BIC	−3.825	−4.288	−4.077	−4.377	−4.684	−4.359

(*Continued*)

Table 3 (*Continued*)

Inclusion of higher moments and the performance of the model

*Panel G: GARCH model with higher moments (only include hyper-skewness and hyper-kurtosis), (HM-GARCH)—*t *distribution*						
Likelihood	431,408	483,624	459,818	493,550	528,261	491,602
Avg. improvement	11%	7%	4%	5%	9%	7%
*Panel H: GARCH model with higher moments included (HM-GARCH) in both mean and variance equations—*t *distribution*						
AIC	−3.676	−4.272	−4.250	−4.419	−4.523	−4.219
BIC	−3.675	−4.271	−4.249	−4.418	−4.522	−4.219
Likelihood	414,488	481,696	479,179	498,200	509,946	475,719
Avg. improvement	7%	7%	8%	6%	6%	4%

We begin with forecasting models, namely, the AR(1) model, Equation (4), and our proposed HM-GARCH model, Equation (5). The models take the following forms:

$$R_t = \alpha + \beta_1 R_{t-1} + \varepsilon_t, \tag{4}$$

$$R_t = \alpha + \beta_1 R_{t-1} + \sum_{j=0}^{n} \gamma_j (HM_{t-j}) + \varepsilon_t. \tag{5}$$

The variables are as previously defined in Equation (1); R_{t-1} is the one-period lag of the exchange rate returns, and *HM* represents the high-order moments.

Following Anatolyev and Gerko's (2005) EP trading strategy, we use the sign of the forecasts as signals for buying and selling. That is, investors go long if the prediction of the next period's return is greater than or equal to 0 and go short otherwise. However, one of the main drawbacks of using directional trading rules, such as the EP trading strategy, is that we only use the sign of the forecasts as a trading signal and the numerical values of the forecasts are not involved in making trading decisions. Therefore, we propose the following modified trading strategy, based on the EP directional trading rule. We first obtain the forecasts of the currency returns from Equations (4) and (5). Then, for those positive forecasts that are numerically greater than or equal to the current actual return (i.e., $\hat{r}_t \geq r_t$), we go long if one of the following conditions are met:[7]

(i) The current cumulative actual return (from the beginning of the sample) becomes less than or equal to the cumulative return of the OOS forecasts. The cumulative return over the OOS forecast shows the future direction of

[7] In other words, we go short when we have a significantly larger negative value for either the predicted or current actual return.

returns, therefore, we buy as long as we expect the future prices (the OOS period) to rise.

(ii) The current actual return becomes less than or equal to the cumulative return of the n-period ahead forecasts.[8] The logic behind this condition is similar to that behind condition (i), and it accounts for short-run expectation of the future prices.

(iii) The current actual return (r_t) becomes greater than the previous period's actual return (r_{t-1}).[9] This condition is simply applied to improve the profitability of our trading strategy by using the last known information in the market only when the basic condition (i.e., $\hat{r}_t + r_t \geq 0$) is met.

Let $Long_t$ be an indicator of the trading position at time t, where r_t is the actual return, $\hat{r}_t$ is the forecast of the return at time t, and n is the number of OOS periods.

$$Long_t = 1 \text{ if } \hat{r}_t \geq r_t \text{ AND } \left(\sum_{t=1}^{n} r_t \leq \sum_{n} \hat{r}_t OR\, r_t \leq \sum_{i=1}^{96} \hat{r}_{t+i} OR\, r_t > r_{t-1} \right).$$
$$Long_t = 0 \text{ otherwise.} \tag{6}$$

The logic behind these restrictions is that they reduce the transaction cost significantly by reducing the chance of changing trading positions, which improves the profitability of the trading strategy. In addition, they allow us to focus more on the predictability performance of our model by including the numerical values of the forecasts in the trading position signals rather than using only the sign of the forecasts.

Transaction costs are an important part of judging the success or otherwise of trading strategies. There are, however, different views in the literature on what exactly should be the transaction cost. Earlier studies, such as Szakmary and Mathur (1997), use a 10-basis point transaction cost. Many recent studies suggest that the realized transaction costs are actually lower than 0.1%. Neely, Weller and Dittmar (1997), and Hsu and Kuan (2005) use five basis points to represent the transaction cost. On the other hand, Neely and Weller (2003) use two basis points, while Neely and Weller (2013) use only one basis point as the transaction cost in the case of intraday currency market trading. In this study, we allow the transaction cost of 10 basis points (0.1%) each time a new position (i.e., long or short) is established. By considering a higher transaction cost, we provide a relatively more stringent test of the economic

[8] We use different periods for "n" such as 4, 32, and 96 periods ahead that correspond to one-hour, eight-hour, and one-day periods ahead, respectively. Here, we only report the 96-period (one-day) ahead forecasts.

[9] We assign the $r_t > r_{t-1}$ and not the $r_t \geq r_{t-1}$ constraint to reduce the chance of changing the trading position when there is no change in the value of the exchange rate return.

significance of our proposed trading strategy. It follows that the adjusted returns, or profits (P_t), are computed as follows:

$$P_t = Long_t r_t + (Long_t - 1)r_t - 0.001\,|Long_t - Long_{t-1}|\,. \tag{7}$$

The first two terms on the right-hand side of the equation represent the raw returns when an investor takes either a long or a short position in the market. The final term in the equation accounts for the transaction cost (0.1%) that is incurred whenever a new position is taken in the market.

Table 4 compares the annualized net profits of our proposed HMTS with those of Anatolyev and Gerko's (2005) EP trading strategy. We present the annualized profits of our proposed strategy for each of the six currencies and for the nine sets of OOS periods. Then, we compare the profitability of our HMTS with that of the EP trading strategy. Based on the HMTS results, we find that the AUD, CHF, and JPY show the highest annualized returns followed by the CAD, EUR, and GBP. Second, we observe that for all currencies, the annualized returns increase as the size of the IS (OOS) period decreases (increases) until the IS period is 40% and the OOS period is 60%.[10] The annualized returns decline thereafter.[11] Finally, we notice that the HMTS offers significantly higher profits for investors (up to 152% for the JPY) than the EP trading strategy.[12] More specifically, using the HMTS, 152%, 130%, 127%, 95%, 85%, and 78% higher profits are achieved when trading the JPY, CAD, CHF, EUR, GBP, and AUD, respectively, in comparison to the EP trading strategy.

Next, we apply the HMTS to test whether the predicted returns obtained from our proposed model (HM-GARCH) generate more profits than the forecasted returns obtained from the TG model.

Table 5 reports the annualized net profits—using the HM-GARCH and TG models. In addition, the Sharpe ratio is used here as a measure of comparative profitability.[13] We present the annualized profits and the Sharpe ratios of both models for each of the six currencies based on each combination of IS and OOS periods. The results reveal the following messages. First, we find that in terms of profitability, the HM-GARCH model outperforms the TG model in 74% of the 54 forecasting models (i.e., nine sets of sample for six currencies) and offers up to 20% higher profits for investors (i.e., JPY horizon 9 [h9]—10% IS and 90% OOS). By contrast, in the remaining cases, the TG model offers 3% higher profits (i.e., CHF horizon 8

[10] This corresponds to the 2004–2008 IS period and 2008–2014 OOS period that contains the largest IS period before the 2008 GFC.

[11] We also notice a very similar trend in the root mean squared error (RMSE) of the forecasts. They increase as the size of the IS (OOS) period decreases (increases) until they reach their maximum level at a 40% IS and 60% OOS period, and they decline afterward. This tells us something about the importance of the sample sizes that are required to generate performance-based statistics.

[12] We do not report net profits obtained from EP trading strategy in Table 4. However, the detailed results on EP trading strategy are available on request.

[13] It is the proportion of the average return of the trading strategy over its standard deviation.

Table 4

Annualized net profits of the HMTS and EP trading strategies

This table report results on out-of-sample predictive evaluation for nine sets of OOS periods (90-10, 80-20, 70-30, 60-40, 50-50, 40-60, 30-70, 20-80, and 10-90), which compares profit obtained from HMTS and EP trading strategy. We report results on Theil U and root mean squared error (RMSE) statistics for each of the nine OOS horizons. Additionally, for each horizon, we report annualized net profits (%) using the HMTS. To compute profits, we consider a transaction cost of 0.1%. Profits are all significantly different from 0 at the 1% level of significance. Profits obtained from the EP trading strategy are not reported here but detailed results are available on request from the corresponding author. The final three rows of this table report the average net profits from our proposed trading rule, HMTS, the average net profit using the EP trading strategy, and the percentage of improvement in profitability when we use HMTS rather than the EP trading strategy, respectively.

	AUD	CAD	CHF	EUR	GBP	JPY
h1: 90% in-sample and 10% out-of-sample						
Theil U	0.567	0.638	0.638	0.598	0.591	0.616
RMSE	0.040	0.025	0.031	0.027	0.026	0.037
Net profit	6.749	3.967	5.244	4.922	4.139	6.460
h2: 80% in-sample and 20% out-of-sample						
Theil U	0.596	0.635	0.645	0.602	0.597	0.618
RMSE	0.037	0.026	0.032	0.029	0.025	0.032
Net profit	6.266	4.519	5.396	4.602	4.036	5.749
h3: 70% in-sample and 30% out-of-sample						
Theil U	0.574	0.639	0.644	0.605	0.601	0.614
RMSE	0.045	0.030	0.036	0.033	0.028	0.035
Net profit	8.345	5.689	6.141	5.783	4.571	6.125
h4: 60% in-sample and 40% out-of-sample						
Theil U	0.606	0.647	0.633	0.616	0.612	0.620
RMSE	0.045	0.033	0.038	0.034	0.031	0.034
Net profit	7.765	6.304	6.504	5.749	5.107	6.329
h5: 50% in-sample and 50% out-of-sample						
Theil U	0.591	0.654	0.645	0.619	0.615	0.615
RMSE	0.054	0.037	0.039	0.036	0.035	0.036
Net profit	10.054	6.880	7.196	6.709	5.735	6.230
h6: 40% in-sample and 60% out-of-sample						
Theil U	0.589	0.646	0.632	0.616	0.604	0.610
RMSE	0.058	0.038	0.041	0.037	0.036	0.038
Net profit	10.298	7.117	7.710	6.826	6.490	7.178
h7: 30% in-sample and 70% out-of-sample						
Theil U	0.602	0.647	0.651	0.616	0.605	0.608
RMSE	0.054	0.038	0.039	0.036	0.035	0.038
Net profit	8.903	7.088	7.232	6.350	6.178	7.079
h8: 20% in-sample and 80% out-of-sample						
Theil U	0.598	0.652	0.657	0.616	0.602	0.610
RMSE	0.055	0.037	0.038	0.035	0.034	0.037
Net profit	8.971	6.941	6.961	6.449	5.445	6.945
h9: 10% in-sample and 90% out-of-sample						
Theil U	0.604	0.664	0.666	0.621	0.606	0.616
RMSE	0.053	0.037	0.038	0.034	0.034	0.036
Net profit	8.269	6.915	6.752	6.223	6.252	6.848
NTS avg. net profit	8.402	6.158	6.571	5.957	5.328	6.549
EP avg. net profit	4.728	2.672	2.891	3.062	2.888	2.595
Improvement (%)	78	130	127	95	85	152

Table 5

Annualized net profits from HM and TG models (transaction cost 0.1%)

In this table, we report annualized net profits using our proposed HMTS. Here, we use both TG and HM-GARCH models to forecast returns for nine sets of OOS periods (90-10, 80-20, 70-30, 60-40, 50-50, 40-60, 30-70, 20-80, and 10-90). Profits reported in this table are all statistically significant at the 1% level. Additionally, we also report the Sharpe ratios in square brackets. Finally, we conduct the Diebold-Mariano (DM) test and report its coefficient and the corresponding *p*-values in parentheses.

	AUD	CAD	CHF	EUR	GBP	JPY
h1: 90% in-sample and 10% out-of-sample						
HM profit	6.749	3.967	5.244	4.922	4.139	6.460
	[0.736]	[0.702]	[0.739]	[0.828]	[0.671]	[0.759]
TG profit	6.266	3.793	4.817	4.895	4.086	6.404
	[0.661]	[0.658]	[0.652]	[0.821]	[0.659]	[0.749]
DM	−12.009	−41.146	−21.830	−28.205	−23.190	−35.513
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
h2: 80% in-sample and 20% out-of-sample						
HM profit	6.266	4.519	5.396	4.602	4.036	5.749
	[0.699]	[0.801]	[0.758]	[0.675]	[0.672]	[0.787]
TG profit	6.087	4.241	5.075	4.541	3.970	5.827
	[0.670]	[0.725]	[0.691]	[0.662]	[0.657]	[0.804]
DM	24.697	−38.684	−42.407	−30.267	−25.752	45.371
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
h3: 70% in-sample and 30% out-of-sample						
HM profit	8.345	5.689	6.141	5.783	4.571	6.125
	[0.871]	[0.893]	[0.763]	[0.780]	[0.684]	[0.892]
TG profit	7.056	5.698	5.873	5.256	4.501	6.135
	[0.673]	[0.896]	[0.713]	[0.676]	[0.669]	[0.895]
DM	−12.769	−36.296	−18.768	−25.974	−21.923	−38.229
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
h4: 60% in-sample and 40% out-of-sample						
HM profit	7.765	6.304	6.504	5.749	5.107	6.329
	[0.701]	[0.911]	[0.782]	[0.713]	[0.700]	[0.886]
TG profit	7.612	6.307	6.235	5.609	4.997	6.397
	[0.681]	[0.912]	[0.733]	[0.688]	[0.678]	[0.903]
DM	−16.971	−40.198	−25.757	−32.646	−25.594	−45.755
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
h5: 50% in-sample and 50% out-of-sample						
HM profit	10.054	6.880	7.196	6.709	5.735	6.230
	[0.863]	[0.896]	[0.871]	[0.851]	[0.711]	[0.773]
TG profit	8.607	6.969	7.085	5.911	5.570	6.193
	[0.680]	[0.917]	[0.848]	[0.698]	[0.681]	[0.765]
DM	−7.534	−43.370	−26.253	−32.883	−22.419	−34.946
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
h6: 40% in-sample and 60% out-of-sample						
HM profit	10.298	7.117	7.710	6.826	6.490	7.178
	[0.804]	[0.891]	[0.925]	[0.840]	[0.822]	[0.894]

(*Continued*)

Table 5 (*Continued*)

Annualized net profits from HM and TG models (transaction cost 0.1%)

	AUD	CAD	CHF	EUR	GBP	JPY
TG profit	9.109	7.225	7.650	6.066	5.717	7.119
	[0.671]	[0.915]	[0.912]	[0.699]	[0.677]	[0.882]
DM	−16.574	−6.584	1.785	−24.000	−18.795	−51.809
	(0.000)	(0.000)	(0.074)	(0.000)	(0.000)	(0.000)
h7: 30% in-sample and 70% out-of-sample						
HM profit	8.903	7.088	7.232	6.350	6.178	7.079
	[0.675]	[0.908]	[0.885]	[0.805]	[0.804]	[0.899]
TG profit	8.866	7.118	7.275	6.494	5.494	7.047
	[0.671]	[0.914]	[0.894]	[0.835]	[0.673]	[0.892]
DM	−1.105	−37.311	−28.782	−19.651	−30.890	−57.966
	(0.269)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
h8: 20% in-sample and 80% out-of-sample						
HM profit	8.971	6.941	6.961	6.449	5.445	6.945
	[0.719]	[0.907]	[0.862]	[0.861]	[0.687]	[0.896]
TG profit	8.505	6.970	7.189	6.537	5.371	6.898
	[0.666]	[0.914]	[0.912]	[0.882]	[0.674]	[0.885]
DM	−21.104	14.998	−45.522	−33.880	−26.431	−31.816
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
h9: 10% in-sample and 90% out-of-sample						
HM profit	8.269	6.915	6.752	6.223	6.252	6.848
	[0.665]	[0.915]	[0.828]	[0.824]	[0.869]	[0.891]
TG profit	8.135	6.940	6.424	5.520	6.233	5.722
	[0.650]	[0.921]	[0.765]	[0.685]	[0.865]	[0.674]
DM	−10.296	−45.508	−31.718	−30.773	−24.575	−38.421
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)

[h8]—20% IS and 80% OOS). Second, the HM-GARCH model contains a higher Sharpe ratio (shown in square brackets) in 74% of the models, with a value up to 32.2% higher than the TG model (i.e., JPY horizon 9 [h9]—10% IS and 90% OOS). By contrast, in the remaining cases, the TG model shows a maximum ratio of 5.5% (i.e., CHF horizon 8 [h8]—20% IS and 80% OOS). Third, the results of the Diebold and Mariano (DM; 1995) test imply that the null hypothesis of equal predictive accuracy is rejected at the 1% level of significance in 96% of the forecasting models.[14] In other words, the observed differences in performance between the two forecasting models are statistically different from 0. According to the DM statistics (based on the absolute-error loss[15]), the forecasting error of the TG model is significantly

[14] Following Harvey, Leybourne and Newbold (1997), we compare the statistic with the Student's t distribution rather than the standard normal distribution.

[15] We also evaluate the forecasting performance of the models based on the squared error loss and find similar results.

higher than that of the HM-GARCH model in 93% of the models, suggesting that the HM-GARCH model significantly outperforms the TG model in terms of the forecasting accuracy.

In short, we conclude that (1) the HM-GARCH model outperforms the TG model in terms of forecasting accuracy and offers the highest returns to investors, (2) the annualized returns increase as the size of the IS (OOS) period decreases (increases) until the IS period is 40% and the OOS period is 60%, and (3) our findings regarding profitability (economic significance) are generally consistent with our statistical evidence of predictability.

4.4. Robustness tests

4.4.1. Are results sensitive to the GFC?

In this section, we aim to confirm the robustness of our findings by checking its sensitivity with respect to the GFC. The GFC marked an important phase in the currency market. As argued in Melvin and Taylor (2009), for instance, the onset of the GFC instigated massive unwinding of carry trade. This led to massive losses for currency market investors. Due to this loss, August 2007 is officially associated as the start of currency market crisis (see Melvin and Taylor, 2009). The currency crisis contributed to currency market volatility that brings into the picture the role of higher order moments. Given this, it is imperative to account for the GFC in testing our hypotheses. We test whether the superior performance of our proposed model holds when we apply it to different sample periods. To achieve this goal, we conduct similar analyses with three subsamples (i.e., pre-GFC, GFC, and post-GFC) of data.

There are many possible criteria for identifying the GFC's timeline. Melvin and Taylor (2009) suggest the early summer (July and August) of 2007, as the beginning of the crisis period in equity and currency markets. In addition, Frankel and Saravelos (2012) study different crisis measures, such as percentage change in nominal currency, equity market returns, and real gross domestic product. Based on their indicators, they declare the second quarter of 2009 as the start of the recovery period (the end of financial crisis).[16] Therefore, in this study, we define our three subsample periods as follows: the pre-GFC sample covers the period from January 1, 2004 to June 30, 2007; the GFC sample covers the period from July 1, 2007 to May 31, 2009; and the post-GFC sample covers the period from June 1, 2009 to January 1, 2014. Table 6 compares the statistical performance (i.e., according to both the AIC and BIC) and forecasting accuracy between the TG model and our proposed HM-GARCH model. Consistent with our findings from the full sample period of data, we find that over all three subsample periods (i.e., pre-GFC, GFC, and post-GFC), the

[16] This is consistent with what the National Bureau of Economic Research reported in September 2010 (http://www.nber.org/cycles/sept2010.html), determining that the recession ended and the recovery began in June 2009.

Table 6

Performance of the models under different subsamples

In this table, we compare the performance of TG model with HM-GARCH model by dividing data into three subsample periods. These three subsample are pre-GFC (January 1, 2004 to June 30, 2007), the GFC (July 1, 2007 to May 31, 2009), and the post-GFC (June 1, 2009 to January 1, 2014). To facilitate the comparison between TG (see Panel A) and HM-GARCH (see Panel B) based models, here we report AIC and BIC statistics for each of the six currencies over each of the three subsample periods. Finally, in Panel C we report the DM test results to examine the null hypothesis of equal predictive accuracy over the three subsample periods. Here, p-values are reported in parentheses.

	Pre-GFC (January 2004 to June 2007)		GFC (July 2007 to May 2009)		Post-GFC (June 2009 to January 2014)	
	AIC	BIC	AIC	BIC	AIC	BIC
Panel A: Traditional GARCH model (TG)						
AUD	−3.775	−3.774	−2.882	−2.879	−3.604	−3.603
CAD	−4.251	−4.249	−3.522	−3.519	−4.185	−4.184
CHF	−4.170	−4.169	−3.673	−3.671	−3.953	−3.951
EUR	−4.457	−4.456	−3.917	−3.914	−4.220	−4.219
GBP	−4.540	−4.539	−3.825	−3.823	−4.394	−4.393
JPY	−4.261	−4.260	−3.509	−3.506	−4.137	−4.136
Panel B: GARCH model with higher moments included (HM-GARCH)						
AUD	−4.299	−4.297	−3.011	−3.008	−4.052	−4.051
CAD	−4.688	−4.686	−3.743	−3.741	−4.563	−4.562
CHF	−4.596	−4.594	−4.056	−4.054	−4.367	−4.366
EUR	−4.869	−4.868	−4.238	−4.235	−4.588	−4.587
GBP	−4.942	−4.940	−4.125	−4.122	−4.736	−4.734
JPY	−4.757	−4.755	−3.939	−3.936	−4.611	−4.610

Panel C: Forecasting accuracy of TG and HM-GARCH under three subsamples

		DM	
AUD	−21.919	−3.356	−18.348
	(0.000)	(0.000)	(0.000)
CAD	−23.749	−12.822	2.034
	(0.000)	(0.000)	(0.000)
CHF	−19.125	−15.009	−21.428
	(0.000)	(0.000)	(0.000)
EUR	−15.115	−15.328	−21.439
	(0.000)	(0.000)	(0.000)
GBP	−15.879	−7.124	−20.674
	(0.000)	(0.000)	(0.000)
JPY	−15.734	−16.859	−21.644
	(0.000)	(0.000)	(0.000)

HM-GARCH model outperforms the TG model. Results presented in Panels A and B of Table 6 imply that, in all cases, the AIC and BIC significantly improve when we include higher moments in our model.

In addition, the DM test results (Panel C) reveal that the null hypothesis of equal predictive accuracy is rejected at the 1% level of significance in all cases. In other words, the observed differences in the performance between the two forecasting models are significant. According to the DM statistics (based on the absolute error loss), the forecasting error of the TG model is significantly higher than that of the HM-GARCH model in 94% of the models, suggesting that regardless of the sample periods considered, the HM-GARCH model significantly outperforms the TG model in terms of forecasting accuracy.[17] The main implication of these results is that our key findings are insensitive to the use of different sample periods, implying that our findings—both statistical and trading strategy based—are robust.

4.4.2. *Are trading strategy profits sensitive to the GFC?*

Likewise with the analysis in Section 4.4.1., it is possible that profits are sensitive to the financial crisis. This is almost expected given the crisis. However, the key question here whether the profitability of our HMTS holds. Our findings (untabulated) suggest two things. First, as expected profits are lowest during the crisis period and highest in the pre-crisis period. Second, that regardless of the subsample period considered profits tend to maximize with HMTS.

4.4.3. *Are results sensitive to different assumptions about error distribution?*

Our earlier results assumed that the GARCH errors were Student's *t* distributed. We check whether results are sensitive to a different error distribution assumption. We report results from the normal and generalized error distributions (GED) and report results in Table 3 (Panels C and D), respectively. The results are consistent in the sense they suggest that the HM-GARCH model (regardless of the error distribution properties) beats the TG model.

4.4.4. *Are results sensitive to different model specifications?*

In the next set of robustness test, we check whether a model different to GARCH leads to the same conclusions about the role of higher order moments. In this regard, we use an EGARCH model (HM-EGARCH) and a Power-GARCH model

[17] We also evaluate the forecasting performance of the models based on the squared error loss and find similar results.

Table 7

Model performance based on different data frequencies

This table compares the statistical performance of the TG model with our proposed HM-GARCH model. The lag lengths used in Equations (1) and (2) according to lag length selection criterion are two lags for skewness and hyper-skewness, one lag for kurtosis and hyper-kurtosis, and one lag for ARCH and GARCH processes. For the model based on 30-minute (hourly) data, four (two) lags of the dependent variable are used. To select the best model between TG and HM-GARCH, here we report values on Akaike information criterion (AIC), Bayesian information criterion (BIC), and the log-likelihood.

	AUD	CAD	CHF	EUR	GBP	JPY
Panel A: Traditional GARCH model (TG)—30-minute data—t distribution						
AIC	−3.137	−3.756	−3.680	−3.861	−3.943	−3.793
BIC	−3.135	−3.754	−3.680	−3.860	−3.942	−3.792
Likelihood	176,837	211,725	207,484	217,641	222,264	213,820
Panel B: Traditional GARCH model (TG)—hourly data—t distribution						
AIC	−3.112	−3.673	−3.597	−3.787	−3.856	−3.728
BIC	−3.111	−3.672	−3.596	−3.786	−3.855	−3.726
Likelihood	87,730	103,550	101,409	106,755	108,686	105,077
Panel C: GARCH model with higher moments included (HM-GARCH)—30-minute data—t distribution						
AIC	−3.395	−4.172	−4.104	−4.119	−4.258	−4.255
BIC	−3.394	−4.171	−4.103	−4.118	−4.257	−4.254
Likelihood	191,399	235,234	231,393	232,211	240,046	239,904
Panel D: GARCH model with higher moments included (HM-GARCH)—hourly data—t distribution						
AIC	−3.592	−3.945	−4.066	−3.799	−3.988	−3.877
BIC	−3.590	−3.942	−4.063	−3.797	−3.986	−3.875
Likelihood	101,251	111,197	114,607	107,093	112,431	109,289

(HM-Power-GARCH). The results are reported in Panels E and F of Table 3. Overall, the results confirm that the higher order moments regardless of the type of GARCH model employed beats the TG model.

So far in testing the performance of the higher order moments, we included all higher order moments in the model. We check whether only including the two highest order moments—namely, hyper-skewness and hyper-kurtosis—in the GARCH model would beat the TG model. The results reported in Table 3 (Panel G) again confirm that a hyper-kurtosis and hyper-skewness based GARCH model comfortably beats the TG model.

In our results so far, we included the higher order moments only in the mean equation. We now include all moments in both the mean and the variance equations of the GARCH model and test how well a HM-GARCH model performs vis-à-vis a TG model. The results reported in Panel H confirm that the HM-GARCH model comfortably beats the TG model.

4.4.5. Are results sensitive to different data frequencies?

Our next set of robustness test deals with data frequency. So far, we have utilized 15-minute data and showed that higher order moment-based GARCH type models beat the TG model. We check whether the role of higher order moments is sensitive to different data frequencies. In this regard, we use 30-minute data and hourly data. A summary statistic of these two data frequencies is presented in Table 1 (see Panels B and C).

The results are reported in Table 7. Panels A and B contain results from the TG model while Panels C and D contain results from the HM-GARCH model. Our main finding is that the test statistics are minimized when using the HM-GARCH model compared to the TG model. This is true regardless of whether we use 30-minute or hourly data. The main conclusion, therefore, is that the importance of higher order moments holds regardless of data frequency.

5. Concluding remarks

In this paper, we examine the role of high-order moments of exchange rate returns, not previously considered in this literature, in forecasting exchange rate returns. The proposed HM-GARCH model's performance is compared with the TG model of exchange rate returns. An extensive IS and OOS forecasting analysis applied to 15-minute returns of the six most liquid currencies (i.e., the AUD, CAD, CHF, EUR, GBP, and JPY) against the USD over the period 2004–2014 reveals two new insights on exchange rate behavior. First, the inclusion of higher moments—odd (even) moments in modeling returns (variance) improves the statistical performance of the HM-GARCH model in forecasting exchange rate returns compared to a TG model that simply ignores the high-order moments. Second, we devise a range of trading strategies that generate buy and sell trading signals based on forecasted exchange rate returns from the HM-GARCH and TG models. We find that profits are maximized when trading signals are based on exchange rate return forecasts generated using the HM-GARCH models. Both results, we show, are robust to different samples of data. We test the robustness of our results along multiple fronts and find that the performance of the HM-GARCH model holds regardless of (1) the GFC, (2) the assumptions about GARCH errors, (3) model specifications, and (4) data frequencies.

Appendix

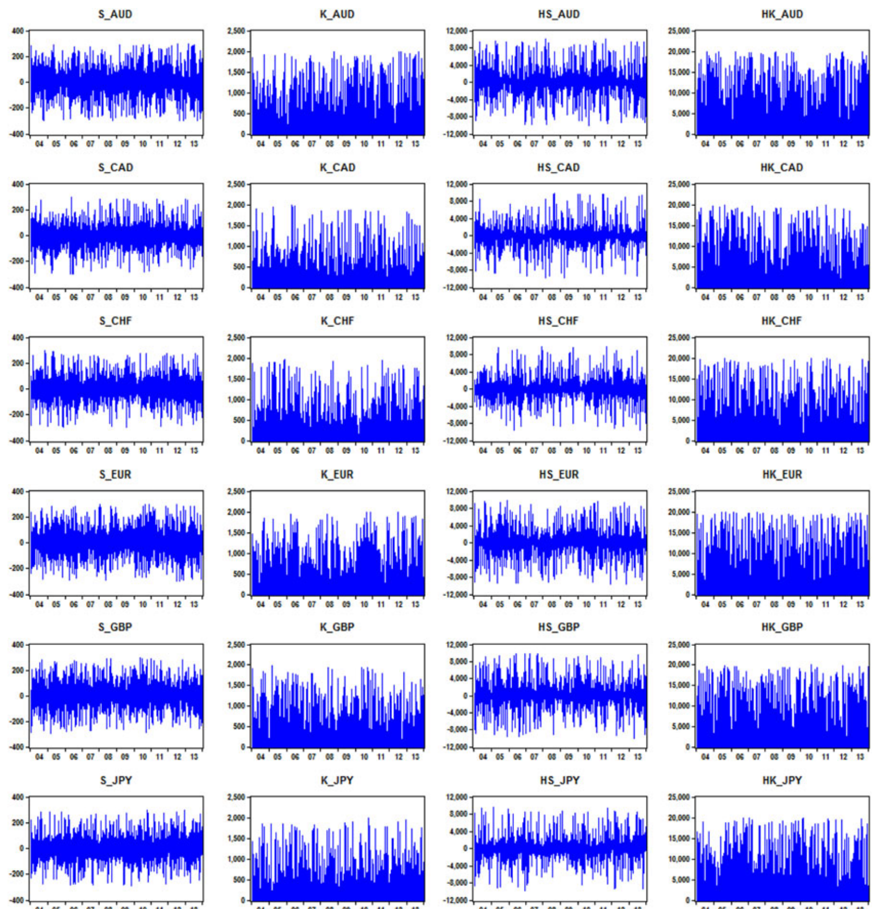

Figure A1

A plot of higher moments

This figure presents a plot of the higher moments (i.e., skewness [S], kurtosis [K], hyper-skewness [HS], and hyper-kurtosis [HK]) for each of the six currencies. We estimate the parameters using the GARCH model with Student's *t* error distribution to extract the conditional variance and residuals. Then following the definitions (see Equation [3]), we compute each of the conditional higher moments by dividing the estimated residual at time *t* by the variance at time *t* – 1 to the power of 3, 4, 5, and 6 for the skewness, kurtosis, hyper-skewness, and hyper-kurtosis, respectively.

References

Aggarwal, R., 1990. Distribution of spot and forward exchange rates: Empirical evidence and investor valuation of skewness and kurtosis, *Decision Sciences* 21, 588–595.

Anatolyev, S. and A. Gerko, 2005. A trading approach to testing for predictability, *Journal of Business and Economic Statistics* 23, 455–461.

Andersen, T.G., T. Bollerslev, and F.X. Diebold, 2007. Roughing it up: Including jump components in the measurement, modeling and forecasting of return volatility, *Review of Economics and Statistics* 89, 701–720.

Athayde, G. and R. Flores, 2006. On certain geometric aspects of portfolio optimization with higher moments, in: E. Jurczenko and B. Maillet, eds., *Multi-Moment Asset Allocation and Pricing Models* (John Wiley & Sons, Hoboken, NJ).

Bank for International Settlements, 2013. Triennial Central Bank Survey: Foreign exchange and derivatives market activity in 2013, Basel, Switzerland.

Brooks, C., S.P. Burke, S. Heravi, and G. Persand, 2005. Autoregressive conditional kurtosis, *Journal of Financial Econometrics* 3, 399–421.

Brooks, C., A. Cerny, and J. Miffre, 2012. Optimal hedging with higher moments, *Journal of Futures Markets* 32, 909–944.

Brunnermeier, M.K., S. Nagel, and L.H. Pedersen, 2008. Carry trades and currency crashes, *National Bureau of Economic Research Macroeconomic Annual* 23, 313–347.

Brunnermeier, M.K. and L.H. Pedersen, 2009. Market liquidity and funding liquidity, *Review of Financial Studies* 22, 2201–2238.

Campbell, J.Y. and S.B. Thompson, 2008. Predicting excess stock returns out of sample: Can anything beat the historical average, *Review of Financial Studies* 21, 1509–1531.

Chevallier, J., 2011. Detecting instability in the volatility of carbon prices, *Energy Economics* 33, 99–110.

Choi, J.J., S. Hauser, and K.J. Kopecky, 1999. Does the stock market predict real activity? Time series evidence from the G-7 countries, *Journal of Banking and Finance* 23, 1771–1792.

Conrad, J., R.F. Dittmar, and E. Ghysels, 2008. Skewness and the bubble. *Working paper*, University of Michigan.

Dickey, D.A. and W.A. Fuller, 1979. Distribution of the estimators for autoregressive time series with a unit root, *Journal of the American Statistical Association* 74, 427–431.

Diebold, F.X. and R. Mariano, 1995. Comparing predictive accuracy, *Journal of Business and Economic Statistics* 13, 253–265.

Doskov, N. and L. Swinkels, 2015. Empirical evidence on the currency carry trade, 1900–2012, *Journal of International Money and Finance* 51, 370–389.

Frankel, J. and G. Saravelos, 2012. Can leading indicators assess country vulnerability? Evidence from the 2008–09 global financial crisis, *Journal of International Economics* 87, 216–231.

Genotte, G. and H. Leland, 1990. Market liquidity, hedging, and crashes, *American Economic Review* 80, 999–1021.

Gnabo, J.Y., L. Hvozdyk, and J. Lahaye, 2014. System-wide tail comovements: A bootstrap test for cojump identification on the S&P 500, US bonds and currencies, *Journal of International Money and Finance* 48, 147–174.

Harris, R.D.F. and F. Yilmaz, 2009. A momentum trading strategy based on the low frequency component of the exchange rate, *Journal of Banking and Finance* 33, 1575–1585.

Harvey, C.R., J.C. Liechty, M.W. Liechty, and P. Müller, 2010. Portfolio selection with higher moments, *Quantitative Finance* 10, 469–485.

Harvey, C.R. and A. Siddique, 1999. Autoregressive conditional skewness, *Journal of Financial and Quantitative Analysis* 34, 465–487.

Harvey, D., S. Leybourne, and P. Newbold, 1997. Testing the equality of prediction mean squared errors, *International Journal of Forecasting* 13, 281–291.

Hsu, P.H. and C.M., Kuan, 2005. Reexamining the profitability of technical analysis with data snooping checks, *Journal of Financial Econometrics* 3, 606–628.

Kho, B., 1996. Time varying risk premia, volatility and technical reading profits: Evidence from foreign currency future markets, *Journal of Financial Economics* 41, 249–290.

Kostakis, A., K. Muhammad, and A. Siganos, 2012. Higher co-moments and asset pricing on London stock exchange, *Journal of Banking and Finance* 36, 913–922.

Kraus A. and R. Litzenberger, 1976. Skewness preference and the valuation of risk assets, *Journal of Finance* 31, 1085–1100.

Kraus, A. and R. Litzenberger, 1983. On the distributional conditions for a consumption oriented three moment CAPM, *Journal of Finance* 38, 1381–1391.

LeBaron, B., 1999. Technical trading rule profitability and foreign exchange intervention, *Journal of International Economics* 49, 125–143.

Lee, C.I. and I. Mathur, 1996. Trading rule profits in European currency spot cross-rates, *Journal of Banking and Finance* 20, 949–962.

Levich, R.M. and L.R. Thomas, 1993. The significance of technical trading-rule profits in the foreign exchange market: A bootstrap approach, *Journal of International Money and Finance* 12, 451–474.

Malevergne, Y. and D. Sornette, 2006. Multi moments method for portfolio management: Generalized capital asset pricing model in homogeneous and heterogeneous markets, in: B. Maillet and E. Jurczenko, eds., *Multi-Moment Asset Allocation and Pricing Models* (Wiley & Sons, Hoboken, NJ).

Melvin, M. and M.P. Taylor, 2009. The crisis in the foreign exchange market, *Journal of International Money and Finance* 28, 1317–1330.

Mende, A. and L. Menkhoff, 2006. Profits and speculation in intra-day foreign exchange trading, *Journal of Financial Markets* 9, 223–245.

Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf, 2012. Currency momentum strategies, *Journal of Financial Economics* 106, 660–684.

Mittnik, S. and M.S. Paolella, 2000. Conditional density and value at risk prediction of Asian currency exchange rates, *Journal of Forecasting* 19, 313–333.

Narayan, P.K., S. Mishra, S. Narayan, and K. Thuraisamy, 2015. Is exchange rate trading profitable? *Journal of International Financial Markets, Institutions and Money* 38, 217–229.

Narayan, P.K., S.S. Sharma, and K. Thuraisamy, 2014. An analysis of price discovery from panel data models of cds and equity returns, *Journal of Banking and Finance* 41, 167–177.

Neely, C.J. and P.A. Weller, 2003. Intraday technical trading in the foreign exchange market, *Journal of International Money and Finance* 22, 223–237.

Neely, C.J. and P.A. Weller, 2013. Lessons from the evolution of foreign exchange trading strategies, *Journal of Banking and Finance* 37, 3783–3798.

Neely, C.J., P.A. Weller, and R. Dittmar, 1997. Is technical analysis in the foreign exchange market profitable? A genetic programming approach, *Journal of Financial and Quantitative Analysis* 32, 405–426.

Okunev, J. and D. White, 2003. Do momentum-based strategies still work in foreign exchange currency markets? *Journal of Financial and Quantitative Analysis* 38, 425–447.

Olson, D., 2004. Have trading rule profits in the currency markets declined over time? *Journal of Banking and Finance* 28, 85–105.

Osler, C. and T. Savaser, 2011. Extreme returns: The case of currencies, *Journal of Banking and Finance* 35, 2868–2880.

Osler, C.L., 2003. Currency orders and exchange rate dynamics: Explaining the predictive success of technical analysis, *Journal of Finance* 58, 1791–1819.

Palandri, A., 2015. Do negative and positive equity returns share the same volatility dynamics? *Journal of Banking and Finance* 58, 486–505.

Plantin, G. and H.S. Shin, 2007. Carry trades and speculative dynamics. *Working paper*, Princeton University.

Premaratne, G. and A.K. Bera, 2000. Modeling asymmetry and excess kurtosis in stock return data. *Illinois Research and Reference Working paper no. 00–123*.

Raj, M., 2000. Transactions data tests of efficiency: An investigation in the Singapore futures markets, *Journal of Futures Markets* 20, 687–704.

Rubinstein, M., 1973. The fundamental theorem of parameter-preference security valuation, *Journal of Financial and Quantitative Analysis* 8, 61–69.

Szakmary, A.C. and I. Mathur, 1997. Central bank intervention and trading rule profits in foreign exchange markets, *Journal of International Money and Finance* 16, 513–535.

Taylor, M.P. and H. Allen, 1992. The use of technical analysis in the foreign exchange market, *Journal of International Money and Finance* 11, 304–314.

Taylor, S.J., 1994. Trading futures using a channel rule: A study of predictive power of technical analysis with currency examples, *Journal of Futures Markets* 14, 215–235.

Westerfield, J.M., 1977. An examination of foreign exchange risk under fixed and floating rate regimes, *Journal of International Economics* 7, 181–200.